

Astrophysics and cosmology with the scattering transform

by
Sihao Cheng

A dissertation submitted to The Johns Hopkins University in conformity
with the requirements for the degree of Doctor of Philosophy

Baltimore, Maryland
August, 2021

© 2021 Sihao Cheng
All rights reserved

Abstract

Quantifying textures and patterns in physical data is an important but challenging task. Recently, a novel statistic borrowing ideas from convolutional neural networks (CNNs), called the scattering transform, has shown its great potential. It is a sweet spot between the power spectrum and CNNs: it yields a compact set of summary statistics with a simple definition while sharing the high performance of CNNs. In this thesis I provide intuitive understandings and interpretations for the scattering transform. I also discuss its connection to and advantages over other common statistics, in particular the N -point functions. Then, I show promising examples of its applications in physics research, including rigid parameter inference tasks in weak lensing cosmology and exploratory data analysis in astronomy and oceanography.

Thesis Readers

Dr. Bhuvnesh Jain

Professor, Department of Physics and Astronomy
University of Pennsylvania

Dr. Marc Kamionkowski

Professor, Department of Physics and Astronomy
Johns Hopkins University

Dr. Stéphane Mallat

Professor, Applied Mathematics
Collège de France

Dr. Brice Ménard (Primary Advisor)

Professor, Department of Physics and Astronomy
Johns Hopkins University

Dr. Alex Szalay

Professor, Department of Physics and Astronomy
Johns Hopkins University

Dedicated to my parents, brother, and Siyu

Acknowledgements

Forewords

During the Ph.D. study, I have been working on two topics: one in cosmology and the other in stellar physics. The main narrative of this thesis will focus on the cosmology project. Nevertheless, I will briefly introduce both projects below and provide a copy of my work about white dwarf stars in the Appendices.

1) Cosmology: In the past century, cosmology has developed from a highly speculative subject to a mature field with precise models, shifting the focus of observational cosmology largely into parameter inference. Our universe in the ‘early time’, based on observations of the cosmic microwave background, is quite homogeneous. The matter distribution can be well modelled by a Gaussian random field, for which the power spectrum is the sufficient statistic to describe it. Afterwards, gravitational interaction of matter leads to both a growth and a non-linear evolution of fluctuations, producing all kinds of non-Gaussian structures in the universe such as halos, filaments, walls, and voids, which we see in the current ‘late-time’ universe. Several large-scale cosmology surveys costing billions of dollars have been built or planned to measure the late-time cosmic matter distribution through weak lensing and galaxy clustering. In order to better extract information from these non-Gaussian structures, a statistical tool beyond the power spectrum is desired.

In the cosmology project, I explored the application of a novel statistical estimator called the scattering transform to extract information from the cosmic density field. By the time I started my PhD, the scattering transform had never been mentioned in the astrophysics literature. Borrowing ideas from convolutional neural networks, this estimator is highly efficient in describing complex structures in physical fields. Moreover, it is robust, compact, and interpretable, which is important for practical analysis. My work demonstrated that the scattering transform extracts as much information from weak lensing maps as what convolutional neural networks can do, and its sampling distribution is much Gaussian than traditional non-Gaussian statistics. One of my papers won the outstanding paper award of the International Astrostatistics Association in 2020.

In this thesis, I shall first discuss in general the task of extract information from grid data (such as a field or a time series) and the most widely used mathematical tool, the power spectrum in chapter 1. Then, in chapter 2, I shall introduce the scattering transform in the language of physicists and astronomers. After summarising its key features, I shall discuss the understanding of its internal operations and the interpretation of its coefficients. I argue that the compactness and robustness of scattering coefficients make them particularly advantageous for probing the non-Gaussian structures in physics. In chapter 3, I shall show its application to regression problems in cosmology, and in chapter 4, applications to exploratory data analysis in other areas such as oceanography. Finally, I shall summarize the main conclusions of this thesis in chapter 5 and present a code for efficiently calculating the translation-invariant scattering coefficients in an appendix.

2) White dwarfs: In the appendix, I shall also present a copy of my papers about white dwarfs.

For this project, I studied the physics and evolution of white dwarf stars, the destiny of most stars in the universe. White dwarfs have the mass of the sun but the size of the earth. They are extremely compact objects and ideal laboratory for dense plasma physics. In my work, I measured the age of white dwarfs with a novel way, using knowledge from Galactic kinematics. Then I compared the ages measured from the traditional ‘photometric’ and the novel ‘kinematic’ methods and see if they are consistent.

I first applied this idea to study white dwarf cooling. As star remnants, white dwarfs have no nuclear reaction inside and just cool down in a cosmic time scale, similar to the cooling of a cup of hot tea on the table. Cooling models have been established and elaborated during the past fifty years to a level many believed to be mature. However, by incorporating the kinematic information obtained from the Gaia satellite, I found that a population of white dwarfs must experience an extreme cooling delay, unexpected by any current cooling models. Going deeper, I showed that the missing physics may be the sink of ^{22}Ne , an important isotope provides in stellar evolution with more neutrons than protons and thus feeling more downward gravity than upward pressure in a white dwarf. This explanation has now been accepted by the community. Interestingly, it is also related to the behavior of freezing (liquid–solid phase transition) of dense plasma, which can only be probed in the extreme condition of white dwarfs. This discovery is considered one of the major discoveries of Gaia.

I also applied my method to white dwarf mergers. Measuring the merger rate sets strong constraints on binary evolution models, and interestingly, white dwarf mergers are also suspected to be the trigger of type-Ia supernova explosions, which are among the most energetic events in the universe and have been used as an anchor for cosmological distance measurement. I found that a significant amount (about 20%) of high-mass white dwarfs do not have consistent ages and must originate from mergers of to lighter white dwarfs. This fraction and the corresponding merger rate have been difficult to measure for decades, and my novel method provides the most precise estimate of them.

Publications during the Ph.D. study

First-author publications:

[A Cooling Anomaly of High-mass White Dwarfs](#)

Discovered a new population of white dwarfs with an unexpected, extremely long cooling delay using Gaia data

Sihao Cheng, Jeffrey D. Cummings, and Brice Ménard
2019, *ApJ*, 886, 100

[Two delays in white dwarf evolution revealed by *Gaia*](#)

Sihao Cheng

2019, *Proceedings of IAU*, 15 (S357), 175

[Double White Dwarf Merger Products among High-mass White Dwarfs](#)

Measured the merger rate of white dwarf with unprecedented high precision using a novel kinematic method

Sihao Cheng, Jeffrey D. Cummings, Brice Ménard, and Silvia Toonen
2020, *ApJ*, 891, 160

[A new approach to observational cosmology using the scattering transform](#)

Introduced a new statistic inspired by convolutional neural nets to observational cosmology, and demonstrated that it outperforms classic estimators and is on a par with neural nets

Sihao Cheng, Yuan-Sen Ting, Brice Ménard, and Joan Bruna
2020, *MNRAS*, 499, 5902

[Weak lensing scattering transform: dark energy and neutrino mass sensitivity](#)

Visualised what the scattering transform sees from a lensing map and emphasised its advantage of Gaussianizing quickly

Sihao Cheng and Brice Ménard
2021, arXiv:2103.09247, accepted to *MNRAS*

A guide of scattering transform for physicists

explaining and visualising the scattering transform using the language of physicists

Sihao Cheng and Brice Ménard
2021, in preparation

Weak lensing scattering transform: application to observational data

Sihao Cheng et al.
2021, in preparation

Contributing-author publications:

[Carbon star formation as seen through the non-monotonic initial–final mass relation](#)

Marigo, P. **et al.**

2020, *Nature Astronomy*

[A Gravitational Redshift Measurement of the White Dwarf Mass–Radius Relation](#)

Used populational gravitational redshift to probe the white dwarf mass–radius relation over a wide mass range

Vedant Chandra, Hsiang-Chih Hwang, Nadia L. Zakamska, and **Sihao Cheng**

2020, *ApJ*, 899, 146

[Multi-Gigayear White Dwarf Cooling Delays from Clustering-Enhanced Gravitational Sedimentation](#)

Evan B. Bauer, Josiah Schwab, Lars Bildsten, and **Sihao Cheng**

2020, *ApJ*, 902, 93

[An Increase in Small-planet Occurrence with Metallicity for Late-type Dwarf Stars in the Kepler Field and Its Implications for Planet Formation](#)

Cicero X. Lu, Kevin C. Schlaufman, and **Sihao Cheng**

2020, *AJ*, 160, 253

[Forever young white dwarfs: when stellar ageing stops](#)

Camisassa, M. **et al.**

2021, *A&A Letters*, 649, 7

Contents

Abstract	ii
Dedication	iii
Acknowledgements	iv
Forewords	v
Contents	ix
List of Figures	xi
Chapter 1 Introduction	1
1.1 Goals and challenges of information extraction	2
1.2 The power spectrum $P(k)$	3
1.3 Can we do things differently?	4
Chapter 2 The scattering transform	6
2.1 Definition	6
2.1.1 Extending the power spectrum	6
2.1.2 Simplifying a CNN	7
2.1.3 Formalism	9
2.2 Generative models	11
2.3 Key properties	13
2.4 Understanding the operations	14
2.4.1 An overview on translation-invariant descriptors	14
2.4.2 The role of wavelet convolution and modulus	15
2.4.3 Hierarchy: information extraction beyond the power spectrum	18
2.4.4 Scattering operations in Fourier space	19
2.5 Interpretability	20
2.5.1 First-order coefficients	20
2.5.2 Second-order coefficients	22
2.6 Relation to other methods	26
2.6.1 Difference from wavelet transform	26
2.6.2 Relation to higher-order statistics	26
2.6.3 Scale interactions, cross-correlations and phase harmonic estimator	30

Chapter 3 Applications in cosmology	32
3.1 Parameter inference in cosmology	32
3.2 Weak lensing cosmology	34
3.3 Forecast for clumpiness σ_8 and matter density Ω_m	39
3.3.1 Simulated lensing maps	39
3.3.2 Statistical descriptors	40
3.3.3 Cosmological inference framework	41
3.3.4 Results	42
3.3.5 Discussion	47
3.3.6 Conclusions	50
3.4 Forecast for dark energy parameters and neutrino mass M_ν	51
3.4.1 Simulated lensing maps	51
3.4.2 Summary statistics	53
3.4.3 Likelihood and parameter inference	54
3.4.4 Sampling distribution of different summary statistics	55
3.4.5 Results	58
3.4.6 Conclusions	60
3.5 Application to observational data	60
3.5.1 HSC weak lensing data	61
3.5.2 Mock data	63
3.5.3 From shear catalog to mass maps	63
3.5.4 Weak lensing scattering transform	64
Chapter 4 Applications in exploratory research	67
4.1 Sea temperature field	67
4.1.1 Data and preprocessing	67
4.1.2 The scattering representation	69
4.2 Galaxy morphology	71
Chapter 5 Summary	76
Software implementation of the scattering transform	80
Appendix II: Merger products among massive white dwarfs	81
Appendix III: A cooling anomaly of white dwarfs	93
References	121
Curriculum vitae	134

List of Figures

Figure 1-1	Two fields with the same power spectrum but different textures. The left panel shows a simulated matter density field of late-time universe, which is highly non-Gaussian. The right panel is a Gaussian random field.	2
Figure 2-1	Upper panel: comparison between Fourier and localised kernels in 1D. Solid and dashed curves represent the real and imaginary part of the profiles, respectively. Lower panel: comparison between Fourier convolution and convolution with localised kernels.	8
Figure 2-2	The scattering transform’s multi-layer structure and similarity to CNNs. The scattering operation is composed of a local (wavelet) convolution and a pixel-wise modulus.	9
Figure 2-3	Image generation using the scattering statistics. Upper panels are fields from various disciplines in physics: Turing pattern, Ising model, turbulence revealed by the sea surface temperature, turbulence revealed by Solar ultraviolet emission, and gravitational collapsing of matter in the Universe. Lower panels are randomly generated images which have the same scattering coefficients as the upper counterparts.	12
Figure 2-4	<i>Upper panel:</i> profile of a Morlet wavelet ($j = 6, l = 0$, image size 512×512 pixels) in the real space and another one ($j = 1, l = 1$) in Fourier space. The centre of the Fourier space represents zero frequency. <i>Lower panel:</i> radial frequency profiles of a family of wavelets. Dilating/contracting (by factor of 2) and rotating (by π/L) one wavelet give the whole family of wavelets used in the scattering transform.	16
Figure 2-5	The scattering coefficients and sparsity. Given the same power spectrum, spatially-spread textures (Gaussian textures) have most energy in S_1 coefficients, spatially-sparse textures leak energy from S_1 to higher order scattering coefficients. This provides an intuitive understanding of the reduced coefficients $s_{21} \equiv S_2/S_1$: the sparsity of textures.	21
Figure 2-6	The similarity and difference between power spectrum and 1st-order scattering coefficients. Here we use 2D fields as examples. It can be seen that when the field structures are not extremely sparse in real space, the s_1 coefficients basically follow the square root of the (logarithmically binned) power spectrum, such as in the cases of Gaussian fields and random points in large scales. However, when the structures are very sparse, the s_1 coefficients significantly deviate from the square root of power spectrum.	22

Figure 2-7	The s_1 coefficients of random point process and Gaussian white noise, as a function of scale combinations (j_1, j_2)	23
Figure 2-8	Sparsity and structures in a field.	23
Figure 2-9	Visual interpretation of scattering coefficients s_{21} , generated by starting from a non-Gaussian field (right panel) and requiring the s_{21} coefficients to decrease while keeping the power spectrum unchanged. It can be seen that the s_{21} measures the spatial sparsity of structures in a field at various scales.	24
Figure 2-10	Visual interpretation of scattering coefficients s_{22} , generated by requiring the s_{22} coefficients changing from below unity to above unity, while keeping the s_1 and s_{21} coefficients as a particular function of scales. It can be seen that the s_{22} coefficients contain information about structures described by combinations of orientations. They differentiate straight lines from swirls.	25
Figure 2-11	Comparison between the moment approach with scattering approach: calculating moments is equivalent to amplifying the tail of $p(x)$ and then integrating over x . When the tail is heavy, it takes the risk of divergence. In contrast, the scattering approach is equivalent to folding the core of $p(x)$	27
Figure 2-12	The fraction of Fisher information contained in moment and scattering statistics, respectively, of log-normal distributions. The x axis is defined as $\sigma_\delta \equiv \sqrt{\exp(\sigma^2) - 1}$, where σ is the variance of the corresponding Gaussian variable, and σ_δ is also the ratio of the standard deviation to the mean of the log-normal distribution. The folding (scattering) strategy works much better in the heavy-tail regime than traditional moments.	28
Figure 2-13	Relations between different translation-invariant descriptors, which are on the route of deriving the scattering transform from the power spectrum.	29
Figure 2-14	The corresponding 2-point and 4-point functions of scattering coefficients when replacing the modulus operation by modulus squared.	30
Figure 3-1	The evolution of cosmic density field in the universe in the Millennium Simulation. At early times (left panel), the density fluctuations are small and Gaussian. As the universe evolves, the fluctuations grow and becomes non-Gaussian at late times (right panel, the current universe).	33
Figure 3-2	The on-going and planned major weak lensing surveys in the next ten years.	35
Figure 3-3	Left: The 4-meter telescope used for the Dark Energy Survey. Right: The CCD array contains 74 pieces of CCDs, each being similar to but much more sensitive than the detector used in digital cameras. The array can cover an area of about 14 moons on the sky.	35
Figure 3-4	The scattering transform of three fields (I_0) with indistinguishable power spectra. Row 1 shows a realization of convergence maps in cosmology $(\Omega_m, \sigma_8) = (0.292, 0.835)$, row 2 shows cosmology $(\Omega_m, \sigma_8) = (0.566, 0.520)$, row 3 is for a Gaussian random field with the same (2D) power spectrum as row 1. It can be seen by eye that the average intensity of the 2nd-order scattering fields (the last column), which corresponds to an s_2 coefficient and measures the clustering strength of structures highlighted by I_1 , is significantly different from each other, while their power spectra (the first column) are indistinguishable.	37

Figure 3-5	Images generated with selected sets of summary statistics, including the power spectrum, bispectrum, and scattering coefficients (S_1, S_2), showing that the scattering coefficients characterize the field better. As the image generation is random, one should compare textures, instead of the exact positions of features, between the input and generated maps.	38
Figure 3-6	<i>Upper-left Panel:</i> The fiducial cosmology (black) and two other cosmologies on the (Ω_m, σ_8) plane. <i>Upper-right Panel:</i> Illustration of reduced scattering coefficients $s_1(j_1)$ and $s_2(j_1, j_2)$ for a single j_1 scale. <i>Lower Panel:</i> The power spectrum and scattering coefficients for the three cosmologies in noiseless case. The first row presents coefficients of the fiducial cosmology and of Gaussian random fields with the same power spectrum, and the second row shows changes of coefficients (Δ coef.) when we move from the fiducial cosmology to the other two. Error bars and gray shaded regions show cosmic variance, i.e., the variability among realizations. The 1st-order scattering coefficients behave similarly to the power spectrum, while the 2nd-order scattering coefficients can break the Σ_8 degeneracy, along which non-Gaussianity of weak lensing field changes.	43
Figure 3-7	The 1σ Fisher forecast of cosmological parameters from a $3.5 \times 3.5 \text{ deg}^2$ noiseless convergence map with 0.41 arcmin per pixel resolution. The de-correlated 2nd-order scattering coefficients s_2/s_1 provide critical information to break the Σ_8 degeneracy along which the power spectrum cannot distinguish, therefore drastically improve the constraint.	45
Figure 3-8	The 1σ Fisher forecast of cosmological parameters (Ω_m and σ_8) from different descriptors of a $3.5 \times 3.5 \text{ deg}^2$ convergence map smoothed with $\sigma = 1'$ Gaussian filter. The scattering coefficients have comparable performance as a state-of-the-art CNN [58] at all noise levels, and 3–5 times better than the power spectrum depending on the noise level.	46
Figure 3-9	Dependence of the (Ω_m, σ_8) constraints with different methods on galaxy shape noise. The figure of merit (FoM) is defined as the 1σ confident area on the (Ω_m, σ_8) plane. Note that the CNN result [58] is reported in terms of posterior, while others are Fisher forecast. For noisy cases the difference between scattering transform and CNN is not intrinsic but due to the difference between posterior and Fisher forecast. Peak count’s performance does not increase as fast because it is more sensitive to smoothing scale than the other methods.	47
Figure 3-10	Scale coverage of wavelets used by the scattering transform. We also show the profile of a 1 arcmin Gaussian filter in frequency space for comparison.	53
Figure 3-11	Skewness distribution of the (log) scattering and (binned) bispectrum coefficients, measured from 10,000 realizations of convergence maps at the fiducial cosmology of MassiveNuS simulation set. If the PDF of these statistical coefficients are Gaussian, then the skewness should be zero.	56

Figure 3-12	Fisher forecast with noiseless κ maps (95% confidence ellipses). The scattering coefficients have remarkably high constraining power for cosmological parameters. <i>Left:</i> Forecast for dark energy equation of state index $w(a) = w_0 + (1 - a)w_a$, from 20,000 deg ² noiseless map with l nominally up to roughly 30,000 and source redshift $z = 2$. <i>Right:</i> Forecast for neutrino mass M_ν , from 20,000 deg ² noiseless map with source redshift $z = 1$	58
Figure 3-13	Forecast of cosmological parameters (95% confidence contours) for a Rubin-observatory-like survey with 20,000 deg ² field of view, with M_ν , Ω_m , and A_s (or σ_8) to be constrained. Colours represent results of different sets of summary statistics of lensing convergence (κ). Improvement from adding redshift tomography is also shown.	59
Figure 3-14	Left: the design and size of the HSC camera. Right: The CCD array of the Hyper Suprime-Cam on the 8.2m Subaru telescope. It has 870 Mega pixels in total and covers 1.5 degree field-of-view in diameter, equal to the diameter of 3 full moons.	61
Figure 3-15	A small patch of the HSC survey. Each blob in the figure is a galaxy far away. Their shapes are measured and averaged to estimate the gravitational lensing effect of the foreground dark matter. Regions near bright stars (the bright objects with spikes) are masked because of their severe influence on shape measurements.	62
Figure 3-16	A set of 100 N -body simulations sampled in the $\Omega_m - \sigma_8$ plane to build the likelihood emulator for scattering statistics.	63
Figure 3-17	Upper: a mock convergence map for the field WIDE12H without galaxy shape noise, redshift bin 0.6–0.9. Lower: an in-painted field of the upper panel. The in-painting is conducted from the corresponding masked shear map.	65
Figure 3-18	The scattering coefficients measured on real data and on a realization of the corresponding mock data in the fiducial cosmology. It can be seen that the simulation based forward-modelling approach can successfully recover the behavior of scattering statistics in real data, which includes a number of observational effects.	66
Figure 4-1	An example of ‘cutouts’ of the sea surface temperature dataset from NASA Aqua satellite. The left panel shows the image from the original Level-2 data product, the right panel shows the result after in-painting masked pixels and flat field correction.	68
Figure 4-2	Randomly selected sea surface temperature fields (upper) and a set arranged based on their 2nd-order scattering coefficients (lower). The x-axis of the right panel represents an averaged measure of the feature sparsity (s_{21}), and the y-axis represents that of the shape indicator (s_{22}).	70
Figure 4-3	An interesting sea surface temperature field with an eddy (left) and a gallery of similar fields selected based on the s_{22} scattering coefficients (the shape indicator).	71

Figure 4-4	The distribution of the $\sim 100,000$ sea surface temperature cutout fields on four reduced dimensions: the amplitude index, s_1 (power spectrum) power index, and two non-Gaussianity indices based on s_2 scattering coefficients. Colour in the left panel represents the feature sparsity index (the x-axis of the right panel), and the colour in the right panel represents s_1 power index (the x-axis of the left panel).	72
Figure 4-5	The non-Gaussianity distribution (feature sparsity s_{21} and texture shape s_{22}) for images with similar s_1 power index (i.e., similar shape of the power spectrum). The shown cutouts are randomly chosen from the corresponding cells. The s_1 power index is between -0.1 and 0.1	72
Figure 4-6	Same as the previous figure, but with the s_1 power index between 0.35 and 0.45	73
Figure 4-7	Same as the previous figure, but with the s_1 power index between 0.75 and 0.85	73
Figure 4-8	A gallery of randomly selected galaxy images (left) and a set arranged based on the 2nd-order scattering coefficients (right). The x- and y-axes of the right panel represent the first two UMAP projection of the coefficients, which is similar to a non-linear version of the principal component analysis (PCA).	75
Figure 5-1	Different approaches to extract information from a field. Summary statistics can serve as the representation, while convolutional neural networks (CNNs) usually hide such representation inside. The scattering transform borrows the ideas of local convolutions, ‘low-order’ operations, and the hierarchical design, which have proven powerful in CNNs. But, it belongs to the statistical approach and needs no training.	77
Figure V-2	H–R diagram of WDs in <i>Gaia</i> DR2. We show the 250 pc sample of WDs with high-quality measurements and a grid of WD masses m_{WD} and photometric ages τ_{phot} derived from the combined O/Ne- and C/O-core WD cooling model. WDs evolve along their cooling tracks, i.e., the constant-mass curves. The red region includes 1395 nearby, high-mass, hot WDs selected in Section 5.	83
Figure V-3	Delay-time distributions of double-WD mergers used in our model. These distributions are generated from binary population synthesis. The x -axis is the delay time of binary evolution, i.e., $\tau - t_{\text{cool}}$ for the resulting white dwarf. The y -axis is in linear scale and normalized to their maximum values. We input to our model the shapes of these five distributions as probability distribution and do not use the information from their normalization. . . .	87

Figure V-4	Velocity distribution of our white dwarf sample. We show the sample of white dwarfs from all five mass bins ($0.8\text{--}1.3 M_{\odot}$) as an example. v_l and v_b in the left and right panel of the figure means the Galactic longitude and latitude components of the transverse velocity. We present the observed histograms of the absolute values of v_l and v_b in 20 bins between 0 and 100 km s^{-1} and Poisson errors. We also show the theoretical velocity distributions for $f_m = 0, 1$, and the average of best-fitting values weighted by the sample size in each mass bin, which is about 0.2. Note that the y -axes are in logarithmic scale.	88
Figure V-5	Our estimates for the fraction of double-WD merger products among high-mass white dwarfs, in five bins of white dwarf photometric mass. The sample sizes in these mass bins, from lower to higher masses, are 408, 431, 323, 176, and 57, respectively.	88
Figure V-6	A comparison of the observed and simulated double-WD merger rate. The red data points with error bars are our observational estimates based on double-WD merger products. The histogram shows binary population synthesis results. Other data points show estimates in the literature based on the orbital distribution of observed double-WD systems: the light-grey one is an estimate for all double-WD mergers [124], and the green one is for systems with at least one extremely low-mass (ELM) WD [125], which provides a lower limit of the merger rate. The blue data point shows the observed SN Ia rate. Comparisons between a data point and the histogram should be made in terms of the area under the horizontal ‘error bar’ of the data point and the area under the histogram in the same mass range. . . .	90
Figure V-7	The H–R diagram of WDs selected from Section 5. In the top right panel we use the 150 pc sample to show the number-density distribution with a higher contrast. The Q branch is marked by the red arrow. In the main panel, we show our main 250 pc sample color-coded with transverse velocities v_T . Fast WDs ($v_T > 70 \text{ km s}^{-1}$) are emphasized by large symbols, and high-mass WDs ($> 1.08 M_{\odot}$) are emphasized by high symbol opacity. The grid of WD mass and photometric age is also plotted (using the O/Ne model for high-mass WDs). For the mass range marked by dark blue texts, the first (second) number corresponds to the O/Ne (C/O) model.	97
Figure V-8	The normalized photometric age τ_{phot} distribution of high-mass WDs in three consecutive mass ranges. The mass-dependent peaks trace the position of the Q branch. Crystallization should not produce any peak on this plot, because the τ_{phot} is calculated from a model including crystallization effects; the completeness stays high for at least 1 Gyr after the peaks in each mass range, so these mass-dependent peaks cannot be explained by a peak in the star formation history or by incompleteness. Therefore, there must be an extra cooling delay piling up WDs on the Q branch.	98

Figure V-9	The fraction of fast WDs in different mass ranges (<i>left panel</i>) and for different velocity cuts (<i>right panel</i>). There are significantly more fast-moving WDs on the Q branch than both before and after it in terms of photometric age. According to the age–velocity–dispersion relation (AVR), fast WDs are old. As argued in Section 5, this high fraction of fast WDs on the Q branch can only be explained by a subset of WDs experiencing an extra cooling delay on the Q branch.	99
Figure V-10	An illustration of the two-population scenario of the Q branch: a normal-cooling population (blue dots), and a population with the extra cooling delay (red dots). The number density of WDs on the H–R diagram is determined by the cooling rate, because WDs accumulate where the cooling rate is low. Here, we use the best-fitting values ($f_{\text{extra}} = 6\%$ and $t_{\text{extra}} = 8$ Gyr) from our later analysis to generate this mock H–R diagram. An animated version of this figure is available, where blue dots move with the normal cooling rate, while red dots (the extra-delayed population) move slowly on the Q branch. Each second in the animation corresponds to 1 Gyr in physical time, and the duration of this animation is 11 seconds. More animations can be found on the website: https://sihaocheng.github.io/Qbranch	101
Figure V-11	A sketch of the extra-delay scenario. The axes are the same as in Figure V-8. Some WDs are subject to the extra cooling delay and the rest of them are not, corresponding to the extra-delayed and normal-cooling population in the figure. We also illustrate the quantities of equation V.23.	102
Figure V-12	The comparison of AVRs constrained by this work and in the literature. The shaded regions show the 16th and 84th percentiles of the AVR posterior constrained by our high-mass WD sample. Symbols with error bars show the AVR measured for main-sequence stars by GCS and RAVE [204, 205]. In our model, the $\tau < 3.5$ Gyr part of the AVR is mainly constrained by the normal-cooling WDs (population ‘s’ and ‘m’), and the older part by the extra-delayed WDs. Note that the turnings at 7 and 10 Gyr reflect our parameterization of the AVR.	106
Figure V-13	The corner plot of the posteriors of the AVR parameters. We use flat priors for these parameters within the ranges shown on this figure. We have checked that there are no correlations between these parameters and the three main parameters f_{extra} , t_{extra} , and f_{m} , and the three components of solar motion in our model.	110
Figure V-14	The posterior distribution of the main parameters for setup 1. f_{extra} is the fraction of extra-delayed population, t_{extra} is the length of the extra cooling delay, and f_{m} is the fraction of normal-cooling double-WD merger products. Note that in setup 1, it is $f_{\text{extra}} + f_{\text{m}}$ rather than f_{m} that is the total fraction of merger products.	111

- Figure V-15** The observed and modeled velocity distributions. v_l and v_b are the Galactic longitude and latitude components of the transverse velocity. For the observed distribution, we present the histogram between -150 and 150 km s⁻¹ with 23 bins and the Poisson error of each bin. Note that the y-axis is in logarithmic scale. The solid and dashed curves, which are not very different, are the velocity distributions of the best-fitting models under setup 1 and 2, respectively. Both models fit the observations quite well. The dotted curves are the velocity distributions when no white dwarf has the extra cooling delay or merger delay. Its discrepancy to observed histograms shows the necessity of the two delays. 113
- Figure V-16** The effective zone of ²²Ne settling for C/O-core DA WDs with $D/D_s = 3.5$, assuming no suppression from crystallization. ²²Ne settling significantly delays the WD cooling between the solid and dashed curves, when $L_{\text{extra}}^{\text{Ne}}$ is close to L_{surf} . The colors represent two ²²Ne abundances of WDs, corresponding to $[M/H] = 0$ and 0.15 in their progenitor stars. We observe that the position, trend, and narrowness of the effective zone of ²²Ne settling match the Q branch quite well. 114
- Figure V-17** A part of the H-R diagram showing the spectroscopically verified WDs, with Q-branch DQ and hot-DQ white dwarfs marked by red and green open circles. The dots without circles are mostly DA white dwarfs. We estimate that half of the extra-delayed white dwarfs are DAs, because half of the fast-moving white dwarfs on the Q branch are DAs. We also mark the known magnetic DQs with larger black circles. Note that the mass range here ($> 0.9 M_{\odot}$) is larger than for the sample in our main analysis (1.08–1.23 M_{\odot}). 117

Chapter 1

Introduction

In attempting to understand the laws of Nature, most physicists face the challenge of extracting relevant information from data, produced by sensors or computer simulations. Different sets of data analysis tools are used in different disciplines, sometimes motivated by the properties of the object of study (symmetries, invariants), aspects of the data (noise level) or simply convenience or habit. Among them, a few are used ubiquitously, such as the power spectrum and correlation functions. However, being second-order moment analyses, they unfortunately lack the ability to characterise non-Gaussian structures, which are present in a wide variety of interesting fields. Figure 1-1 shows a comparison of non-Gaussian and Gaussian field textures with the same power spectrum. To distinguish them, one needs to go beyond the power spectrum. Over the past decade, a new type of estimator has gained popularity in virtually all fields of science: convolutional neural networks (CNNs), a novel paradigm remarkably efficient at extracting certain types of information from pixelized data but whose properties are not yet fully understood. Compared to the traditional mathematical tools, neural networks lack transparency, stable mathematical properties or interpretability, which are crucial to scientific research.

With growing depth, neural networks enjoy a fantastic level of expressivity, capable of capturing the highly complex sets of varying features produced by the biological world, from the appearance of cats to human speech. However, when considering a wide range of physical fields, we are in a different regime. We deal with a level of complexity often substantially lower than that involved in the typical images or sounds considered in ‘deep learning’ applications. In this thesis, I advocate for the use of the *scattering transform*, introduced by Mallat [1] to extract information from physical fields. I shall show that the scattering transform provides an approach to data analysis that in many ways conveniently stands *in between* the power spectrum and CNNs.

The scattering transform has many attractive properties. In particular, it can efficiently extract information from complex signals while being fully deterministic and not requiring any training. In addition, understanding the properties of the scattering transform allows one to decipher many key aspects of the inner workings of CNNs.

The scattering transform was originally introduced in the mathematics literature [1] with follow-up works that appeared in the signal processing and computer science literature. So far, the scattering transform has been used primarily in audio/visual signal processing and has been mostly discussed in the mathematics and computer science literature [e.g., 2–5]. It has already been used in a number of scientific applications: intermittency in turbulence [6], quantum chemistry and material science [7–9], graph-structured data [10], geography [11], astrophysics [12–14], and cosmology [15,

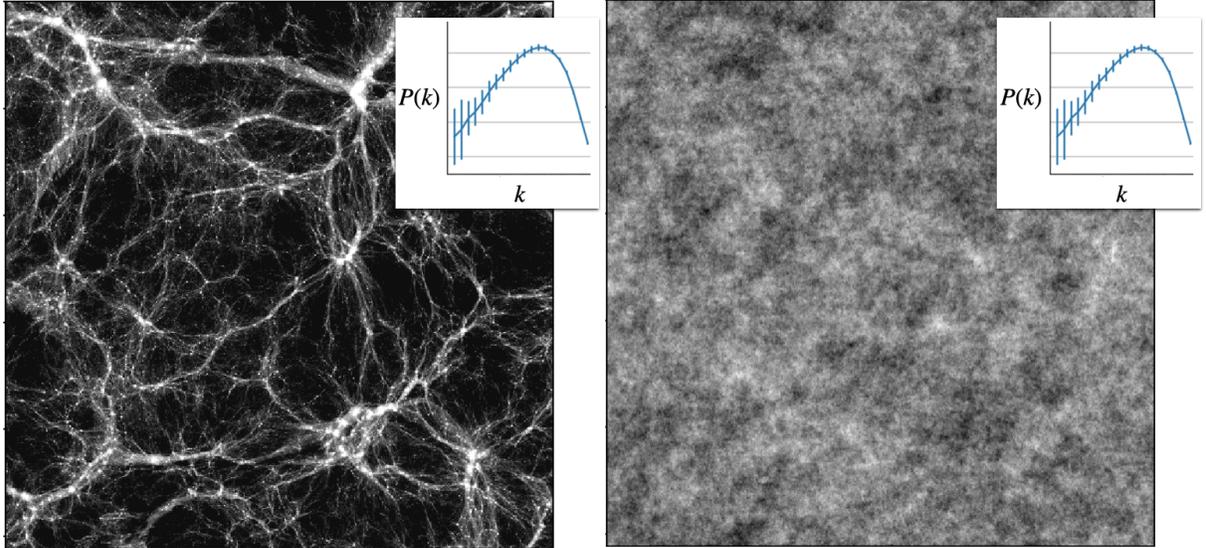


Figure 1-1. Two fields with the same power spectrum but different textures. The left panel shows a simulated matter density field of late-time universe, which is highly non-Gaussian. The right panel is a Gaussian random field.

16]. In several of these applications, the scattering transform reached state-of-the-art performance compared to the CNNs in use at the time. It has been typically used in the context of classification tasks but, as we will discuss, if a model is available, it is also powerful in the context of parameter estimation or regression.

1.1 Goals and challenges of information extraction

Extracting physical information from data means a mapping from the extremely high-dimensional function space of data to a low-dimensional space that corresponds to a few classes or physical parameters. To do so, it is convenient and often necessary to first describe the field with a mathematical vocabulary, which aims at discarding irrelevant variabilities and concentrating relevant information into a smaller set of descriptors or summary statistics. Finding the best language or description to extract physical information from data is often the key challenge of data analysis.

In physics, a number of properties are fundamental. Understanding how energy is distributed in a system is often an important goal. In addition, symmetries and invariants play a major role. For example, many physical laws are invariant under translation in space and/or time. When relevant, such properties should be incorporated into the chosen representation to discard the irrelevant variability. Ideally, the summary statistic should also be robust and compact, which are necessary requirements for interpretability. Robustness means the stability or continuity of the descriptors with respect to physically meaningful perturbations of the field, such as additive noise, geometric (or temporal) deformation, change of physical parameters of the field, potential distortions from the detector, etc.

In the community of computer vision, a lot of attention has been focused on classification problems. In physics, parameter inference (regression) and exploratory data mining are equally

common and generic as classifications, if not more so, Nevertheless, it is interesting to note that the requirements for good statistical descriptors turn out to be similar for all those three purposes. Especially, stability and robustness are extremely important for practical data analysis in all kinds of research. Computer scientists have put great efforts to guarantee stability when designing a descriptor. But unfortunately, this point is often less discussed or even omitted in the community of physicists when we choose our statistical descriptors.

In summary, to extract information from a field, the challenge is often reduced to finding a set of (statistical) descriptors that are invariant to translation and/or rotation, robust, and compact, while being physically informative. To better understand how to obtain such properties in a statistical estimator, we begin by discussing the key aspects and limitations of the most commonly used summary statistic in physics: the power spectrum. In order to simplify the formalism and discussion, we restrict our domain to stationary ergodic processes. In other words, we will primarily characterize textures. We will discuss extensions beyond that domain towards the end of the paper.

1.2 The power spectrum $P(\mathbf{k})$

The power spectrum is a ubiquitous translation-invariant descriptor in physics and beyond. It is used from the most formal theoretical studies to exploratory data analyses. It quantifies the variance of a field as a function of scale. For many fields, it is related to the physical energy of the system: for example, for fluids, the variance of velocity field is the kinetic energy; for fields described by wave equations, the variance of field intensity is the potential energy.

Given a field $I(\vec{x})$, the power spectrum is defined through the Fourier transform of the field, $\tilde{I}(\vec{k})$, as

$$P(\vec{k}) \equiv \tilde{I}(\vec{k})\tilde{I}(-\vec{k}) = |\tilde{I}(\vec{k})|^2, \quad (1.1)$$

where \vec{k} is a frequency. Its generic use is motivated by a number of properties:

- **Translation-invariance:** Defined from Fourier modulus, the power spectrum is translation-invariant and therefore immune to the irrelevant variability introduced by translations.
- **Energy extraction and scale separation:** The power spectrum is a partition of the variance or energy density of the field as a function of scale. This is often an important quantity to measure to extract characteristic scales of a system, the range of scales above/below which noise contributions dominate and its decay rate as a function of frequency informs on the regularity of the fluctuations.
- **Dimensionality reduction:** The power spectrum itself does not compress the data. However, two common binning schemes can substantially reduce the number of coefficients without losing much information. 1) Statistical isotropy in many fields allows for an average over orientations and reducing the power spectrum into a 1D object. 2) In many cases, for example when long-range correlation is not too strong, the power spectrum turns out to be a smooth function of scale, enabling a binning in neighbouring scales.
- **Asymptotic normality:** To perform precise parameter inference, one needs to know the likelihood. It is desirable to use summary statistics whose PDF can be well approximated by a Gaussian distribution, because it facilitates likelihood parametrisation. As the binned power spectrum can be seen as averages of locally measured power spectra at different \vec{x} , one expects it to Gaussianize according to the central limit theorem under certain conditions.

- **Theoretical predictions:** Defined through only addition and multiplication, the mathematics of power spectrum are familiar to theoretical physicists. Especially, when the field is in perturbative regime, the power spectrum can be accurately predicted from the field equation.

Unfortunately, power spectrum-based analyses have a serious limitation. Though it is a sufficient statistic for Gaussian random fields, it cannot characterise more complicated structures, the so-called non-Gaussianities or interactions among scales. Such information is largely carried by the Fourier phases of the field, which do not contribute to the power spectrum. It can be retrieved by measuring higher-order statistics of the field, but unfortunately, higher-order statistics are plagued with two main issues:

- **Non-robustness/divergence:** higher-order statistics are obtained by *multiplying* groups of random variables. This process amplifies the distribution tail of the input signal, increasing the variance and all higher-order moments of the statistics, which reduces the speed towards asymptotic normality and causes non-robustness of the statistics. In more extreme cases, they can even diverge.
- **Information dilution:** the number of higher-order statistics required to describe all possible configurations increases steeply with the order n , leading to a cumbersome description with relevant information diluted in many coefficients.

We will elaborate in section 2.6.2 when comparing them to the scattering transform. Both of those issues make higher-order statistics inefficient in concentrating information and hard to use in practice. As a result, it remains a fundamental challenge to find a good strategy to extract information beyond the power spectrum.

1.3 Can we do things differently?

Over the past decades, it has become clear that convolutional neural networks (CNNs) are very efficient in extracting targeted information from a field. This approach differs from traditional statistics in two ways. Firstly, it uses different operations. Especially, when compressing data and constructing representation, CNNs use operations other than the *multiplication* operation used in moment-based statistics as its non-linearity. These operations never elevate the power of the input data. Secondly, CNN’s data representation is usually implicit and is learned from a training set, not pre-determined. Nevertheless, it is instructive to examine for structural connections between CNNs and the power spectrum. CNNs involve a series of convolutions with localised kernels (learned during a training procedure) followed by a non-linear function and a ‘pooling’, i.e., an average or a Max of neighbouring pixels. The power spectrum, under a revisit, can be defined in real space, equivalently to its traditional definition, by:

$$P(\mathbf{k}) \propto \langle |I \star e^{-i\mathbf{k}\cdot\mathbf{x}}|^2 \rangle, \tag{1.2}$$

which involves a convolution, a non-linear function, and an average – the three key operations used in CNNs. Therefore, the power spectrum can be seen as a 1-layer CNN with pre-determined kernels.

A closer look at their differences can inform us on how to push the power spectrum toward a CNN: Can we use convolution with localised kernels instead of Fourier modes? Can we drop the square, since modulus is already a non-linearity? Can we add more layers? It can also inform us on

how to simplify a CNN: Can we use pre-determined kernels without learning? The answer to all these questions is: yes, and the scattering transform can be naturally obtained along these lines of thinking.

Chapter 2

The scattering transform

The scattering transform was originally proposed by Mallat [1] as a tool for signal processing. It has started to reveal its great potential in many areas, including audio signal processing [2, 5], image classification [3], texture classification [4], intermittency in turbulence and finance data [6], material science [7–9], graph-structured data [10], and astrophysics [12–14] and cosmology [15, 16]. Several of these examples reached state-of-the-art performance compared to the CNNs in use at the time. In contrast to neural networks, the scattering transform comes with provable mathematical properties. Interestingly, the scattering transform has also provided key insights into deciphering the remarkable behaviour and performance of CNNs [3]. A perhaps counter-intuitive feature of CNNs is that the convolution, though restricting the flexibility of the neural network, dramatically boosts its performance on many types of data. In addition, a successful CNN architecture can often be re-purposed for very different tasks. These facts suggest that a certain mathematical structure enables efficient information extraction from a wide range of complex data. Understanding this structure may dramatically simplify the costly training process required when using neural networks.

2.1 Definition

Before formally introducing the scattering transform, I shall first present two lines of intuition leading to its core mathematics: (i) by re-orienting the power spectrum approach and (ii) by simplifying the principle behind a CNN.

2.1.1 Extending the power spectrum

Instead of defining the power spectrum as $P(k) \equiv \langle \tilde{I}(k)\tilde{I}(-k) \rangle$, which emphasises its being a two-point multiplication in Fourier space, we can express the power spectrum in real space as Eq. 2.23. Furthermore, the binned power spectrum can be expressed similarly in real space by replacing the Fourier kernel with localised kernels ψ_k :

$$\begin{aligned} P_{\text{binned}}(\vec{k}) &\equiv \int P w \cdot d\vec{k} \\ &\propto \langle |I \star \psi_k|^2 \rangle, \end{aligned} \tag{2.1}$$

where $w(\vec{k})$ is the binning window in Fourier space, and the convolution kernel ψ_k is defined by $\tilde{\psi}_k = w(\vec{k})^{0.5}$, with the tilde denoting Fourier transform. A comparison between Fourier kernels

and localised kernels is illustrated in the upper panel of Figure 2-1. As a result, the binned power spectrum can be seen as a one-layer CNN with local kernels, as illustrated in the lower panel of Figure 2-1

Such re-expressions in real space can be interpreted as follows: the convolution selects fluctuations and features *around* a scale; the non-linearity – modulus squared – converts the selected fluctuations into their *local strength*; and the average measures the overall intensity, which is a translation-invariant quantity. In this ‘feature extraction’ interpretation of the power spectrum, two improvements can be made. First, one can simply drop the square after the modulus to obtain a *lower-order* statistic, which has a similar meaning in real space:

$$\langle |I \star \psi_k| \rangle. \quad (2.2)$$

Second, we can repeat the operation $I \rightarrow |I \star \psi_k|$ to go beyond the single-scale analysis and construct statistics for more complicated structures:

$$\begin{aligned} S_1(k_1) &\equiv \langle |I \star \psi_{k_1}| \rangle \\ S_2(k_1, k_2) &\equiv \langle ||I \star \psi_{k_1}| \star \psi_{k_2}| \rangle \\ &\dots \end{aligned} \quad (2.3)$$

This is a new approach to extend the power spectrum and characterise complicated structures in a field, compared to the higher-order approach. Like the power spectrum, the statistics S_n are all *translation-invariant*. The locality of kernels is crucial for this extension, because the higher-order coefficients with $n > 1$ only probe the interactions of Fourier modes *within* the window set by the previous kernel ψ_{k_1} . As illustrated in Figure 2-1, if the kernels are Fourier modes, the field after convolution and modulus will become constant everywhere, and the ‘scattering’ strategy would not work anymore.

2.1.2 Simplifying a CNN

The scattering transform can also be obtained by simplifying certain aspects of a convolutional neural network. The main modification is to use pre-determined kernels for the convolutions instead of learning them from the training set through the challenging and costly optimisation process. As pointed out by Bruna & Mallat [3], the key elements found in CNNs can be simplified or replaced as follows:

- learnable convolutions \rightarrow pre-defined wavelet convolutions
- ReLU \rightarrow modulus
- pooling \rightarrow average
- multi-layer \rightarrow hierarchical structure

With such simplifications in mind, the forward-propagation operation done in the first layer of a CNN can be written as a ‘scattering operation’:

$$I_0 \rightarrow I_1 \equiv |I_0 \star \psi|, \quad (2.4)$$

where ψ is a pre-determined convolution kernel. An activation amplitude over the entire field can then be obtained by $S_1 = \langle I_1 \rangle$. The connection of two layers extracting information from two

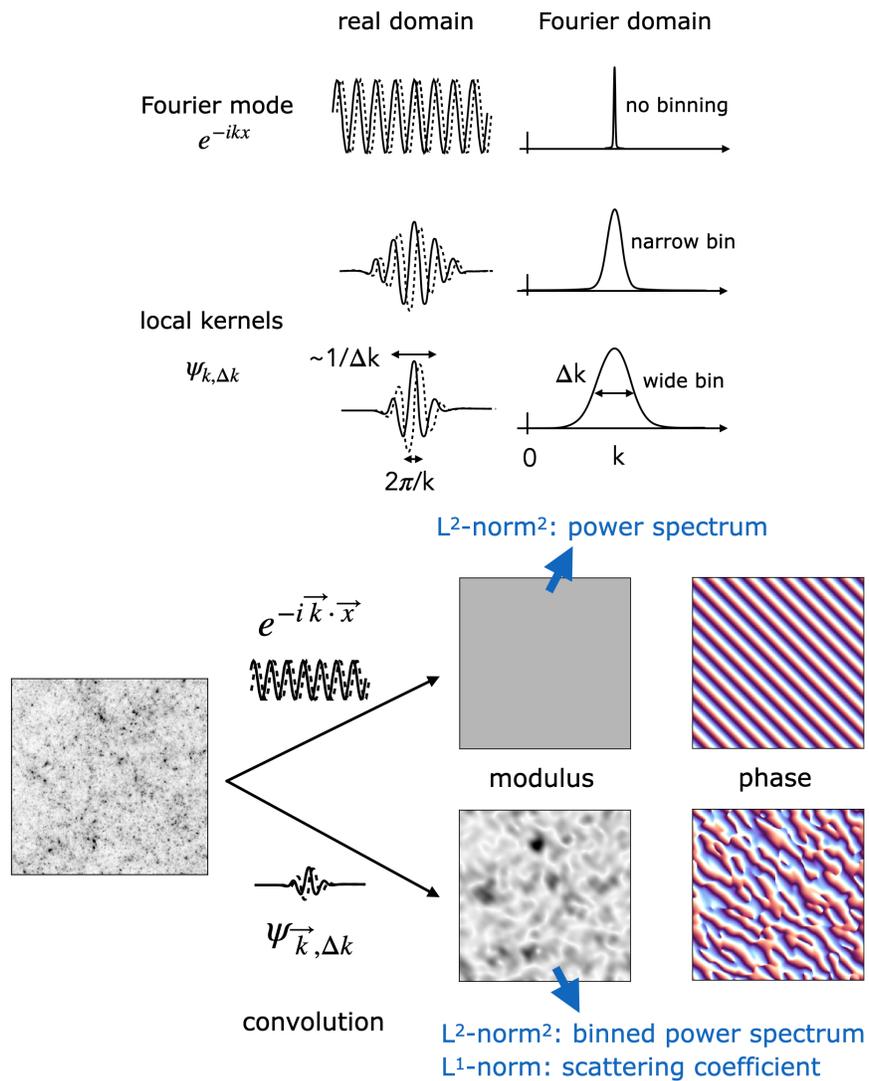


Figure 2-1. Upper panel: comparison between Fourier and localised kernels in 1D. Solid and dashed curves represent the real and imaginary part of the profiles, respectively. Lower panel: comparison between Fourier convolution and convolution with localised kernels.

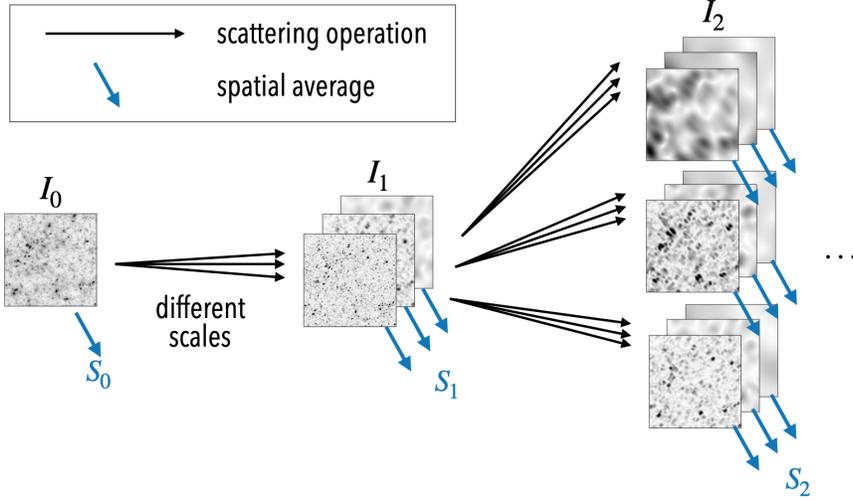


Figure 2-2. The scattering transform’s multi-layer structure and similarity to CNNs. The scattering operation is composed of a local (wavelet) convolution and a pixel-wise modulus.

different scales can be obtained by the successive application of the scattering operation for two kernels

$$I_2 \equiv ||I_0 \star \psi_1| \star \psi_2| \quad (2.5)$$

and an activation amplitude over the entire field can similarly be obtained by $S_2 = \langle I_2 \rangle$. Generalizing this to n layers:

$$I_n \equiv ||I_0 \star \psi_1| \star \psi_2| \dots \star \psi_n|, \quad (2.6)$$

results in a convolution tree or planar network which outputs a set of coefficients $S_n = \langle I_n \rangle$, the scattering coefficients. Forward-propagation into such a simplified CNN is therefore similar to calculating scattering coefficients. With the insight provided by both the expansion of the power spectrum and the simplification of CNNs, we now present the full formalism of the scattering transform.

2.1.3 Formalism

Conceptually, the scattering transform is composed of wavelet convolutions, modulus, hierarchy, and average. It yields translation-invariant descriptors $S_n(j_1, \dots, j_n)$ from an input field $I_0(x)$ by recursively applying the following operations:

$$I_{n-1} \rightarrow I_n \equiv |I_{n-1} \star \psi^j| \quad (\text{the scattering operation}) \quad (2.7)$$

$$S_n \equiv \langle I_n \rangle, \quad (2.8)$$

where ψ^j stands for a wavelet indexed by j , its logarithmic scale. When the signal $I(\vec{x})$ is higher than one dimension, an index l for the orientation of wavelet should also be added: $j \rightarrow j, l$. Considering stationary ergodic processes, the ensemble average can be estimated from an average over the extent x of a realization.

As illustrated in Figure 2-2, successive applications of the scattering operation form a tree structure, i.e. a planar multi-layer network, with the scattering fields $I_n(x)$ at its nodes. Each I_n is the intensity map of around a scale in the previous-order field I_{n-1} , which is similar to a local power spectrum analysis of the previous-order field. The average operation at each node is used to extract a translation-invariant scattering coefficient. It is similar to the pooling operation in convolutional neural networks. The 0th-, 1st-, and 2nd-order scattering coefficients can be written explicitly as:

$$S_0 \equiv \langle I_0 \rangle \quad (2.9)$$

$$S_1(j_1) \equiv \langle I_1^{j_1} \rangle = \langle |I_0 \star \psi^{j_1}| \rangle \quad (2.10)$$

$$S_2(j_1, j_2) \equiv \langle I_2^{j_1, j_2} \rangle = \langle \left| |I_0 \star \psi^{j_1}| \star \psi^{j_2} \right| \rangle. \quad (2.11)$$

As mentioned before, when $I_0(\vec{x})$ is higher than one dimension, such as an image, the scale index j becomes a scale index j plus an orientation index l , and the kernels ψ^j (or $\psi^{j,l}$) are chosen to be a family of wavelets, which are produced by dilating (and rotating) a mother wavelet.

Number of scattering coefficients

The number of scattering coefficients is determined by the number of wavelet combinations. It is usual to consider a dyadic sequence of scales: 2^j , with integer $1 \leq j \leq J$ which cannot exceed the length of the signal 2^J . With J choices of scales for each wavelet, there are J^n available combinations at the n th order. However, only a subset carries relevant information. As we will explain in section 2.4.1, the modulus operation extracts the envelope of the signal which scatters information and energy only into larger scales. As a result, only combinations with $j_2 > j_1$ are significant, which reduces the number of informative coefficients by a factor of 2^{n-1} . Thus, the number of useful scattering coefficients at each order is:

$$\begin{aligned} 0^{\text{th}} \text{ order: } & 1 \text{ coefficient,} \\ 1^{\text{st}} \text{ order: } & J \text{ coefficients,} \\ 2^{\text{nd}} \text{ order: } & J(J-1)/2 \text{ coefficients.} \end{aligned}$$

Because J is the logarithm of the dynamical range of scales in the field, the number of scattering coefficients increases slowly with the field size. They form a relatively compact set of descriptors.

Similar to the power spectrum, the scattering transform also applies to data in higher-than-one-dimension space, such as images or 3D physical fields. The only change of scattering transform's formalism is that the wavelet would not only be labelled by its scale j , but also its orientation l . As a result, all the scale indexes should be replaced by that of oriented scales: $j \rightarrow j, l$. The angular size and angular sampling rate remain a free choice. For example, for the two-dimension case, one may probe L orientations by using wavelets with angular size π/L in Fourier space, whose position angles are $\pi l/L$, with $0 \leq l < L$. Compared to 1D cases, there are L^n times more coefficients at the n th order. Fortunately, in many applications, this original set of scattering coefficients can be reduced.

Reduction of orientations

When considering isotropic fields, there are different levels of orientation reductions. To the simplest extreme, one can averages over all orientation indexes, which reduces the number of coefficients by

an order of L^n and creates a much more compact set of statistical descriptors:

$$s_1(j_1) \equiv \langle S_1(j_1, l_1) \rangle_{l_1} \quad (2.12)$$

$$s_2(j_1, j_2) \equiv \langle S_2(j_1, l_1, j_2, l_2) \rangle_{l_1, l_2}, \quad (2.13)$$

where $\langle \cdot \rangle_l$ denotes an average over orientation indices. As an example, for an image with a 256×256 pixels size, there are only 29 reduced scattering coefficients (S_0, s_1, s_2) when the full range of scales ($J = 7$) is probed.

However, those reduced coefficients $s_n(j_1, \dots, j_n)$ only depend on scales and do not provide the ‘shape’ information of the non-Gaussianity, which depends on angles. A less aggressive reduction for the 2nd-order coefficients is to keep the angular dependence on $l_2 - l_1$ and average only over l_1 , which reduce the number of coefficients by L . An even more informative reduction is to apply the scattering idea to orientations [4].

Normalization

At all orders, the scattering coefficients possess the units of the input field I_0 . It is sometimes convenient to manipulate dimensionless statistics. In addition, as the coefficients S_n are proportional to their previous-order field I_{n-1} , they are correlated. In order to deal with unitless coefficients, de-correlate them, and ease their interpretation, one can use the normalized scattering coefficients given by

$$s_n = \frac{S_n}{S_{n-1}}, \quad (2.14)$$

[6]. This is similar to the normalization commonly found in the context of moment statistics, such as the dimensionless skewness and kurtosis. Note that in the above equation S_n and S_{n-1} should belong to the same branch of the scattering tree, i.e., having the same $\{j_1, \dots, j_{n-1}\}$. The 1st-order ratio S_1/S_0 is not always necessary, depending on the physical meaning of the field.

2.2 Generative models

To demonstrate the power of the scattering transform, we first show some results of a stochastic generative model using the scattering coefficients (see Figure 2-3). Technically, we randomly generate fields with the same *translation-invariant* scattering coefficients as the input field, which is equivalent to drawing realizations of images from a micro-canonical ensemble defined by the particular summary statistics. Using only about 100 scattering coefficients, one can largely reproduce the non-Gaussian textures present in a wide variety of input fields. For comparison, using the power spectrum would typically result in Gaussian textures as in the right panel of Figure 1-1.

When it comes to non-Gaussian fields, there is no universal metric to compare different statistical estimators. In order to show the ability of the scattering transform to characterise a field, it is informative to build a generative model and check its performance.

A generative model can visualise what an input field looks like in the eye of a summary statistics. It randomly generates new fields with the same summary statistics as the input field, which form a ‘microcanonical ensemble’ of fields given the particular values of the statistics [see, e.g., 17]. If the generated fields have similar textures to the input field, or in other words, if the input field is a representative and typical one in the ensemble, then it means that the statistics capture the main features of the input field.

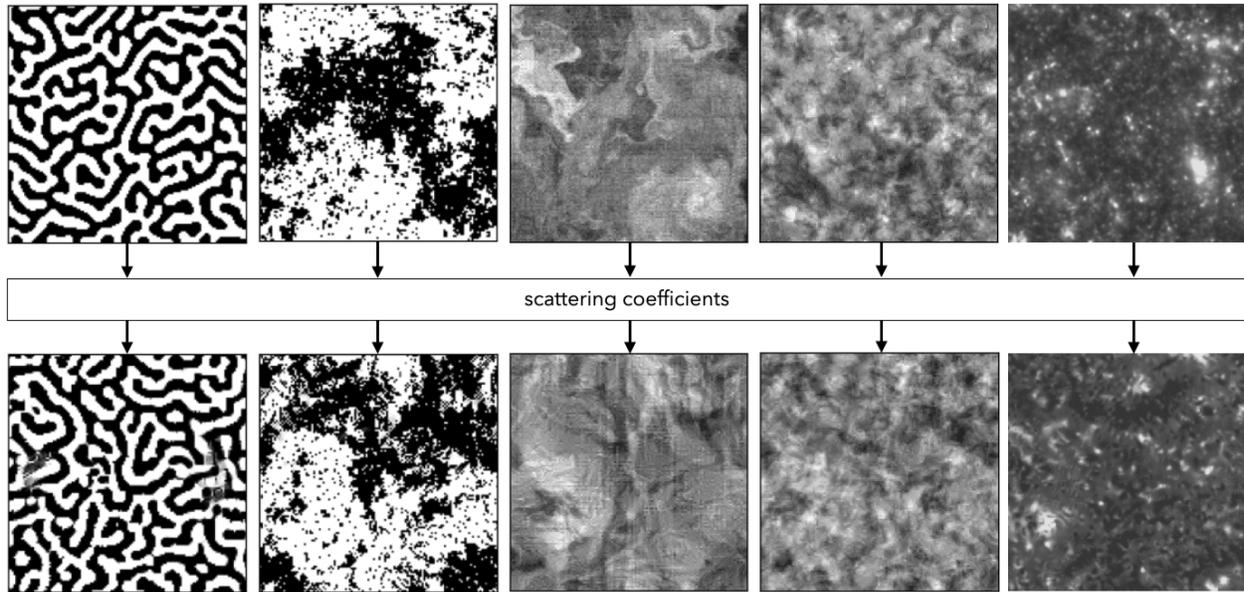


Figure 2-3. Image generation using the scattering statistics. Upper panels are fields from various disciplines in physics: Turing pattern, Ising model, turbulence revealed by the sea surface temperature, turbulence revealed by Solar ultraviolet emission, and gravitational collapsing of matter in the Universe. Lower panels are randomly generated images which have the same scattering coefficients as the upper counterparts.

As a familiar example, fields generated with the power spectrum statistics are Gaussian random field. That is to say, fields generated using the power spectrum for the left panel of Figure 1-1 would be like its right panel. For comparison, in Figure 2-3 we show the field generation results using the scattering statistics for different fields chosen from various disciplines of physics. The striking similarity between the input and generated fields is strong evidence of the great power of the scattering statistics to characterise various realistic non-Gaussian textures in physics.

Technically, we built those generative models by starting from a random field and then modifying it in order to minimise the difference between its summary statistics and that of the input field, until they became the same. We used the ‘adam’ optimiser in the python package `torch.optim` to implement the minimisation. Codes are publically available online¹.

There are several caveats about the generative models to keep in mind. For example, the image quality somewhat depends on the optimisation’s initial condition and the choice of loss function. More importantly, the concept of information must always be related to a particular task. In the image generation case, this task is ‘to distinguish fields under the metric of human eyes and brains’, which can be different from the task of ‘inferring the underlying physical parameters’. Nevertheless, in practice, the generation and physical inference abilities are often closely relative, which may interestingly suggest that the human visual system is a well-evolved and optimised to extract information from the physical world.

¹https://github.com/SihaoCheng/scattering_transform

2.3 Key properties

The scattering transform generates a statistical description of a field with a number of attractive properties for extracting physical information from data. It shares all the desirable properties listed for the power spectrum in section 1.2, namely: translation invariance, rotation invariance (after a reduction), dimensionality reduction, scale separation, energy extraction. As we will discuss in section 2.5, given the interpretability of the scattering coefficients, links to theoretical models are also possible. Here, we highlight additional properties of this estimator which offer key advantages for the analysis of physical fields:

- **Extraction of non-Gaussian information:** The raison d'être of the scattering transform is its ability to extract non-gaussian information and thus characterize some of the morphology of structures present in a field. It does so by being able to capture interactions between neighbouring Fourier modes, in contrast to the power spectrum. The fast convergence of the energy/information extraction enables one, for a wide range of physical fields, to quantify the relevant morphological information using only a second-order scattering transform. This is illustrated in Figures 1-1 and 2-3.
- **A novel way to describe energy distribution:** energy-preserving representations are often important to physicists. Similar to the power spectrum, the scattering transform provides an energy partition of the input field, but it does so with *non-linear* transforms which enables the extraction of both energy and non-Gaussian information at the same time. When the admissible wavelets (meaning that their squared profiles in Fourier space add up to one at all frequencies) are used, the energy (L^2 -norm²) is exactly partitioned into scattering coefficients [1]:

$$\langle |I_0|^2 \rangle = \|S\|_2^2 = S_0^2 + \sum_{j_1} S_1^2 + \sum_{j_1, j_2} S_2^2 + \dots$$

where S is the set of all scattering coefficients of I_0 . Note that with S_0 excluded, the others add up to the total variance of I_0 .

At each order of the scattering transform, energy is moved towards lower frequency by the modulus and some of it is extracted by the average over the field. The remaining part is further shifted towards lower frequencies by a new application of the modulus of the wavelet transform. Repeating this process guarantees the extraction of all the energy of the input field.

Strictly speaking, the partition in Eq. 2.15 conserves energy for an infinite expansion. In practice, one can only work with a truncated expansion but it is interesting to point out that one can decompose the energy into a finite set of low-order scattering coefficients together with the next-order scattering *field* I_{n+1} :

$$\langle |I_0|^2 \rangle = S_0^2 + \sum_{j_1} S_1^2 + \dots + \sum_{j_1 \dots j_n} S_n^2 + \sum_{j_1 \dots j_n} \langle |I_{n+1}|^2 \rangle. \quad (2.15)$$

It is also important to point out that the energy of scattering coefficients with order at least n comes almost only from higher frequency ranges.

- **Stability and asymptotic normality:** multiple stability properties make the scattering transform a robust estimator, which is crucial for both regression and classification problems.

Ideally, a small perturbation of the input field should yield a small change in the descriptors. The scattering transform has proven stability to adding noise and to geometrical deformations, thanks to the ‘low-order’ non-linear operation, modulus, and the logarithmic binning of scales performed by wavelets. For comparison, higher-order statistics are not stable to additive noise, which means they are sensitive to outliers. This property is not met by Fourier coefficients, unless one averages them using logarithmic bins of scales and regular frequency kernels as opposed to sharp bins using k_{\min}, k_{\max} values. In addition, by using a low-order non-linearity (modulus), the scattering transform never amplifies the tail of the field pdf. The average over the field thus gaussianizes quicker than the average involved in the power spectrum, and much quicker than the higher-order moment-based statistics. This property facilitates likelihood parameterisation and improves inference robustness.

- **Fast convergence and compactness:** an extremely useful property of the energy partition is that, for a wide range of physical fields, the norm of scattering coefficients decays exponentially fast when their order increases. More precisely, this exponential decay of the coefficients is guaranteed as soon as the Fourier transform of I decays at least as fast as $O(|k|^{-1})$, as proved by Waldspurger [18]. As a result, in many practical applications, *the leading orders of the scattering transform is sufficient to extract relevant information* from a field. One can thus work with a compact yet powerful set of summary statistics. This was shown by Andén & Mallat with a database of audio signals and Bruna & Mallat [3] with the Caltech-101 image texture database, applications of the scattering transform indicate that only a sub-percent fraction of the energy remains to be extracted by orders greater than three. In this paper, we therefore focus on only second-order scattering transforms. The visual examples displayed in Figure 2-3, which are all computed using only a second-order scattering transform, convincingly show that the relevant perceptual information is sufficiently well captured.

2.4 Understanding the operations

2.4.1 An overview on translation-invariant descriptors

Translation invariance is common in physical laws. But, incorporating such invariance into a descriptor turns out to be non-trivial. In this section, we step back and describe the big picture behind translation invariant descriptors, which helps to understand the scattering transform and many other statistical descriptors.

First, one would like to use *linear* operations alone to construct translation invariant descriptors, given their simplicity and interpretability. Unfortunately, only one such quantity can be constructed, that is the global mean of the field. The reason is as follows. Let us consider a field $I(x)$ as a vector in function space. Translations are orthogonal linear operations of I . They share the same set of eigen vectors: the Fourier modes e^{-ikx} , with eigen values $e^{-ik\Delta x}$. Translation invariance requires an eigen value of unity, which can only be obtained for $k = 0$, corresponding to the mean of the field over x .

Therefore, to extract translation-invariant information beyond the mean, an estimator must involve *non-linearity*. A typical paradigm is thus to first non-linearly transform the field and then

take the mean² over x :

$$\text{descriptor} \equiv \langle \text{transformed field} \rangle, \quad (2.16)$$

where the role of the non-linear transform is to convert high-frequency information into the $k = 0$ mode. In addition, the transform must not interfere with the ‘translation invariant’ property, in the sense that it should *commute* with translations. For example, *pointwise* non-linear transforms, such as the modulus used in the scattering transform and the activation functions in CNNs, satisfy the commutation requirement. Carefully designed *multiplications*, such as in the case of higher-order moment-based statistics, also satisfy this requirement. The power spectrum can be interpreted in both ways.

As non-linearity is usually awkward to analyse and interpret, one would still hope to use linear operations as much as possible, so that the non-linear operator is kept as simple as possible:

$$\text{descriptor} \equiv \langle \text{linearly and nonlinearly transformed field} \rangle. \quad (2.17)$$

Convolutions, which can be found in the power spectrum, scattering transform, CNNs, etc., are the only allowed linear operations in the transform because the commutation condition requires such linear operators be diagonal in Fourier domain. Intuitively, convolutions are used to linearly decompose the signal or select features in the signal in a translation-invariant way.

2.4.2 The role of wavelet convolution and modulus

The core operation $I \rightarrow |I \star \psi^{j,l}|$ employed by the scattering transform comprises two steps: a convolution by a complex-valued wavelet and a modulus operation. In short, the wavelet convolution selects scales, and the modulus converts fluctuations into their local strength.

Wavelets

Let us discuss the wavelet convolution first. As a wavelet is a band-pass filter, the wavelet convolution selects Fourier modes around a central frequency and coarsely separates information of different scales (see Figure 2-4). The wavelet framework provides a regular and systematic way to organise such kernels with symmetry and other provable mathematical properties.

Locality: Wavelet transform (convolutions) is a linear transform with properties *in between* the Fourier and the pixel decomposition, i.e., with partial locality in both real space and Fourier space. For example, its reaction to translations smaller than the wavelet scale resembles Fourier coefficients (pure phase shift), whereas for large displacements, it behaves like pixels. These localised, multi-scale kernels are not an arbitrary choice, but required by the properties of fields in our physical world. Such wavelet-like kernels have also been found in animal’s vision systems [19] and in the first layers of CNNs regardless of the training set [20], suggesting a fundamental role played by wavelets. This is because fields and signals in physical world often involves sparse superposition of *local* structures such as peaks, edges, eddies, filaments, patches, etc, or transient features in time series signals. These structures are revealed much clearer and simpler in the wavelet representation (i.e., the results of wavelet convolution) than in the Fourier or pixel representations. It has been recognised since 1990s that *wavelets* in general provide a set of basis that sparsify the representation [e.g., 21]. We

²Alternatively, one can use non-linear translation-invariant operations such as global maximum or minimum after a transform, in which case the transform itself need not be non-linear anymore.

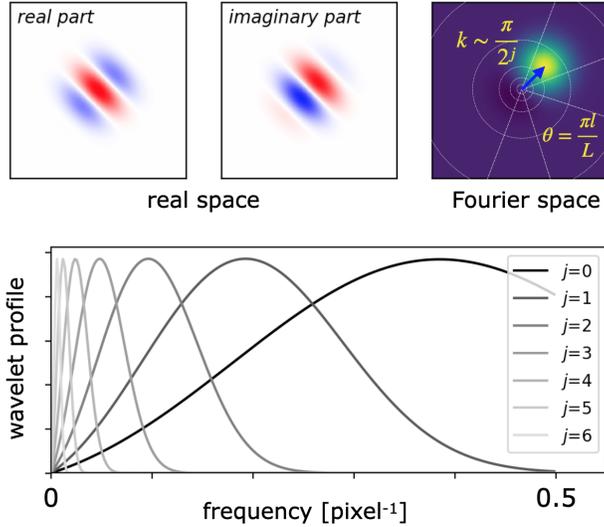


Figure 2-4. *Upper panel:* profile of a Morlet wavelet ($j = 6, l = 0$, image size 512×512 pixels) in the real space and another one ($j = 1, l = 1$) in Fourier space. The centre of the Fourier space represents zero frequency. *Lower panel:* radial frequency profiles of a family of wavelets. Dilating/contracting (by factor of 2) and rotating (by π/L) one wavelet give the whole family of wavelets used in the scattering transform.

will come back to sparsity when discussing the interpretation of S_2 coefficients. In addition, as shown in Figure 2-1, localised kernels make a wide selection of scales and leave rich structures in real space, which makes the iteration in the scattering transform and CNNs possible.

Regularity: Wavelets in a family have the same shape but different sizes and orientations, which is a desired regularity for physics research. It thus provides a wide, *logarithmic* ‘tiling’ of the Fourier space and allows for a logarithmic sub-sampling of scales, which significantly reduces the number of scattering coefficients to be calculated. Wavelet convolution itself is just a linear decomposition, but its wide logarithmic tiling of scales provides the foundation of deformation stability of the scattering coefficients.

A common choice is to use wavelets with dyadic scales, whose width in Fourier space is roughly equal to its central frequency and can be sampled at factor-of-two scales (Figure 2-4). In general, they provide a good compromise between the need for separating scales and probing scale interactions in the scattering framework. When the field is higher than 1D, there is usually also a uniform tiling of orientations.

Also, the scattering transform uses complex-valued wavelets, so after such convolutions, the field becomes complex-valued. Complex-valued wavelets lie in one side of the Fourier space. Their relation to real-valued wavelets is similar to that of the complex to real Fourier modes (e^{-ikx} vs. $\sin(kx)$). For on real-valued fields, they do not miss any information, but react much more smoothly to translations and modulus.

Morlet wavelets

In the framework of scattering transform, it is free to choose which family of wavelet to use. Among many, the Morlet wavelets provide a good balance for the locality in both real space and Fourier

space. Morlet wavelets basically have Gaussian profiles. If we simply use a Gaussian envelope to modulate a plane wave, then we obtain a Gabor function,

$$G(\mathbf{x}) = \frac{1}{\sqrt{|\boldsymbol{\Sigma}|}} e^{-\mathbf{x}^T \boldsymbol{\Sigma}^{-1} \mathbf{x} / 2} e^{i \mathbf{k}_0 \cdot \mathbf{x}}, \quad (2.18)$$

where $\boldsymbol{\Sigma}$ is the covariance matrix describing the size and shape of the Gaussian envelope, and \mathbf{k}_0 determines the frequency of the modulated oscillation. To keep maximum symmetry, usually $\boldsymbol{\Sigma}$ is selected to have only 1 eigen-value different from the others, and \mathbf{k}_0 to be along that eigen-direction. Thus we denote the eigen-value along \mathbf{k}_0 by σ^2 and the other eigen-value by σ^2/s^2 . The parameter s is also the ratio of transverse to radial width of the wavelet in Fourier space.

The Fourier transform of a Gabor function is simply a Gaussian filter centred at \mathbf{k}_0 ,

$$\tilde{G}(\mathbf{k}) = e^{-(\mathbf{k}-\mathbf{k}_0)^T \boldsymbol{\Sigma} (\mathbf{k}-\mathbf{k}_0) / 2}. \quad (2.19)$$

Wider envelope in real space makes narrower filter in Fourier space. Note that the product $k_0 \sigma$ determines the number of oscillations within $\pm \pi \approx 3$ standard deviation of the Gaussian envelope and allows for a trade off between spatial and frequency resolution.

Unfortunately, a Gaussian profile in Fourier space does not go to zero at zero frequency. This contradicts the admissibility of wavelet which requires wavelets to strictly be band-pass filters, not low-pass filters. Therefore, a small correction is required. A simple solution is to introduce an offset, β , before the Gaussian modulation. In Fourier space this is equivalent to subtracting another Gaussian profile centred at 0 to cancel out the 0-frequency contribution. Families of wavelets created in this way are called Morlet wavelets. Formally,

$$\psi(\mathbf{x}) = \frac{1}{\sqrt{|\boldsymbol{\Sigma}|}} e^{-\mathbf{x}^T \boldsymbol{\Sigma}^{-1} \mathbf{x} / 2} \left(e^{i \mathbf{k}_0 \cdot \mathbf{x}} - \beta \right), \quad (2.20)$$

where $\beta = e^{-\mathbf{k}_0^T \boldsymbol{\Sigma} \mathbf{k}_0 / 2}$ is determined by the admissibility criterion. Its Fourier transform is

$$\tilde{\psi}(\mathbf{k}) = \tilde{G}(\mathbf{k}) - \beta e^{-\mathbf{k}^T \boldsymbol{\Sigma} \mathbf{k} / 2}. \quad (2.21)$$

For example, a set of parameters for 2D dyadic wavelets is given by the ‘kymatio’ package (<https://www.kymatio.io>),

$$\begin{aligned} \sigma &= 0.8 \times 2^j \\ k_0 &= \frac{3\pi}{4 \times 2^j} \\ s &= 4/L, \end{aligned} \quad (2.22)$$

where σ is in unit of pixels, j is an integer starting from 0, and k_0 is between 0 and 2π . This choice allow a family of Morlet wavelets best covers the whole Fourier space with a dyadic sequence of scales (2^j). Examples of the Morlet wavelets we use are shown in Figure 2-4. Within the wavelet envelope, there are about 2 cycles of oscillations, because $k_0 \sigma \approx 2$.

Modulus

One key idea of the scattering transform is to generate ‘first-order’ statistics, in contrast to higher-order moments, which multiply an increasing number of field intensities and cause instability to

outliers. Being a linear operator, the wavelet convolution certainly keeps the ‘first-order’ property. However, for a homogeneous random field, convolution alone cannot extract information beyond the mean of the original field $\langle I \rangle$, because the expected value operator commutes with all linear transformations. Extracting more information requires non-linear operations. For example, in N -point functions, the multiplication of field intensities plays the role. The scattering transform, on the other hand, employs the *modulus* operation, which is a natural choice to preserve the desired property of working with first-order statistics.

Pointwise modulus is a non-linear operation resulting in a non-negative field. It has two important properties:

1) intuitively, the modulus of complex-values (also called analytical) wavelets converts fluctuations into their *envelop*, i.e., their *local strengths*. The envelop has in general lower frequencies than the original fluctuations. So, modulus has the property that it ‘scatters’ information and energy from high frequencies to low frequencies. This naturally fits into the paradigm of constructing translation-invariant descriptors for high-frequency information (see section 2.4.1).

2) it keeps norms of the field, which allows the scattering transform to provide an energy (variance) partition of the field (see section 2.3). This is desired both in terms of physical interpretations and of stability and robustness.

As the modulus is taken in the real space and is non-linear, its behaviour in Fourier space is not simple. Nevertheless, we collect some intuitive understandings and present them in section 2.4.4 for interested readers.

2.4.3 Hierarchy: information extraction beyond the power spectrum

There are a number of similarities between the power spectrum and each single iteration of the scattering transform. Indeed, the power spectrum can be defined using the formalism of the 1st-order scattering coefficients $S_1 = \langle |I_0 \star \psi| \rangle$:

$$P(\mathbf{k}) \propto \langle |I_0 \star \psi'|^2 \rangle \text{ with } \psi' = e^{-i\mathbf{k}\cdot\mathbf{x}} . \quad (2.23)$$

The differences between the two estimators S_1 and $P(k)$ are the choice of convolution kernels (wavelets ψ or Fourier modes ψ') and that of the norm (L^1 versus L^2). Therefore, the 1st-order scattering coefficients have similarity to the power spectrum. Both of them characterize the strength of fluctuations (or clustering) as a function of scale.

However, in the case of the power spectrum, the convolution kernel ($\psi' = e^{-i\mathbf{k}\cdot\mathbf{x}}$) is completely de-localized in real space. Thus, the power spectrum’s version of I_1 fields ($|I_0 \star \psi'|^2$) lose all spatial information. In contrast, the use of localized wavelets in the scattering transform allows I_1 to preserve spatial information, as shown in Figure 2-1. According to the analogy with the power spectrum, the mean of an I_1 field characterizes the average amplitude of Fourier modes selected by the wavelets, whereas the spatial distribution of fluctuations in I_1 , missing in the power spectrum analogue, in turn encodes the phase interaction between those Fourier modes. This information can be extracted by applying the scattering operations once again, $I_1 \rightarrow I_2 = |I_1 \star \psi_2| = ||I_0 \star \psi_1| \star \psi_2|$, and then measuring the mean of I_2 , i.e., 2nd-order scattering coefficients S_2 .

According to the power spectrum analogy, S_2 coefficients resemble the power spectrum of I_1 fields and measure clustering properties on I_1 . Because I_1 fields highlight the regions where fluctuations around a scale are stronger, the 2nd-order coefficients can be understood as measuring the clustering of structures highlighted in I_1 , i.e., the ‘clustering of (clustered) structures’.

This leads to an interesting intuition: we need two points to describe the scale of one structure and an additional two points for another one. Therefore, the 2nd-order scattering coefficients S_2 , which measure the clustering of sized structures, include information up to about 4-point. In general, an n th-order scattering coefficient S_n will contain information up to about 2^n -point function of the input field. By this ‘hierarchical clustering’ design, the scattering-transform expansion quickly includes information from higher-order statistics.

However, it should be noted that there is still a fundamental difference between the scattering transform and N -point functions. There are mainly two difficulties associated with N -point functions to characterize a random field: the failure to describe distribution tails and the huge number of configurations. The first difficulty, related to the multiplication of multiple random variables, leads to high variances and also prevents the extraction of information from fields whose pdf has a tail [22]. The scattering transform, which uses modulus and does not enhance the tail, can significantly alleviate this problem. The second difficulty may be alleviated by an efficient binning. For example, the hierarchical wavelet transform used in the scattering transform *is* a binning strategy that can also be applied to N -point functions. A detailed comparison to higher-order statistics is discussed in section 2.6.2.

2.4.4 Scattering operations in Fourier space

It is enlightening to collect some intuition of the scattering transform in the Fourier domain. In general, as a non-linear operator, a modulus in real space will mix Fourier modes and scatter information among different frequencies. In particular, taking the modulus of $I \star \psi$, where ψ has a single peak in Fourier space, will re-express I ’s information around ψ ’s frequency in forms of lower frequencies. In other words, the typical frequency of $|I \star \psi|$ is lower than $I \star \psi$.

Intuitively, this is because the modulus is converting complex-valued oscillations into its local strength, namely its envelope. Formally, this can be revealed by first writing $|I \star \psi|$ as $\sqrt{(I \star \psi)(I \star \psi)^*}$, where $*$ stands for complex conjugate, and then Taylor expanding the square root in terms of $(I \star \psi)(I \star \psi)^* - C$, where C is the mean of $(I \star \psi)(I \star \psi)^*$ over all pixels [23]. The leading term of the Taylor expansion is proportional to $(I \star \psi)(I \star \psi)^* - C$ itself, which corresponds to $I \star \psi$ ’s auto-correlation in Fourier space. When the power spectrum of I is a smooth function, the frequency distribution of $I \star \psi$ is similar to ψ . For the Morlet wavelets used in the scattering transform, the central wavenumber of the wavelet ψ is roughly k_0 (as defined in section 2.4.2), and its half-width in Fourier space around $1/\sigma$. So, its auto-correlation will have an half-width around $\sqrt{2}/\sigma$ and a centroid at 0. As $\sqrt{2}/\sigma \approx 0.75k_0 < k_0$ (Equation 2.22), this means that the typical frequency of $|I \star \psi|$ is lower than $I \star \psi$. Therefore, the core operation $I \rightarrow |I \star \psi|$ re-expresses high frequency information of I_n in terms of lower frequency modes including the 0-frequency component in the next-order fields I_{n+1} . As the 0-frequency component is translation invariant, it can be directly used as a statistical descriptor of the original field.

Writing the modulus $|x|$ as $\sqrt{|x|^2} = \sqrt{x \cdot x^*}$ brings an interesting question: what happens if we replace each modulus by modulus squared? It can be shown that, in this case, the n th-order scattering coefficients will exactly become some averaged 2^n -point-spectra weighted (binned) by wavelets. Nevertheless, they are not equivalent to any degenerate case of 2^n -point functions in either real or Fourier domain. For example, at the 2nd order, these ‘pseudo’ scattering coefficients become $\iiint \tilde{I}_0(\mathbf{k}_1)\tilde{I}_0(-\mathbf{k}'_1 - \mathbf{k}_2)\tilde{I}_0(\mathbf{k}'_1)\tilde{I}_0(-\mathbf{k}'_1 + \mathbf{k}_2) \cdot W \cdot d\mathbf{k}_1 d\mathbf{k}'_1 d\mathbf{k}_2$, where the weight is determined by the wavelets: $W = \tilde{\psi}_1(\mathbf{k}_1)\tilde{\psi}_1(\mathbf{k}_1 + \mathbf{k}_2)\tilde{\psi}_1(-\mathbf{k}'_1)\tilde{\psi}_1(-\mathbf{k}'_1 + \mathbf{k}_2)\tilde{\psi}_2^2(\mathbf{k}_2)$, and the tilde sign denotes Fourier

conjugate. Although these ‘pseudo’ coefficients may help us understand the connection between scattering transform and N -point functions in terms of how they organise spatial configurations, the genuine scattering transform is fundamentally different from N -point functions, because it generates ‘first-order’ estimators, which alleviates the problem of classic moments described in Carron [22] when dealing with heavy tailed probability distribution. We will discuss this further in section 2.6.2.

2.5 Interpretability

Being mathematically well-defined, the scattering coefficients can be interpreted. As we will describe below, they can be related to interpretable quantities such as scale, energy, feature sparsity and shapes. A finite scattering transform expansion does not uniquely determine a signal but characterizes some of its statistical properties. After having computed the scattering coefficients of an input field, for example an image, it is not possible to exactly recover it from the coefficients. However, a perceptually plausible sample of the same texture can be reconstructed, as illustrated in figure 2-3.

2.5.1 First-order coefficients

The S_1 coefficients are qualitatively similar to the power spectrum amplitudes. Both of them characterize fluctuation strength as a function of scale. They differ only in two aspects: the scattering transform selects fluctuations using a family of dilated overlapping wavelets rather than potentially arbitrary and narrow bins in scale. Second, while the power spectrum uses the L^2 norm of the convolved field, the scattering transform uses the L^1 norm which does not amplify the extreme values of the fluctuations.

In Figure 2-5, we apply the scattering transform to a series of two-dimensional random point processes with densities ranging from very dense to very sparse, with the dense end asymptotically corresponding to a Gaussian random field. We point out that the majority of physical fields, after a convolution by a localized kernel representing the typical features, would lie in the intermediate regime. The upper panels illustrate realizations of these fields and the bottom panels show how the energy is partitioned through the first, second and higher-order scattering coefficients.

As expected, in the case of a Gaussian random field, most of the energy is carried by the first-order coefficients. Higher-order coefficients carry only a negligible fraction of the total energy. In addition, being generated from random points, all these fields possess a power spectrum consistent with white noise. As a result, the scattering coefficients do not depend on scale. When the density of the Poisson process decreases, the sparsity of the field increases. The partition of energy now spreads towards second- and higher-order scattering coefficients. This effect is scale-dependent and more pronounced on small scales. Indeed, considering large scales is similar to smoothing the field on those scales, thus reducing the level of shot noise and making the fluctuations closer to Gaussian. Interestingly, we observe that, in this regime, the fraction of energy carried by scattering coefficients at orders higher than two is relatively small. It also means the s_1 coefficients are approximately just the square root of a logarithmically binned power spectrum, because of the energy conservation of the scattering transform:

$$\int P(k) \tilde{\psi}^{j_1}(k) dk = (S_1^{j_1})^2 + \sum_{j_2} (S_2^{j_1, j_2})^2 + \dots \quad (2.24)$$

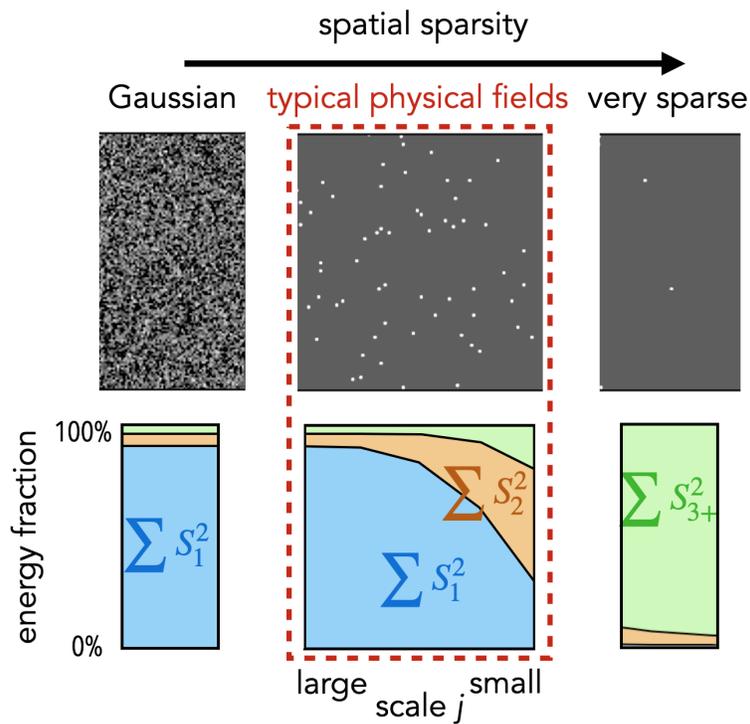


Figure 2-5. The scattering coefficients and sparsity. Given the same power spectrum, spatially-spread textures (Gaussian textures) have most energy in S_1 coefficients, spatially-sparse textures leak energy from S_1 to higher order scattering coefficients. This provides an intuitive understanding of the reduced coefficients $s_{21} \equiv S_2/S_1$: the sparsity of textures.

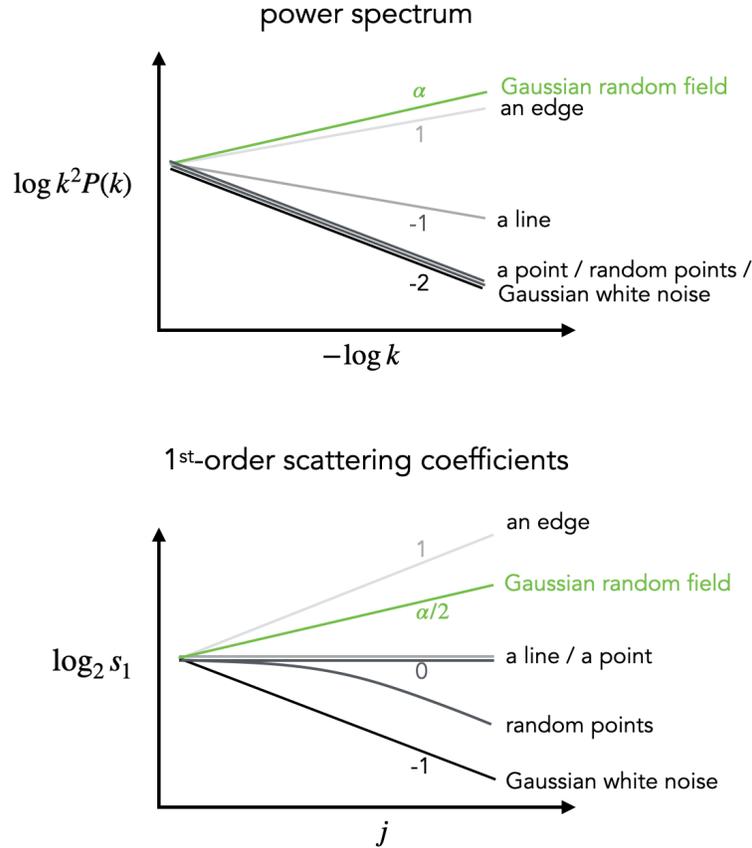


Figure 2-6. The similarity and difference between power spectrum and 1st-order scattering coefficients. Here we use 2D fields as examples. It can be seen that when the field structures are not extremely sparse in real space, the s_1 coefficients basically follow the square root of the (logarithmically binned) power spectrum, such as in the cases of Gaussian fields and random points in large scales. However, when the structures are very sparse, the s_1 coefficients significantly deviate from the square root of power spectrum.

For fields with significant sparsity, the s_1 coefficients may deviate from the power spectrum, as shown in Figure 2-6. For a very sparse field, most energy will be stored in coefficients with higher scattering orders.

2.5.2 Second-order coefficients

Second-order scattering coefficients offer deeper insight into some of the statistical properties of a field. They characterize transient phenomena such as localized structures in space, rapid changes in time or amplitude modulation. Substantial non-Gaussian information emerges from these coefficients. In particular, they provide co-occurrence information at the scales j_1 and j_2 and thus capture interferences of the field between features selected with two successive wavelets ψ_{j_1} and ψ_{j_2} . This is why they are called scattering coefficients.

The second-order coefficients are obtained after applying the scattering operation to the transformed field I_1 . They characterize the assembly or clustering of *patterns* (with a given size or scale j_1) by quantifying the strength of their fluctuations mapped in $I_1^{j_1}$ as a function of scale $j_2 > j_1$.

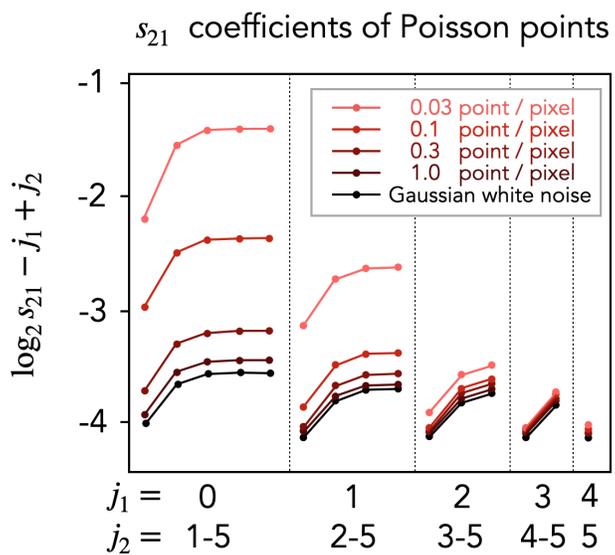
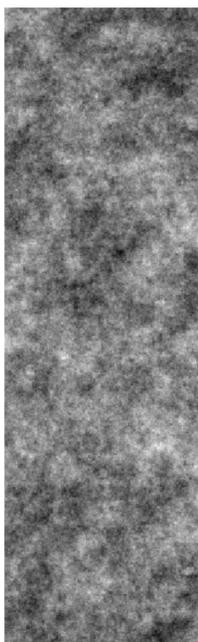


Figure 2-7. The s_1 coefficients of random point process and Gaussian white noise, as a function of scale combinations (j_1, j_2) .

Gaussian field:
not spatially sparse

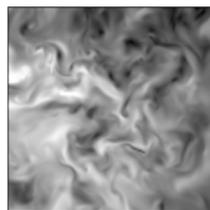


fields with more features:
sparse in some representation

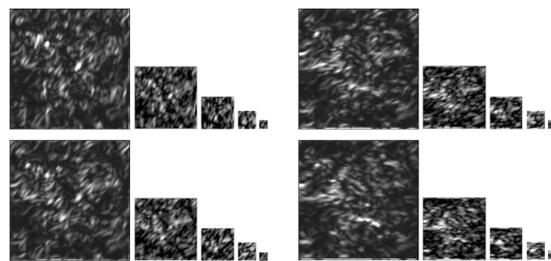


points

spatially sparse



typical physical fields



spatially sparse after wavelet convolution

Figure 2-8. Sparsity and structures in a field.

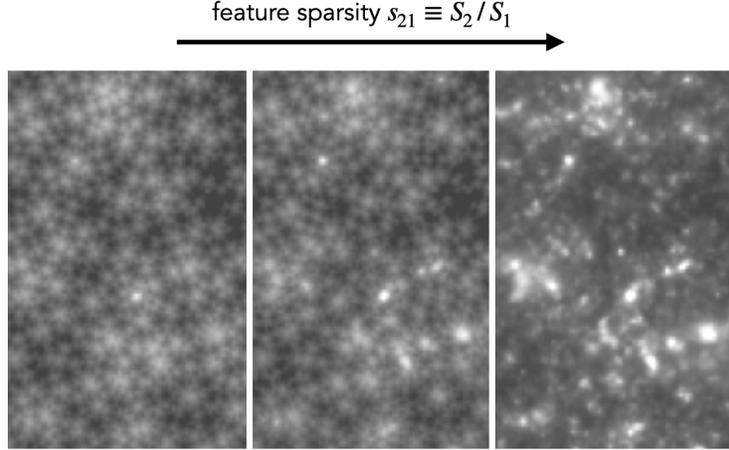


Figure 2-9. Visual interpretation of scattering coefficients s_{21} , generated by starting from a non-Gaussian field (right panel) and requiring the s_{21} coefficients to decrease while keeping the power spectrum unchanged. It can be seen that the s_{21} measures the spatial sparsity of structures in a field at various scales.

In other words, this corresponds to the clustering, on a scale j_2 , of structures on scale j_1 . Being a function of the two scales j_1 and j_2 and two orientations l_1 and l_2 , the second-order coefficients are more numerous than at first-order. As mentioned above, as they are significantly correlated, dimensionality reduction techniques can be used to construct a more compact set of summary statistics, if need be. However, in order to directly interpret some of the coefficients, one can focus on two particularly meaningful combinations:

- **Feature sparsity s_{21} :** the ratio between S_2 and S_1 coefficients can be interpreted as an estimate of feature *sparsity* at various scales, as illustrated in figure 2-9. Intuitively, sparsity indicates whether fluctuations or structures are concentrated at a few positions or widely spread. If not relevant orientation information is expected, it is useful to consider the orientation-averaged ratio

$$s_{21} \equiv \langle S_2 / S_1 \rangle_{l_1, l_2} . \quad (2.25)$$

This quantity directly informs us on departures from Gaussian random fields. It is sensitive to the strength of structure or localized features present in the field. If the density of such features were to increase and start to overlap to the point of becoming ambiguous, the central limit theorem will apply and the field will Gaussianize. During this process s_{21} will decrease and approach unity. Sparsity and structures are related to lower entropy, compared to the level reachable with the same energy but through a Gaussian random field. Unlike the estimation of L^2 -norm or energy, sparsity estimates are not generic and depend on the chosen basis. Certain manifestations of sparsity will only be fully revealed by the use of certain bases. For the ensemble of physical fields, the proper basis set that sparsifies the representation is usually composed of wavelets. That is why the scattering transform is also based on wavelet convolutions.

- **Shape s_{22} :** When angular information becomes non-negligible, a less aggressive orientation reduction is needed. As an example, we present here one simple option that also has interesting interpretations. In the 2nd order, coefficients depend on two orientations l_1 and l_2 . Information

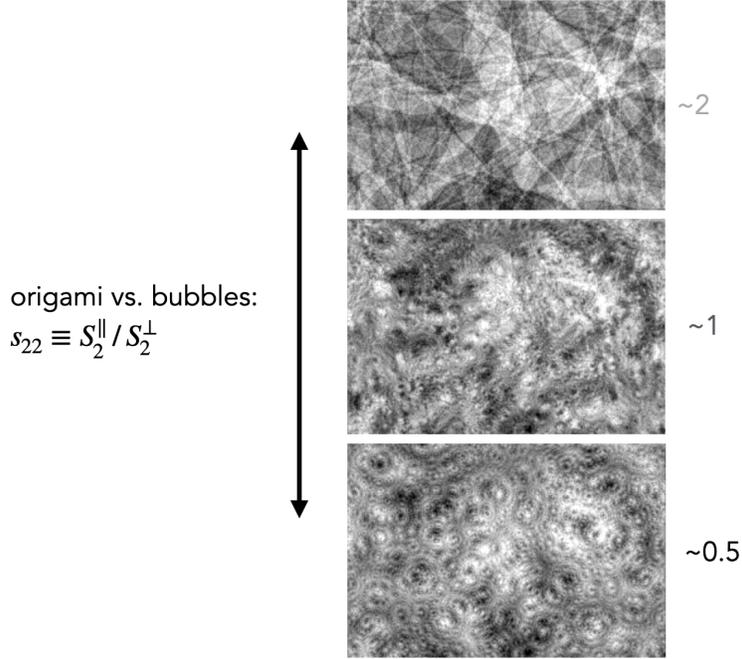


Figure 2-10. Visual interpretation of scattering coefficients s_{22} , generated by requiring the s_{22} coefficients changing from below unity to above unity, while keeping the s_1 and s_{21} coefficients as a particular function of scales. It can be seen that the s_{22} coefficients contain information about structures described by combinations of orientations. They differentiate straight lines from swirls.

about the shape of features or fluctuations can be an important aspect of a field. This is usually the property that allows the identification of textures around us, whose aspects can vary due to variations in angle, illumination, scale, etc. Shape information is captured by a number of the $S_2(j_1, j_2, l_1, l_2)$ coefficients but it is possible to introduce a convenient reduction which carries valuable information. If rotation invariance applies, the global orientation can be averaged out but the relative angle between l_1 and l_2 remains informative. We can probe such a dependence by selecting two key orientations – parallel and perpendicular – and consider the ratios of their respective scale-dependent coefficients as a new group of reduced coefficients, which we tentatively call as

$$s_{22} \equiv \langle S_2^{l_1=l_2} / S_2^{l_1 \perp l_2} \rangle_{l_1}. \quad (2.26)$$

This quantity is easily interpretable, as illustrated in Figure 2-10. For $s_{22} < 1$ corresponds to curvy, soft-look patterns like bubbles and swirls, as the anisotropic fluctuations at one scale (j_1) are mainly distributed along the orthogonal direction at larger scales (j_2). For $s_{22} > 1$ corresponds to straight, hard-look patterns like long lines, filaments, and origami textures, as the anisotropic fluctuations are mainly distributed along the same direction.

The visualisation in Figure 2-10 is achieved by randomly generating images with the same S_1 and s_{21} , i.e., similar power spectrum and feature sparsity, but different s_{22} coefficients. Technically, we set the scale dependence of S_1 and s_{21} as power laws ($\propto 2^{aj_1+bj_2}$). In each panel, we set no scale dependence for s_{22} to acquire the basic intuition. In real world, these s_{22} coefficients can of course depend on the scale combinations (j_1, j_2), resulting in rich possibilities of textures. Note that similar to s_{21} (a measure of sparsity), s_{22} also depends only on the two scales j_1 and j_2 . So, after including s_{22} , the set of reduced scattering statistics becomes even

more powerful while remaining compact.

2.6 Relation to other methods

2.6.1 Difference from wavelet transform

Wavelet transform (an organised set of wavelet convolutions) serves as a crucial part of the scattering transform. However, the two have different properties and goals, and should not be confused. Mainly, wavelet transform is a linear transform. Its coefficients are *not* translation-invariant. In a sense, it is similar to the Fourier transform. In contrary, the goal of the scattering transform is to generate translation-invariant coefficients, and non-linearity (modulus) is a crucial part of the scattering transform. In a sense, it is similar to the power spectrum, though it has multiple layers of non-linearity to capture characterise complex features.

The advantage of the linear bases (or ‘frames’, when being redundant) provided by wavelets is that it sparsifies the representation of natural images and physical fields, which means structures and features in those fields are most clearly seen in such representation. This property of wavelets transform (wavelet convolutions) is used in the scattering transform. Moreover, it has also been used to obtain effective approximations, such as in tasks of compression, reconstruction, and de-noising [e.g., 24].

2.6.2 Relation to higher-order statistics

Higher-order statistics such as N -point functions, moments, and poly-spectra are multiplications of field variables $I(x_2)I(x_2)..I(x_n)$ and their linear combinations. Despite its close connection to perturbation theory, it is a beast for data analysis of non-Gaussian fields. It has two practical problems: 1) the multiplications amplify ‘outliers’ (distribution tail), making the statistics non-robust; 2) the number of statistics explodes with N , so information is diluted quickly.

The scattering transform can be seen as a low-order version of the N -point functions. Both of them have the ability of separating scales and probing their interactions. Nevertheless, there are two main differences: 1) the scattering transform replaces the multiplication by a ‘first-order’ non-linearity (modulus), which makes the coefficients stable and robust, and 2) it bins the N -point configurations by wavelets, which significantly reduces the number of coefficients without losing much relevant information.

Difference: Don’t amplify the tail. Fold the core.

Higher-order statistics and scattering coefficients employ different non-linear operations to extract information: successive multiplications versus successive modulus operations. As we will show below, the former amplifies the tail of the underlying distribution whereas the latter folds the core. While the former leads to instabilities or even divergence, the latter guarantees stability and convergence. To compare the two approaches, we use a simple example, using a one-dimensional single-variable probability distribution $p(x)$. Instead of a specific wavelet convolution, we will simply consider taking the mean. The insights gained from this idealized experiment will also apply to generic high-dimensional random fields.

To probe the outskirts of a distribution, moment-based statistics use products of random variable(s)

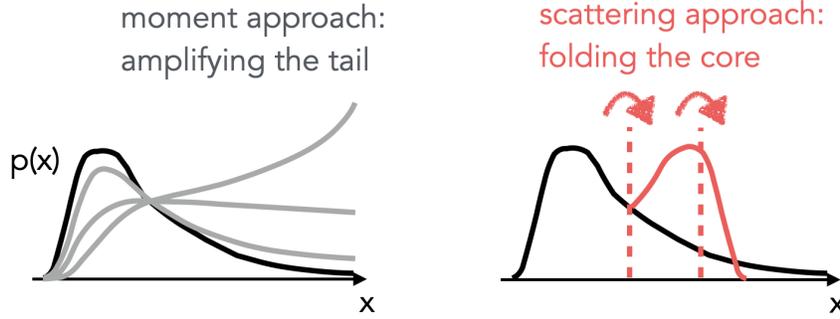


Figure 2-11. Comparison between the moment approach with scattering approach: calculating moments is equivalent to amplifying the tail of $p(x)$ and then integrating over x . When the tail is heavy, it takes the risk of divergence. In contrast, the scattering approach is equivalent to folding the core of $p(x)$.

to down-weight the distribution's core: it *amplifies* the tail (see the left panel of Figure 2-11), causing potential non-robustness and divergence problems. The scattering transform, however, follows another approach to probe the distribution tail: with repeated modulus, it *folds* the core instead of amplifying the tail, as shown in the right panel of Figure 2-11. In single-variable cases, wavelet convolution is reduced to just a subtraction of the mean, which is similar to the idea of central moments. Indeed, the single-variable analogue of ‘scattering transform’ can be defined by first repeatedly transforming the sample through subtracting the sample mean and taking modulus $x_n \rightarrow x_{n+1} \equiv |x_n - \bar{x}_n|$:

$$\begin{aligned}
 x_0 &= \text{original sample} \\
 x_1 &= |x_0 - \bar{x}_0| \\
 x_2 &= |x_1 - \bar{x}_1| \\
 &\dots
 \end{aligned}
 \tag{2.27}$$

and then taking the mean of each transformed set $S_n \equiv \bar{x}_n$ as the n th-order ‘scattering coefficients’. Similar to eq. 2.15, these statistics provide a partition of the sample variance, which indicates stability and robustness.

Below we compare the amount of information extracted by the moment and scattering approaches, quantified by the Fisher information. With a finite sample drawn from a distribution, there is a limit of how well one can constrain the distribution parameters, formalised by the Fisher information matrix and Cramér-Rao inequality. In practice, one usually uses summary statistics instead of the original sample data points to constrain distribution parameters. The question here is, can these summary statistics squeeze out all available information and reach the limit set by the Fisher information matrix? In contrast to what many physicists believe, moments can fail thoroughly even if all high orders are taken into account. To explore the transition from Gaussian to non-Gaussian regimes relevant to cosmology, Carron [22] studied a family of log-normal distributions³. He showed that moments fail quickly and completely when it becomes highly non-Gaussian, as shown in Figure 2-12, a problem well-known to statisticians but overlooked by physicists at times. In fact, for any distributions with a tail decaying slower than exponential, moments do not form a

³ $p(x|\mu, \sigma) = \frac{1}{x\sigma\sqrt{2\pi}} \exp(-\frac{(\ln x - \mu)^2}{2\sigma^2})$. When σ is small, it tends to Gaussian; when it is large, the tail becomes heavier.

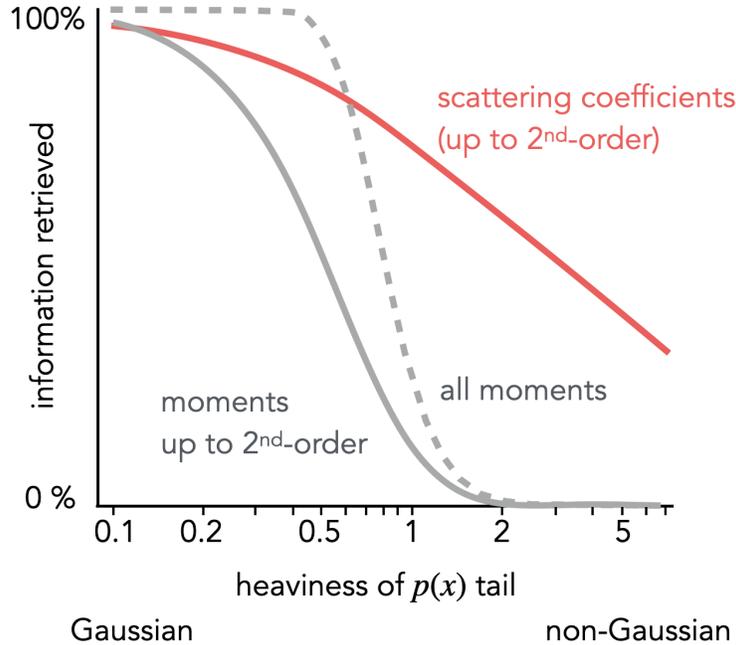


Figure 2-12. The fraction of Fisher information contained in moment and scattering statistics, respectively, of log-normal distributions. The x axis is defined as $\sigma_\delta \equiv \sqrt{\exp(\sigma^2) - 1}$, where σ is the variance of the corresponding Gaussian variable, and σ_δ is also the ratio of the standard deviation to the mean of the log-normal distribution. The folding (scattering) strategy works much better in the heavy-tail regime than traditional moments.

complete set anymore, even if they all exist and are finite. Moreover, the information stored in the first several orders is diluted quickly [see Figure 2 of 22].

What about the scattering approach? In Figure 2-12, we show results of the same experiment with the one-variable scattering coefficients defined in eq. 2.27. These repeated-modulus statistics retrieve much more information in the highly non-Gaussian regime⁴. Although they do not form a sufficient set of statistics, either, they fail much slower than higher-order moments, because the the modulus does not amplify the distribution’s tail.

Similarity

In section 2.1.1 we obtain the scattering transform by three modifications of the power spectrum: dropping out the square of modulus, using local kernels, and repetitions (going to deeper layers). If one insists on using squared modulus while keep the other two modifications, one will obtain exactly some 2^n -point functions. This can be understood by writing the modulus squared as $|I|^2 = I \cdot I^*$, which is a 2-point correlator. This subset of 2^n -point functions serve as a bridge between the

⁴Different from Carron [22], we obtained the derivatives in Fisher matrix numerically instead of using analytical expressions, as no such expressions exist for the scattering statistics yet. However, we did validate our method by applying it to moments and confirmed its consistency with the analytical method. There is another difference between moments and repeated modulus statistics. The repeated modulus statistics is similar to the central moments such as the variance: the higher-order ones are defined from sample statistics of the lower-order ones. As a result, the expectation of sample statistics may depend on sample size. However, this is in general not a worry for the scattering coefficients calculated from a field, because the analogue to ‘sample size’, the number of pixels, is usually huge.

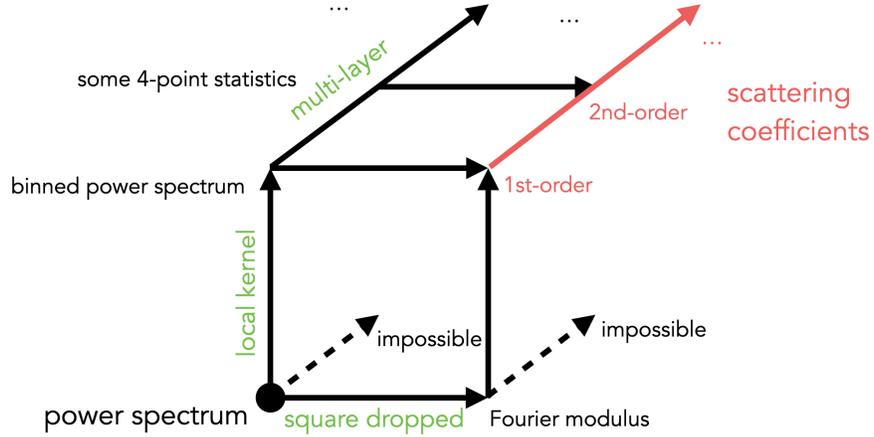


Figure 2-13. Relations between different translation-invariant descriptors, which are on the route of deriving the scattering transform from the power spectrum.

scattering transform and n -point functions to explain their similarities and differences. The relations among these translation-invariant descriptors are illustrated in Figure 2-13. We note that in neither real domain nor Fourier domain, these corresponding 2^n -point functions are equivalent to any ‘squeezed’ limit of higher-order functions. Instead, they are 2^n -point functions weighted by the wavelets ψ_1, ψ_2 , etc. used at each order.

Below I present the detailed derivation. For convenience, let’s call the coefficients obtained from squared modulus as ‘pseudo scattering coefficients’, and our task is to show that they can be written as N -point function. First, we note that the pseudo scattering coefficients at each order are defined from the previous order by

$$\begin{aligned}
 S_{n+1}^{\text{pseudo}} &\equiv \langle I_{n+1}^{\text{pseudo}} \rangle \\
 &\equiv \int |I_n \star \psi_{n+1}|^2 d\mathbf{x} \\
 &= \int |\tilde{I}_n \cdot \tilde{\psi}_{n+1}|^2 d\mathbf{k} \\
 &= \int P_n(\mathbf{k}) \cdot \tilde{\psi}_{n+1}(\mathbf{k})^2 d\mathbf{k}, \tag{2.28}
 \end{aligned}$$

where $P_n(k)$ is the power spectrum of I_n . This equation means that *each pseudo scattering coefficient is an averaged power spectrum of the previous-order pseudo scattering field*. The derivation makes use of three properties: (1) Fourier transform is ortho-normal; (2) a convolution in real space is a multiplication in Fourier space; and (3) wavelets ψ_{n+1} are real-valued band-pass filters in Fourier space.

Applying the above statement to the first order, we obtain that S_1^{pseudo} are averaged (binned) power spectrum (2-point function) of the input field, weighted by wavelets:

$$\begin{aligned}
 S_1^{\text{pseudo}} &= \int P_0(\mathbf{k}) \cdot \tilde{\psi}_1(\mathbf{k})^2 d\mathbf{k} \\
 &= \text{averaged (binned) power spectrum of } I_0. \tag{2.29}
 \end{aligned}$$

Applying it to the second order, S_2^{pseudo} are binned power spectrum of the intermediate scattering

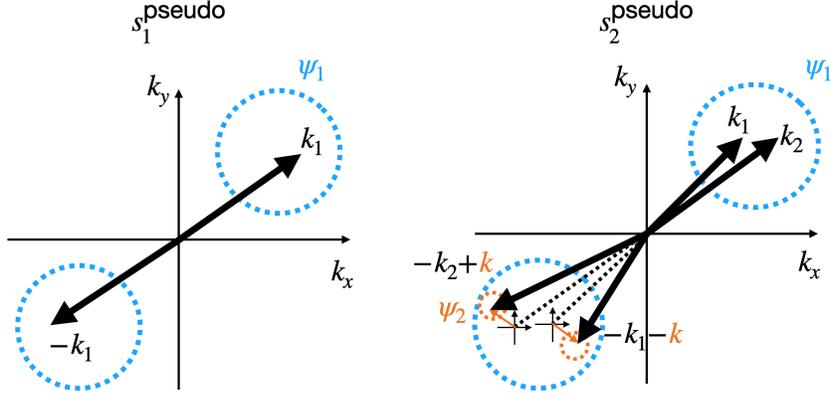


Figure 2-14. The corresponding 2-point and 4-point functions of scattering coefficients when replacing the modulus operation by modulus squared.

fields I_1^{pseudo} , whose Fourier transform is

$$\begin{aligned} \widetilde{I_1^{\text{pseudo}}}(\mathbf{k}) &= (I_0 \star \psi_1) (\widetilde{I_0 \star \psi_1})^* \\ &= \int \tilde{I}_0(\mathbf{k}') \tilde{I}_0(\mathbf{k}' + \mathbf{k})^* \cdot [\tilde{\psi}_1(\mathbf{k}' + \mathbf{k}) \tilde{\psi}_1(\mathbf{k}')] \cdot d\mathbf{k}' . \end{aligned} \quad (2.30)$$

Again, because wavelets are real-valued in Fourier space ($\tilde{\psi} = \tilde{\psi}^*$), we have

$$\begin{aligned} S_2^{\text{pseudo}} &= \int P_1 \cdot \tilde{\psi}_2^2 d\mathbf{k} \\ &= \int \widetilde{I_1^{\text{pseudo}}} \cdot \widetilde{I_1^{\text{pseudo}}}^* \cdot \tilde{\psi}_2^2 d\mathbf{k} \\ &= \iiint \tilde{I}_0(\mathbf{k}_1) \tilde{I}_0(\mathbf{k}_1 + \mathbf{k})^* \tilde{I}_0(-\mathbf{k}_2)^* \tilde{I}_0(-\mathbf{k}_2 + \mathbf{k}) \\ &\quad \cdot [\tilde{\psi}_1(\mathbf{k}_1) \tilde{\psi}_1(\mathbf{k}_1 + \mathbf{k}) \tilde{\psi}_1(-\mathbf{k}_2) \tilde{\psi}_1(-\mathbf{k}_2 + \mathbf{k}) \tilde{\psi}_2(\mathbf{k})^2] \\ &\quad \cdot d\mathbf{k}_1 d\mathbf{k}_2 d\mathbf{k} \\ &= \text{averaged tri-spectrum of } I_0 . \end{aligned} \quad (2.31)$$

The integrand is the product of four Fourier coefficients of the input field. Each coefficient is weighted by the first wavelet $\tilde{\psi}_1$, and their difference \mathbf{k} is weighted by the second wavelet $\tilde{\psi}_2$. An illustration of these 4-point configurations in Fourier space is shown in Figure 2-14.

Higher-order cases can be derived in a similar way, revealing that the n th-order pseudo scattering coefficients S_2^{pseudo} are exactly some averaged 2^n -point functions. Note that the ‘binning’ or averaging weights are defined by the wavelets $\psi_1, \psi_2, \dots, \psi_n$ used at each order, so they are not equivalent to any squeezed limit in either real or Fourier domain.

2.6.3 Scale interactions, cross-correlations and phase harmonic estimator

The scattering transform probes scale interactions within the frequency range of a single wavelet. There is a natural question regarding interactions of very different frequencies that are not located in the range of a single wavelet? Another related question is: in addition to two scales, can the

scattering transform describe interactions of two fields, in the flavor of the cross-spectrum and cross-correlation function?

In the original design, the scattering transform is only applicable to a single field. Nevertheless, it may be adapted to two-field applications, which is related to another recently proposed set of descriptors called the *phase harmonic* statistics [25].

Both the cross-spectrum $\Gamma(k) \equiv \langle \tilde{I}(k)\tilde{I}'(-k) \rangle$ and cross-correlation function $\gamma(\Delta x) \equiv \langle I(x)I'(x + \Delta x) \rangle$ (where I and I' are two fields to be studied) have the same limitation as the power spectrum $P(k) \equiv \langle \tilde{I}(k)\tilde{I}(-k) \rangle$: they treat scales separately and cannot probe scale interactions (because of the translation-invariance requirement). In the moment-based approach, one has to add more points to the multiplication (such as 3-point or 4-point statistics) in order to probe scale interactions.

The scattering transform represents a different approach for scale interaction: manipulating the phase in real space after a local-kernel convolution. In particular, it sets all phases to zero by taking the modulus, $|I \star \psi_k|$. This idea can also be used when cross-correlating two fields at different scales: $\langle |I \star \psi_{k_1}| \cdot |I' \star \psi_{k_2}| \rangle$, which avoids the requirement that $k_1 + k_2 = 0$ in cross-spectrum. Such technique of cross-correlating the wavelet modulus of different frequency bands has also been used to characterise complex textures [26].

The phase harmonics extend the way of manipulating phases. In addition to using modulus to set phases to zeros (in real space), one can also multiply the phases of k_1 -band convolution $I \star \psi_{k_1}$ by k_2/k_1 to match the phase changing rate in the k_2 -band convolution $I \star \psi_{k_2}$. This is another way to ‘cross-correlate’ fluctuations in different frequency bands, for the same field or for two different fields.

There are several differences between the scattering transform and phase harmonics. Most importantly, the number of scattering coefficients, especially after the reduction, is much smaller than that of phase harmonic statistics, which make the scattering transform more compact and convenient to use. In addition, the scattering coefficients are ‘first-order’ statistics, whereas the phase harmonics are second-order. Nevertheless, phase harmonics have the advantage that they can directly be used to explore the relation between two fields.

Chapter 3

Applications in cosmology

This section is largely based on three papers:

Cheng, Ting, Ménard, & Bruna, 2020, MNRAS, 499, 5902

Cheng & Ménard, 2021, arXiv:2103.09247 (accepted for publication in MNRAS)

Cheng et al., 2021 (in preparation).

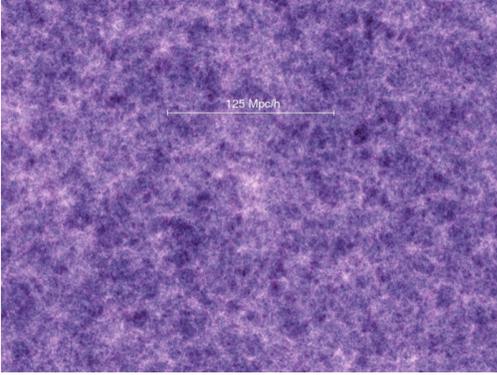
Non-Gaussian fields are ubiquitous in astrophysics and cosmology but challenging to analyze, as no generic tool such as the power spectrum exists as a sufficient statistic, and therefore the dimensionality of their description can be arbitrarily high. In addition, there is usually little guidance on which statistical estimator will be most appropriate for parameter inference. In this chapter, we explore the application of scattering transform for the analysis of the matter distribution in the Universe, a highly studied non-Gaussian field in cosmology.

3.1 Parameter inference in cosmology

Our imagination and speculation on the universe has a history as long as human civilization. However, it is in the past century that we finally found the physical laws that can describe the evolution of our the universe. After decades of development, cosmology has become a mature subject with precise models that can explain both the expansion of the universe as a whole and the formation of structures in it. Accordingly, observational cosmology has also entered an era of percent-level precision measurements and shifted its focus into precise parameter inference, using the statistical properties of cosmic density field to infer parameters about the initial condition, the expansion history, and the composition of our universe [27].

Observations of the cosmic microwave background show that our universe started from a small, hot, and uniform status with tiny density fluctuations. The cosmic density field at this ‘early-time’ can be well described by a Gaussian random field, for which the power spectrum is the only relevant statistic. Thus, power spectrum analysis has led to huge successes in cosmology [e.g., 28, 29]. As the universe evolves, the mean density decreases because of the expansion of space, but the fluctuations of density grow non-linearly due to local gravitational attraction, producing the structures we see in the current ‘late-time’ universe, such as halos, filaments, walls, and voids in the density field in small scales. As a result, cosmological information starts to elude Gaussian statistics such as the power spectrum and correlation function [30–33]. Figure 3-1 illustrates the formation of all these non-Gaussian structures in the Millenium Simulation [34]. In the past two decades, large surveys have revolutionize observational astronomy and cosmology by creating an explosion of homogeneous

$t = 0.21 \text{ Gyr}$



$t = 13.6 \text{ Gyr}$

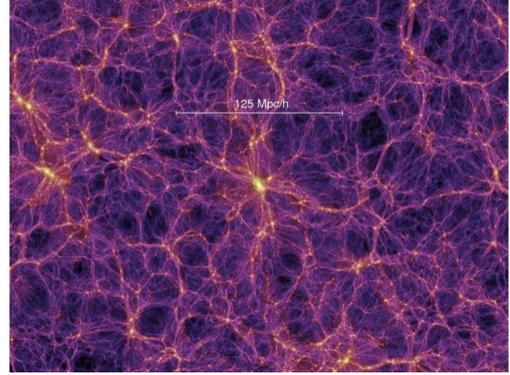


Figure 3-1. The evolution of cosmic density field in the universe in the Millennium Simulation. At early times (left panel), the density fluctuations are small and Gaussian. As the universe evolves, the fluctuations grow and becomes non-Gaussian at late times (right panel, the current universe).

observational data suitable for statistical analysis. Thanks to these large-scale, deep surveys, we are able to probe the structures in the cosmic density field from the observational side and compare them with simulations. Currently, several new survey instruments are also in construction, such as the Rubin Observatory, Euclid space mission, and Roman space telescope (Figure 3-2), with billions of dollars invested. Their unprecedented depth and area will allow the revelation of more non-Gaussian aspect of the cosmic density field. In this context, it is certainly desired that the statistical analysis of observational data can extract as much information as possible, i.e., a robust non-Gaussian statistic that suits for the properties of the ‘late-time’ cosmic structures and optimally extracts information is desirable.

In many areas of astrophysics and, in particular, in cosmology, extracting non-Gaussian information has been attempted through N -point correlation functions [e.g., 35–38, for weak lensing applications] and polyspectra, their Fourier equivalents [e.g., 39]. Correlation functions are convenient for theoretical predictions and for measuring weak departures from Gaussianity. However, being higher powers of the input field, these statistics suffer from an increasing variance and are not robust to outliers in real data, making them gradually less informative [40]. If the distribution of field intensity has a long tail, the amount of information accessible to N -point functions will quickly decrease [22]. In addition, the number of configurations to consider for N -point functions explodes with the number of points used. As a result, information is highly diluted among coefficients, and it becomes a challenge to efficiently extract information with N -point functions. Other methods, including performing a non-linear transformation before calculating correlation functions [41–44], using topological properties such as Minkowski functionals [45–48], and using biasing properties such as counts of clusters, peaks, and voids [49–54], have also been considered. However, in the cosmological context, these excursions into non-Gaussian signal analyses have had limited impact in improving existing constraints on cosmological parameters so far.

Recently, convolutional neural networks [CNNs, e.g., 55] have claimed supremacy in a wide variety of applications aimed at extracting information from complex data. They have also shown promises to efficiently retrieve cosmological information well beyond second-order statistics [e.g., 56–58]. While the potential of this method is enormous, it also comes with a number of issues. To precisely and robustly estimate cosmological parameters, CNNs require a large training set. In

addition, when applied to real data, systematic errors not included in the training process of CNN can hardly get checked and controlled, whereas for traditional statistics, a simple χ^2 test can do so. As such, the use of CNNs in real data comes with limitations regarding interpretability and validity.

In this chapter, we advocate using the scattering transform to efficiently and robustly extract statistical information from non-Gaussian cosmological fields. As discussed earlier in this thesis, the operations and structure of the scattering transform has close similarities with those built in CNNs, but the scattering transform does not require any training, and like traditional statistics, it generates coefficients with proved properties. It can therefore hopefully overcome the aforementioned limitations encountered with CNNs. I shall show that the scattering statistics outperform the power spectrum, bispectrum, and peak counts, and is on par with the state-of-the-art CNNs. I shall also show its application to the real weak lensing data from the Subaru HSC survey.

3.2 Weak lensing cosmology

The cosmic density field carries information about the evolution and composition of the universe. However, based on various observational evidence including the cosmic microwave background and galaxy kinematics, we now believe that most matter in our universe is ‘dark matter’, which are not composed of atoms (baryons) and generally does not emit photons. Fortunately, their gravity causes distortion of light emitted from background galaxies. This effect is weak for the general dark matter density fluctuation, thus is called weak lensing. Locally the distortion create two effects: magnification and shape distortion. Weak magnification is in general more difficult to measure, because the intrinsic luminosity galaxies ranges many orders of magnitude. Fortunately, the shape distortion effect, usually called cosmic shear, can be precisely measured by averaging galaxy shapes. The lensing distortion effect is related to the cosmic density field by mere gravitational theory. Therefore, cosmic shear is a neat probe of the cosmic density field, and weak lensing has become a main focus of modern cosmology. Currently, there are three large on-going weak lensing programs: the Kilo-Degree Survey (KiDS)¹ with the 2.6m VLT survey telescope in Chile, the Dark Energy Survey (DES)² with the 4-meter Victor M. Blanco telescope in Chile, and Hyper Suprime-Cam (HSC) survey³ with the 8-meter Subaru telescope in Hawaii. Going deeper and wider, the next generation surveys include the Vera C. Rubin Observatory (LSST)⁴, the *Euclid*⁵ space mission, and the Nancy Grace Roman Space Telescope (*WFIRST*)⁶. Weak lensing cosmology is one of their key science goals of those billion-dollar level instruments.

The lensing field is measured in the following way. From a deep galaxy survey, the shape (ellipticity) of background galaxies are carefully measured. Assuming that there is no preferential orientations in the Universe, these ellipticities statistically reflects the gravitational shear caused by the foreground matter. In practice, binning and averaging are usually carried out to reduce data from large samples of noisy galaxy shapes. One may bin the distance between galaxy pairs to calculate the 2-point correlation function, or one may bin the position of galaxies and create a shear γ map. The shear map is a spin-2 field which can be represented by complex numbers. In absence of artefacts such as masks, the shear map has a simple relation to a real-valued scalar field called

¹<http://kids.strw.leidenuniv.nl>

²<https://www.darkenergysurvey.org>

³<http://hsc.mtk.nao.ac.jp/ssp/>

⁴<https://www.lsst.org>

⁵<https://sci.esa.int/euclid>

⁶<https://roman.gsfc.nasa.gov>

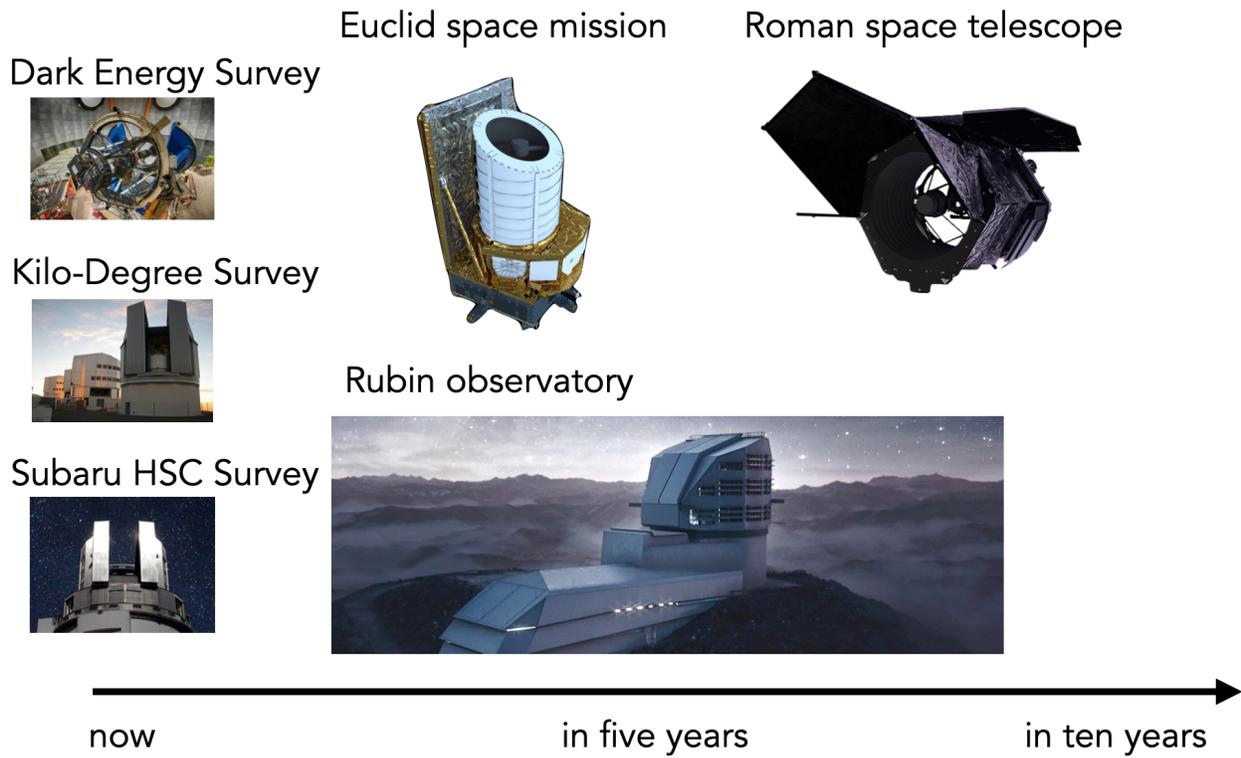


Figure 3-2. The on-going and planned major weak lensing surveys in the next ten years.



Figure 3-3. Left: The 4-meter telescope used for the Dark Energy Survey. Right: The CCD array contains 74 pieces of CCDs, each being similar to but much more sensitive than the detector used in digital cameras. The array can cover an area of about 14 moons on the sky.

convergence κ , which is a weighted projection of 3D matter density field along the line of sight.

Formally, the distortion effect of the background due to lensing of foreground material can be described locally by the 2D distortion matrix (the Jacobian matrix of the mapping from lensed to unlensed background image):

$$A_{ij} = \frac{\partial \beta^i}{\partial \theta^j} \equiv \begin{pmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_2 \end{pmatrix} \quad (3.1)$$

where β is the original un-lensed position of a background source, θ is the observed position of the background source, κ is the convergence, and (γ_1, γ_2) are the two components of shear. Intuitively, the convergence κ describes the magnification lensing effect, and the shear γ describes the shape deformation effect. However, there is only one independent scalar field in the lensing physics, therefore κ , γ_1 , and γ_2 can be unified through a lensing potential Φ . Their relations in Fourier space is particularly simple:

$$\tilde{\kappa}(\mathbf{l}) = \frac{-l^2}{2c^2} \tilde{\Phi}(\mathbf{l}) \quad (3.2)$$

$$\tilde{\gamma}_1(\mathbf{l}) = \frac{-l_x^2 + l_y^2}{2c^2} \tilde{\Phi}(\mathbf{l}) \quad (3.3)$$

$$\tilde{\gamma}_2(\mathbf{l}) = \frac{-l_x l_y}{c^2} \tilde{\Phi}(\mathbf{l}), \quad (3.4)$$

therefore measuring shears and convergence are in general equivalent. When κ and γ are much smaller than unity, the lensing potential can be well approximated as the gravitational potential integrated along the line of sight [59, 60]. Thus, the convergence for a single source can be written as an weighted integral of the matter density fluctuation $\delta(x, t) = \rho(x, t)/\bar{\rho}(t) - 1$ along the line of sight:

$$\kappa(\theta, \chi_s) = \int_0^{\chi_s} d\chi q(\chi, \chi_s) \delta(\chi, r(\chi)\theta), \quad (3.5)$$

where χ is the co-moving distance to the foreground lens, χ_s is the comoving distance to the source, $q(\chi, \chi_s)$ is called the lensing kernel, which peaks roughly in the midpoint between the source and observer, and $r(\chi)$ is the angular diameter distance of lens. The lensing kernel q can be expressed explicitly as:

$$q(\chi, \chi_s) = \frac{3\Omega_m H_0^2}{2c^2 a(\chi)} \frac{r(\chi)r(\chi_s - \chi)}{r(\chi_s)}, \quad (3.6)$$

where Ω_m is the ratio of current matter density to critical density $\frac{8\pi G\rho_0}{3H_0^2}$, H_0 is the current expansion rate of the universe $\frac{da(z)/dt}{a(z)}|_{z=0}$, $a(\chi)$ is the scale factor (the size) of the universe at the time when light travels and reaches the position of the lens at distance χ . More details about cosmology with gravitational lensing can be found in the reviews written by Kilbinger [61] and Mandelbaum [62].

To acquire an intuitive sense of lensing fields, we show in Figure 3-4 an example of simulated lensing convergence fields [63], together with the transformed fields after the scattering operations. The two lensing fields and a Gaussian random field all have the same power spectrum but differ in non-Gaussianity. It is clear that after two interactions of scattering operation, the non-Gaussianity is manifested in the mean (i.e., the 2nd scattering coefficients), which can be trivially extracted.

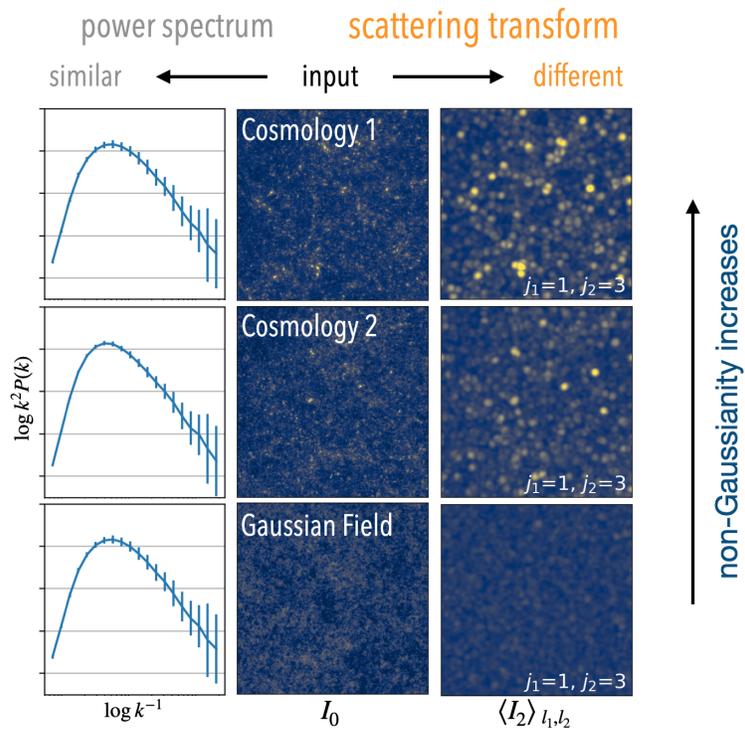


Figure 3-4. The scattering transform of three fields (I_0) with indistinguishable power spectra. Row 1 shows a realization of convergence maps in cosmology $(\Omega_m, \sigma_8) = (0.292, 0.835)$, row 2 shows cosmology $(\Omega_m, \sigma_8) = (0.566, 0.520)$, row 3 is for a Gaussian random field with the same (2D) power spectrum as row 1. It can be seen by eye that the average intensity of the 2nd-order scattering fields (the last column), which corresponds to an s_2 coefficient and measures the clustering strength of structures highlighted by I_1 , is significantly different from each other, while their power spectra (the first column) are indistinguishable.

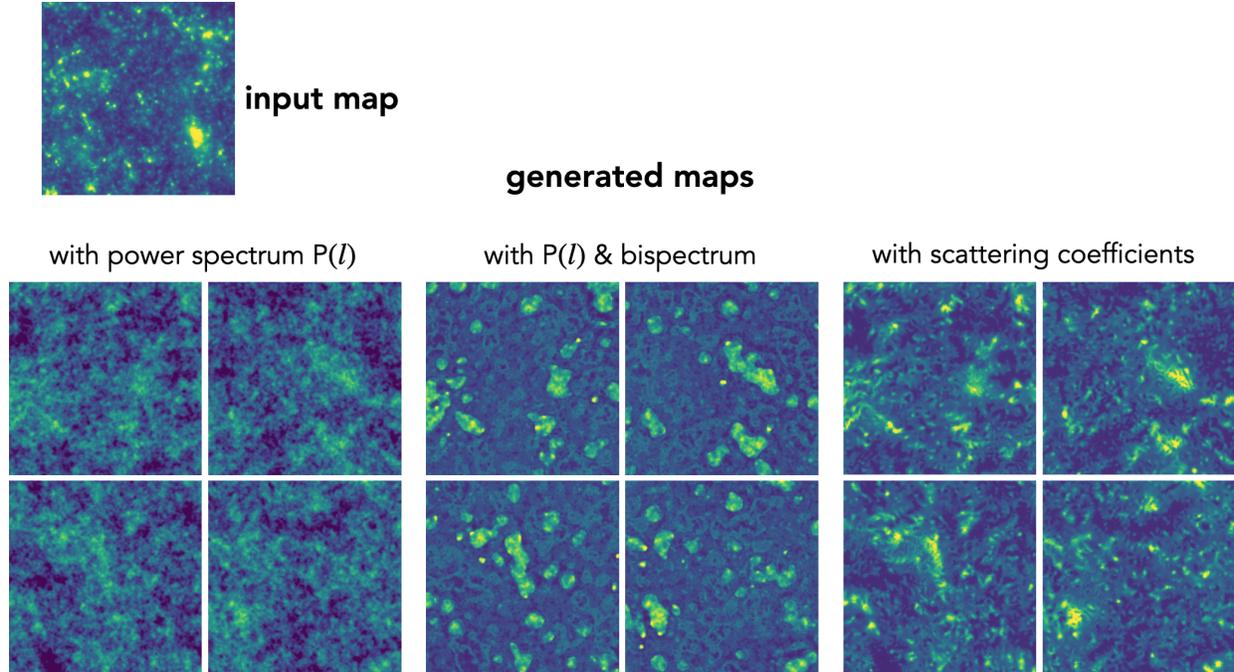


Figure 3-5. Images generated with selected sets of summary statistics, including the power spectrum, bispectrum, and scattering coefficients (S_1 , S_2), showing that the scattering coefficients characterize the field better. As the image generation is random, one should compare textures, instead of the exact positions of features, between the input and generated maps.

Before going to the details of cosmological inference, let’s first play again with the image generation technique introduced in Chapter 2.2 to see how well the scattering statistics perform on weak lensing convergence maps. Here we add two other statistics, the power spectrum and the bispectrum, for comparison. As a reminder, the idea of is to randomly generate new images (or, in other words, to sample from the ensemble of images) that have the same summary statistics as the target one, and then visually check the *texture* of the generated images [see, e.g., 17, 64].

To implement it, we start from a random image and modify it in order to minimize the difference of the summary statistics between the generated and the target images. Technically, we start from a Gaussian random field that has the same power spectrum as the target image, and then use the ‘adam’ optimizer in the python package `torch.optim` to minimize a loss function defined from the difference of summary statistics between the generated and target image. In the scattering coefficient and bispectrum cases, we also minimize the difference of their L_1 norms and set a lower bound for all pixels. Our code for image generation is available online⁷.

In Figure 3-5, we show the generated images using power spectrum P , bispectrum B , and the scattering coefficients S . For each set of statistics, four realizations are shown to illustrate the sampling variance. Compared to the power spectrum and bispectrum, the results from scattering coefficients look much more similar to a real lensing convergence map (target image), especially in the textures created by halos on all scales.

⁷https://github.com/SihaoCheng/scattering_transform

3.3 Forecast for clumpiness σ_8 and matter density Ω_m

We now show that the scattering transform can be a powerful tool in observational cosmology to extract non-Gaussian information from the matter density field. To illustrate this point, we consider an application with 2-dimensional fields: we show how well cosmological parameters can be constrained using the scattering coefficients of weak lensing convergence maps $\kappa(\vec{\theta})$ or, equivalently, measurements of cosmic shear. Being projections of the density field along the line-of-sight, these maps present an appreciable level of non-gaussianities on scales smaller than a few degrees, reflecting the non-linear growth of matter fluctuations.

We explore the use of our reduced scattering coefficients on simulated weak lensing convergence maps to infer Ω_m and σ_8 and compare their performance with that of the power spectrum. We also compare our results with that of a state-of-the-art CNNs by Ribli *et al.* [58] and peak count statistics.

3.3.1 Simulated lensing maps

We use mock convergence maps generated by the Columbia Lensing team⁸ and described in Zorrilla Matilla *et al.* [63] and Gupta *et al.* [56]. The maps are produced through ray-tracing to redshift $z = 1$ in the output of dark-matter-only N -body simulations for a set of Λ -CDM cosmologies. Each simulation is run in a $240 h^{-1}$ Mpc box with 512^3 particles. The cosmologies differ only in two parameters: the present matter density relative to the critical density Ω_m , and a normalization of the power spectrum σ_8 . Other cosmological parameters are fixed: baryon density $\Omega_b = 0.046$, Hubble constant $h = 0.72$, scalar spectral index $n_s = 0.96$, effective number of relativistic degrees of freedom $n_{\text{eff}} = 3.04$, and neutrino masses $m_\nu = 0.0$. The dark energy density is set so that the universe is spatially flat, i.e., $\Omega_\Lambda = 1 - \Omega_m$. For each cosmology, 512 convergence maps with 3.5×3.5 deg² field of view are generated from the simulations, allowing us to sample cosmic variance. The corresponding scales are well suited to probing the non-Gaussianities of the convergence field [61]. These maps were also used by Ribli *et al.* [58]. To compare our results to Ribli *et al.* [58], we use the same resolution as theirs, down-sampling the original 1024^2 pixel maps to a 512^2 resolution with 0.41 arcmin per pixel.

In practice, convergence or shear estimates are obtained from measurements of galaxy shapes, with a level of noise that depends on the galaxy ellipticity distribution and their number density on the sky. To first order, background galaxies used for shear measurements have a wide range of redshifts and are not correlated. The noise can be well approximated as Gaussian white noise. Its contribution to the convergence maps can be modelled [65] as

$$\sigma_{\text{noise}}^2 = \frac{\sigma_\epsilon^2}{2n_g A_{\text{pix}}}, \quad (3.7)$$

where σ_ϵ^2 is the intrinsic variance of ellipticity of galaxies, which is taken to be 0.4^2 , n_g is the number density of background galaxies, A_{pix} is the area per pixel, which is 0.1682 arcmin². For some on-going surveys such as KiDS, DES, and Subary HSC survey, n_g is 10 - 20 arcmin⁻² [66–68]; for some upcoming surveys we expect substantially higher densities: $n_g \sim 25$ arcmin⁻² for the survey at Vera C. Rubin Observatory (LSST), $n_g > 30$ arcmin⁻² for *Euclid*, and $n_g \sim 50$ - 75 arcmin⁻² for the planned survey with Nancy Grace Roman Space Telescope (*WFIRST*).

⁸<http://columbialensing.org>

After adding noise, we also smooth the maps. As the power of Gaussian white noise is distributed mostly at high frequencies, smoothing the convergence maps can help to increase the signal-to-noise of specific estimators. By default, we do not smooth the noiseless maps, and we perform a $\sigma = 1$ arcmin (2.44 pixel) Gaussian smoothing on noisy maps.

3.3.2 Statistical descriptors

Scattering coefficients: For each 3.5×3.5 deg² convergence field in each cosmology, we apply the scattering transform up to 2nd order using the ‘kymatio’ python package⁹ [69] and then calculate the reduced coefficients (s_0, s_1, s_2) as defined in Section 2.1.3. To probe the available range of scales, we set $J = 8$ and $L = 4$ in the scattering transform, i.e., we use 8 scales spaced logarithmically with central wavelengths between 1.2 arcmin and 75 arcmin and 4 azimuthal orientations, resulting in 32 different wavelets used in total.

By default, the ‘kymatio’ package only calculate the 2nd-order coefficients with $j_2 > j_1$, because the coefficients with $j_2 \leq j_1$ is mainly determined by the property of wavelets but not the input field, as illustrated by the upper-right sketch of Figure 3-6. Intuitively, this is because structures of a particular size, say j_1 , do not have meaningful clustering at scales smaller than their own size. A mathematical reasoning for this property can also be found in section 2.4.4. To demonstrate these coefficients’ behaviour, we modified the ‘kymatio’ code to calculate them, and show them together with the coefficients with $j_2 > j_1$ in Figure 3-6. Nevertheless, we checked that they do not contribute to constraining cosmological parameters, and therefore in our inference analysis, we only use 2nd-order coefficients with $j_2 > j_1$, which yields an even more compact set of $1 + 8 + 28 = 37$ scattering coefficients used for our cosmological inference.

Power spectrum: For the same set of input fields, we also compute the power spectrum and peak count statistics using the publicly available ‘LensTools’ python package¹⁰ [70]. The power spectrum is calculated within 20 bins in the range $100 \leq l \leq 37500$ (corresponding to 0.58–216 arcmin) with logarithmic spacing, following the setting adopted in Ribli *et al.* [58].

Peak count: In our analysis, a peak is defined as a pixel with higher convergence (κ) than its eight neighbors. Then, peaks are binned by their κ values and counted in each bin. We adopt a binning similar to that in Liu *et al.* [52]. We use 20 bins in total, including 18 bins linearly spaced between $\kappa = -0.02$ and 0.12, one bin for peaks below -0.02 , and one bin above 0.12. For reference, $\kappa = 0.12$ corresponds to a significance of peak $\nu \equiv \kappa/\sigma_{\text{noise}}$ around 7 when $n_g = 30$. Although using more bins for very high peaks ($\kappa > 0.12$) may enhance the constraining power of the peak count method, we do not use them in this study, because the count distribution of these rare peaks can no longer be approximated by Gaussian distribution [e.g., 71].

To obtain constraints on the cosmological parameters, we use the Fisher inference framework [72, 73], in which we assume the probability distribution of statistical descriptors is a multivariate Gaussian distribution for a given cosmology. The mean vector and covariance matrix of this Gaussian distribution are dependent on cosmological parameters and estimated from the 512 realizations of each cosmology in simulations. Details of our cosmological inference framework are described in Appendix 3.3.3. Because s_1, s_2 , and power spectra must be positive for a non-trivial field, we consider their logarithm to better satisfy a multivariate Gaussian likelihood. To perform the cosmological inference analysis with the three methods introduced above, we use

⁹<https://www.kymatio.io>

¹⁰<https://lenstools.readthedocs.io>

- 37 scattering coefficients
- 20 power spectrum coefficients
- 20 peak count coefficients.

3.3.3 Cosmological inference framework

Here we describe the Fisher forecast formalism used to infer the cosmological parameters in this study. According to the Cramér–Rao inequality, the variance of any unbiased estimator $\hat{\boldsymbol{\theta}}$ for model parameters $\boldsymbol{\theta}$ cannot be smaller than the inverse of the Fisher information matrix $\mathbf{I}(\boldsymbol{\theta})$ of the model:

$$\text{cov}(\hat{\boldsymbol{\theta}}) \geq \mathbf{I}(\boldsymbol{\theta})^{-1}. \quad (3.8)$$

Elements of the Fisher matrix is defined as

$$I_{m,n}(\boldsymbol{\theta}) \equiv \left\langle \frac{\partial \ln p(\mathbf{x}|\boldsymbol{\theta})}{\partial \theta_m} \frac{\partial \ln p(\mathbf{x}|\boldsymbol{\theta})}{\partial \theta_n} \right\rangle, \quad (3.9)$$

where \mathbf{x} is the observable, p is the likelihood function, $\langle \cdot \rangle$ is the expectation over \mathbf{x} . In our cosmological case, $\boldsymbol{\theta}$ represents cosmological parameters, $\boldsymbol{\theta} = (\Omega_m, \sigma_8)$, and \mathbf{x} represents the statistical descriptors such as the scattering coefficients. The function $p(\mathbf{x}|\boldsymbol{\theta})$ is called the likelihood of $\boldsymbol{\theta}$ when \mathbf{x} is fixed, and is called the probability density function (PDF) of \mathbf{x} when $\boldsymbol{\theta}$ is fixed.

In our study, we assume that given any cosmology $\boldsymbol{\theta}$, the PDF of statistical descriptors \mathbf{x} is Gaussian:

$$p(\mathbf{x}|\boldsymbol{\theta}) \propto \frac{1}{\sqrt{|\mathbf{C}|}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right], \quad (3.10)$$

where $\mathbf{C}(\boldsymbol{\theta})$ and $\boldsymbol{\mu}(\boldsymbol{\theta})$ are the mean and covariance matrix depending on the cosmological parameters $\boldsymbol{\theta}$. Thus, elements of the Fisher matrix can be written as

$$I_{m,n} = \frac{\partial \boldsymbol{\mu}^T}{\partial \theta_m} \mathbf{C}^{-1} \frac{\partial \boldsymbol{\mu}}{\partial \theta_n} + \frac{1}{2} \text{tr}(\mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial \theta_m} \mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial \theta_n}), \quad (3.11)$$

where the first and second items describe the information from cosmological dependence of $\boldsymbol{\mu}$ and \mathbf{C} , respectively. To obtain these items for arbitrary cosmology, we first calculate the sample mean and covariance matrix of the 512 realizations of each cosmology in the simulations (Section 3.3.1). The sample mean is an unbiased estimator of the real mean vector, but to unbiasedly estimate the inverse of covariance matrix, \mathbf{C}^{-1} , a correction factor is needed [74]:

$$\widehat{\mathbf{C}}^{-1} = \frac{N - D - 2}{N - 1} \widehat{\mathbf{C}}^{-1}, \quad (3.12)$$

where $\widehat{\mathbf{C}}^{-1}$ is the unbiased estimator in the inverse, N is the number of independent sample used for the estimation, D is the dimension of each data vector, and $\widehat{\mathbf{C}}$ is the sample covariance before Bessel’s correction. Then, with a further assumption that $\boldsymbol{\mu}$ and \mathbf{C} have smooth cosmological dependence, we use 3rd-order polynomials to fit for the cosmological dependence of $\boldsymbol{\mu}$ ’s elements and use 2nd-order polynomials for \mathbf{C} ’s elements.

3.3.4 Results

In this section, we examine the distribution and cosmological sensitivity of scattering coefficients, and present their constraining power for two cosmological parameters, Ω_m and σ_8 . We show that the scattering coefficients provide substantially more information than the power spectrum and is on a par with CNN.

Cosmological sensitivity of the scattering coefficients

In Figure 3-6, we present the distributions of reduce scattering transform in the noiseless case together with the power spectrum. In the first row, we show the values for a fiducial cosmology that has the Planck cosmology of $\Omega_m = 0.309$ and $\sigma_8 = 0.816$ [75]. The expected values of these descriptors are estimated by averaging over different realizations of a given cosmology. Error bars, which are the sample standard deviations of realizations, represent the cosmic variance in this noiseless case. We can see the similarity between the power spectrum and s_1 coefficients, as they have similar physical meanings. We can also see the different behaviours of s_2 coefficients for $j_2 < j_1$ and $j_2 > j_1$, as discussed in Section 3.3.2.

Then, we investigate the cosmological sensitivity of the power spectrum and scattering coefficients. The power spectrum is known to be mostly sensitive to one combination of the cosmological parameters, namely

$$\Sigma_8 \equiv \sigma_8 \left(\frac{\Omega_m}{0.3} \right)^a, \quad (3.13)$$

with a around 0.6 [e.g., 61], but can hardly distinguish cosmologies with the same Σ_8 , as illustrated in the upper panel of Figure 3-6. Breaking this degeneracy requires the extraction of non-Gaussian information from lensing maps.

In the second row of Figure 3-6, we show the response of coefficients as cosmological parameters move along (orange curves) and across (blue curves) the Σ_8 degeneracy. Gray areas indicate cosmic variance of the fiducial cosmology. As expected, the 1st-order scattering coefficients show a cosmological sensitivity similar to that of the power spectrum, because both of them measure the strength of fluctuations as a function of scale.

The 2nd-order scattering coefficients, on the other hand, characterize the spatial distribution of sized fluctuations. To make the 2nd-order scattering coefficients less correlated with the 1st-order ones, here we present de-correlated 2nd-order coefficients s_2/s_1 , as each $s_2(j_1, j_2)$ is proportional to the corresponding $s_1(j_1)$ according to their definitions [6]. These s_2/s_1 exhibit particularly high sensitivity to cosmological change along the Σ_8 degeneracy. In addition, they are indifferent to the other direction of cosmological change, which means they provide a piece of information roughly orthogonal to that carried by the 1st-order coefficients s_1 or the power spectrum. In noisy cases, though the information from s_2/s_1 is not orthogonal to s_1 anymore, we have checked that they still provide substantial sensitivity along the Σ_8 degeneracy. Due to this additional sensitivity, the scattering transform can be used to better constrain cosmological parameters than the power spectrum.

Constraining cosmological parameters

We now present the cosmological constraints set by the scattering coefficients measured from a single $3.5 \times 3.5 \text{ deg}^2$ field. For reference, we note that LSST will generate about 2,000 times more

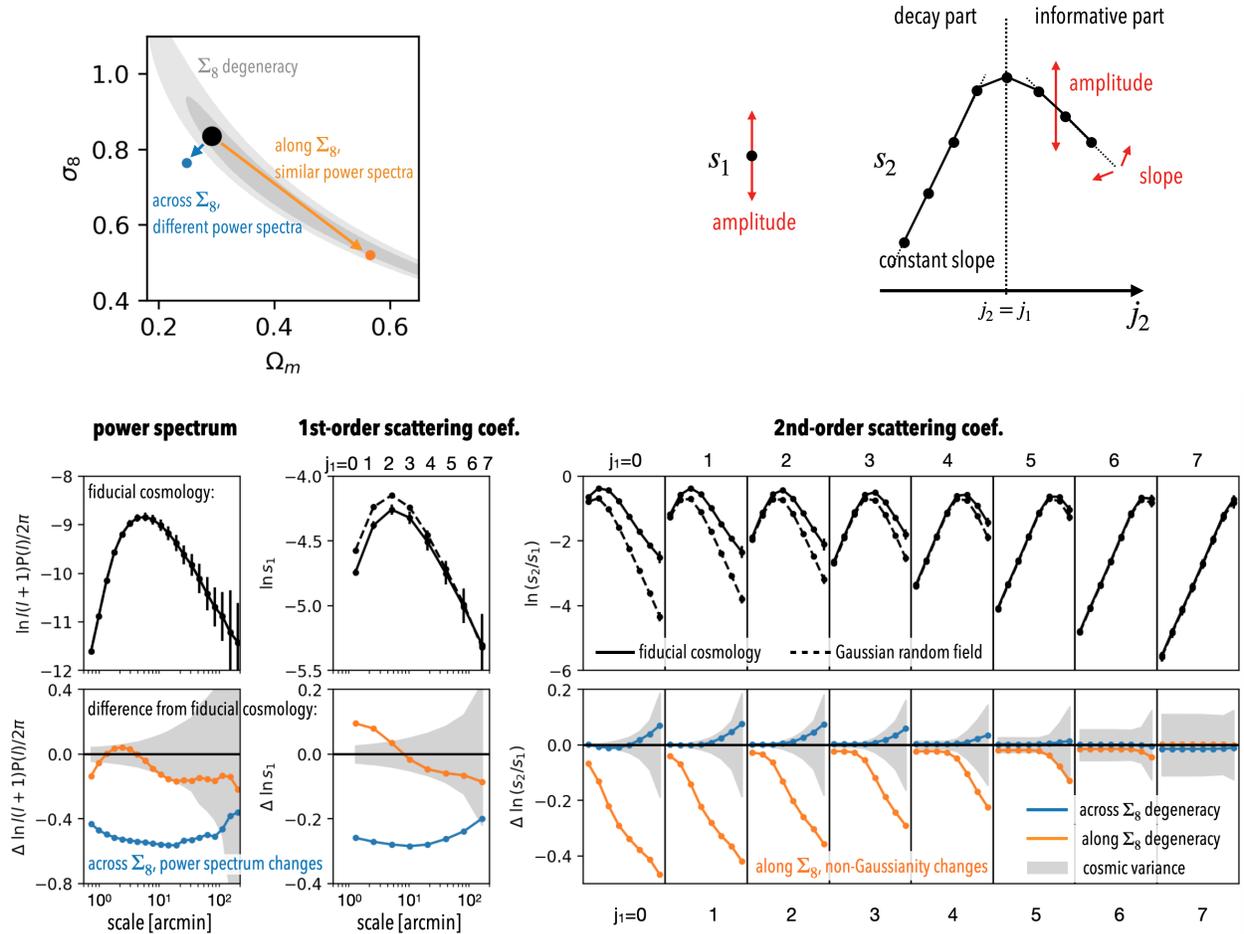


Figure 3-6. *Upper-left Panel:* The fiducial cosmology (black) and two other cosmologies on the (Ω_m, σ_8) plane. *Upper-right Panel:* Illustration of reduced scattering coefficients $s_1(j_1)$ and $s_2(j_1, j_2)$ for a single j_1 scale. *Lower Panel:* The power spectrum and scattering coefficients for the three cosmologies in noiseless case. The first row presents coefficients of the fiducial cosmology and of Gaussian random fields with the same power spectrum, and the second row shows changes of coefficients (Δ coef.) when we move from the fiducial cosmology to the other two. Error bars and gray shaded regions show cosmic variance, i.e., the variability among realizations. The 1st-order scattering coefficients behave similarly to the power spectrum, while the 2nd-order scattering coefficients can break the Σ_8 degeneracy, along which non-Gaussianity of weak lensing field changes.

Table 3-I. Comparison of the constraining power for (Ω_m, σ_8) between different methods, with a single $3.5 \times 3.5 \text{ deg}^2$ convergence map. The figure of merit is defined as the reciprocal of the 1σ confident area based on Fisher matrix (or the 68% posterior contour, in parentheses) on the (Ω_m, σ_8) plane. The convergence maps are smoothed with $\sigma = 1'$ Gaussian filter except for the case shown in the last column with no smoothing. The galaxy density n_g is in unit of arcmin^{-2} .

Methods	Ω_m - σ_8 Figure of Merit				
	$n_g = 10$	30	100	noiseless	no smoothing
$s_0 + s_1 + s_2$	50	140	329	1053	3367
$s_0 + s_1$	21	55	133	492	565
$s_1 + s_2$	39	91	181	446	1720
power spectrum $P(l)$	20	40	67	104	253
peak count	30	89	162	170	667
CNN [58]	(44)	(121)	(292)	(1201)	(-)

data, leading to constraints about 40 times tighter than the numbers presented below. In this study, we only probe the constraints on Ω_m and σ_8 and leave the work of using scattering coefficients to constrain the dark energy equation of state parameter w or neutrino mass M_ν to future study. Cosmological inference is just another aspect of the cosmological sensitivity problem examined in the previous subsection. The Fisher inference formalism we use in this study is described in section 3.3.3.

We first present results in the noiseless case. In Figure 3-7 we demonstrate the 1σ Fisher forecast of Ω_m and σ_8 using all scattering coefficients (red ellipse) and power spectrum (gray ellipse). The scattering coefficients provide a dramatically tighter constraint than the power spectrum. We also show a break-down of this constraining power into contributions from 1st-order (blue ellipse) and 2nd-order (orange ellipse) coefficients alone. As expected, the 1st-order coefficients (s_1) and power spectrum set similar constraints. The slight difference of ellipse orientation originates from the difference between the L^1 and L^2 norms used by the scattering transform and the power spectrum. The de-correlated 2nd-order scattering coefficients (s_2/s_1) provide a strong constraint along the Σ_8 degeneracy, consistent with our cosmological sensitivity discussion in Section 3.3.4.

The 0th-order coefficient s_0 is the mean of the $3.5 \times 3.5 \text{ deg}^2$ field. While its expectation value over the sky is zero, it does carry relevant information on those scales by capturing larger-scale modulations of the convergence field. We also note that it has strong correlations with other scattering coefficients (and the power spectrum), which is a sign of being in the non-linear regime of cosmology [e.g., 76]. Therefore, although the expected value of s_0 is identically zero in all cosmology, combining s_0 with other coefficients helps to substantially tighten the constraints on cosmological parameters. However, this piece of information may not scale as fast with the increasing field of view as the small-scale information, because in real data each patch of $3.5 \times 3.5 \text{ deg}^2$ fields on the sky are not independent. The mass sheet degeneracy [e.g., 77] is another problem for using s_0 , though the s_0 of small patches may be obtained by inheriting the zero-point solution of the whole survey. We find that including s_0 only improves the constraint of Σ_8 , consistent with the understanding that it is a leakage of larger-scale fluctuation. Similar improvement is also found when combining s_0 with the power spectrum.

To be more quantitative, we compare different methods using the reciprocal of the area of their 1σ Fisher forecast ellipses on the (Ω_m, σ_8) plane as the figure of merit (FoM). In the noiseless case,

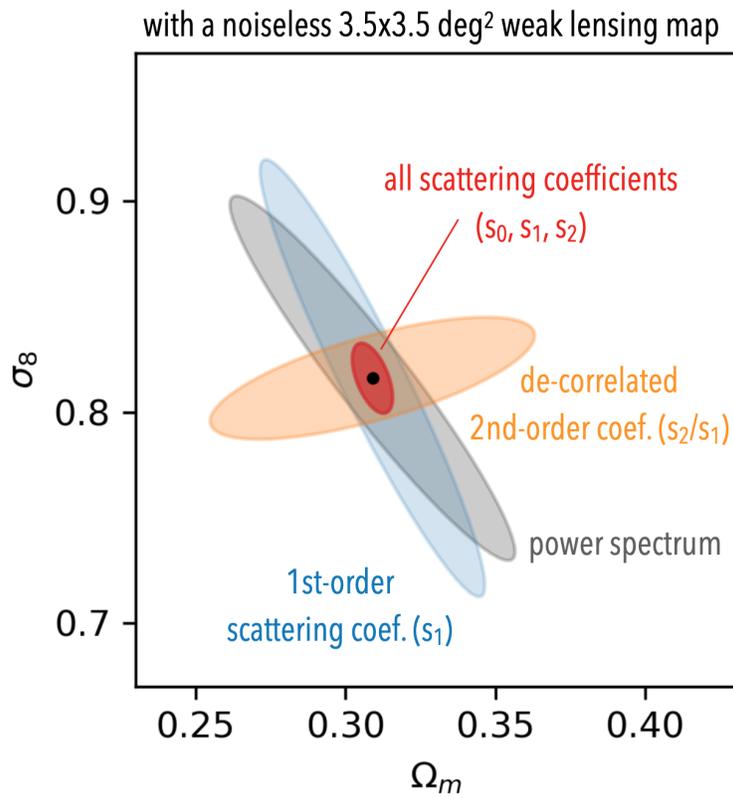


Figure 3-7. The 1σ Fisher forecast of cosmological parameters from a $3.5 \times 3.5 \text{ deg}^2$ noiseless convergence map with 0.41 arcmin per pixel resolution. The de-correlated 2nd-order scattering coefficients s_2/s_1 provide critical information to break the Σ_8 degeneracy along which the power spectrum cannot distinguish, therefore drastically improve the constraint.

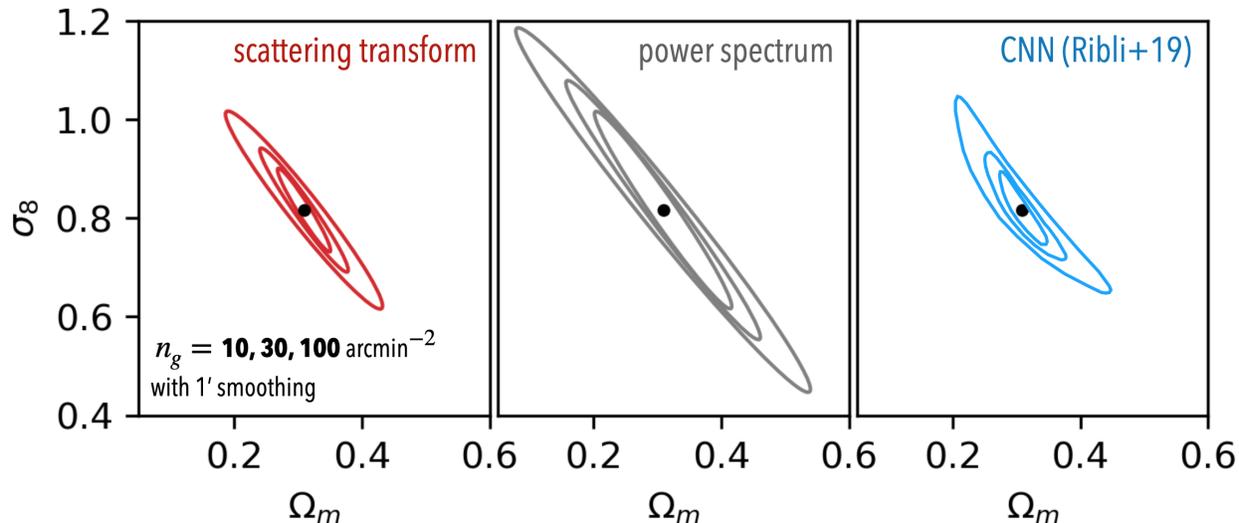


Figure 3-8. The 1σ Fisher forecast of cosmological parameters (Ω_m and σ_8) from different descriptors of a $3.5 \times 3.5 \text{ deg}^2$ convergence map smoothed with $\sigma = 1'$ Gaussian filter. The scattering coefficients have comparable performance as a state-of-the-art CNN [58] at all noise levels, and 3–5 times better than the power spectrum depending on the noise level.

combining all scattering coefficients (s_0, s_1, s_2) leads to a constraint that is 14 times tighter than that of the power spectrum, 5 times tighter than peak count statistics, and 3.3 times tighter than the joint constraint from power spectrum and peak count.

We then compare the performance of the scattering transform to a state-of-the-art CNN analysis by Ribli *et al.* [58]. To perform a meaningful comparison, we follow Ribli *et al.* [58] to use noiseless convergence maps smoothed with a $\sigma = 1$ arcmin Gaussian filter. Interestingly, we find that the scattering coefficients extract a similar amount of cosmological information to the CNN trained in Ribli *et al.* [58]. The corresponding figures of merit are shown in Table 3-I.¹¹

We now consider convergence fields in the presence of galaxy shape noise. As the noise level increases, small-scale structures, which carry plenty of cosmological information, get erased. As a result, the constraining power of the scattering coefficients (as well as other methods) degrades. In Figure 3-8 we show the Fisher forecast of Ω_m and σ_8 from a $3.5 \times 3.5 \text{ deg}^2$ convergence map under three noise levels, using the scattering coefficients and the power spectrum. We also show the posterior constraints from CNNs trained by [58] on the same simulations. The figures of merit for these methods, together with the peak count method, are listed in Table 3-I. Again, we find that the scattering transform not only outperforms the power spectrum and peak count, but also provides cosmological constraints on a par with state-of-the-art CNNs.

To summarise, we have demonstrated the power of the scattering transform for cosmological parameter inference with weak lensing data. For simplicity, we focused on the convergence field but a similar analysis can also be performed on the shear field. In Figure 3-9, we present quantitative comparisons between the four techniques discussed in our study. It shows the high performance of the scattering transform over a wide range of noise levels. We therefore advocate using this new estimator in the analysis of existing and upcoming weak lensing surveys, in observational cosmology,

¹¹We note that Ribli *et al.* [58] do include the field mean information in their CNN training. So, a fair comparison would be $s_1 + s_2$ versus power spectrum, and $s_0 + s_1 + s_2$ versus CNN.

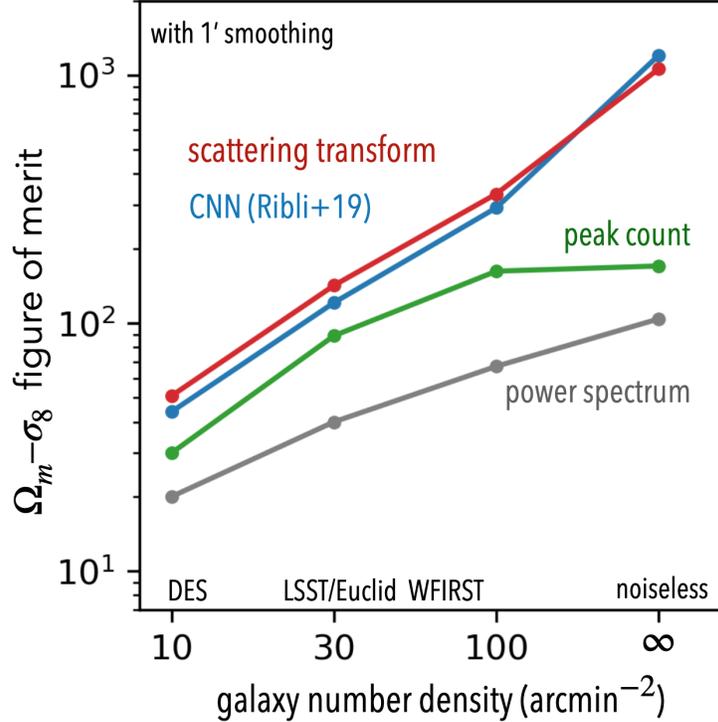


Figure 3-9. Dependence of the (Ω_m, σ_8) constraints with different methods on galaxy shape noise. The figure of merit (FoM) is defined as the 1σ confident area on the (Ω_m, σ_8) plane. Note that the CNN result [58] is reported in terms of posterior, while others are Fisher forecast. For noisy cases the difference between scattering transform and CNN is not intrinsic but due to the difference between posterior and Fisher forecast. Peak count’s performance does not increase as fast because it is more sensitive to smoothing scale than the other methods.

and more generally, in the analysis of stochastic fields encountered in physics.

3.3.5 Discussion

Inference for non-Gaussian fields

In physics, many inference problems concern estimating physical parameters from realizations of random fields. Ideally, one would like to use the likelihood function of the field itself, but this is often out of reach except for several simple cases such as some Gaussian random fields. Therefore, for the inference problem to be feasible, a statistical representation of the data is often used. Statistical descriptors reduce the dimensionality of the data vector and they tend to Gaussianize according to the central limit theorem. Both of these properties help to regularize the likelihood. However, it is still challenging to find a proper representation because in general a random field can be random in too many different ways. In these cases, a useful characterization must be one that makes use of known properties of the field.

Viewed in this direction, traditional statistical approaches with their own representation framework may or may not suit the properties of particular fields. For example, the peak count statistic used in weak lensing cosmology suits the presence of distinct haloes in convergence maps. N -point

functions, closely related to perturbation theory and convenient for analytical prediction, represent the field with a series expansion, which makes them good descriptors for fields slightly deviating from a Gaussian one. A highly non-Gaussian field, however, requires using larger N . As the number of coefficients and the complexity of configurations increase rapidly with N , N -point functions quickly become an inefficient and non-robust representation of the input field. On the other hand, CNNs try to learn the field properties and search for informative representation through a training optimization.

Fortunately, the non-Gaussian fields that originate from physical interactions do often have common properties. Such fields typically display localized, coherent structures in multi-scales, and smaller structures often act as building blocks of larger structures. These properties can be used as the ‘domain knowledge’ to guide our design or choice of the statistical representation in a general sense. As we will explain in the next section, the design and operations of the scattering transform leads to an efficient and robust representation for such fields, because they are tailored for these properties.

Attractive properties of the scattering transform

Efficiency: All the three elements (wavelet convolution, modulus, and the hierarchical design) play essential roles to make the scattering transform efficient. The use of wavelets balances the resolution in real and frequency domain. As a result, the scattering transform can capture localized information from a large range of scales with only a few coefficients, at each order. After selecting structures of scale j_n in one order, the scattering transform then selects structures ‘assembled’ by these j_n -scale structures in the next order. This hierarchical design allows the n th order scattering coefficients to quickly access configurations described by about 2^n points. Moreover, the ‘low-order’ non-linear operator, modulus, helps to collect information even beyond the access of 2^n point functions.

These strategies concentrate relevant information to a limited set of statistical descriptors, which is desirable in terms of compactness of the representation and the signal-to-noise ratio of each estimator. For example, in our case, the scattering transform compresses weak lensing information into 37 coefficients, a number that is much smaller than typical bi-spectrum descriptors, while achieving CNN-like constraint on cosmological parameters.

Robustness: All scattering coefficients are ‘first-order’ statistics in the sense that they are proportional to the input field, and it is proved that the scattering representation is non-expansive, i.e. the distance between two vectors in the scattering representation never exceeds their distance in the original pixel-based representation [1, 3]. Therefore, it does not amplify the process variability. This is in contrast to the N -point correlation function approach, which requires multiplying an increasing number of field fluctuations and causes high variability. As a result, the scattering coefficients are low variance descriptors and insensitive to outliers.

The locality of wavelets, which is related to their logarithmic spacing and widths in frequency space, also introduces stability to deformations [1], which is a desired property of robust descriptors that classical N -point functions do not have.

Interpretability: the scattering coefficients have a simple and intuitive interpretation. They describe clustering properties of the field in the following way:

The 1st-order scattering coefficients are similar to a coarsely binned power spectrum, which characterize the clustering strength at different scales j_1 . As the scattering transform uses an L^1 norm as opposed to an L^2 norm, the ratio between s_1 coefficients and the power spectrum provides

a measure of sparsity of the field. This explains why in Figure 3-7 the constraints from 1st-order coefficients and the power spectrum are slightly different, and just combining these two can also provide a stronger constraint on cosmology than using power spectrum alone.

The 2nd-order scattering coefficients characterize the clustering strength of j_1 -scale structures separated by j_2 -scales. In other words, these coefficients characterize the clustering of structures selected over a given frequency range, or the ‘clustering of clustering’. Their departure from their Gaussian counterparts is a robust measure of the strength of non-Gaussianities. The n th-order scattering coefficients, though not shown explicitly in this study, can in turn be understood as the strengths of n th-order hierarchy of clustering of the field at all different combinations of scales.

Comparison to CNNs

The scattering transform and CNNs share a number of properties. Both of them have hierarchical layers with localized convolution kernels and use a simple non-expansive non-linear operation. Although CNNs are usually trained to directly map a field to physical parameters, their inside can be considered as composed of a convolutional part that extracts spatial features and a second part that learns the mapping from these features to physical parameters. Both parts are trainable and trained together. The scattering transform, on the other hand, uses preset wavelets as convolutional kernels and just a few layers (in our case two layers). So it can be viewed as a non-trainable mini-CNN playing the role of the first part of trainable CNNs. In the scattering transform’s approach, the second part of trainable CNNs is supplanted by using traditional regression techniques.

The trainable kernels make CNNs more flexible and may lead to a higher performance for finer classification problems such as classifying different types of rabbits, but in the mean time this over-parametrization defines a much more brittle statistical model [17, 78]. Our results imply that compared to CNNs, the scattering transform has enough expressiveness to characterize the matter density field in the cosmological context while holding provable stability properties. Indeed, as shown by Ribli *et al.* [58], a CNN trained on convergence maps internally generates kernels similar to (azimuthally averaged) Morlet wavelets. Our results also imply that much of the power of CNNs may be detached from its trainable nature.

Over-parametrized models tend to over-fit, i.e., to ‘remember’ single realizations instead of comprehending the overall property of the whole training set. Thus the over-parametrized CNNs require a large number of simulations as training set to alleviate the over-fitting problem. In contrast, the scattering transform uses preset kernels, thus has no parametrization in the kernels. In addition, the choice of CNN architecture can modify the results substantially, as can be seen in the comparison between results of Ribli *et al.* [58] and Gupta *et al.* [56]. As such, CNNs usually require much, and often ad-hoc, fine-tuning. The scattering transform, on the other hand, is not subject to these sources of variability. It requires the use of simulations only to probe the cosmic variance of the descriptors. Without learning the kernels, the scattering transform also significantly save calculation time.

Another view on the over-fitting problem is given by the framework of maximum-entropy regularized estimation, which looks for the most ‘non-committal’ statistical model under the constraints of a ‘feature vector’ of sufficient statistics [79]. There is thus a tension in the design of such vector of sufficient statistics [17]: On the one hand, the features should be descriptive enough so that they introduce enough constraints, i.e., typical samples from the estimated model should also be typical in the true distribution; On the other hand, one would like the features to be efficiently

estimated from the available samples, so that the corresponding statistical model is robust under resampling. In other words, typical samples from the true distribution should remain typical under the estimated statistical model.

Finally, when applied to observational data, the scattering transform holds another advantage over CNNs, namely the possibility to investigate systematic effects. As traditional statistics, the scattering coefficients can be used to derive not only the best-fitting cosmological parameters, but also an evaluation of the goodness of fit and therefore a sanity check of the result. In contrast, although the internal machinery of CNNs can be roughly divided into a feature extraction part and a regression one, the CNNs are trained as a whole to learn a direct mapping from the data to the physical parameters. Due to the over-parametrization nature, outputs from intermediate layers (i.e. the intermediate abstraction of CNN) do not typically have good statistical properties. Therefore, when using CNN, it is challenging to check for systematic error in real data.

Relation to peak count method

The non-linear gravitational evolution of density fluctuations in the universe gives rise to haloes, which are virialized systems locally bound by gravity. As highlighted by Ribli *et al.* [58] in their Figure 10, a substantial amount of non-Gaussian cosmological information can be extracted from these features. The peak count method directly captures information in the abundance of haloes. However, it does not characterize the spatial information including profiles and positions of these haloes, which is also sensitive to cosmological parameters. The scattering transform implicitly extracts a comprehensive information of the abundance, profile, and distribution of haloes by first highlighting structures of particular scales and then characterizing their clustering at other scales. In the limit of small j_1 and large j_2 , the 2nd-order scattering coefficients can be understood as a measure of the ‘two-halo term’ in the halo model at scale j_2 , weighted by the halo response to the first wavelet with scale j_1 . This response is related to halo profiles. In general, the scattering transform provides a non-parametric description of the one-halo, two-halo, and transitional regime where haloes overlap and form larger haloes.

3.3.6 Conclusions

Characterizing arbitrary non-Gaussian fields is challenging as the dimensionality of their description can be arbitrarily high. The subset of fields relevant in physics, however, tends to be more constrained as they typically display localized, coherent structures. In the cosmological context, the matter density field presents another characteristic property, namely hierarchical clustering. An efficient statistical descriptor of the cosmological density field would ideally make use of these properties.

In this paper, we advocate the use of the scattering transform [1, 3], which generates statistics designed to extract information from complex fields with provable stability properties. It involves operations similar to those found in convolutional neural networks (CNNs): it uses wavelet convolution, which is particularly suitable for characterizing localized structures; it uses modulus as the non-linear operation; and it iterates these operations. However, in contrast to CNNs, the scattering transform does *not* require training. It generates a compact set of robust coefficients, which forms a representation of the input field and can be used as efficient summary statistics for non-Gaussian information.

We applied the scattering transform to a parameter inference problem in the context of weak lensing cosmology. For simplicity, we focused on the convergence field but a similar analysis can

also be performed on the shear field. We used simulated convergence maps generated by ray-tracing N -body simulation results [56, 63] and measured their scattering coefficients to infer the cosmological parameters Ω_m and σ_8 . On maps with and without galaxy shape noise, the scattering transform outperforms the power spectrum and peak counts, and is on par with state-of-the-art CNNs.

As described in section 3.3.5, the scattering transform possesses a series of attractive properties for parameter estimation. It is efficient, robust, and interpretable. Obtained by iteratively applying wavelet convolution and modulus and finally taking the expectation value, the scattering coefficients can be interpreted as the strength of a hierarchy of clustering at various combinations of scales. Different from N -point functions, all scattering coefficients have the welcome property that they remain proportional to the input field, thus avoid instability problems and extract much more information when the field distribution has a long tail. Similar to classic statistical estimators, the scattering transform requires no training or tuning and offers the possibility to investigate systematic errors potentially present with real data.

In this chapter we demonstrated applications of the scattering transform in weak lensing data. Using it with existing and upcoming surveys (e.g. DES, LSST, *Euclid*, *WFIRST*) can be of great interest to improve constraints and provide consistency checks. Based on its properties and design, the scattering transform can also be an attractive approach for many other applications: in observational cosmology, astrophysics, and beyond.

3.4 Forecast for dark energy parameters and neutrino mass M_ν

In the last section, we demonstrated the power of the scattering transform by applying it to simulated weak lensing maps to constrain the Ω_m and σ_8 cosmological parameters. In this section, we extend our analysis to other cosmological dependencies: the dark energy equation of state w_0 , w_a and the neutrino mass sum M_ν . Constraining these parameters is one of the key goals of upcoming weak lensing surveys and they are often used in the figure of merits to compare different survey strategies.

To visualize the power of scattering coefficients, we will first show that one can use them to generate random images with textures very similar to real lensing maps, to a level not achievable by traditional moment-based statistics such as the power spectrum and bispectrum. Then, we will present the constraining power of scattering coefficients for dark energy parameters and neutrino mass.

To explore the constraining power of the scattering coefficients on cosmological parameters, we calculate the Fisher forecast [72, 73, 80] and Bayesian posterior of the cosmological parameters, after building a likelihood emulator for the scattering coefficients using simulated lensing convergence maps.

3.4.1 Simulated lensing maps

To explore the effects of dark energy and neutrino mass on lensing scattering coefficients, we use two sets of simulated convergence maps, which were generously made available by the Columbia lensing team¹². Both datasets were designed for probing the non-Gaussian information in weak lensing cosmology. The convergence maps were produced through ray-tracing N -body simulations to certain source redshifts without the Born approximation, using the `LensTools` python package

¹²<http://columbialensing.org>

[70]. All the cosmologies are spatially flat. The first dataset, described in Zorrilla Matilla *et al.* [63], Petri [70], and Liu *et al.* [81], varies dark energy properties and matter density (w_0 , w_a , and Ω_m). The convergence is traced to one source redshift ($z = 2$). The second dataset, **MassiveNuS** [82], varies neutrino mass, matter density, and primordial fluctuation amplitude (M_ν , Ω_m , and A_s , equivalent to M_ν , Ω_m , and σ_8), and the convergence is traced to five source redshifts ($z = 0.5, 1, 1.5, 2, 2.5$). Both datasets contain multiple realizations of convergence maps in multiple cosmologies, with 3.5×3.5 deg² field of view (where non-linear gravitational evolution becomes important) and 512×512 pixel resolution. Thus these maps cover multipole moments l from 100 to roughly 30,000. Below we describe these datasets in more detail.

Dark Energy simulations: This dataset contains simulations of 7 $w_0 w_a$ CDM cosmologies with different dark energy properties, including the present dark energy density Ω_Λ and the dark energy equation of state index $w(a)$, parametrized by w_0 and w_a through $w(a) = w_0 + (1 - a)w_a$ [83]. Each simulation was run in a 240 Mpc/ h box with 512^3 particles and used to generate 1,024 mock convergence maps with 3.5×3.5 deg² field of view. The lensing sources were set at redshift $z = 2$. We down-sampled the original 2048^2 pixel maps to a 512^2 resolution with 0.41 arcmin per pixel, by averaging adjacent pixels. Besides the dark energy properties, other cosmological parameters are fixed: baryon density $\Omega_b = 0.046$, Hubble constant $h = 0.72$, scalar spectral index $n_s = 0.96$, normalization of fluctuation amplitude $\sigma_8 = 0.8$, effective number of relativistic degrees of freedom $n_{\text{eff}} = 3.04$, total neutrino masses $M_\nu = 0.0$, and temperature of the cosmic microwave background $T_{\text{CMB}} = 2.725$ K. The dark matter density is set so that the universe is spatially flat, i.e., $\Omega_m = 1 - \Omega_\Lambda$. Pipelines for the N -body simulations and ray-tracing are described in Petri [70], Zorrilla Matilla *et al.* [63], and Liu *et al.* [81]. Specifications of the simulations are also listed on the website of the Columbia Lensing team.

Neutrino mass simulations (MassiveNuS): Generated by Liu *et al.* [82], this dataset contains convergence maps for several source redshifts created from 101 $\Lambda\nu$ CDM simulations with different M_ν , Ω_m , and A_s (equivalent to M_ν , Ω_m , and σ_8). Other cosmological parameters are fixed: $\Omega_b = 0.046$, $h = 0.70$, $n_s = 0.97$, $w = -1$. Each simulation was run in a 512 Mpc/ h box with 1024^3 particles and used to generate 10,000 convergence maps with 3.5×3.5 deg² field of view and 512^2 -pixel resolution. The fiducial cosmology has parameters $(M_\nu, \Omega_m, A_s) = (0.1 \text{ eV}, 0.3, 2.1 \times 10^{-9})$. One cosmology with massless neutrinos $(M_\nu, \Omega_m, A_s) = (0 \text{ eV}, 0.3, 2.1 \times 10^{-9})$ is used to calculate the covariance matrix of summary statistics. In this study, we will use the convergence maps for 5 source redshifts ($z = 0.5, 1.0, 1.5, 2.0, \text{ and } 2.5$), in 28 cosmologies that are close to the fiducial cosmology, i.e. in the range of $A_s \in [1.8, 2.7] \times 10^{-9}$, $M_\nu \in [0.06, 0.60] \text{ eV}$, $\Omega_m \in [0.28, 0.32]$. This range of parameters is enough for estimating the scattering coefficients' constraining power for cosmological parameters in a Rubin-observatory-like survey.

In real lensing surveys, the galaxy shape noise will dilute cosmological information on small scales, acting as a scale cut. For simplicity, we approximate such noise by Gaussian white noise on convergence maps [65], $\sigma_{\text{noise}}^2 = \langle \sigma_\epsilon^2 \rangle / n_{\text{gal}} A_{\text{pix}}$, where A_{pix} is the pixel area. We adopt an averaged ellipticity squared $\langle \sigma_\epsilon^2 \rangle = 0.3^2$ and a number density of background galaxies n_{gal} of 44 per arcmin² to simulate the weak lensing survey of the Rubin Observatory [84]. The multiple source redshifts of the **MassiveNuS** convergence maps also allow us to explore the information gain from redshift tomography. We consider two situations: (1) only the convergence maps for source redshift $z = 1$ are used, with $n_{\text{gal}} = 44$ per arcmin²; (2) convergence maps for five source redshifts are used, with $n_{\text{gal}} = 8.83, 13.25, 11.15, 7.36, \text{ and } 4.26$ per arcmin², respectively. The source densities are derived from the expected source redshift distribution for the Rubin Observatory survey: $n(z) \propto z^2 \exp(-2z)$ with a normalization of 50 per arcmin² [84]. We calculate the galaxy number density in each fiducial

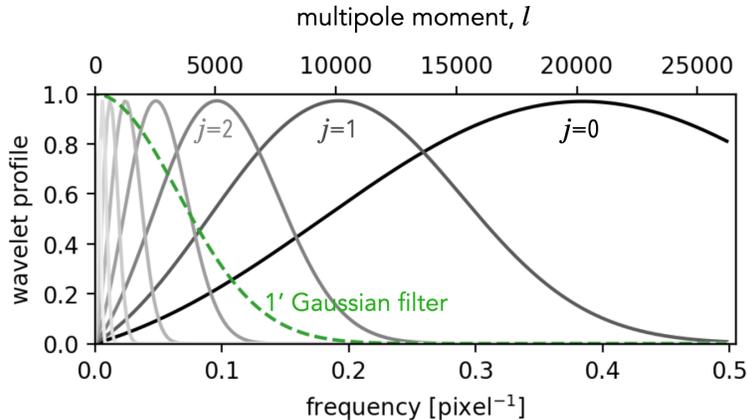


Figure 3-10. Scale coverage of wavelets used by the scattering transform. We also show the profile of a 1 arcmin Gaussian filter in frequency space for comparison.

redshift bins by integrating the distribution from $z = 0.25$ to 2.75 with steps of $\Delta z = 0.5$ as the bin widths. For noisy maps, we also carry out a smoothing with Gaussian kernel $\sigma = 1$ arcmin. Finally, the effects of masks and imperfection in shear-to-convergence conversion [e.g., 85–87], which are important practical problems for weak lensing cosmology, are not taken into account in this analysis. Our primary goal is only to compare the statistical properties of different estimators.

3.4.2 Summary statistics

We calculate three sets of summary statistics from the simulated lensing convergence maps:

- the binned power spectrum $P(l)$ which we sample with 30 linear bins of scale,
- the binned bispectrum $B(l_1, l_2, l_3)$ which we sample with 10 linear bins for each of l_1, l_2 , and l_3 in the bispectrum monopole $B(l_1, l_2, l_3)$, resulting in 125 coefficients, and
- the reduced scattering coefficients s_1^j and $s_2^{j_1, j_2}$.

For the scattering coefficients, we use the same Morlet wavelets as described in Cheng *et al.* [15], with 4 azimuthal orientations ($L = 4$) and 8 dyadic scales ($J = 8$). The scales correspond to 8 logarithmic bins from $l = 100$ to 37,000, as shown in Figure 3-10. We average over the azimuthal orientations and use the 36 reduced scattering coefficients s_1 and s_2 . On noisy maps, the small- j (small scale, high- l) scattering coefficients are dominated by noise, and we verify that abandoning the small- j coefficients does not change the precision of cosmological inference. As a result, the set of informative scattering coefficients on noisy maps is even more compact than on noiseless maps, which is desirable for inference analysis. As these scattering coefficients are non-negative, we also take the logarithm of them. This step helps to Gaussianize their PDF. We have checked the influence of this step to the cosmological constraints, and found negligible effects.

In summary, for each $3.5 \times 3.5 \text{ deg}^2$ lensing map, we calculate 36 scattering coefficients, 30 binned power spectrum coefficients, and 125 binned bispectrum coefficients for the cosmological parameter inference. We have developed a python3 package optimized in speed to calculate the translation-invariant scattering coefficients. The package, together with codes calculating the bispectrum, are

available online¹³. We estimated the power spectrum using the `LensTools` python package [70].

Before entering the next section, we provide some technical discussion about how we calculated the binned bispectrum. The bispectrum is the products of three Fourier coefficients,

$$B(\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3) \equiv \langle \tilde{I}(\mathbf{k}_1) \tilde{I}(\mathbf{k}_2) \tilde{I}(\mathbf{k}_3) \rangle \quad (3.14)$$

Because of symmetry, it is enough to consider only cases with $k_1 < k_2 < k_3$. Under statistical homogeneity, the coefficients vanish unless $\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3 = 0$. To reduce the number of coefficients, one can use the binned bispectrum mono-pole, which averages the bispectrum according to the magnitude of k_1 , k_2 , and k_3

$$B_{i,j,k} = \int_{k_i}^{k_{i+1}} d\mathbf{k}_1 \int_{k_j}^{k_{j+1}} d\mathbf{k}_2 \int_{k_k}^{k_{k+1}} d\mathbf{k}_3 (2\pi)^2 \delta^{(2)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) \cdot \tilde{I}(\mathbf{k}_1) \tilde{I}(\mathbf{k}_2) \tilde{I}(\mathbf{k}_3) \quad (3.15)$$

It is not convenient to select $\delta(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3)$ triplets within each bin. Fortunately, expanding the right-hand side of Eq. 3.15 in real space will simplify the integrant,

$$B_{i,j,k} = \langle (I \star f_i)(I \star f_j)(I \star f_k) \rangle, \quad (3.16)$$

where $f_i(\mathbf{x}) = \int_{k_i}^{k_{i+1}} e^{i\mathbf{k}\cdot\mathbf{x}} d\mathbf{k}$ is the inverse Fourier transform of the binning, an annulus in Fourier space and a wavelet-like profile in real space. This form of expression, similar to the scattering coefficients, means that the bispectrum (and other N -point functions) can also be expressed as the spatial mean of some non-linear transformation of the input field, though it has different properties from the scattering one.

3.4.3 Likelihood and parameter inference

Parameter inference requires a likelihood function $p(\mathbf{x}|\boldsymbol{\theta})$. We first assume that its function form is Gaussian, i.e., the probability distribution function (PDF) of the summary statistics \mathbf{x} is a multivariate Gaussian distribution for a given cosmology $\boldsymbol{\theta}$. We also assume the cosmological dependence of the mean vector of the Gaussian PDF can be approximated by a smooth function in the range of interest. For the dark energy dataset, we use a linear dependence, because there are only 7 cosmologies. For the `MassiveNuS` dataset, we use second-order polynomials. We set the covariance matrix to be independent of cosmology. The latter assumption is conservative but robust when the likelihood is not exactly Gaussian [43]. The cosmological dependence of mean vector and the covariance matrix are estimated using the simulated convergence maps.

Given the likelihood, we use the Fisher matrix and Bayesian posterior to quantify the constraining power of lensing scattering coefficients for cosmological parameters. The Fisher matrix allows us to forecast the constraining power using only on the local cosmological dependence of the likelihood. For the noiseless cases we consider in this paper, it provides almost identical results as the posterior. However, with increasing noise level, the difference may become non-negligible. We will therefore use the Bayesian posterior for the results of noisy data. The mathematical formulations of these two inference frameworks are given below.

¹³https://github.com/SihaoCheng/scattering_transform

For a statistical model of observable \mathbf{x} , the Cramér–Rao inequality sets the lower limit of any unbiased estimators of model parameters obtained from the observable, therefore it can be used to quantify the uncertainty of parameter inference.

$$\text{cov}(\hat{\boldsymbol{\theta}}) \geq \mathbf{I}(\boldsymbol{\theta})^{-1}, \quad (3.17)$$

where $\mathbf{I}(\boldsymbol{\theta})$ is the Fisher information matrix, whose elements are defined as

$$I_{m,n}(\boldsymbol{\theta}) \equiv - \left\langle \frac{\partial \ln p(\mathbf{x}|\boldsymbol{\theta})}{\partial \theta_m} \frac{\partial \ln p(\mathbf{x}|\boldsymbol{\theta})}{\partial \theta_n} \right\rangle. \quad (3.18)$$

In our case, $\boldsymbol{\theta}$ is the cosmological parameters, and \mathbf{x} is the vector of summary statistics such as the scattering coefficients. The likelihood is assumed to be Gaussian,

$$p(\mathbf{x}|\boldsymbol{\theta}) \propto \frac{1}{\sqrt{|\mathbf{C}|}} \exp\left[-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \mathbf{C}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right], \quad (3.19)$$

where $\mathbf{C}(\boldsymbol{\theta})$ is the covariance matrix of summary statistics. We further assume $\mathbf{C}(\boldsymbol{\theta})$ is cosmology independent. Thus, the Fisher matrix becomes

$$I_{m,n} = \frac{\partial \boldsymbol{\mu}^T}{\partial \theta_m} \mathbf{C}^{-1} \frac{\partial \boldsymbol{\mu}}{\partial \theta_n}. \quad (3.20)$$

We use linear or second-order polynomials to fit the cosmological dependence of the mean vector $\boldsymbol{\mu}(\boldsymbol{\theta})$, and use the sample covariance matrix $\widehat{\mathbf{C}}$ estimated from simulations at fiducial cosmologies ($\Omega_m = 0.26$, $w_0 = -1$, $w_a = 0$ for the dark energy dataset and $\Omega_m = 0.3$, $M_\nu = 0$, $A_s = 2.1 \times 10^{-9}$ for the MassiveNuS dataset) to estimate the inverse of covariance matrix [74]:

$$\widehat{\mathbf{C}}^{-1} = \frac{N - D - 2}{N - 1} \widehat{\mathbf{C}}^{-1}, \quad (3.21)$$

where $\widehat{\mathbf{C}}^{-1}$ is an unbiased estimator, N is the number of independent sample used for the estimation (1,000 for the dark energy dataset and 10,000 for the MassiveNuS dataset), D is the number of summary statistics used, which is 36, 30, and 125 per redshift bin, for scattering coefficients, binned power spectrum, and binned bispectrum, respectively.

The Bayesian posterior is another way to quantify the uncertainty of model parameter inference. Given the likelihood $p(\mathbf{x}|\boldsymbol{\theta})$, the posterior probability distribution of model parameter is

$$p(\boldsymbol{\theta}|\mathbf{x}) \propto p(\mathbf{x}|\boldsymbol{\theta})p(\boldsymbol{\theta}), \quad (3.22)$$

where $p(\boldsymbol{\theta})$ is the prior distribution of $\boldsymbol{\theta}$. For the noisy cases with MassiveNuS, we use a flat prior in $\Omega_m \in [0.28, 0.32]$, $M_\nu \in [0.06, 0.6]$ eV, and $A_s \in [1.8, 2.7] \times 10^{-9}$ to avoid extrapolation. The posterior is then sampled using `emcee` [88], a Markov Chain Monte Carlo (MCMC) sampler. We used 32 walkers, ran 1,000 steps to burn-in and then 10,000 steps to sample the posterior. Sampling convergence was checked within and among chains.

3.4.4 Sampling distribution of different summary statistics

Before exploring the constraining power on cosmological parameters, we first present one key advantage for using the scattering coefficients, namely that their probability density distribution (PDF) or sampling distribution is well Gaussianized.

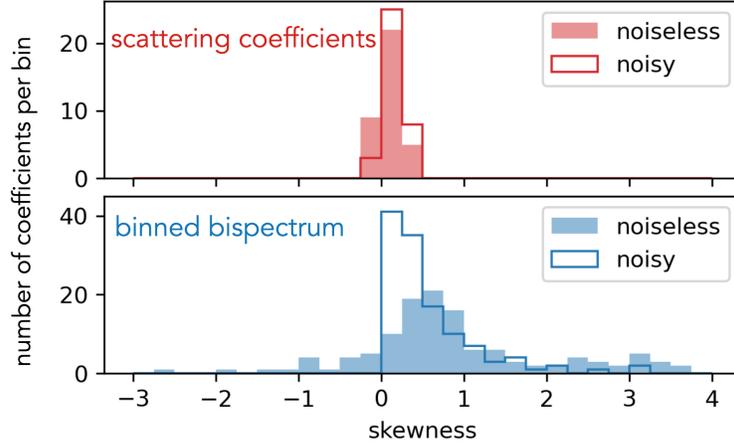


Figure 3-11. Skewness distribution of the (log) scattering and (binned) bispectrum coefficients, measured from 10,000 realizations of convergence maps at the fiducial cosmology of MassiveNuS simulation set. If the PDF of these statistical coefficients are Gaussian, then the skewness should be zero.

Table 3-II. Number of directions that are significantly non-Gaussian in the high-dimension PDF of summary statistics. Values outside and inside parentheses are for noiseless and noisy maps, respectively. Rows are different criteria to identify non-Gaussianity, which systematically suggest that the scattering coefficients have nearly-Gaussian PDF, whereas the bispectrum coefficients do not.

summary statistics	log scattering coefficients	binned bispectrum
number of coefficients	36	125
$ \text{skewness} > 0.5$	0 (0)	87 (49)
$ \text{skewness} > 1$	0 (0)	45 (22)
kurtosis > 0.5	1 (1)	120 (78)
kurtosis > 1	0 (0)	111 (49)
whitened, $ \text{skewness} > 0.5$	0 (0)	18 (4)
whitened, $ \text{skewness} > 1$	0 (0)	9 (2)
whitened, kurtosis > 0.5	0 (0)	124 (25)
whitened, kurtosis > 1	0 (0)	119 (11)

For many cosmological studies, the typical inference framework assumes a Gaussian PDF for the statistical estimator. This is the so-called Gaussian likelihood assumption, though strictly speaking, likelihood is $p(\mathbf{x}|\boldsymbol{\theta})$ with the data \mathbf{x} fixed, while in practice one usually parametrizes and fits the PDF (the same expression with the model parameters $\boldsymbol{\theta}$ fixed) to simulations. So we will call it the Gaussian PDF assumption. This assumption comes from the central limit theorem when the field of view tends to infinity. However, with a finite field of view, there is no guarantee for a Gaussian PDF [e.g., 89–92], and approximating a non-Gaussian PDF with a Gaussian one may introduce bias and/or underestimation of uncertainty to parameter inference. Therefore, summary statistics that Gaussianize quickly are favoured. Note that the non-Gaussianity of the field is not the same as the non-Gaussianity of the summary statistics. Some statistics of a highly non-Gaussian field can be nearly Gaussian distributed due to central limit theorem. On the other hand, statistics of a Gaussian random field can have non-Gaussian sampling distribution. For example, the power of each Fourier mode is expected to have a χ^2 -distribution.

Deviations from the Gaussian PDF approximation are harmful to inference, among which the most problematic is the presence of a heavy tail. To quantify this effect, we will consider both the skewness and kurtosis parameters of the 1D marginal distributions. More precisely: we calculate these parameters for both noiseless and noisy maps, using the fiducial cosmology of MassiveNuS simulation with source redshift $z = 1$. The skewness is defined as $\langle (\frac{x-\mu_x}{\sigma_x})^3 \rangle$ for each summary statistic x , where μ and σ are the sample mean and standard deviation calculated from the 10,000 realizations. It measures the asymmetry and tail heaviness of the one-variable PDFs, which is one possible deviation from a Gaussian distribution. The scattering coefficients¹⁴ have much lower skewness than the bispectrum coefficients. We show the histograms of their skewness parameters in Figure 3-11. Similarly, we measure the kurtosis defined as $\langle (\frac{x-\mu_x}{\sigma_x})^4 \rangle - 3$, which is more sensitive to the tail heaviness. For a 1D Gaussian distribution, given our sample size, the skewness and kurtosis should be 0 ± 0.03 and 0 ± 0.05 , respectively. The measured values for the scattering coefficients are indeed much closer to these expectations than the values obtained from the bispectrum which, in some cases, exceed the upper bound by more than two orders of magnitude. The effect of global whitening (using principle component analysis to remove correlations between dimensions) is also examined. Under all these metrics, the scattering coefficients show a low non-Gaussianity. In contrast, many dimensions of the bispectrum data vector are asymmetric or heavy tailed, even on the noisy maps where the field itself is more Gaussian. This is shown in Table 3-II where we present the number of coefficients exceeding a given skewness or kurtosis threshold for each estimator. Tests using the 95-percentile, yet another measure of tail heaviness, show similar results. Assessing the non-Gaussianity of parameters in high-dimensional spaces can be done from numerous points of view. For example, Sellentin & Heavens [89] proposed to use all pairs of 2D marginal PDFs. Such tests will be valuable in future studies aimed at inferring cosmological parameters.

The difference in the types of distributions found for scattering and bispectrum coefficients can be understood from their respective designs. The bispectrum coefficients are products of three random variables. The multiplication in general yields a distribution with heavier tail than the original variables. In contrast, the scattering transform uses a ‘first-order’ modulus which does not amplify the tail. Therefore, as the field of view increases, the central limit theorem Gaussianizes the scattering coefficients much quicker than bispectrum coefficients. Moreover, the logarithmic binning of scales through wavelets combines more Fourier modes, which also helps the Gaussianization process.

¹⁴We point out that, even without taking the logarithm of the scattering coefficients, the skewness values of the ‘raw’ scattering coefficients are still relatively small, typically between 0 and 1.

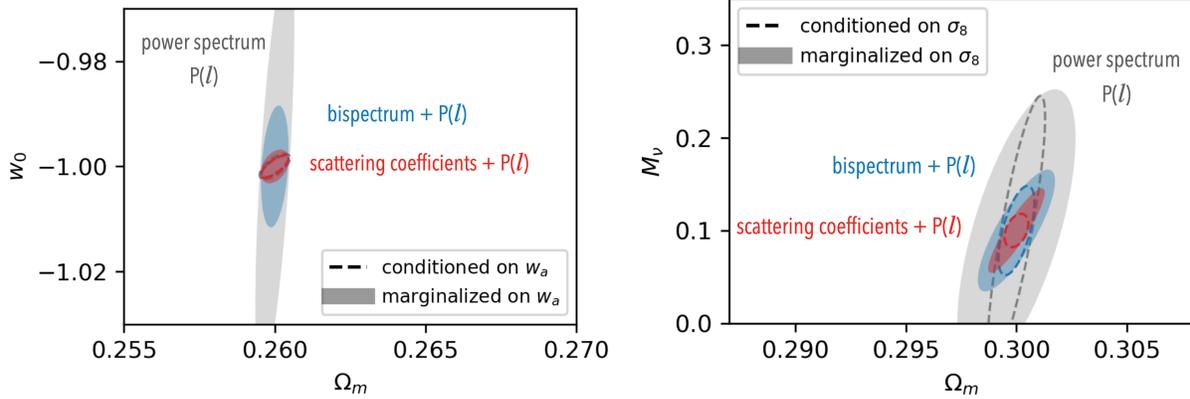


Figure 3-12. Fisher forecast with noiseless κ maps (95% confidence ellipses). The scattering coefficients have remarkably high constraining power for cosmological parameters. *Left:* Forecast for dark energy equation of state index $w(a) = w_0 + (1 - a)w_a$, from 20,000 deg^2 noiseless map with l nominally up to roughly 30,000 and source redshift $z = 2$. *Right:* Forecast for neutrino mass M_ν , from 20,000 deg^2 noiseless map with source redshift $z = 1$.

In summary, we find that the scattering coefficients Gaussianize particularly well, which is desirable and necessary for accurate cosmological inference.

3.4.5 Results

Having checked the Gaussian likelihood assumption for the scattering coefficients, now we compare the cosmological constraints obtained by different summary statistics. These constraints were obtained based on simulated convergence maps with varied cosmological parameters, and 1,000 (for the dark energy simulation set) or 10,000 (for the neutrino mass simulation set `MassiveNuS`) realizations for each cosmology.

Because the convergence maps have a long-tailed PDF (approximately log-normal), we do expect the scattering coefficients to extract more information than the bispectrum. Indeed, we find a significantly tighter constraint using the scattering coefficients on noiseless maps (when all high- l modes are included). In Figure 3-12 we show the Fisher forecast of cosmological parameters using noiseless convergence maps, which corresponds to l from 100 to roughly 30,000. For simplicity, we only show a plane of two parameters for each dataset. In both panels of Figure 3-12, the marginalized ellipses are the projection of the three-parameter Fisher ellipsoids onto the plane, while the conditioned ellipses are their cross-sections.

For dark energy parameters (left panel), the power spectrum constrains w_0 well when w_a is known (conditioned), but constrains poorly if w_a is also to be constrained (marginalized), meaning a strong degeneracy between w_0 and w_a . Compared to the bispectrum, which is known to partly break this degeneracy [36], the scattering coefficients almost eliminate it in the noiseless case. In weak lensing cosmology, there is another well-known degeneracy between Ω_m and σ_8 , which the scattering coefficients can also break [15]. However, the dark energy dataset used in this study has a fixed σ_8 value. So, it remains to be investigated whether or not the scattering coefficients can break both degeneracies at the same time.

On the noiseless maps, scattering coefficients also provide much tighter constraint on the neutrino mass sum M_ν than the power spectrum and bispectrum (see the right panel of Figure 3-12). In the

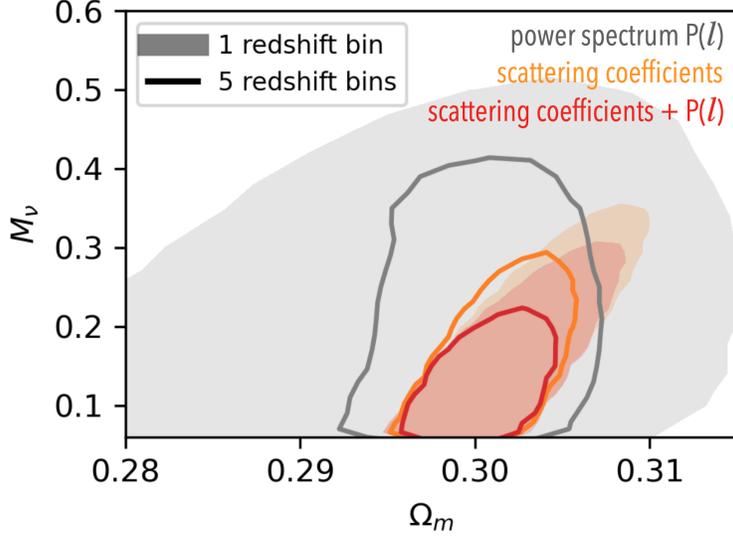


Figure 3-13. Forecast of cosmological parameters (95% confidence contours) for a Rubin-observatory-like survey with 20,000 deg² field of view, with M_ν , Ω_m , and A_s (or σ_8) to be constrained. Colours represent results of different sets of summary statistics of lensing convergence (κ). Improvement from adding redshift tomography is also shown.

parameter space spanned by Ω_m , σ_8 , M_ν , the power spectrum’s constraining power comes mainly from its overall amplitude, which results in a disc-like Fisher ellipsoid. Adding the bispectrum reduces the disc’s area by partly breaking the Ω_m - σ_8 degeneracy, whereas adding the scattering coefficients almost eliminate it, reducing it to a one-dimension instead of two-dimension degeneracy. This is consistent with the finding presented in the figure 5 of [15]. The ratio between the volumes of the three-parameter Fisher ellipsoids (i.e., the square root of determinant of the covariance matrix) shown in the right panel of Figure 3-12 is roughly 20:4:1.

When the galaxy shape noise is taken into account, the constraining power of all summary statistics weakens due to the loss of small scale (high- l) information. In Figure 3-13 we show the forecast of neutrino mass sum (M_ν) for a Rubin-observatory-like survey, using the Bayesian posterior with fiducial cosmology at $(M_\nu, \Omega_m, A_s) = (0.1 \text{ eV}, 0.3, 2.1 \times 10^{-9})$. In spite of noises, there are still substantial non-Gaussian structures in the lensing convergence maps, thus the scattering coefficients still set much tighter cosmological constraints (more than two times in M_ν) than using the power spectrum alone. In this noisy case, the combination of bispectrum and power spectrum sets a similar constraint to the combination of scattering coefficients and the power spectrum (though not shown in the figure). Figure 3-13 also show the effect of redshift tomography. A five-bin redshift tomography improves the scattering coefficients’ constraint for M_ν by about 40%.

With the same dataset (MassiveNuS), similar explorations were performed for other non-Gaussian statistics, including the PDF [93], bispectrum [94], peak count [95], starlet peak count [96], minima count [97], Minkowski functionals [98], and starlet l_1 -norm [99]. They all found similar improvement from redshift tomography for their respective summary statistics.

3.4.6 Conclusions

In cosmology, substantial information is stored in non-Gaussian structures on small scales, which requires statistics beyond the power spectrum to extract. Motivated by ongoing and upcoming deep surveys, we explore the weak lensing application of a novel and powerful non-Gaussian statistics, called the scattering transform. We extend the constraining forecast of Ω_m and σ_8 in Cheng *et al.* [15] to more cosmological parameters including the dark energy parameters w_0 , w_a and neutrino mass M_ν . To do so, we use mock convergence maps from two publicly available simulation suites (the dark energy and MassiveNuS sets) made by the Columbia Lensing team.

We first show that using the scattering coefficients, one can generate random images with textures very close to a simulated lensing map, which cannot be achieved by using the power spectrum and bispectrum coefficients. Then, we show that the scattering coefficients provide significantly better constraint for dark energy parameters w_0 , w_a and neutrino mass M_ν than using power spectrum alone.

For noiseless maps (when high- l modes are accessible), the scattering transform also outperforms the bispectrum and power spectrum. This result can be explained by the ‘first-order’ nature of scattering coefficients which stay proportional to the field intensity, in contrast to higher-order statistics which amplify the distribution tail. This *lower-order* nature together with the wavelet weighting strategy efficiently concentrates cosmological information into a compact set of scattering coefficients. Moreover, the *lower-order* nature makes the distribution of scattering coefficients much more Gaussian and robust than higher-order statistics, which is essential for accurate likelihood parametrization and parameter inference.

We also provide a forecast of M_ν with a noise level of the Rubin observatory survey, where the scattering coefficients set a 2 times tighter constraint than using the power spectrum alone, and redshift tomography improves the constraint by an additional 40%. Similar forecast for the dark energy parameters is left for future study, due to the redshift limitation of the mock convergence maps. For noisy maps, although the constraining power of the scattering coefficients is not significantly better than the combination of bispectrum and power spectrum, we argue that the scattering coefficients are still preferred, because the PDF of scattering coefficients is much more Gaussian than bispectrum coefficients.

The scattering transform yields a compact set of *lower-order* summary statistics that are stable, powerful, and efficient for characterizing non-Gaussian structures. Together with the previous section, we have shown that the scattering transform has great potential to serve as the non-Gaussian summary statistics in observational cosmology.

3.5 Application to observational data

In the previous section, I demonstrated the high constraining power of the scattering transform on cosmological parameters by showing its forecast based on simulations. In this section, I will demonstrate that there is no major difficulty in applying the scattering transform to real weak lensing data. In particular, in real lensing data, the direct observables are the shape of galaxies. In the process of making a mass map from the shear catalog, several aspects can deviate from the simplified model used for our forecasts in the previous sections, such as masks, non-Gaussian galaxy shape noise, non-uniform galaxy density, etc. Nevertheless, these are no fundamental difficulties, because currently all the inference with the scattering coefficients is based on simulations and forward-modelling,

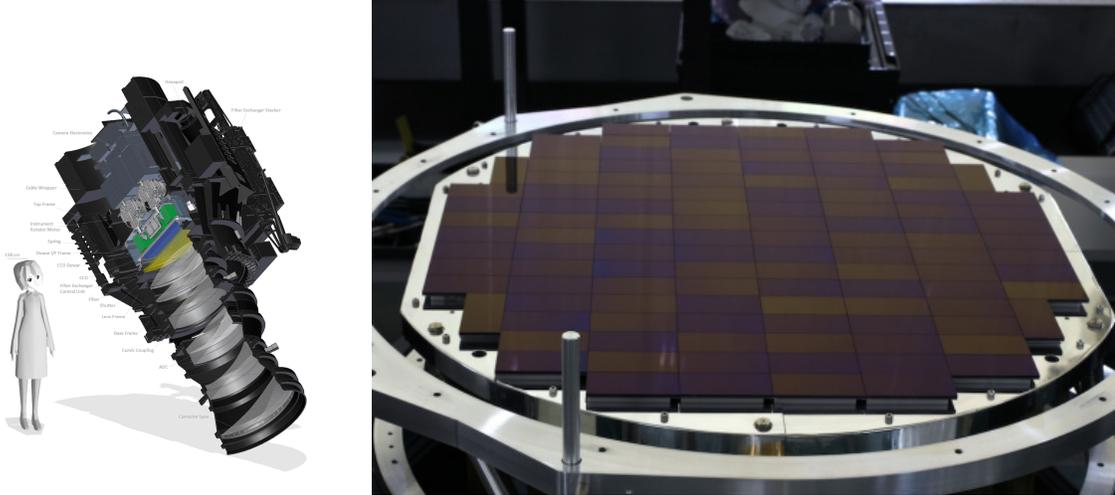


Figure 3-14. Left: the design and size of the HSC camera. Right: The CCD array of the Hyper Suprime-Cam on the 8.2m Subaru telescope. It has 870 Mega pixels in total and covers 1.5 degree field-of-view in diameter, equal to the diameter of 3 full moons.

given that there is not yet analytical predictions for the scattering coefficients. Forward-modelling these effects, compared to the traditional approach of ‘de-biasing’ the imperfections, is relatively straightforward to run and does not increase computation cost substantially, as most of the time is dedicated to running the cosmology simulations. In this section, I shall apply the scattering transform to the data from Subaru Hyper Suprime-Cam (HSC) survey [100]. This is an on-going project with blinding policy to avoid confirmation bias, meaning that project members cannot dig into cosmological constraints until the whole inference pipeline has finished and has been tested. So I will not show constraints on cosmological inference and only present results of one field in the HSC survey for illustration.

3.5.1 HSC weak lensing data

The Subaru Hyper Suprime-Cam (HSC) survey [100] is a wide-field imaging survey on the 8.2m Subaru telescope at Hawaii. The whole survey consists of three layers, with the wide layer planned to cover a 1,400 deg² area over 5-6 years. Figure 3-15 shows the wide-field CCD array of the HSC. We choose HSC because among the major on-going weak lensing surveys, HSC has the deepest photometry (5σ point-source depth of $i \sim 26$) and best seeing (mean i -band seeing of $0''.58$), therefore it is expected to reveal the most non-Gaussian structures in the convergence field. Here we use its first data release [100], referred to as HSC S16A. It covers an area of 136.9 deg² area split into six fields: XMM, GAMA09H, GAMA15H, HECTOMAP, VVDS, and WIDE12H. We used the shear catalog compiled by Mandelbaum *et al.* [68] with a mean galaxy density around 22 arcmin⁻², which is ready for weak lensing cosmology. Because of the blinding policy, I will not analyze the whole dataset but only focus on one of the six fields, the field WIDE12H, for illustration purpose.

Following the pseudo power spectrum analysis performed by the HSC team [101], we conduct a conservative cut of redshift between 0.3 and 1.5, which is best suited for the HSC photometry with a break at 400 nm. The photometric redshifts of source galaxies in HSC survey are estimated based on photometry in its *grizy* bands with several different codes [102], including a classical template-fitting code (Mizuki), a machine-learning code using self-organizing map (MLZ), a neural



Figure 3-15. A small patch of the HSC survey. Each blob in the figure is a galaxy far away. Their shapes are measured and averaged to estimate the gravitational lensing effect of the foreground dark matter. Regions near bright stars (the bright objects with spikes) are masked because of their severe influence on shape measurements.

network code (Ephor AB), and others. Slightly different from Hikage *et al.* [101], we apply the tomographic photo-z cuts on the redshifts determined by the MLZ code, which leads to about 9.5 million galaxies in the six fields and a mean galaxy number density of 20 arcmin^{-2} .

The shape of galaxies in the HSC shear catalog Mandelbaum *et al.* [68] is measured on the *i*-band coadded images with the re-Gaussianization PSF correction method [103], where the galaxy ellipticity is defined as

$$\mathbf{e} = (e_1, e_2) = \frac{1 - (b/a)^2}{1 + (b/a)^2} (\cos 2\phi, \sin 2\phi), \quad (3.23)$$

where a and b are the observed length of major and minor axes, ϕ is the position angle of major axis in equatorial coordinate system. The shear of each galaxy is estimated by

$$\gamma^{(\text{obs})} = \frac{1}{1 + \langle m \rangle} \left(\frac{\mathbf{e}}{2R} - \mathbf{c} \right), \quad (3.24)$$

where R is called responsivity, m the multiplicative bias, and \mathbf{c} the additive bias. The responsivity is given by [104, 105]

$$R = 1 - \langle e_{\text{rms}}^2 \rangle, \quad (3.25)$$

where $\langle e_{\text{rms}}^2 \rangle$ is the weighted average of intrinsic root mean square ellipticity in each direction, and

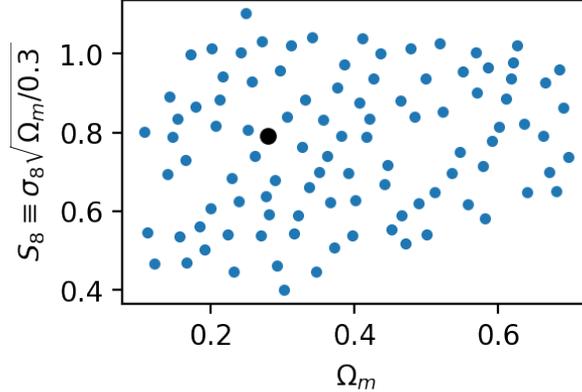


Figure 3-16. A set of 100 N -body simulations sampled in the $\Omega_m - \sigma_8$ plane to build the likelihood emulator for scattering statistics.

the weight w is defined as

$$w = \frac{1}{\sigma_e^2 + e_{\text{rms}}^2}, \quad (3.26)$$

where σ_e is the measurement error of galaxy shape. The shape error and bias are estimated based on simulated HSC images of galaxies observed in the COSMOS survey [68], a deep and high-resolution survey conducted by the Hubble Space Telescope.

3.5.2 Mock data

For theoretical predictions, we use the mock shear catalog generated by Shirasaki *et al.* [106] for the first data release of the HSC survey. This mock catalog was generated based on 108 quasi-independent full-sky N -body simulations [107] with the Λ CDM model and WMAP9 cosmological parameters [108]: $\Omega_{\text{cdm}} = 0.233$, $\Omega_{\text{b}} = 0.046$, $\Omega_{\text{m}} = \Omega_{\text{cdm}} + \Omega_{\text{b}} = 0.279$, $\Omega_{\Lambda} = 0.721$, $h = 0.7$, $\sigma_8 = 0.82$, and $n_s = 0.97$. The sky footprint of the HSC fields was rotated in 21 angles to take full advantage of the large volume of each simulation, resulting in a total of 2,269 realizations of mock HSC lensing maps. Then, mock galaxies were placed at the same positions as the observed galaxies in each field of view and assigned randomly rotated ellipticities. The observationally estimated photometric redshifts were also used for calculating the cosmic shear effect on each mock galaxies. This set of mock data in a fiducial will be used to estimate the cosmic variance of the scattering coefficients and further to determine the uncertainty of inference of cosmological parameters.

To perform cosmological inference, one also need a likelihood function, which concerns different cosmologies. We use a set of cosmic-varying simulations (Shirasaki *et al.* in preparation), which has 50 realizations in each cosmology but samples 100 different cosmologies on the $\Omega_m - \sigma_8$ plane, to build an emulator for the scattering statistics. The sampling is shown in Figure 3-16.

3.5.3 From shear catalog to mass maps

Similar to other non-Gaussian statistics used for weak lensing, the scattering transform calculates statistics from a pixelized map instead of a catalog of objects. Therefore, we first generate a pixelized

shear map by weighted averaging the ellipticity of galaxies within cells of sizes 0.88^2 arcmin², according to the galaxy density of the HSC survey. We abandon cells with galaxy density smaller than half of the mean galaxy density to reduce the influence of their high noise. Thus this shear maps have missing pixels corresponding to both the masked regions near bright stars where the shear measurements are not reliable and their neighborhood. For most non-Gaussian descriptors in weak lensing, such as bispectrum, peak count, Minkowski functional, etc, the next step is to convert the shear maps into convergence maps, because shear maps are spin-2 maps with redundant information for weak lensing (the corresponding B-mode), while the convergence maps are scale fields directly corresponding to the projected mass distribution. Interestingly, the scattering transform can also actually be applied directly to shear maps, as a spin-2 field can be expressed by a field of complex numbers. Moreover, in cases with no masks or edge effects, the scattering coefficients of a shear map is almost the same as the corresponding convergence map. Nevertheless, to be consistent with other estimators, we sticks to convergence maps in this study. A traditional way to perform the conversion is by using their simple point-wise relation in Fourier space (eq. 3.4), which is called the Kaiser-Squires method [85]. This method is ideal without masks or edges but suffers from edge effect in presence of masks and finite field of view. In principle, as long as one stick to the forward-modelling approach as treat the mock data in the same way as the real data, there is no bias introduced by the imperfection of mass map conversion. However, edge effects do reduce the signal to noise of some descriptors on the mass map if they are not carefully treated. Therefore, several improved methods have been proposed accordingly to reduce the edge effect by making use of prior knowledge about the properties of the field. In this study we use an in-painting technique based on the sparsity property of in discrete cosine transform and suggested for weak lensing mass maps [86]. An example of in-painted noiseless convergence map is shown in Figure 3-17.

3.5.4 Weak lensing scattering transform

Having obtained the pixelized and in-painted lensing mass map for both observations and simulations, we then calculate the scattering coefficients of these mass (convergence) maps, which is the first application of the scattering transform to real cosmological data.

Figure 3-18 shows the reduced scattering coefficients of mass maps in the field WIDE12H. The scale index j is related to the exponent of wavelet central wavelengths. The scale $j=0$ roughly corresponds to a wavelength between 2 to 4 pixels, i.e., a multipole moment l between 1000 and 2000 on the sky. Each increase of j by unity corresponds to a larger scale and thus smaller l by a factor of 2. The drop of s_1 coefficients at small scales is due to a 1 arcmin Gaussian smoothing on the mass map. The s_1 coefficients are the orientation average of the original S_1 coefficients, which are similar to the square root of power spectrum in logarithmically-spaced bins. The orientation-reduced coefficients s_{21} and s_{22} are defined in chapter 2. They characterize non-Gaussian properties of the mass map.

It can be seen that the curves representing observations and a random realization of simulated data almost overlap with each other, for both the 1st-order and 2nd-order coefficients, meaning that the behavior of scattering coefficients in real data is basically expected and under control. There is an offset of several percent in s_1 coefficients at all scales, which should not be the problem of the scattering coefficients because we confirm that a similar offset is also present in the power spectrum, which is a well-tested estimator in cosmology. This offset may be caused by cosmic sampling variance, e.g., the presence of a rare cluster in the field of view, or by the fact that our universe actually has a different parameter from the fiducial cosmology we used for the simulations. To comply with the

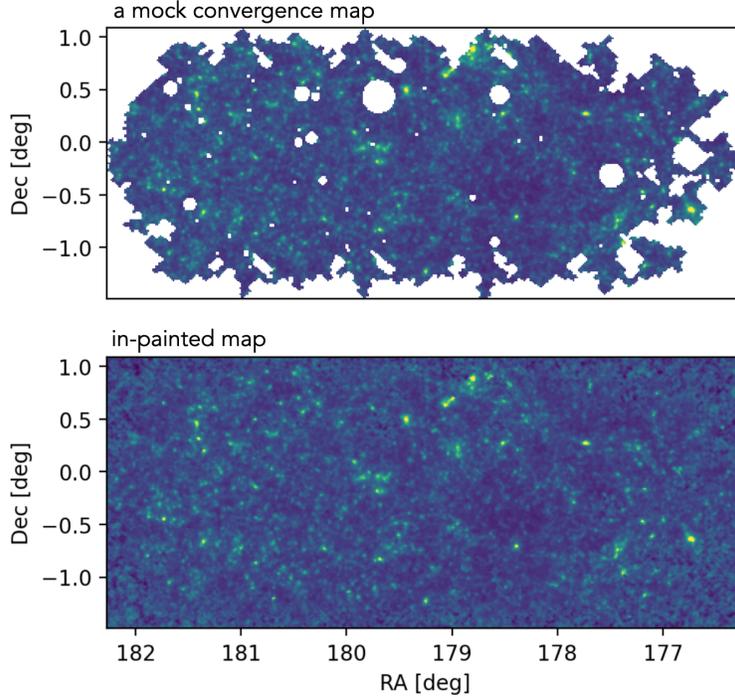


Figure 3-17. Upper: a mock convergence map for the field WIDE12H without galaxy shape noise, redshift bin 0.6–0.9. Lower: an in-painted field of the upper panel. The in-painting is conducted from the corresponding masked shear map.

blinding policy, I will come back to this question only after the likelihood emulator and pipeline of systematic checks have been built and tested. It is worth mentioning that this offset, either due to cosmic variance or cosmological dependence of the coefficients, is much smaller than the expected variance of higher-order statistics. So, based on the consistency of observational data with simulations, we conclude that the scattering coefficients have expected and controlled behaviors. In the next step, we will use a set of simulations sampled at various positions in the parameter space of σ_8 and Ω_m to build a likelihood emulator and conduct cosmological inference, which is similar to the procedures described in the previous two sections.

To summarize, I presented the first application of the scattering transform to observational weak lensing data, which differ from the simplified situations assumed for a fast forecast in the previous sections in terms of various observational effects. Nevertheless, as we follow a forward-modelling approach to build up the inference framework, there is no major difficulty in including these factors in our model and predicting the behavior of scattering coefficients accordingly. I have shown that the scattering transform calculated on real data are qualitatively consistent with expectation. Moreover, its robustness and compactness make it ideal for practical analysis. Right now, we are working on building and testing the likelihood emulator for the HSC survey, and the first cosmological inference with the scattering statistics is expected to be published in the near future.

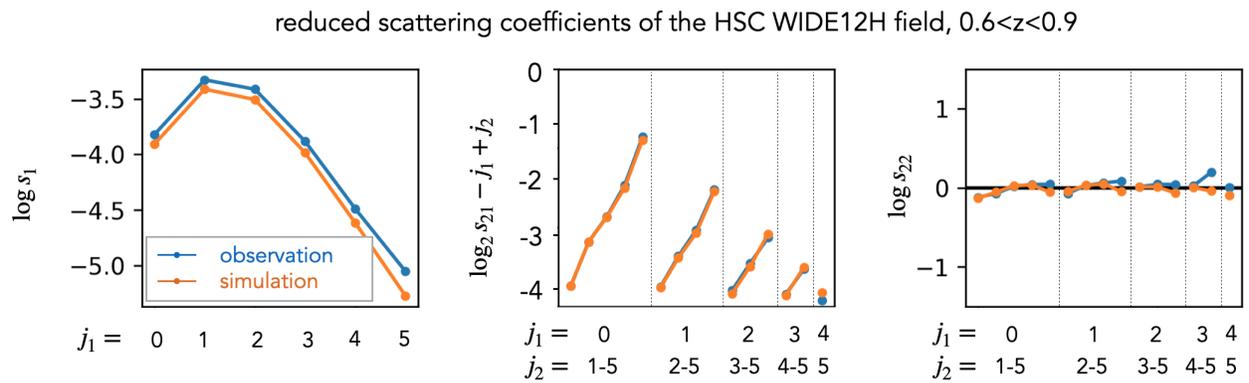


Figure 3-18. The scattering coefficients measured on real data and on a realization of the corresponding mock data in the fiducial cosmology. It can be seen that the simulation based forward-modelling approach can successfully recover the behavior of scattering statistics in real data, which includes a number of observational effects.

Chapter 4

Applications in exploratory research

In the literature, the majority of applications of the scattering transform has been targeted to classification problems, partly because of its computer-vision origin. However, it is by no means limited to classification. Being a natural extension of the power spectrum, it is suitable for a variety of problems, including parameter inference (regression) and exploratory data analysis – two types that are commonly met by physicists.

In the previous section, I showed applications of the scattering transform in regression problems in the context of cosmology. In this section, I shall show examples where the scattering transform is used to facilitate data mining and exploration. Interestingly, in many cases the scattering transform can be used to achieve similar goals to CNNs but with much less computational cost and more interpretability.

4.1 Sea temperature field

Satellite-borne sensors have been collecting high quality data for oceanography studies since 1990s. In this section, I shall show an application of the scattering transform to remote sensing oceanographic images. In particular, I demonstrate how useful it can be to organize cutouts of the ocean temperature field taken at different place and time.

The sea surface temperature field contains information about the dynamics of ocean currents. Similar to the cosmic matter density field, it is also non-Gaussian, therefore the power spectrum alone does not capture the full information stored in the field. Recently, convolutional neural networks have been trained to find outliers (regions showing field textures significantly different from the majority of the sample) from the sea surface temperature dataset [109], in addition to other machine learning applications in oceanographic data [e.g., 110–113]. In this section, I shall show that the scattering transform can be used to achieve similar goals with much less computational cost and more interpretability than convolutional neural networks.

4.1.1 Data and preprocessing

We use the nighttime MODerate-resolution Imaging Spectroradiometer (MODIS) Level-2 (L2) ¹ sea surface temperature (SST) dataset obtained from the NASA spacecraft Aqua in the year of

¹<https://www.ghrsst.org/ghrsst-data-services/products/>

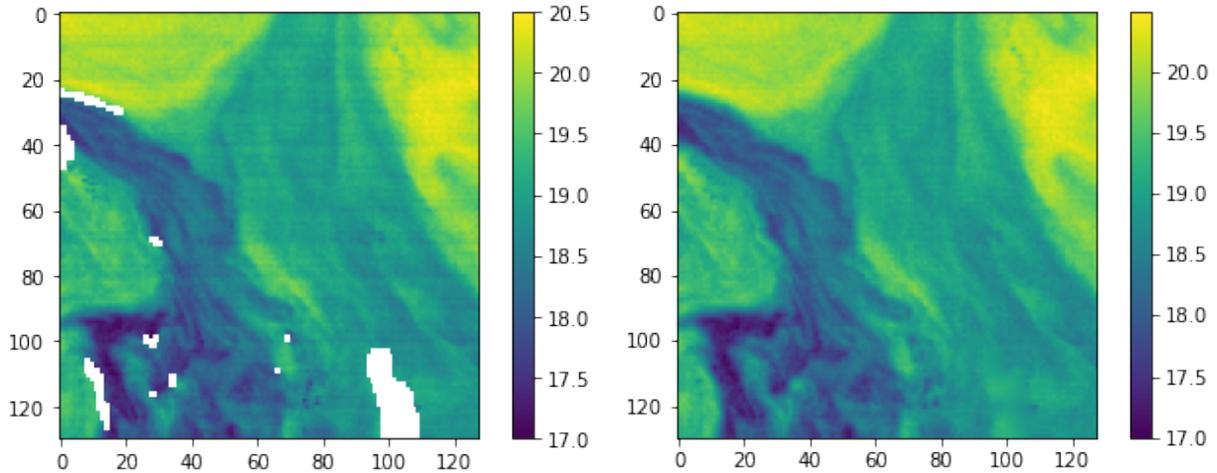


Figure 4-1. An example of ‘cutouts’ of the sea surface temperature dataset from NASA Aqua satellite. The left panel shows the image from the original Level-2 data product, the right panel shows the result after in-painting masked pixels and flat field correction.

2015. The NASA Aqua satellite runs in a near polar orbit and works in a double-scanning mode: a south–north scanning by the orbiting motion of the satellite, and a west–east scanning by a rotating mirror in the satellite. As a result, with relatively small number of pixels in the detector (10×1 to 40×1 pixels for each band), the satellite can scan a large area within a short period. Its viewing swath width (roughly along circles of latitudes) is 2,330 km, which covers the entire surface of the Earth every one to two days. The MODIS detector contains 36 bands with wavelengths spanning from $0.4 \mu\text{m}$ to $14 \mu\text{m}$ (visible light to mid-infrared). The Level-2 sea surface temperature dataset is mainly based on the 4 and $11 \mu\text{m}$ infrared emission measurement, with a spatial resolution of 1 km per pixel.

In the Level-2 data product ², pixels are kept as on the scanning track and packed into 5-minute segments, referred as ‘granules’. Each granules corresponds to a 2000×1350 km field of view and usually includes large fraction of pixels corresponding to lands and clouds, where no meaningful sea surface temperature can be retrieved. As clouds and lands would introduce patterns not related to the physics of sea temperature, we follow Prochaska *et al.* [109] to cut the granules into 128×128 pixel (128×128 km at the nadir) ‘cutout’ images and use only the cutouts with less than 5 % of pixels masked as clouds or lands (defined by a threshold of data quality greater than ‘2’). Also, cutout centres are required to be within 480 pixels from the nadir, so that the geometric deformation along the scanning direction (roughly west–east) remains less than a factor of two. Slightly different from Prochaska *et al.* [109], the cutouts in this study are sampled by first making a square tiling of the granules and then ‘drizzle’ along both x- and y-axes by 64 pixels to over-sample, instead of randomly located within the granules. This procedure yields 104,134 cutouts through out the year of 2015.

Moreover, masked pixels are in-painted using the biharmonic algorithm in the `skimage.restoration` python package, which comply with the Navier-Stokes equation. As shown in the left panel of Figure 4-1, there are clear horizontal patterns present in the field, along the scanning direction

²<https://oceancolor.gsfc.nasa.gov/data/aqua/>

of the satellite. This is due to issues of the calibration of pixel sensitively, similar to the effect of ‘flat field’ in astronomy. Unfortunately, experiments shows that such flat field of the 10×1 pixel detector varies with time, so it can only be estimated and corrected locally (within each cutout). Nevertheless, local de-flatting turns out to significantly reduce the stripes without introducing noticeable noise. By successfully de-flatting the data, we avoid the smoothing and down-sampling adopted in Prochaska *et al.* [109], which loses small-scale information. The right panel of Figure 4-1 shows the result of in-painting and the flat field correction.

Turbulence fields, including the sea temperature fields, often have a red power spectrum. In such cases, the spectrum leakage of low-frequency power due to the window function (field of view) may cause a severe problem for multi-scale analysis such as the power spectrum and scattering transform. It can also be understood as the edge effect. To reduce the such effect, we fit the 2D field with a linear function and remove the slope.

4.1.2 The scattering representation

After pre-processing the data, we now discuss the physical structures in the sea surface temperature fields. Being a tracer of turbulence, these fields contain rich non-Gaussian structures and morphological information, which is well suited for the scattering transform. To simplify the story, we can further reduced the scattering coefficients into four indices:

1. s_1 power index: the power-law index of s_1 as a function of scale. It is derived from the orientation-averaged s_1 coefficients at scales $j = 1$ and $j = 4$;
2. amplitude index $s_1^{j=1}$: the logarithm of S_1 coefficient at scale $j = 1$ (around 6-pixel wavelength), averaged over orientations l ;
3. feature sparsity index: the averaged logarithmic s_{21} coefficients over all scale combinations (j_1, j_2) except those with $j_1 = 0$ (to reduce the influence of white noise at small scale);
4. texture shape index: the averaged logarithmic s_{22} coefficients over all scale combinations (j_1, j_2) except those with $j_1 = 0$ (to reduce the influence of white noise at small scale).

The first two s_1 -related indices are basically the same as the power spectrum at 6-pixel scale and the power index of power spectrum, which is expected for field with not extremely sparse features (as discussed in section 2.5.1) and is confirmed for these sea surface temperature fields. The non-Gaussian information, on the other hand, is mainly extracted by the second-order scattering coefficients, which can be further reduced to the feature sparsity index s_{21} and shape index s_{22} with interesting intuitive meanings. To demonstrate the power of these quantities, we show in Figure 4-2 some example fields arranged according to these two indices. Fields with different features are automatically grouped and separated. For example, different regions on the figure correspond to quiet fields, big eddies, cloud contamination, front of cold and warm currents, etc.

Interestingly, the scattering representation also provides a measure of similarities between field structures, which is useful for data exploration. For example, Figure 4-3 shows a gallery of cutout fields (right panel) selected according to the similarity to a given field with a big eddy (left panel). The scattering transform accurately recognizes the swirl feature in the fields. The principal component analysis (PCA) of the fields themselves, in contrast, does not work because PCA is not translation invariant.

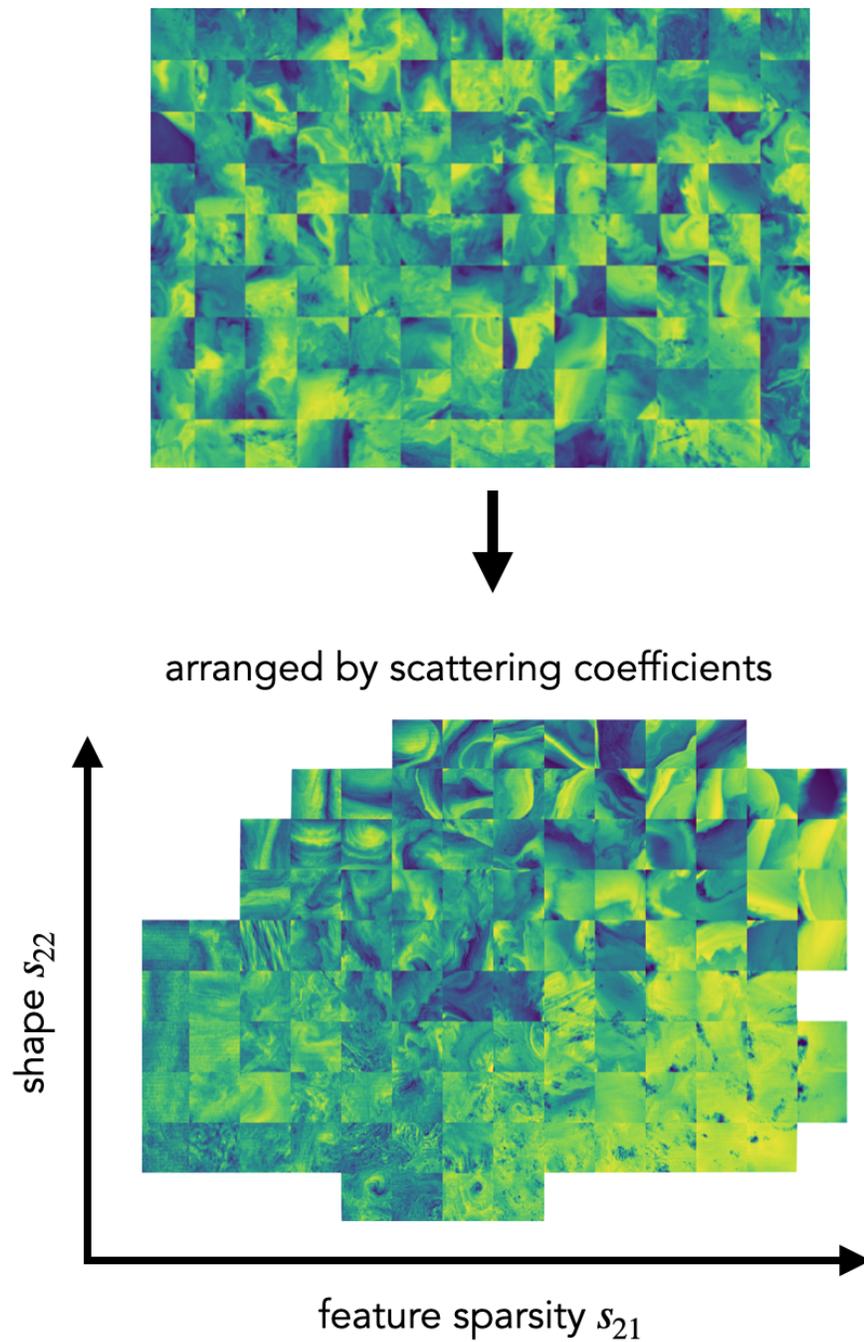


Figure 4-2. Randomly selected sea surface temperature fields (upper) and a set arranged based on their 2nd-order scattering coefficients (lower). The x-axis of the right panel represents an averaged measure of the feature sparsity (s_{21}), and the y-axis represents that of the shape indicator (s_{22}).

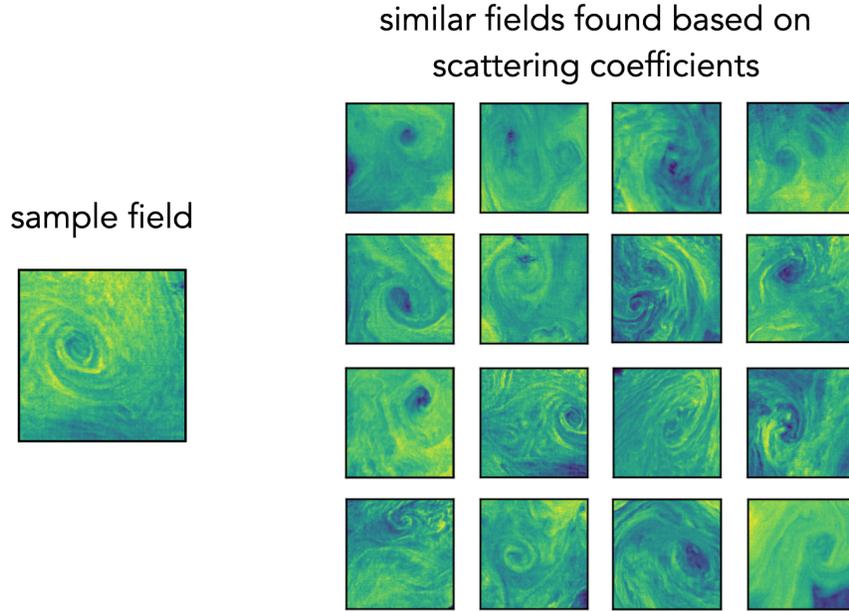


Figure 4-3. An interesting sea surface temperature field with an eddy (left) and a gallery of similar fields selected based on the s_{22} scattering coefficients (the shape indicator).

Combining the 1st- and 2nd-order scattering coefficients makes it possible to separate the information contribution from only Gaussian and non-Gaussian structures. Figure 4-4 shows the distribution of the aforementioned four indices from 1st- and 2nd-order coefficients. It can be seen that there is some correlation between the 1st- and 2nd-order coefficients. Nevertheless, using the 1st-order coefficients (or the power spectrum alone) still cannot separate many structures. Figures 4-5 to 4-7 show the additional non-Gaussian information provided by the s_2 -related indices. They are similar to the right panel of Figure 4-2, but here the fields in each figure are required to have similar s_1 power index, which is equivalent to the power index of power spectrum. These figures clearly show that with similar power spectrum, the morphology of the sea surface temperature field can still vary a lot, corresponding to potentially different underlying physics.

Moreover, in addition to the scale separation of energy, the scattering coefficients provide yet another dimension to describe the energy distribution in the turbulence field, which distinguishes the spatially concentrated and sparsely distributed energy in the field from the more wide-spread and uniformly distributed energy in the field.

This set of result is expected to be published in the oceanography literature in the near future.

4.2 Galaxy morphology

In addition to stationary ergodic fields, the scattering transform is also suitable for characterizing objects in a background. In this section, we discuss the potential of scattering transform to be used in galaxy morphology classification, a long-standing problem in astronomy. As shown in Figure 4-8, galaxies look different. In general, they can be grouped into several classes according to their morphological properties, such as the presence of disk, spiral arms, bar, etc. The development of galaxy morphology classification is basically in line with all other fields that require vision

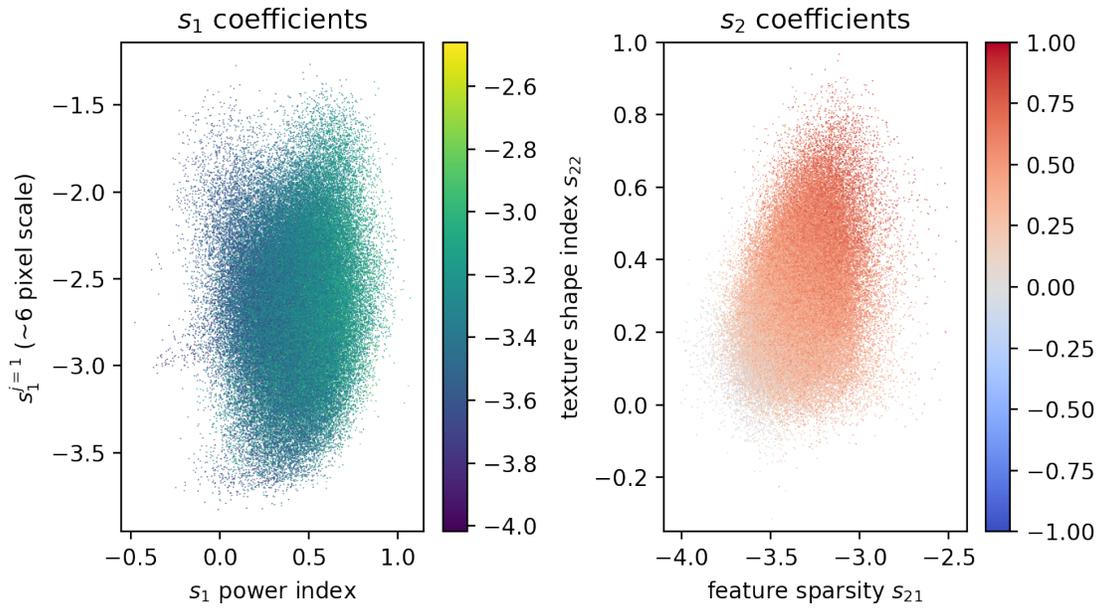


Figure 4-4. The distribution of the $\sim 100,000$ sea surface temperature cutout fields on four reduced dimensions: the amplitude index, s_1 (power spectrum) power index, and two non-Gaussianity indices based on s_2 scattering coefficients. Colour in the left panel represents the feature sparsity index (the x-axis of the right panel), and the colour in the right panel represents s_1 power index (the x-axis of the left panel).

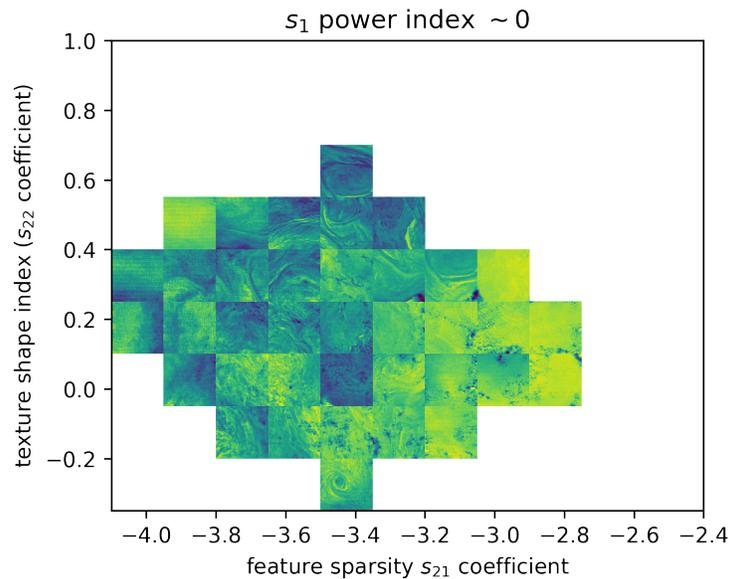


Figure 4-5. The non-Gaussianity distribution (feature sparsity s_{21} and texture shape s_{22}) for images with similar s_1 power index (i.e., similar shape of the power spectrum). The shown cutouts are randomly chosen from the corresponding cells. The s_1 power index is between -0.1 and 0.1 .

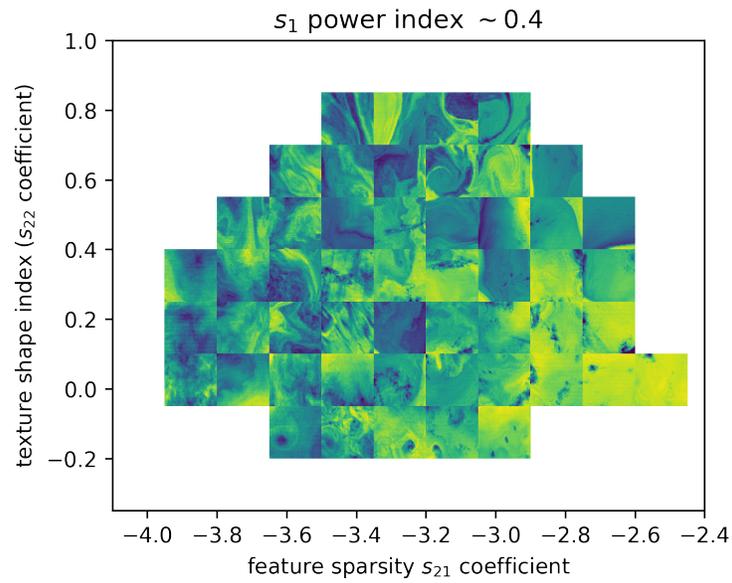


Figure 4-6. Same as the previous figure, but with the s_1 power index between 0.35 and 0.45.

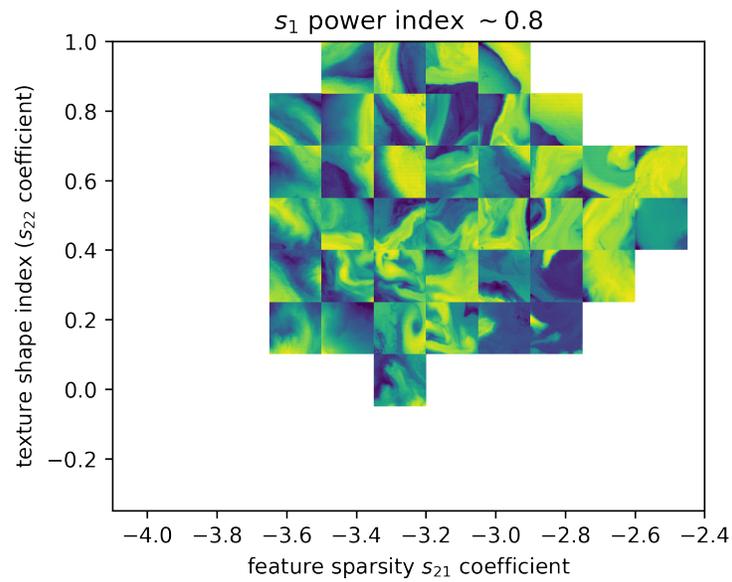
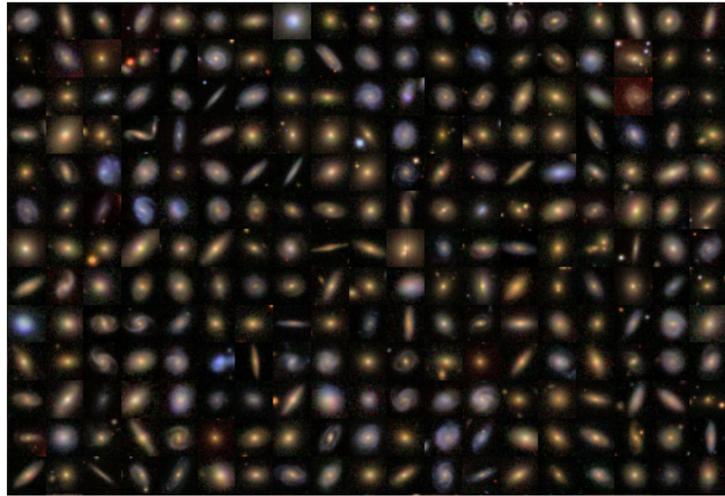


Figure 4-7. Same as the previous figure, but with the s_1 power index between 0.75 and 0.85.

recognition. The early work of morphological classification of galaxies dates back to about a century ago, when Hubble [114] visually inspected 400 galaxies and proposed the famous ‘Hubble sequence’ classification, which is still used today. For a long time, galaxy classification is performed more or less unchanged in this laborious and somewhat subjective manner. Later in 1990s, as large extra-galactic surveys started to revolutionize astronomy by creating exploding amount of data, computer vision techniques were introduced to automatize galaxy classification, but not at the level of human accuracy. In recent years, following the success of convolutional neural networks in computer vision, deep learning was also introduced to this area [115].

However, there are a number of issues with using deep neural networks, e.g., they are prone to over-fitting and usually require a large training set. On the other hand, it is interesting to note that many heavy, deep networks including a 50-layer residual network (ResNet-50) and a 42-layer Inception V3 [116] have almost the same performance as a 7-layer network of Dieleman *et al.* [115] with built-in rotation invariance. It shows how adding physically-motivated properties can simplify a network and save computation. It also suggests that the intrinsic complexity of galaxy morphology may be much lower than those only solvable by deep CNNs and motivates the use of the scattering transform.

To probe the potential of scattering transform in galaxy morphology classification, we randomly selected 100,000 images from the database of the Sloan Digital Sky Survey, some of which are shown in the upper panel of Figure 4-8. Similar to the application to sea surface temperature, we use the scattering coefficients of galaxy images as a low-dimension representation of the galaxies. In the lower panel of Figure 4-8, we arrange these galaxies based on their 2nd-order scattering coefficients. Slightly different from the sea surface temperature application, we use an algorithm called UMAP [Uniform Manifold Approximation and Projection 117] to further compress the coefficients into 2 dimensions. Interestingly, galaxies are automatically separated and clustered according to their morphological properties, as the UMAP space also inherits the morphological meanings and interpretability of the scattering coefficients. Different from CNNs, there is no training of a single parameter involved in the scattering transform. For obtaining this morphological arrangement alone, we do not even need labels from human classification. Therefore, by using the scattering transform, we can bypass CNNs’ heavy heavy computational requirement and the reliance on large training sets.



arranged by scattering coefficients

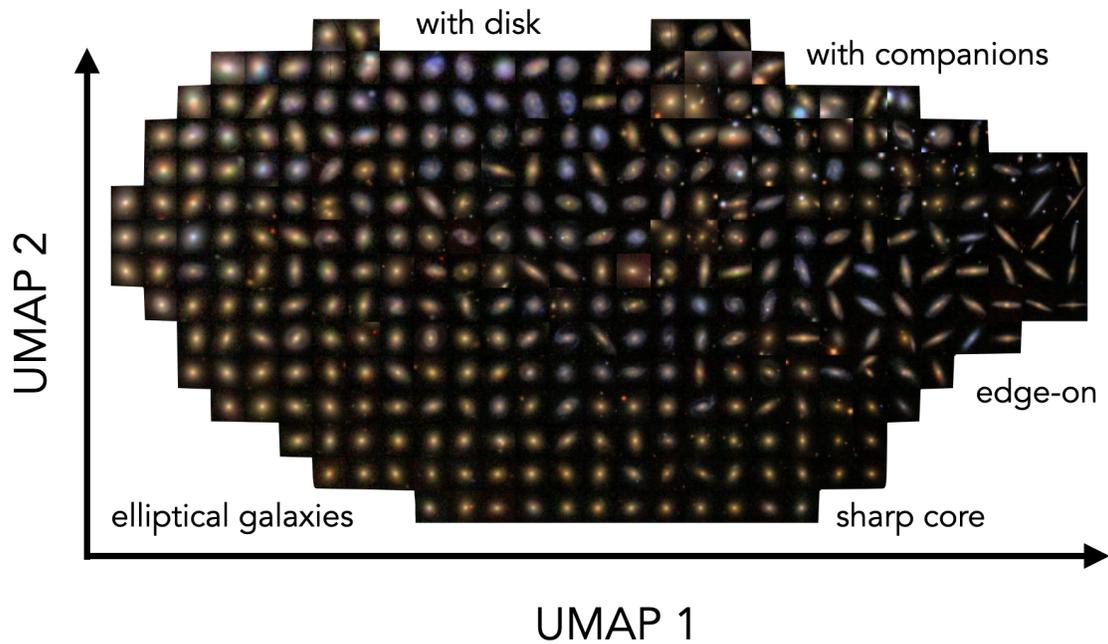


Figure 4-8. A gallery of randomly selected galaxy images (left) and a set arranged based on the 2nd-order scattering coefficients (right). The x- and y-axes of the right panel represent the first two UMAP projection of the coefficients, which is similar to a non-linear version of the principal component analysis (PCA).

Chapter 5

Summary

The complexity of the stochastic fields typically studied by physicists tends to lie between simple Gaussian random fields and the elaborate systems found in the biological world whose complexity appears boundless. This intermediate regime calls for an appropriate data analysis tool which goes beyond a simple power spectrum but does not necessarily motivates the use of deep learning. Deep learning is certainly a powerful method to extract information from complex data, but it comes at a high cost: it requires large amounts of data and computation for training, designing deep convolutional neural networks is a tedious task mainly done empirically, and the networks after training does not have provable properties or guarantee of regularity.

When extracting information from data, it seems desirable to use a tool or an estimator whose expressivity is on par with what is present in the data. In this thesis, I advocate for the use of the *scattering transform* to extract information from physical fields. This estimator, introduced by Mallat [1] and Bruna & Mallat [3] in the mathematics and signal processing literature, provides an approach to data analysis that, in many ways, conveniently stands *in between* the power spectrum and CNNs. Interestingly, it possesses the required expressivity to capture the statistical and morphological properties of a wide range of physical fields.

I aim at motivating this approach to the physics community interested not only in classification tasks but in the full spectrum of data analysis problems, from exploratory data analysis all the way to precise regression problems when models are available. I thus present the formalism of the scattering transform by introducing it from the power spectrum, which most physicists are familiar with, and only focusing on the aspects of potential interest to physicists. After having demonstrated the power of this estimator with a number of synthesis examples, I explain its design and discuss the role of each operation involved in the transform, from both real and Fourier space points of view.

In brief, the scattering transform is a well-defined estimator and a new way extend the power spectrum. It defines a mathematically-motivated network architecture in contrast to classical deep convolutional neural networks. It is similar to a shallow CNN with pre-defined kernels. The adoption of wavelet filters allows a scattering network to have explicit and physically-motivated properties. With pre-defined kernels, the estimator can be applied without any training, making it directly usable for a range of tasks: regression, classification, and exploratory analysis. It enjoys a number of attractive properties for the analysis of physical fields, especially in the context of stationary ergodic fields:

- like the power spectrum, it is invariant to translation (and possibly rotation) and it preserves energy.

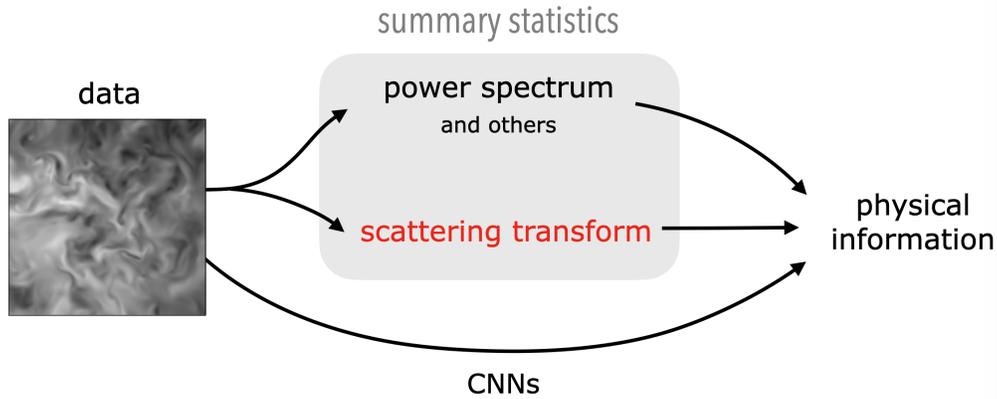


Figure 5-1. Different approaches to extract information from a field. Summary statistics can serve as the representation, while convolutional neural networks (CNNs) usually hide such representation inside. The scattering transform borrows the ideas of local convolutions, ‘low-order’ operations, and the hierarchical design, which have proven powerful in CNNs. But, it belongs to the statistical approach and needs no training.

- in contrast to the power spectrum, it is capable of extracting non-Gaussian or morphological information. Also, in many cases, a scattering transform up to the second order is sufficient to extract the relevant morphological information.
- It produces a compact set of informative coefficients or summary statistics which are stable and robust. These statistics also Gaussianize quickly according to the central limit theorem.
- Last but not least, its coefficients are interpretable. In particular, we showed how two sets of reduced scattering coefficients are particularly informative:
 - a feature sparsity index $s_{21}(j_1, j_2)$ which indicates whether fluctuations or structures are concentrated at certain positions or widely spread.
 - a shape index $s_{22}(j_1, j_2)$ (defined for two-dimensional fields) which describes the level of straight versus curved features present in the data.

Using real data from various areas in physics, I showed how useful the scattering coefficients can be in representing data in a lower-dimensional space, for applications from exploratory data analysis to parameter estimation. In particular, I explored in depth the application of scattering transform in cosmology. After development over nearly a century, modern cosmology has entered an era of high precision, with parametrized models and a large amount of data. Accordingly, a large fraction of its focus has been paid to inference of cosmological parameters in the model, which is essential for both understanding the nature of dark energy and dark matter, which have been parametrized in current models, and discovery of possible new physics not yet included. I argue that cosmology is a perfect area where the scattering transform can be applied. In particular, I showed that in the weak lensing context, the scattering transform extracts information as much as state-of-the-art CNNs, while it can be used just as an ordinary statistical estimator: not only can we avoid the need for training a huge neural network, but also we are guaranteed with provable mathematical properties and a mathematical understanding of where the constraints and information come from. Among different choices of non-Gaussian statistics, the scattering transform is especially

informative, compact, and robust. These properties are crucial for the simulation-based forward modelling approach, which is almost unavoidable in cosmology as we enter the highly non-linear regime with the increasingly deeper surveys.

Beyond cosmology, I also showed other applications of the scattering transform to physical fields, such as the temperature fields in oceanography and galaxy images. They are just two examples of a long list of interesting fields where the scattering transform can benefit relevant research, as non-Gaussian fields are ubiquitous in physics. In recent years, we have observed a fashion of applying heavy CNNs in research of complex data. Indeed, when the data complexity exceeds what simple tools can reach, we have only limited options to go forward, and it is tempting to turn to CNNs given their success in computer vision. However, there is also a growing danger of blindly using, if not abusing, CNNs. because the goal and situation in computer vision are usually not the same as that in scientific research. For example, scientists care a lot about understanding, while neural networks usually provide little understanding of where the information comes from. Also, not every area in scientific research has the luxury of building large training set as required by deep networks. The scattering transform may significantly improve such a situation by providing a new powerful tool to characterize complex structures, and, as shown in my thesis, in many such applications the scattering transform can reach the level of CNNs, because the complexity of physical field is usually not high enough to require a specific vocabulary to describe them. Perhaps the scattering transform cannot classify two sub-types of rabbits well, but the physical fields we usually meet are much simpler. Therefore, I believe wide variety of studies, including cosmology, interstellar medium, asteroseismology, galactic dynamics, and studies beyond astrophysics, will benefit from the adoption of scattering transform.

We are living in a period of revolution. Computational power has been growing exponentially for decades. So does the amount of data that we create and analyze, in almost all scientific fields and in human activities as a whole. And, ten years ago, we experienced the revolution of the convolutional neural net. In such a period, fashion, belief, or even value, are often changing. What most of us are familiar with today may not prevail in the future. It is interesting to speculate for a while. How would we do science in the future? In ten years, are we going to continue focusing on higher-order statistics? Perhaps not as much, because for many cases we have a much better tool now. Will neural networks dominate and become the optimal choice for most topics of scientific research? Perhaps not. It seems to me neural networks represent a different paradigm. They are well suited for efficiently solving many engineering problems, but hardly provide understanding. To me, they are like well-trained dogs good at doing a particular job, whereas mathematical tools used in science, including the scattering transform, are the language we use to think and to communicate our ideas. However, neural network is indeed an interesting object of scientific research, helping us to understand more about ‘learning’. What about the scattering transform? Will it in ten years become as popular as the power spectrum today? Perhaps yes. To me, the scattering transform has all the good reasons to be used in scientific research as a generic mathematical language, just as the power spectrum. It is powerful, interpretable, and easy to use. However, when I started my Ph.D., the scattering transform was not even known to the astronomy and cosmology community. After my work and others, it has gradually collected attentions and attracted interests in the community. And now, several teams have started working on it. But, it is still far from being ‘known’, and a large fraction of the community still has to choose between the power spectrum and a black box when encountering a somewhat complex dataset. For example, to characterize the fluctuation texture of galaxy surface brightness, extra-galactic astronomers redirect a deep neural network trained for recognizing balloons and traffic lights by cutting off the final layers of the network. This application

of CNN is conceptually identical to using the scattering coefficients, except that the scattering transform is defined in a much transparent way. Given all these arguments, it is not unreasonable to imagine the scattering transform will become widely adopted in the future.

There are also several interesting open questions about the theoretical aspect of the scattering transform, in the context of physics. For example, in physics, we care about analytical predictions. Is it possible to derive analytical predictions for particular physical fields such as the cosmic density field under gravitational interaction or the turbulence field? It might be challenging because the modulus, which is the key of robustness in practical analysis of observational data, is unfamiliar to most physicists. Nevertheless, it is not impossible. Another question unanswered to me is the deeper reason behind the informativeness of the scattering transform. This is apparently related to the sparse representation of wavelet transform, but to me, there is still a missing connection to the equation of physical fields. Viewed from another direction, the scattering transform is an aggressive data compression from pixels to a few coefficients. The decrease of dimensionality here is much more than what geometric invariances require. So, there should be other principles guiding the information concentration performed by the scattering transform. It seems most of the compression is applied to high-frequency modes, due to the wavelet binning of high-frequency information. This logarithmic binning of scale is motivated by some scale-invariance argument about information distribution in a field. To what extent can we connect this to the equations governing the physical field (including daily-life images in a broad sense)? I believe the endeavors to address these questions will bring interesting improvement in our understanding for both physics and computer vision theory.

Appendix I: Software implementation of the scattering transform

In both python and Matlab, there have been packages implementing the scattering transform. Especially, the “Kymatio” package (<https://www.kymatio.io>) has been developed to take advantage of GPU acceleration. For a wide variety of applications of the scattering transform, the “Kymatio” package is a good choice.

In the implementation, the scattering transform uses fast Fourier transform (FFT) to perform convolution, whose computation complexity is only $O(N \log N)$ with the number of pixel N . The number of FFT required is roughly equal to that of the scattering coefficients (around $J^2 L^2$ to the 2nd order). With GPU, the computing time often becomes negligible compared to the data transfer time.

Despite the convenience of the Kymatio software, it has several limitations. So, we wrote our own code for the scattering transform as made it a publicly available package with example usage and instructions (https://github.com/SihaoCheng/scattering_transform), which also support GPU implementation. Compared to the Kymatio package, our package have the following advantages: it can deal with 2D images with arbitrary size, in addition to base-two side lengths (32, 64, 128, etc. pixels); for the global (translation-invariant) scattering coefficients, which are more relevant for physicists, the calculation speed is much faster (depending on the size of data; about a factor of 5 for 256×256 images); it provides the option to change kernel profiles; and, the code is based purely on numpy and pytorch and is written in a simpler way, easier to modify and to customize. Other differences between the two packages can be found on my web page. We believe these packages can serve as a convenient tool for many physicists.

Appendix II: Merger products among white dwarfs

This appendix is a copy of my paper [118] published in *Astrophysical Journal*.

Abstract

Double white dwarf (double-WD) binaries may merge within a Hubble time and produce high-mass WDs. Compared to other high-mass WDs, the double-WD merger products have higher velocity dispersion because they are older. With the power of *Gaia* data, we show strong evidence for double-WD merger products among high-mass WDs by analyzing the transverse-velocity distribution of more than a thousand high-mass WDs ($0.8\text{--}1.3 M_{\odot}$). We estimate that the fraction of double-WD merger products in our sample is about 20%. We also obtain a precise double-WD merger rate and its mass dependence. Our merger rate estimates are close to binary population synthesis results and support the idea that double-WD mergers may contribute to a significant fraction of type Ia supernovae.

Introduction

During the last few decades, there has been increasing evidence showing that a large number of double white dwarf (double-WD) systems should merge within a Hubble time [e.g., 119–125]. Many double-WD mergers are believed to produce new white dwarfs with higher masses [e.g., 126]. So, a fraction of high-mass white dwarfs in the solar neighborhood are expected to be double-WD merger products [e.g., 127, 128]. To verify the existence of these merger products, some investigators have looked for an excess of high-mass white dwarfs [129–131], and others have searched for kinematic signatures of these merger products [132, 133]. The kinematic method makes use of the following facts: high-mass double-WD merger products are in general older than singly-evolved white dwarfs because of their binary evolution, and according to the age–velocity–dispersion relation (AVR) of the Milky-Way disc [e.g., 134], these older double-WD merger products have higher velocity dispersion. The former method, based on number counts, is influenced by large systematic errors from the adopted initial–final-mass relation of white dwarfs and the sample completeness. In contrast, the kinematic method is less influenced by systematic errors, but it was limited by the sample size of white dwarfs with kinematic measurements.

Thanks to the European Space Agency *Gaia* mission [135], the number of stars with precise kinematic measurements has been enlarged drastically. Cheng *et al.* [136] selected a deep, homogeneous sample of white dwarfs in a narrow mass range ($1.08\text{--}1.23 M_{\odot}$) from *Gaia* Data Release 2

[DR2; 137] to investigate the ‘Q branch’, an overdensity of white dwarfs on the Hertzsprung–Russell (H–R) diagram, which is caused by a cooling anomaly. As a byproduct of their kinematic analysis, the fraction of double-WD merger products among white dwarfs in their mass range were inferred to be about 22%, and they reserved the task of conducting an analysis optimized for detecting double-WD merger products and the discussion on this topic to this paper.

In this paper, we extend the kinematic analysis of high-mass white dwarfs to a wider mass range and adopt a more realistic delay-time distribution for binary evolution. We estimate the fractions of double-WD merger products as a function of mass and calculate the corresponding merger rates. We then compare our results to predictions from binary population synthesis. We also discuss the implication of our results for the progenitor problem of type Ia supernovae (SNe Ia), as the double-WD merger is a promising scenario of type Ia supernova explosions [e.g., 138, 139].

Data

Gaia DR2 provides accurate astrometric [140] and photometric [141, 142] measurements for more than one billion of stars. To search for the kinematic signature of double-WD merger products efficiently, we select nearby, high-mass, hot white dwarfs with precise astrometric and photometric measurements from the *Gaia* DR2 white dwarf catalog compiled by Gentile Fusillo *et al.* [143]. Below, we introduce in detail our sample selection and the derivation of white dwarf parameters.

We first impose the same quality cuts as equations (1)–(5) in Cheng *et al.* [136] and a distance cut

$$1/\varpi < 250 \text{ pc} \tag{V.1}$$

to select white dwarfs with high-precision astrometric and photometric measurements. These cuts do *not* introduce any *explicit* kinematic biases to our white dwarf sample.

Then, as the kinematic signature of double-WD merger products most outstanding among high-mass, hot (young-photometric-age) white dwarfs, we carry out selections on the photometric mass (m_{WD}) and age (τ_{phot}) of white dwarfs. These cuts, equivalent to cuts on the H–R diagram, are designed to both maximize the sample size and minimize contamination from the standard-mass helium-atmosphere white dwarfs [the ‘B branch’, 144]:

$$\begin{aligned} 0.8 M_{\odot} < m_{\text{WD}} < 0.9 M_{\odot}, \\ 0.1 \text{ Gyr} < \tau_{\text{phot}} < 0.7 \text{ Gyr}; \\ \text{or} \\ 0.9 M_{\odot} < m_{\text{WD}} < 1.28 M_{\odot}, \\ 0.1 \text{ Gyr} < \tau_{\text{phot}} < 1.0 \text{ Gyr}. \end{aligned} \tag{V.2}$$

The white dwarf parameters m_{WD} and τ_{phot} are derived in the following way. First, we define the absolute magnitude M_G as $M_G = G + 5 \log(\varpi/\text{mas}^{-1}) - 10$, where G and ϖ are the G -band magnitude and parallax. Then, we convert the H–R diagram coordinate into m_{WD} and cooling time t_{cool} by interpolating a grid of cooling tracks for C/O-core DA (hydrogen atmosphere) white dwarfs [145] and synthetic colors [146–148].¹² For white dwarfs heavier than $1.07 M_{\odot}$ we use cooling tracks

¹<http://www.astro.umontreal.ca/~bergeron/CoolingModels/>.

²We made a python 3 package for this kind of transformation publicly available on https://github.com/SihaoCheng/WD_models.

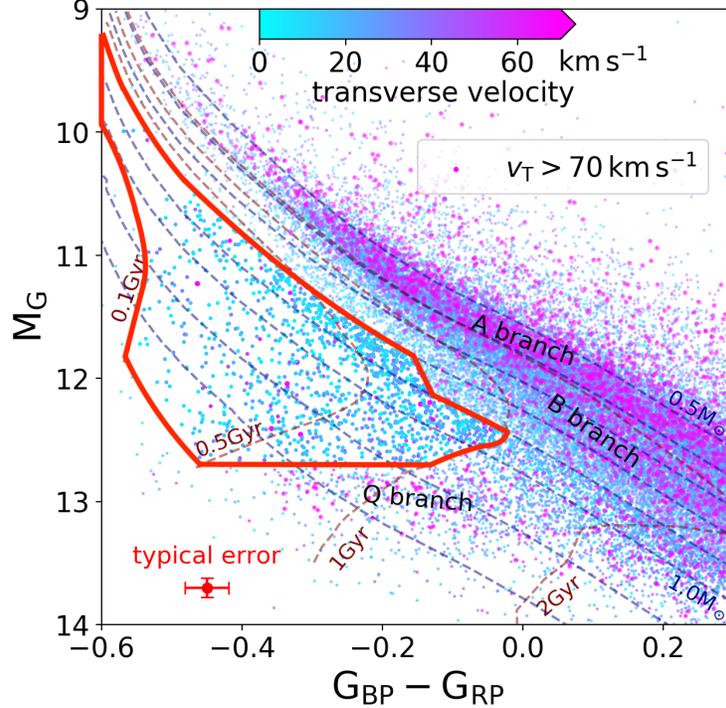


Figure V-2. H–R diagram of WDs in *Gaia* DR2. We show the 250 pc sample of WDs with high-quality measurements and a grid of WD masses m_{WD} and photometric ages τ_{phot} derived from the combined O/Ne- and C/O-core WD cooling model. WDs evolve along their cooling tracks, i.e., the constant-mass curves. The red region includes 1395 nearby, high-mass, hot WDs selected in Section 5.

of O/Ne white dwarfs [149]. Finally, the photometric age τ_{phot} is obtained by adding the cooling time t_{cool} to the main-sequence age, which we calculate based on an initial–final mass relation [150] and the relation between pre-cooling time and main-sequence mass from Choi *et al.* [151] for non-rotating, solar-metallicity stars.

As shown by Cheng *et al.* [136], the ‘Q branch’ on the H–R diagram is produced by an anomalous cooling behavior: some white dwarfs stop cooling and stay on the branch for several billion years, which creates both an overdensity and a high-velocity excess. To avoid modelling the influence of this cooling delay on the velocity distribution and only focus on the binary-evolution delay for double-WD mergers, we exclude the ‘Q branch’ region on the H–R diagram by the cut

$$M_G < 12.7. \quad (\text{V.3})$$

We also apply a color cut at the blue end to control the uncertainty of photometric mass and age determination

$$G_{\text{BP}} - G_{\text{RP}} > -0.6. \quad (\text{V.4})$$

The selection region on the H–R diagram and 1395 selected white dwarfs³ are shown in Figure V-2.

We divide our selected sample into five mass bins, based on the aforementioned photometric mass (assuming C/O-core below $1.07 M_{\odot}$ and O/Ne-core above it). The edges of bins are 0.8, 0.9,

³A catalog of all selected white dwarfs is available on Vizier and on the website: <https://sihaocheng.github.io/DWDmerger>.

1.0, 1.1, 1.2, and 1.28 M_{\odot} . If some white dwarfs heavier than 1.07 M_{\odot} are believed to still hold C/O cores instead of O/Ne cores, such as massive double-WD merger products [e.g., 152], then for those white dwarfs the mass bins correspond to 0.8, 0.9, 1.0, 1.14, 1.24, and 1.32 M_{\odot} . The sample sizes in these mass bins are 408, 431, 323, 176, and 57, respectively. Because of the absolute-magnitude cut and blue color limit, the photometric-age ranges for the five samples are different. We estimate them to be 0.42, 0.82, 0.86, 0.66, and 0.42 Gyr, respectively.

We estimate the completeness of our sample with the completeness–magnitude relation, $c(G)$, derived in Sollima [153] by randomly selecting 180 regions in the sky and comparing *Gaia* DR2 and PanSTARRS DR2 catalogs. His estimate is similar to that given by comparing *Gaia* DR2 and *Hubble Space Telescope* images around globular clusters [154]. So, we adopt the relation in Figure 1 of Sollima [153] and calculate for each mass bin the average completeness, $1/c^{-1}(G_i)$, where G_i represents the G -band magnitude of a single star. The resulting completeness is 88%, 80%, 79%, 79%, and 80%, respectively. In addition, we find that our quality cuts only exclude less than 5% of the objects, given the distance and H–R diagram cuts in Equations V.1–V.4.

Finally, we derive the kinematics of white dwarfs, which are related to the true ages of white dwarfs through the age–velocity–dispersion relation. Because *Gaia* does not provide radial velocity information for white dwarfs due to the narrow wavelength coverage of its spectrometer [135], we focus on the two components of transverse velocity $\mathbf{v}_T = (v_l, v_b)$:

$$v_l = \frac{\mu_l - (A \cos 2l + B) \cos b}{\varpi}, \quad (\text{V.5})$$

$$v_b = \frac{\mu_b + A \sin 2l \sin b \cos b}{\varpi}, \quad (\text{V.6})$$

where μ_l and μ_b are the proper motion in the Galactic longitude and latitude directions, and A and B are the Oort constants taken from Bovy [155].

Model

Our goal is to measure the amount of double-WD merger products among high-mass white dwarfs using the kinematic information. According to the age–velocity–dispersion relation, a group of stars with older true stellar age (τ) has higher velocity dispersion. On the other hand, one can derive the photometric isochrone age (τ_{phot}) of white dwarfs from the H–R diagram. If a white dwarf evolves in isolation, τ_{phot} should be equal to τ , whereas if it originates from a double-WD merger event, then an age discrepancy,

$$\Delta t \equiv \tau - \tau_{\text{phot}}, \quad (\text{V.7})$$

will be created from binary evolution. In general, a white dwarf produced from binary evolution may have positive or negative Δt , but for double-WD mergers with high total mass, the discrepancy is almost always positive [128]. So, for a given τ_{phot} , double-WD merger products are older and have higher velocity dispersion.

Following Wegg & Phinney [132] and Cheng *et al.* [136], we assume that the double-WD merger ‘resets’ the white dwarf back to a sufficiently high temperature, so that the real cooling time is equal to the photometric cooling time. Then, Δt can also be expressed as the difference of pre-cooling times between the two evolutionary scenarios, $\Delta t = (\tau - t_{\text{cool}}) - (\tau_{\text{phot}} - t_{\text{cool}})$, where t_{cool} is the cooling time, and the first item $\tau - t_{\text{cool}}$ is sometimes called the ‘delay time’ of double-WD merger. It has been widely used that the distribution of the age discrepancy, $p(\Delta t)$, for double-WD

Table V-I. Measurements of the Double-WD (DWD) Merger Rate and Type Ia Supernovae (SNe Ia) Rate

Reference	Event Rate ($10^{-13} M_{\odot}^{-1} \text{yr}^{-1}$)	Event Type	Based on
This work	0.20 ± 0.06	DWD mergers with 0.8–0.9 M_{\odot}	Merger products among single WDs
	0.23 ± 0.06	DWD mergers with 0.9–1.0 M_{\odot}	
	0.32 ± 0.06	DWD mergers with 1.0–1.14 M_{\odot}	
	0.28 ± 0.07	DWD mergers with 1.14–1.24 M_{\odot}	
	0.07 ± 0.04	DWD mergers with 1.24–1.32 M_{\odot}	
This work (total)	1.1 ± 0.3	DWD mergers with 0.8–1.32 M_{\odot}	same as above
Maoz <i>et al.</i> [124]	6.3 ± 1.0	All DWD mergers	DWD orbits
Brown <i>et al.</i> [125]	0.3 ± 0.2	DWD mergers with ELM WD	DWD orbits
Li <i>et al.</i> [159]	1.0 ± 0.3	SNe Ia in Milky-Way-like galaxies	Extra-galactic SNe

mergers with high total masses is approximately a power law, i.e., $p(\Delta t) \approx \Delta t^{-1}$ [156], because the binary delay time $\tau - t_{\text{cool}}$ is dominated by the double-WD phase when the orbit shrinks due to gravitational-wave emission, and the single-star pre-cooling time $\tau_{\text{phot}} - t_{\text{cool}}$ is negligible. However, in our mass ranges, none of the two statements are valid. So, we use more realistic distributions for Δt , with the binary delay-time distribution, $p(\tau - t_{\text{cool}})$, obtained from binary population synthesis (see 5 for details) and the values of $\tau_{\text{phot}} - t_{\text{cool}}$ from Section 5.

We consider our white dwarf sample as a mixture of two populations: singly-evolved white dwarfs and double-WD merger products⁴, with fractions $1 - f_m$ and f_m , respectively. If f_m is higher, the tail of the velocity distribution will also be higher, because the double-WD merger products are on average older. For the velocity distribution, we assume that stars with the same true age τ have a Gaussian velocity distribution $\mathbf{v} \sim \mathcal{N}(\mathbf{v}_0(\tau), \Sigma(\tau))$ relative to the Sun [e.g., 157]. The size of this Gaussian distribution is determined by τ through the age–velocity–dispersion relation, and the center of this Gaussian distribution is determined by the solar motion and the age-dependent asymmetric drift $(\sigma_U(\tau)^2/80 \text{ km s}^{-1})$.

We use the same Bayesian framework as constructed by Cheng *et al.* [136] to infer the fraction f_m from the photometric age τ_{phot} and transverse velocity \mathbf{v}_T of each white dwarf. This model uses the conditional probability of \mathbf{v}_T given τ_{phot} and other observables as the likelihood function, and thus it eliminates spatial selection biases. In the model, we set f_m and the solar motion as free parameters and adopt the best-fitting age–velocity–dispersion relation in Cheng *et al.* [136], a flat star-formation history in our sample volume, and the delay-time distribution of double-WD mergers shown in the 5. We do not need to model the ‘extra cooling delay’ included in Cheng *et al.* [136] because this delay has no effect in our selected region. We do not include white dwarf kick effects in our model, because for single-evolved white dwarfs, the kick velocity during the white dwarf formation is less than 1 km s^{-1} [158], and for double-WD mergers, the kick velocity during merger is a few km s^{-1} [e.g., 152], which have only a tiny contribution to the increase of velocity dispersion compared to the contribution from the binary-evolution delay.

⁴The high-mass white dwarfs originating from other types of mergers such as main-sequence and giant star mergers have much shorter age discrepancy Δt than that of double-WD mergers. So, in terms of kinematics, we treat the merger products of other types the same as singly-evolved white dwarfs.

The binary population synthesis

Here we describe the binary population synthesis that we use in this paper to derive the delay-time distribution. The models are synthesized using the binary population synthesis code `SeBa` [160, 161]. The models are identical to the default models used in [127]. We have adopted a Kroupa initial mass function [162] and a uniform mass ratio distribution between 0 and 1 [163, 164]. Furthermore, we assume a uniform distribution in the logarithmic semi-major axis up to $10^6 R_\odot$ [165], and a thermal distribution of eccentricities between 0 and 1 [166].

One of the main sources of uncertainty in the synthetic populations [167] is a phase of unstable mass transfer, i.e., the common-envelope (CE) phase [for a review, see 168]. Similar to [127], we apply the ‘ $\gamma\alpha$ ’ model. This model reproduces the mass ratio distribution [161] and number density [127] of double-WD systems best. In the ‘ $\gamma\alpha$ ’ model, we apply the classical (α -)CE that is based on energy conservation [139], and the (γ -)modelling that is based on a balance of angular momentum [169]. Regarding the former the parameters $\alpha\lambda = 2$ describe how efficient orbital energy can be used to unbind the envelope and how strong the envelope is bound to the donor star, and regarding the latter the parameter $\gamma = 1.75$ describes the efficiency of angular momentum usage. The γ -modelling is applied unless the binary contains a compact object or the CE is triggered by a tidal instability. We note that for our purpose to compare the merger rate, the delay-time distribution of the ‘ $\gamma\alpha$ ’ model does not significantly differ from that of the model that exclusively adopts the α -CE with $\alpha\lambda = 2$ [see 161].

Figure V-3 shows the delay-time distributions in five mass bins, which are used in our kinematic analysis. For the synthesized merger rates shown in Figure V-6, we in addition assume a 50% binary fraction of all stars [see also 164, 170].

Results and Discussions

Constraints on the fraction of merger products

With a thirty-time larger sample of high-mass white dwarfs selected from *Gaia* DR2, we are allowed to go beyond Wegg & Phinney [132] and set strong constraints on f_m . Figure V-4 illustrates the transverse-velocity distribution of white dwarfs in our sample. The clear velocity excess is strong evidence for the existence of double-WD merger products. For clarity we only show the distribution of the whole sample, i.e., the combination of all five mass bins, but similar results can be found in each single mass bin, too. In Figure V-5, we show our estimate of f_m in each mass bin. We find that the fraction of double-WD mergers in our mass range of 0.8–1.3 M_\odot varies from 10% to 35%, with an average of about 20%. This fraction is roughly constant as a function of mass, though declines at the two end are suggested.

To test the robustness of our results, we check for the influence of sample selection, the adopted star-formation history, and the adopted age–velocity–dispersion relation in our model. We found that a different distance cut such as 200 or 300 pc cuts and a linearly decreasing star-formation history with 5 times higher rate at 11 Gyrs ago have less than 20% *fractional* influence on the estimate of f_m . For the influence of the age–velocity–dispersion relation, our results are mostly influenced by the 0–4 Gyr part, where the delay-time distribution is peaked. Adopting the high velocity dispersion from Just & Jahreiß [171] as used by Wegg & Phinney [132] will reduce f_m by a factor of 2, but given the observational constraints from both main-sequence stars [e.g., 134]

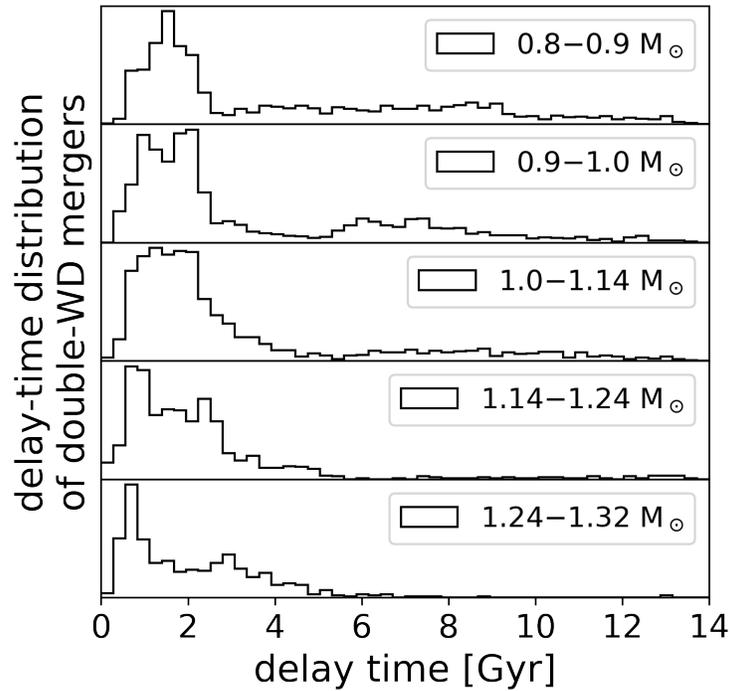


Figure V-3. Delay-time distributions of double-WD mergers used in our model. These distributions are generated from binary population synthesis. The x -axis is the delay time of binary evolution, i.e., $\tau - t_{\text{cool}}$ for the resulting white dwarf. The y -axis is in linear scale and normalized to their maximum values. We input to our model the shapes of these five distributions as probability distribution and do not use the information from their normalization.

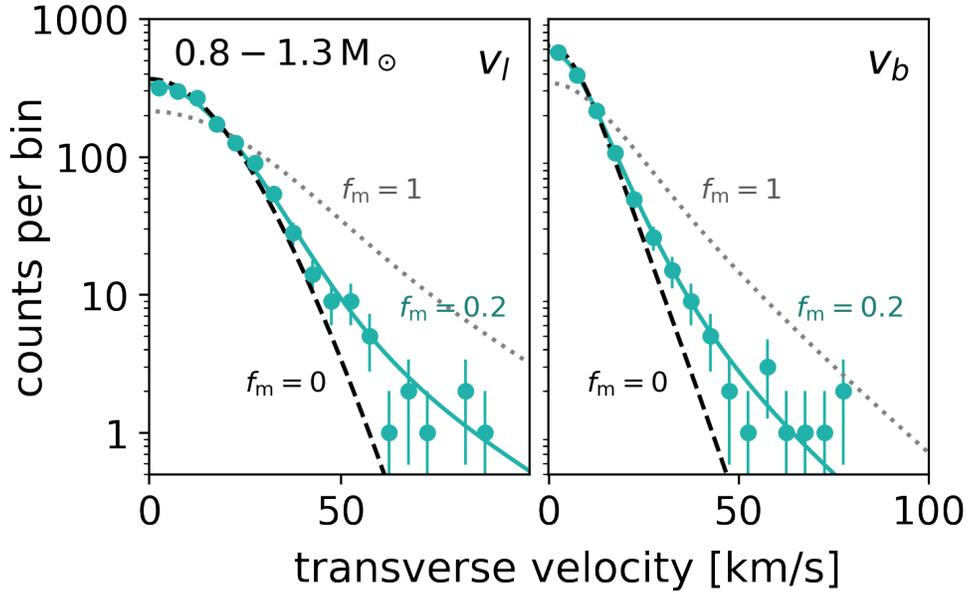


Figure V-4. Velocity distribution of our white dwarf sample. We show the sample of white dwarfs from all five mass bins ($0.8\text{--}1.3 M_{\odot}$) as an example. v_l and v_b in the left and right panel of the figure means the Galactic longitude and latitude components of the transverse velocity. We present the observed histograms of the absolute values of v_l and v_b in 20 bins between 0 and 100 km s^{-1} and Poisson errors. We also show the theoretical velocity distributions for $f_m = 0, 1$, and the average of best-fitting values weighted by the sample size in each mass bin, which is about 0.2. Note that the y -axes are in logarithmic scale.

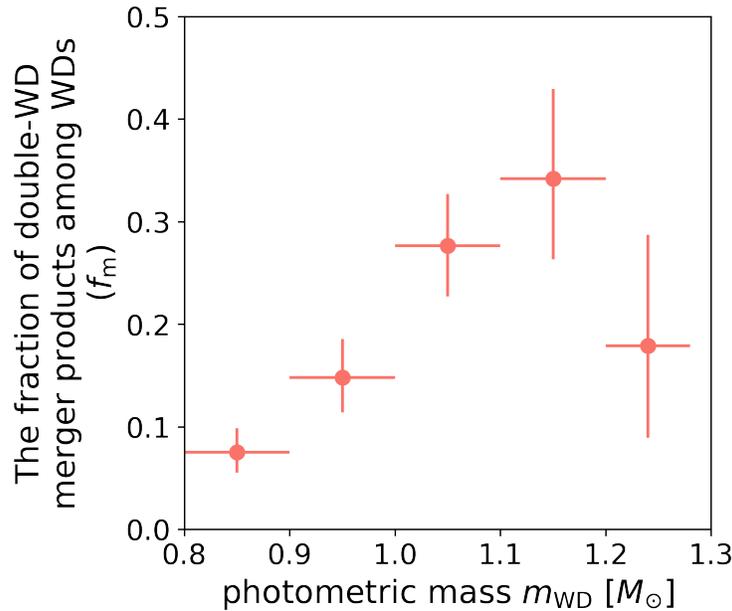


Figure V-5. Our estimates for the fraction of double-WD merger products among high-mass white dwarfs, in five bins of white dwarf photometric mass. The sample sizes in these mass bins, from lower to higher masses, are 408, 431, 323, 176, and 57, respectively.

and white dwarfs [136], such high values of velocity dispersion are unlikely. The effect of adopted delay-time distribution can be seen from the comparison between our results in the mass range of 1.08–1.23 M_\odot and that of Cheng *et al.* [136]: adopting a power-law delay-time distribution leads to a result about 30% lower. So, we estimate the *fractional* systematic error of our results as 30% (a factor of 0.7–1.3). Our estimate of f_m is consistent with population synthesis results [128].

Double-WD merger rate

The fraction of double-WD merger products (f_m) obtained in Section 5 can be translated into double-WD merger rates. Because our sample is nearly volume limited, the merger rate in each mass bin can be estimated by

$$\text{merger rate} = \frac{f_m \cdot N}{m_\star \cdot \Delta\tau_{\text{phot}} \cdot c}, \quad (\text{V.8})$$

where N is the sample size of each mass bin (listed in the caption of Figure V-5), m_\star the stellar mass of the Milky Way within 250 pc. $\Delta\tau_{\text{phot}}$ and c are the photometric-age range and sample completeness of each mass bin, which are estimated in Section 5. The stellar mass m_\star is estimated to be $4.1 \times 10^6 M_\odot$, using the local stellar mass density $\rho_\star = 0.083 M_\odot \text{pc}^{-3}$ [172] and a scale-height 300 pc of the disc. We list our estimate of the current double-WD merger rate in each mass bin in Table V-I. The total merger rate in our mass range (0.8–1.3 M_\odot) amounts to $1.1 \times 10^{-13} M_\odot^{-1} \text{yr}^{-1}$. To make comparison with other measurements easier, we also divide these values by their corresponding mass ranges and show the results in Figure V-6. Note that the mass range of each bin here is slightly different from that in Figure V-5, because we adopt the photometric masses derived from CO white dwarf models (see Section 5 and Table V-I for details).

There are several factors that may lead us to underestimating the merger rate within our mass range. For example, we take the mass of the merger products as the total mass of the original double-WD binary. This is true for CO-CO white dwarf mergers [e.g., 152] but not for He-CO white dwarf mergers, which may lose significant amount of mass during the R Coronae Borealis phase and produce white dwarfs with 0.6–0.7 M_\odot [173]. So, we are likely to underestimate the merger rate of systems with original mass below 1.0 M_\odot , where He-CO white dwarf mergers become important. Similarly, we will miss explosive and disruptive merger events if there be any such events in our mass range, and events that result in extremely magnetic and faint white dwarfs [174] if there be such objects produced.

In Figure V-6 we also show the merger rate from binary population synthesis, with a flat star-formation history assumed. If a decreasing star-formation rate was assumed, as the delay-time distribution is low for mergers with a long delay time, the synthesis would predict a lower current merger rate than plotted in Figure V-6. Details of the population synthesis are shown in 5. We find that the synthesized merger rates are close to our observational estimates without any tuning of parameters. Note that in our analysis of *Gaia* white dwarfs, we only use the *distribution* of the delay time but never use the total merger rate information from the population synthesis. So, the match between the observed and synthesized merger rate is not a circular argument but rather a validation of our understanding of binary evolution.

Then, we compare our results with other estimates of the double-WD merger rate in the literature. While we count the products of mergers, other estimates are obtained by observing pre-merger systems and predicting the merger rate. Maoz *et al.* [124, 175], Badenes & Maoz [123], and Maoz & Hallakoun [176] extrapolate the orbital distribution of double-WD binaries to estimate

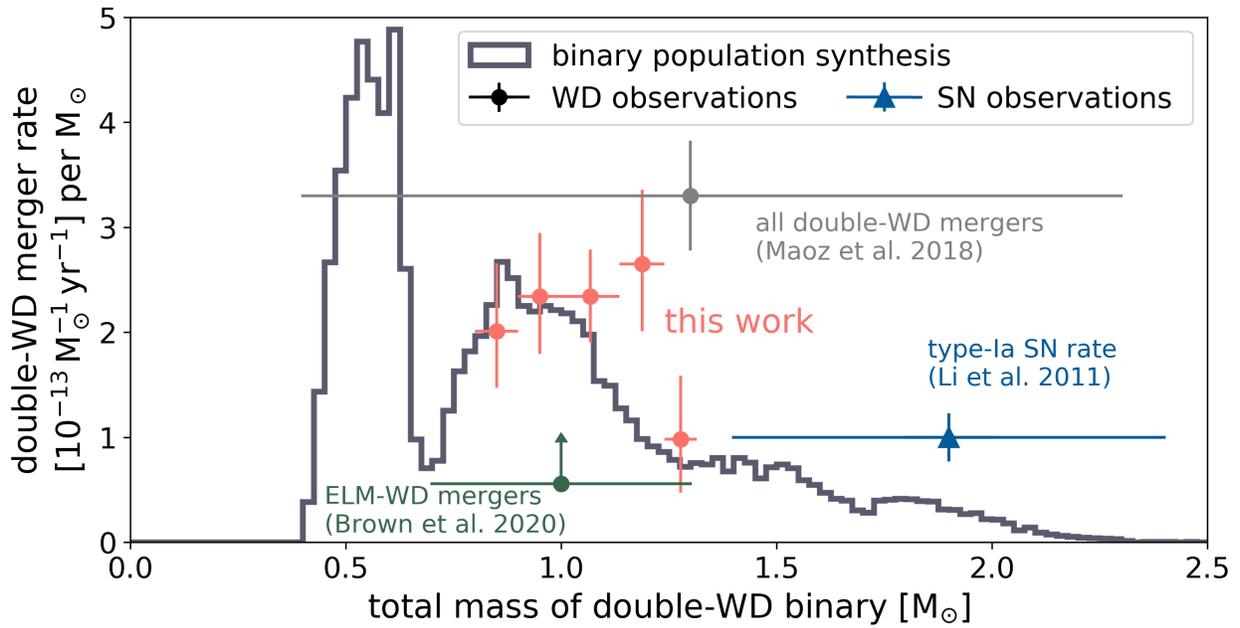


Figure V-6. A comparison of the observed and simulated double-WD merger rate. The red data points with error bars are our observational estimates based on double-WD merger products. The histogram shows binary population synthesis results. Other data points show estimates in the literature based on the orbital distribution of observed double-WD systems: the light-grey one is an estimate for all double-WD mergers [124], and the green one is for systems with at least one extremely low-mass (ELM) WD [125], which provides a lower limit of the merger rate. The blue data point shows the observed SN Ia rate. Comparisons between a data point and the histogram should be made in terms of the area under the horizontal ‘error bar’ of the data point and the area under the histogram in the same mass range.

the total double-WD merger rate, with an up-to-date estimate being $(6.3 \pm 1.0) \times 10^{-13} M_{\odot}^{-1} \text{yr}^{-1}$. Brown *et al.* [125, 177] estimate the merger rate of double-WD binaries with at least one extremely low-mass (ELM; $< 0.3 M_{\odot}$) white dwarf to be $2 \times 10^{-3} \text{yr}^{-1}$ in the Milky Way, corresponding to $0.3 \times 10^{-13} M_{\odot}^{-1} \text{yr}^{-1}$, with a 110% uncertainty including 70% statistical uncertainty. In Table V-I we list these values. In Figure V-6 we assign reasonable mass ranges to these measurements and present the results. For the result from Maoz *et al.* [124], we assign $0.4\text{--}2.3 M_{\odot}$ according to the mass distribution in our binary population synthesis, and for the result from Brown *et al.* [125], we assign $0.7\text{--}1.3 M_{\odot}$ according to the mass distribution of ELM binaries [178]. All data points on Figure V-6 should be understood as the averaged merger rate within the assigned mass range. As these measurements address the merger rates of systems in different mass ranges, one cannot compare them directly. But, if we are allowed to use the mass distribution from binary population synthesis to scale these estimates, we will find that the merger rate obtained by Maoz *et al.* [124] is about 2–3 times of our estimates, and the estimate from Brown *et al.* [125] is consistent with our results, as illustrated in Figure V-6. As discussed in Maoz *et al.* [124], if the merger rates are as high as their estimate, almost all high-mass white dwarfs will need to be double-WD merger products, which is hard to believe given the velocity distribution we observe. Nevertheless, it is noticeable that the observational constraints of the double-WD merger rate from different methods have converged to within a factor of a few.

In summary, our estimates of the double-WD merger rate add significant precision and mass resolution to our knowledge of the double-WD merger rate and provide a validation for current binary population synthesis.

Implication for type Ia supernovae

Type Ia supernovae are important distance indicators, element factories, interstellar medium heaters, and cosmic-ray accelerators, but their progenitors remain unclear [e.g., 179]. The double-WD merger is a promising scenario of type Ia supernova [e.g., 138, 139, 156, 180–186]. The comparison between double-WD merger rate and the type Ia supernova rate is a critical test from this scenario. When a flat star-formation history is assumed, our population synthesis (5) provides a merger rate of about $0.3 \times 10^{-13} M_{\odot}^{-1} \text{yr}^{-1}$ for super-Chandrasekhar double-WD systems, which is about 1/7 of the total synthesized double-WD merger rate and consistent with previous studies [e.g., 187, 188]. For the D⁶ (dynamically driven double-degenerate double-detonation) scenario [e.g., 184], a lower rate is obtained, because it requires in general higher total mass of the system [see figure 2 of 189].

On the other hand, the observed type Ia supernova rate for a Milky-Way-like galaxy (Sb-Sbc type) is $(1.0 \pm 0.3) \times 10^{-13} M_{\odot}^{-1} \text{yr}^{-1}$ [159], or $1.3 \times 10^{-3} M_{\odot}^{-1}$ in terms of a time-integrated rate [190]. This is close to, though 2–3 times higher than, the synthesized rate for the super-Chandrasekhar and D⁶ double-WD merger scenario. As discussed in Section 5, our estimates of the double-WD merger rate within $0.8\text{--}1.3 M_{\odot}$ are in agreement with population synthesis results. So, if we are allowed to extrapolate the merger rate to higher masses according to the mass distribution of mergers from simulations, then

- our measurements support the idea that double-WD mergers can contribute a significant fraction to type Ia supernovae;
- if all type Ia supernovae come from double-WD mergers, it seems that there exist other explosion mechanisms whose requirement on the total mass of the binary is lower than that of the Chandrasekhar- and D⁶ explosion models.

Conclusion

The merger of two white dwarfs in a close binary system may result in a new white dwarf with higher mass. Therefore, among the high-mass white dwarfs observed today, a fraction should come from double-WD mergers. Experiencing binary evolution, these merger products have older true ages than their photometric isochrone ages. According to the age–velocity–dispersion relation in the Milky-Way disc, older stars have higher velocity dispersion. So, the fraction of double-WD merger products (f_m) can be estimated from the velocity distribution of high-mass white dwarfs.

We select a homogeneous sample of high-mass white dwarfs ($0.8\text{--}1.3 M_\odot$, $d < 250$ pc) from *Gaia* DR2, which includes 1395 objects. Our sample is about thirty times larger than that of a previous study with a similar idea [132]. We infer f_m in five mass bins using a Bayesian model of white dwarf transverse velocities. We find

1. about 20% of $0.8\text{--}1.3 M_\odot$ white dwarfs originate from double-WD mergers;
2. the corresponding double-WD merger rates in our mass range add up to $1.1 \times 10^{-13} M_\odot^{-1} \text{yr}^{-1}$.

We show f_m and the merger rate as a function of mass in Figures V-5 and V-6, respectively. We estimate our systematic error to be within 30%, i.e., a factor of 0.7–1.3. Our results are in good agreement with the predictions from binary population synthesis (see 5 for setting details).

Our estimates add significant precision and mass resolution to our knowledge of the double-WD merger rate. If it is allowed to extrapolate the estimates to a higher mass range, our results suggest that double-WD mergers can contribute to a significant fraction of type Ia supernovae.

In a few years, the increasing astrometric and photometric precision provided by future *Gaia* data releases and the radial velocity measurements of white dwarfs by future surveys such as SDSS-V [e.g., 191] will enlarge the available sample size of high-mass white dwarfs and allow for even tighter constraints. We are starting to be able to reliably and precisely compare the observed double-WD merger rates with binary population synthesis, which will shed light upon the progenitor problem of type Ia supernovae.

Appendix III: A cooling anomaly of white dwarfs

This appendix is a copy of my paper [136] published in the *Astrophysical Journal*.

Abstract

Recently, the power of *Gaia* data has revealed an enhancement of high-mass white dwarfs (WDs) on the Hertzsprung–Russell diagram, called the Q branch. This branch is located at the high-mass end of the recently identified crystallization branch. Investigating its properties, we find that the number density and velocity distribution on the Q branch cannot be explained by the cooling delay of crystallization alone, suggesting the existence of an extra cooling delay. To quantify this delay, we statistically compare two age indicators – the dynamical age inferred from transverse velocity, and the photometric isochrone age – for more than one thousand high-mass WDs ($1.08\text{--}1.23 M_{\odot}$) selected from *Gaia* Data Release 2. We show that about 6% of the high-mass WDs must experience an 8 Gyr extra cooling delay on the Q branch, in addition to the crystallization and merger delays. This cooling anomaly is a challenge for WD cooling models. We point out that ^{22}Ne settling in C/O-core WDs could account for this extra cooling delay.

Introduction

Until recently, explorations of the white dwarf region in the Hertzsprung–Russell (H–R) diagram were severely limited by the number of objects with available distance estimates. The European Space Agency *Gaia* mission [135] has changed this situation drastically. *Gaia* is an all-sky survey of astrometry and photometry for stars down to 20.7 magnitude. The H–R diagram of white dwarfs generated by *Gaia* Data Release 2 (DR2) reveals three branch-like features, called the A, B, and Q branches⁵ in Figure 13 of Gaia Collaboration *et al.* [137]. The A and B branches have been understood as standard-mass white dwarfs ($m_{\text{WD}} \sim 0.6 M_{\odot}$) with hydrogen-rich and helium-rich atmospheres, respectively [e.g., 192]. However, the Q branch, as an enhancement of high-mass white dwarfs ($m_{\text{WD}} > 1.0 M_{\odot}$), is still not fully understood. This is a challenge to current white dwarf evolutionary models and an opportunity for studying high-mass white dwarfs.

On the H–R diagram, white dwarfs evolve along their cooling tracks. Unlike the A and B branches, the Q branch is not aligned with any cooling track or isochrone, suggesting that it is

⁵Named after the presence of DA, DB, and DQ white dwarfs [137], respectively. DA white dwarfs have hydrogen lines in their spectr, DB and DQ white dwarfs have helium and carbon lines, respectively.

caused by a delay of cooling instead of a peak in mass or age distribution. This cooling delay makes white dwarfs pile up on the Q branch. The Q branch coincides with the high-mass region of the crystallization branch identified by Tremblay *et al.* [193]. As a liquid-to-solid phase transition in the white dwarf core, crystallization releases energy through latent heat [e.g., 194] and phase separation [e.g., 195–197], which can indeed create a cooling delay. However, the observed pile-up on the Q branch is higher and narrower than expected from the standard crystallization models [193, figure 4], suggesting that an extra cooling delay may exist in addition to crystallization.

In this paper, we investigate this cooling anomaly using kinematic information of high-mass white dwarfs in *Gaia* DR2. In Section 5 we describe our white dwarf sample; in Section 5 we show strong evidence for the existence of an extra cooling delay on the Q branch; in Section 5 we build a model for the white dwarf velocity distribution and use our *Gaia* sample to constrain the properties of this cooling anomaly; in Section 5 we present the best-fit values of these properties and as a byproduct of our model of our analysis, the fraction of double-WD merger products among high-mass white dwarfs; in Section 5 we show that ^{22}Ne settling in massive C/O-core white dwarfs is a promising physical origin of this extra cooling delay; in Section 5 we examine other aspects of the Q branch, which provide evidence that the extra delayed white dwarfs are also double-WD merger products; and in Section 5 we conclude on our findings.

Data

We use data from *Gaia* DR2 [144], which for the first time provides parallaxes ϖ and proper motions μ that are derived purely from *Gaia* measurements [140]. *Gaia* DR2 also provides Vega magnitudes of three wide passbands [141, 142]: the *G* band spans from 350 to 1000 nm, and the G_{BP} and G_{RP} bands are mainly the blue and red parts of the *G* band, separated at the $\text{H}\alpha$ transition [135].

Quality cuts

Gentile Fusillo *et al.* [143] have compiled a catalog of *Gaia* DR2 white dwarfs based on the *G*-band absolute magnitude, *Gaia* color index, and some quality cuts. To select white dwarfs with high-precision measurements, we further apply the following quality cuts:

$$\sigma_{G_{\text{BP}}-G_{\text{RP}}} < 0.10, \quad (\text{V.9})$$

$$\sigma_{\mu}/\varpi < 2 \text{ km s}^{-1}, \quad (\text{V.10})$$

$$\text{parallax_over_error} > 8, \quad (\text{V.11})$$

$$\text{astrometric_excess_noise} < 1.5, \quad (\text{V.12})$$

$$\text{phot_bp_rp_excess_factor} < 1.4, \quad (\text{V.13})$$

$$1/\varpi < 250 \text{ pc}, \quad (\text{V.14})$$

where the color error $\sigma_{G_{\text{BP}}-G_{\text{RP}}}$ is the combined photometric errors in G_{BP} and G_{RP} bands, the proper motion error σ_{μ} is the combined error originating from its two components, and ϖ is the parallax from *Gaia* DR2. These cuts are designed to balance data quality and sample size. They do not introduce *explicit* kinematic biases, which is necessary for our analysis below. While the main sample used in our study uses white dwarfs within 250 pc, we will also occasionally use a subsample of white dwarfs located within 150 pc to clearly show number density enhancements.

Kinematic and physical parameters of white dwarf

Our analysis requires white dwarf absolute magnitude, color index, and the two components of transverse velocity $\mathbf{v}_T = (v_l, v_b)$ in Galactic longitude and latitude directions. Except for the color index $G_{BP}-G_{RP}$, which is directly read from the `bp_rp` column in *Gaia* DR2, we derive the other quantities in the following way:

$$M_G = G + 5 \log(\varpi/\text{mas}) - 10, \quad (\text{V.15})$$

$$v_l = \frac{\mu_L - (A \cos 2l + B) \cos b}{\varpi}, \quad (\text{V.16})$$

$$v_b = \frac{\mu_B + A \sin 2l \sin b \cos b}{\varpi}, \quad (\text{V.17})$$

where G and ϖ/mas are read from *Gaia* DR2 columns `phot_g_mean_mag` and `parallax`, μ_L and μ_B are converted from columns `ra`, `dec`, `pmra`, and `pmdec` with the coordinate conversion function in the `ASTROPY` package [198, 199], and A and B are the Oort constants taken from Bovy [155]. We do not correct for extinction because within the distance cut, extinction is in general tiny and there is no accurate estimate for it. To avoid the influence of hyper-velocity white dwarfs, we further impose a velocity cut:

$$v_T = \sqrt{v_l^2 + v_b^2} < 250 \text{ km s}^{-1}. \quad (\text{V.18})$$

We point out that *Gaia* does not provide any useful radial velocity information for white dwarfs as they have no spectral lines in the 845–872 nm wavelength range of *Gaia*’s spectrometer [135].

We then derive white dwarf photometric isochrone ages and masses from the H–R diagram coordinates:

$$(G_{BP}-G_{RP}, M_G) \rightarrow (\tau_{\text{phot}}, m_{\text{WD}}), \quad (\text{V.19})$$

based on a single-star evolution scenario and white dwarf cooling models⁶. We estimate the main-sequence (MS) ages with an initial–final mass relation [150] and the relation between pre-WD time and main-sequence mass from Choi *et al.* [151] for non-rotating, solar-metallicity stars. For high-mass white dwarfs, the pre-WD ages are negligible. As for white dwarf cooling, we use a table of synthetic colors for pure-hydrogen atmosphere [146–148] and a grid of cooling tracks for C/O-core white dwarfs with “thick” hydrogen layers [145]⁷. In order to convert any H–R diagram coordinate into τ_{phot} and m_{WD} , we linearly interpolate between grid points. Stellar models show that in the single-star-evolution scenario, white dwarfs with a mass higher than about 1.05–1.10 M_\odot have oxygen+neon (O/Ne) cores [e.g., 200, 201]. So, we combine the cooling tracks of $m_{\text{WD}} \leq 1.05 M_\odot$ C/O white dwarfs with the four cooling tracks of $m_{\text{WD}} \geq 1.10 M_\odot$ O/Ne white dwarfs [149].

The O/Ne white dwarf model only gives slightly lower mass estimates than the C/O white dwarf model (e.g., 1.08–1.23 M_\odot in the combined model corresponds to 1.10–1.28 M_\odot in the C/O-only model), and their estimates of the photometric ages τ_{phot} are similar for the white dwarfs we are interested in ($\tau_{\text{phot}} < 3.5$ Gyr). Switching between thick-hydrogen, thin-hydrogen, and helium atmosphere [202] models does not significantly change the photometric-age estimate of our sample either.

⁶We have made this tool a publicly available python 3 module on https://github.com/SihaoCheng/WD_models

⁷<http://www.astro.umontreal.ca/~bergeron/CoolingModels/>.

Mass, age, and Q-branch selection

In Figure V-7 we show the selected white dwarfs on the H–R diagram. In the top right panel, we use the 150 pc sample to show the density distribution with a higher contrast. The Q branch is a factor-two enhancement at around $-0.4 < G_{\text{BP}} - G_{\text{RP}} < 0.2$ and $M_G = 13$. In the main panel, we show our main sample within 250 pc, color-coded by their transverse velocities v_T with respect to the local standard of rest. We adopt $(U_\odot, V_\odot, W_\odot) = (11, 12, 7)$ km s⁻¹ [203] to correct for the solar reflex motion. We emphasize the fast white dwarfs ($v_T > 70$ km s⁻¹) in Figure V-7 with larger dots: they are very likely thick-disk stars. We also plot a grid of photometric age τ_{phot} and mass m_{WD} derived from the combined O/Ne-core and C/O-core white dwarf cooling model. Cooling tracks are the curves with constant m_{WD} . White dwarfs with different birth times form a ‘white dwarf cooling flow’ on the H–R diagram as they move along their cooling tracks.

We focus on the mass range where the Q branch is most prominent. To maximize sample size and minimize the contamination from standard-mass helium-atmosphere white dwarfs (the B branch), we impose the following photometric age and mass cuts:

$$0.1 \text{ Gyr} < \tau_{\text{phot}} < 3.5 \text{ Gyr}, \quad (\text{V.20})$$

$$1.08 M_\odot < m_{\text{WD}} < 1.23 M_\odot, \quad (\text{V.21})$$

where m_{WD} is derived from the combined cooling model for O/Ne and C/O white dwarfs. This mass range corresponds to 1.10–1.28 M_\odot in the C/O-only cooling model. In total, 1070 white dwarfs are selected by these criteria⁸. In this region, the Q branch divides the white dwarf cooling flow into three segments: the early, Q-branch, and late segments, as shown in Figure V-7. We define the Q-branch segment by

$$|M_G - 1.2 \times \text{bp_rp} - 13.2| < 0.2 \quad (\text{V.22})$$

in addition to the previous photometric-age and mass cuts.

An extra cooling delay on the Q branch

Evidence from the photometric-age distribution

As argued by Tremblay *et al.* [193], an enhancement not aligned with mass or age grid, such as the Q branch, should be produced by a slowing down (and therefore a delay) of white dwarf cooling. Such a cooling delay creates a ‘traffic jam’ in the white dwarf flow, and the Q branch is a snapshot of this traffic jam. *Is crystallization alone enough to explain the cooling delay on the Q branch?* If it is, then the distribution of photometric age τ_{phot} derived from a cooling model including crystallization effects should no longer carry signatures of the Q branch. However, observations lead to the antithesis. In Figure V-8 we show the distribution of τ_{phot} in three mass ranges: there is a mass-dependent enhancement tracing the Q branch, which is consistent with the observation by Tremblay *et al.* [193] that the pile-up is higher and narrower than what the standard cooling model predicts. Evolutionary delays from binary interactions or a peak in star formation rate cannot explain this mass-dependent enhancement either. Therefore, an extra cooling delay in addition to crystallization effects (latent heat and phase separation) must exist.

⁸A catalog of all selected white dwarfs is available on the website: <https://sihaocheng.github.io/Qbranch>

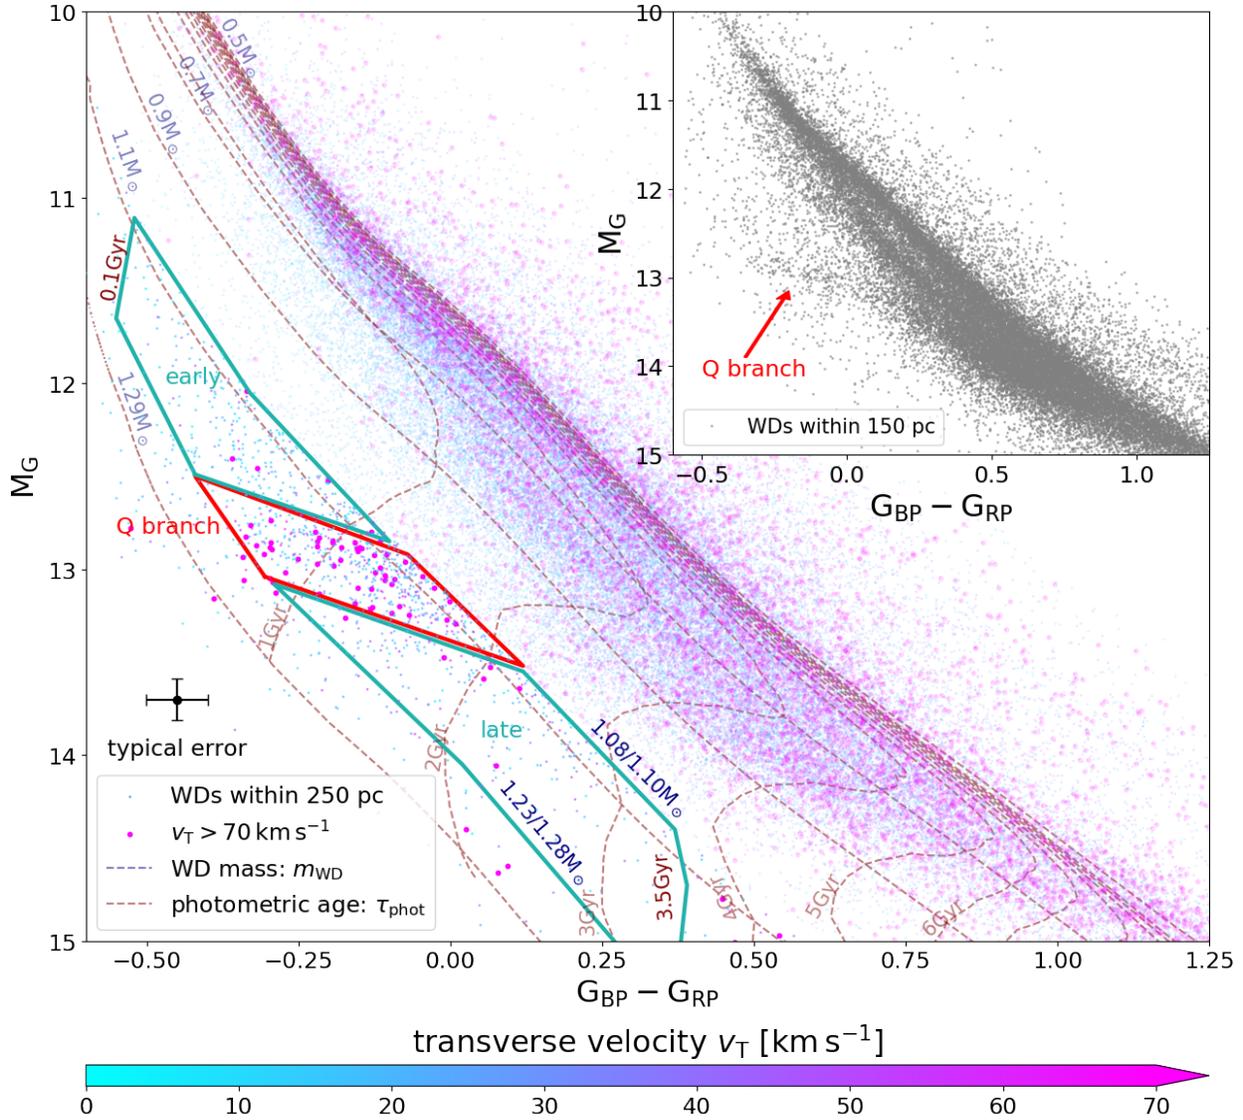


Figure V-7. The H–R diagram of WDs selected from Section 5. In the top right panel we use the 150 pc sample to show the number-density distribution with a higher contrast. The Q branch is marked by the red arrow. In the main panel, we show our main 250 pc sample color-coded with transverse velocities v_T . Fast WDs ($v_T > 70 \text{ km s}^{-1}$) are emphasized by large symbols, and high-mass WDs ($> 1.08 M_\odot$) are emphasized by high symbol opacity. The grid of WD mass and photometric age is also plotted (using the O/Ne model for high-mass WDs). For the mass range marked by dark blue texts, the first (second) number corresponds to the O/Ne (C/O) model.

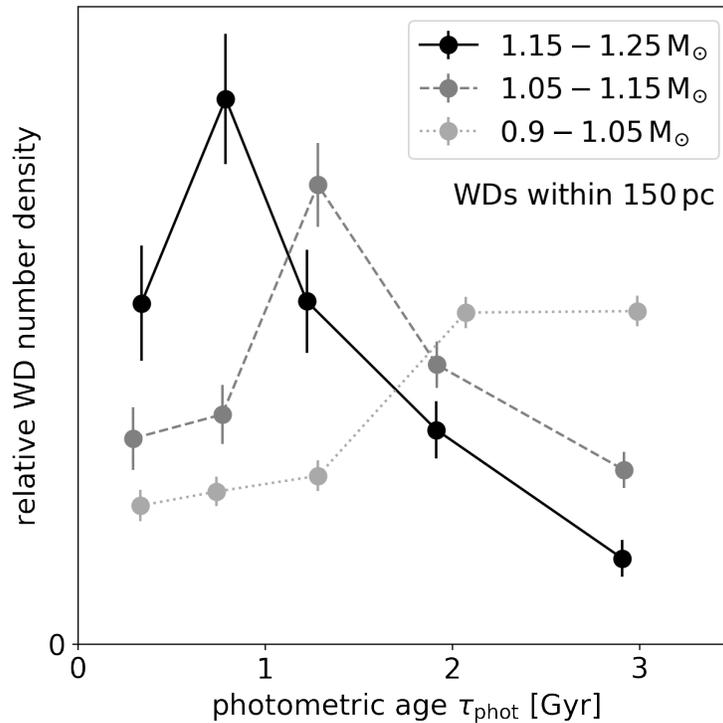


Figure V-8. The normalized photometric age τ_{phot} distribution of high-mass WDs in three consecutive mass ranges. The mass-dependent peaks trace the position of the Q branch. Crystallization should not produce any peak on this plot, because the τ_{phot} is calculated from a model including crystallization effects; the completeness stays high for at least 1 Gyr after the peaks in each mass range, so these mass-dependent peaks cannot be explained by a peak in the star formation history or by incompleteness. Therefore, there must be an extra cooling delay piling up WDs on the Q branch.

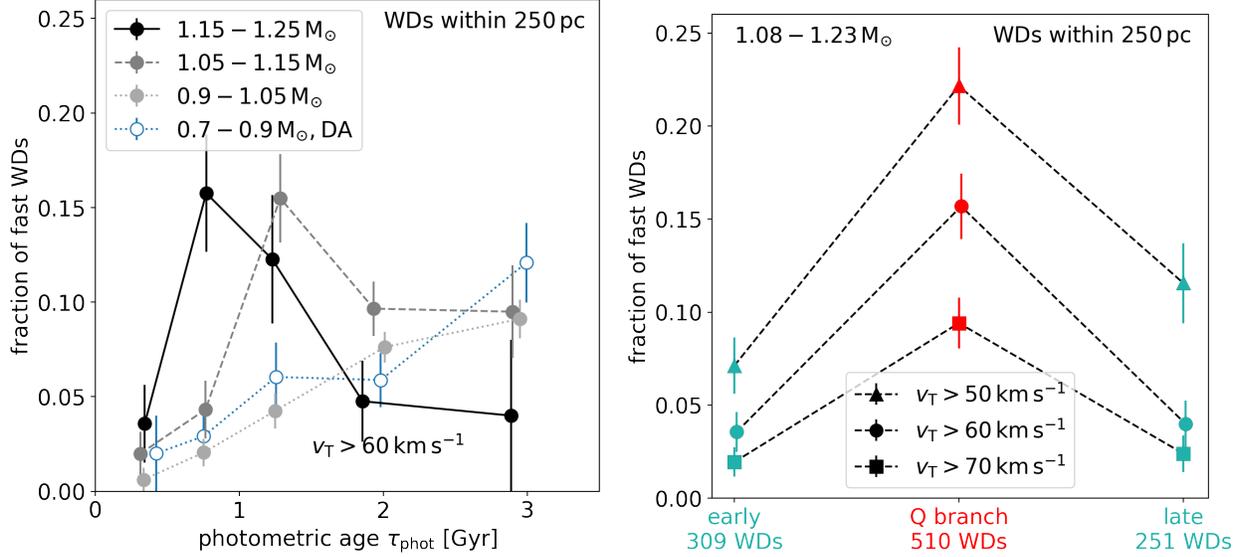


Figure V-9. The fraction of fast WDs in different mass ranges (*left panel*) and for different velocity cuts (*right panel*). There are significantly more fast-moving WDs on the Q branch than both before and after it in terms of photometric age. According to the age–velocity–dispersion relation (AVR), fast WDs are old. As argued in Section 5, this high fraction of fast WDs on the Q branch can only be explained by a subset of WDs experiencing an extra cooling delay on the Q branch.

Evidence from the velocity distribution

Observations show that the velocity dispersion of disk stars in the Milky Way is related to the stellar age τ : older stars have higher velocity dispersion than younger stars [e.g., 204]. So, the transverse velocity v_T derived from *Gaia* DR2 can be used as a ‘dynamical’ indicator of the true stellar age τ . For the Milky Way thin disk, the dispersion of transverse velocity approximately follows a power law increasing from about 25 km s^{-1} at 1.5 Gyr to 55 km s^{-1} at around 6–8 Gyr [e.g., 204]; for the thick disk, the dispersion is about 65 km s^{-1} [e.g., 205]. Given this age–velocity–dispersion relation (AVR), we observe two anomalous things in the velocity distribution of the Q branch:

- There is a strong excess of white dwarfs with $v_T > 70 \text{ km s}^{-1}$ in the Q-branch segment, as shown in Figure V-7. According to the age–velocity–dispersion relation (AVR) mentioned above, these fast white dwarfs should be old stars. Given that the photometric age on the Q branch is only 0.5–2 Gyr, these white dwarfs must have experienced an extra cooling delay for several billion years. In the left panel of Figure V-9 we show that the excess of fast white dwarfs in the Q-branch segment is clear for $m_{\text{WD}} > 1.05 M_{\odot}$; in the right panel we show that this excess is observed for a variety of velocity cuts.
- The fraction of fast white dwarfs in the *late* segment is *lower* than that in the Q-branch segment. This is anomalous, because white dwarfs in the late segment should be older than those in the Q-branch segment, as long as all white dwarfs follow the same cooling law. The only way to create such a reverse of fraction is to have more than one white dwarf population with distinct cooling behaviors.

A two-population scenario of the extra cooling delay

The simplest scenario that can explain both the number enhancement and velocity anomaly on the Q branch is to have an extra-delayed population of white dwarfs in addition to a normal-cooling population. This scenario requires only two free parameters:

- The fraction f_{extra} of the extra-delayed population, and
- The length t_{extra} of the extra cooling delay on the branch.

In Figure V-10 we illustrate this scenario by showing the normal-cooling and extra-delayed populations on the H–R diagram. Before the Q branch, the extra-delayed population has no difference from the normal-cooling population. On the branch, the extra-delayed population has a slower cooling rate, which causes two effects: (1) its members pile up there, creating a number-density enhancement, and (2) the photometric ages τ_{phot} of its members start to seem younger than their true ages τ , creating an age discrepancy. After the branch, the number-density enhancement disappears, but the age discrepancy remains. A detailed parameterization of this scenario is presented in the next section. To create the observed reversal of fast white dwarf fraction in the Q-branch and late segment, the extra-delayed population must have a high number-density contrast between these two regions, which requires that the population fraction f_{extra} be small and the delay time t_{extra} long.

Parameterization of the extra cooling delay

Figure V-11 shows the parameterization of our extra-cooling-delay scenario on the Q branch. The observed enhancement A_{obs} in the photometric-age distribution can be expressed as

$$A_{\text{obs}} \equiv \frac{n_{\text{on branch}}}{n_{\text{off branch}}} - 1 = f_{\text{extra}} A = \frac{f_{\text{extra}} t_{\text{extra}}}{\Delta t_{\text{branch}}}, \quad (\text{V.23})$$

where A is the intrinsic enhancement for the extra-delayed population itself. Δt_{branch} is the average width of the Q branch in terms of photometric age, i.e., the time for a normal-cooling white dwarf to pass through the branch. We directly measure this width from Figure V-7 with just the photometric ages grid and our definition of the Q-branch segment, finding an average value $\Delta t_{\text{branch}} = 0.74$ Gyr. White dwarfs with the extra cooling delay will spend much more time passing the branch.

When only the enhancement A_{obs} is used to investigate the Q branch, there is clearly a degeneracy between f_{extra} and t_{extra} in this two-population scenario. Velocity information helps to break this degeneracy.

Quantitative analysis

Having shown qualitatively the existence of an extra cooling delay on the Q branch, we now attempt to quantify its properties. We build a statistical model that (i) includes double-WD mergers, (ii) uses an anisotropic AVR, and (iii) makes use of the full constraining power of the observations.

Merger products among high-mass white dwarfs

Simulations of binary evolution show that double-WD merger products may account for a considerable fraction of high-mass white dwarfs [e.g., 127, 128]. These merger products also have

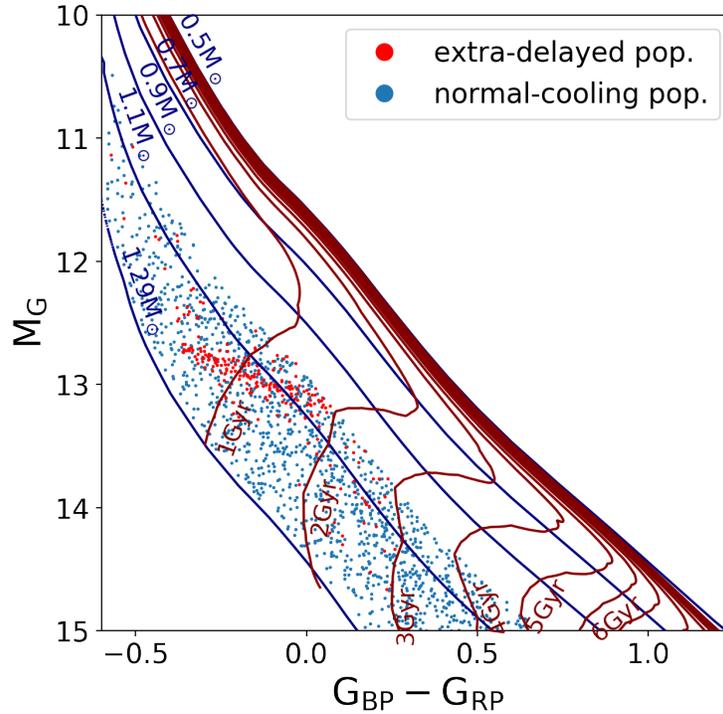


Figure V-10. An illustration of the two-population scenario of the Q branch: a normal-cooling population (blue dots), and a population with the extra cooling delay (red dots). The number density of WDs on the H–R diagram is determined by the cooling rate, because WDs accumulate where the cooling rate is low. Here, we use the best-fitting values ($f_{\text{extra}} = 6\%$ and $t_{\text{extra}} = 8$ Gyr) from our later analysis to generate this mock H–R diagram. An animated version of this figure is available, where blue dots move with the normal cooling rate, while red dots (the extra-delayed population) move slowly on the Q branch. Each second in the animation corresponds to 1 Gyr in physical time, and the duration of this animation is 11 seconds. More animations can be found on the website: <https://sihaocheng.github.io/Qbranch>.

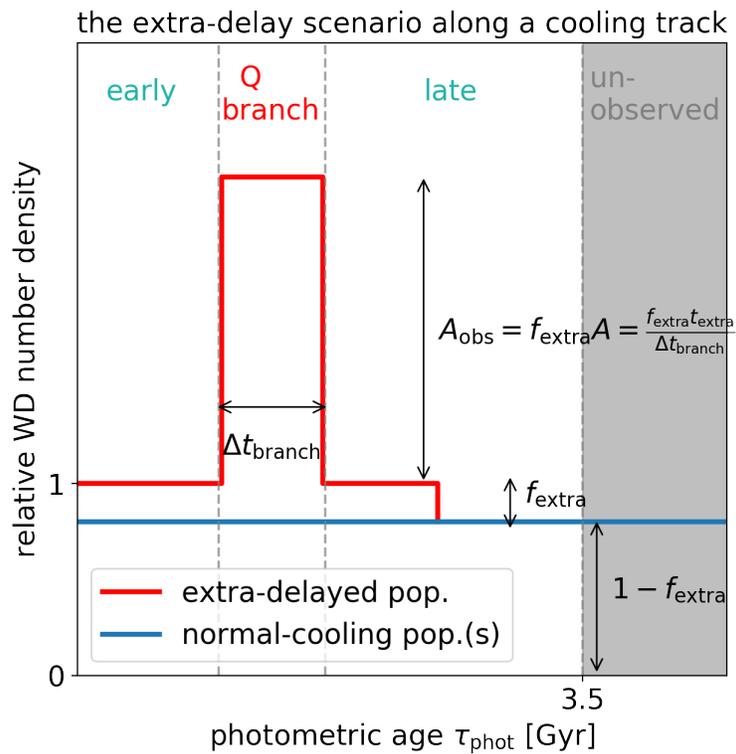


Figure V-11. A sketch of the extra-delay scenario. The axes are the same as in Figure V-8. Some WDs are subject to the extra cooling delay and the rest of them are not, corresponding to the extra-delayed and normal-cooling population in the figure. We also illustrate the quantities of equation V.23.

Table V-II. Delay scenarios of the three populations. For each population, the delay types are shown in the upper part of the table, and the total delay time $\Delta t = \tau - \tau_{\text{phot}}$ for each segment is shown in the lower part. Δt , Δt_{merger} , and Δt_{extra} are not single numbers but random variables following their distributions. They are used to calculate the distributions of true ages τ from photometric ages τ_{phot} .

Population (abbreviation)	single-star evolution (s)	extra-delayed (extra)	merger, no delay (m)
merger delay	no	yes or no (setup 1 or 2)	yes
extra cooling delay	no	yes	no
early	0	Δt_{merger} or 0	Δt_{merger}
Q branch	0	$(\Delta t_{\text{extra}} + \Delta t_{\text{merger}})$ or Δt_{extra}	Δt_{merger}
late	0	$(t_{\text{extra}} + \Delta t_{\text{merger}})$ or t_{extra}	Δt_{merger}

a discrepancy between their true ages and photometric ages due to binary evolution. Therefore, in order to use the velocity distribution to quantitatively constrain the extra cooling delay, the merger population must be modeled simultaneously. Constraining the merger fraction is also of great interest as its value is still a matter of debate [e.g., 124, 129–132]. Therefore, we include the double-WD merger products in our model and set their fraction as a free parameter.

Description of the model

In our model, we consider two evolutionary delays: the extra cooling delay, and the merger delay. Accordingly, we consider three populations of white dwarfs with different combinations of the two delays:

- A generic population of singly evolved white dwarfs that follows normal cooling, denoted by ‘s’;
- A double-WD merger population⁹ with systematic age offsets due to the merger delay and with a normal cooling, denoted by ‘m’;
- A population with the extra cooling delay, denoted by ‘extra’.

Their delay scenarios are listed in Table V-II. For simplicity, we only explore the two extreme situations for the extra-delayed population, where

- Setup 1: *all* members of the extra-delayed population also have the merger delay;
- Setup 2: *no* members of the extra-delayed population have the merger delay.

The distribution function $p(\mathbf{y})$ of observables \mathbf{y} for all white dwarfs can be written as a weighted average of the distribution for each population:

$$p(\mathbf{y}) = f_s p_s(\mathbf{y}) + f_m p_m(\mathbf{y}) + f_{\text{extra}} p_{\text{extra}}(\mathbf{y}), \quad (\text{V.24})$$

⁹We only consider the double-WD mergers because in our mass range, other merger products such as those from MS–RG, MS–MS, and MS–WD mergers usually only have < 0.2 Gyr delay and therefore are indistinguishable from the singly evolved white dwarfs in terms of kinematics.

where the weight f denotes the fraction of each population, satisfying $f_s + f_m + f_{\text{extra}} = 1$.

Our goal is to use observations to constrain two independent population fractions and the delay time of the extra cooling delay:

$$f_m, f_{\text{extra}}, \text{ and } t_{\text{extra}}, \quad (\text{V.25})$$

the last of which is encoded in the distribution $p_{\text{extra}}(\mathbf{y})$. We have two sets of observables \mathbf{y} : the transverse velocities \mathbf{v}_T , and the photometric ages τ_{phot} . They are connected by the AVR $p(\mathbf{v}|\tau)$ and the delay scenario of each population (listed in Table V-II). The delay

$$\Delta t \equiv \tau - \tau_{\text{phot}} \quad (\text{V.26})$$

includes contributions from the extra-cooling and/or the merger delays. We build a Bayesian model based on Equation V.24 to constrain the aforementioned parameters. Our model is similar to that of Wegg & Phinney [132], but we include the extra-delayed population and use a much larger sample. In addition, to avoid the need for modeling selection effects, we derive our constraints from the velocity distribution conditioned on observables other than velocity:

$$p(\mathbf{v}|\text{other observables}) = p(\mathbf{v}_T|\tau_{\text{phot}}, m_{\text{WD}}, l, b). \quad (\text{V.27})$$

The details of this statistical technique and the Bayesian framework of our model are shown in 5.

The free parameters in our model include the population fractions f_m and f_{extra} , the extra delay time t_{extra} , parameters for the AVR, and solar motion. Although constraints on the AVR and solar motion already exist, treating them also as free parameters can avoid potential systematic errors, and the comparison of our best-fitting values with the existing values allows us to check the validity of our method.

Below, we list the main assumptions and simplifications in our model:

1. We assume that upon entering the Q-branch segment, all members of the extra-delayed population suddenly slow down their cooling by a constant factor, and upon leaving the branch, the cooling rates suddenly resume, so that this extra cooling delay can be parameterized by just its length t_{extra} and population fraction f_{extra} . The resulting delay-time distribution is described in Section 5.
2. The velocity distribution of white dwarfs is a superposition of 3D Gaussian distributions as a function of age τ , i.e. $p(\mathbf{v}|\tau) = \mathcal{N}(\mathbf{v}_0(\tau), \Sigma(\tau))$. The details of this Gaussian velocity model are shown in 5.
3. The true-age distribution of high-mass white dwarfs within 250 pc is uniform up to 11 Gyr, i.e. $\tau \sim U[0, 11 \text{ Gyr}]$.
4. For the double-WD merger products, we follow Wegg & Phinney [132] and assume that the resulting white dwarf is reheated enough that its cooling age after the merger is almost equal to the photometric cooling age. We also assume a fixed delay-time distribution for double-WD mergers (see Section 5) and a parameterization of the AVR (see Section 5).

Delay-time distributions

The three white dwarf populations in our model are defined by their different delay signatures Δt , which concern the extra cooling delay Δt_{extra} and double-WD merger delay Δt_{merger} . The delay scenario of each population in each segment is listed in Table V-II.

The extra cooling delay Δt_{extra} is built up on the Q branch. We adopt a uniform distribution $\Delta t_{\text{extra}} \sim U[0, t_{\text{extra}}]$ of this delay for white dwarfs in the Q-branch segment and a constant value in the late segment. Note that Δt_{extra} is a random variable with a probability distribution, whereas t_{extra} , as a model parameter to be constrained, is the upper limit of Δt_{extra} . In the Q-branch segment, we do not further distinguish if a white dwarf has just started or is about to complete their extra cooling delay, because the uncertainty of H–R diagram coordinate due to different atmosphere types and astrometric/photometric error is comparable to the width of the Q branch on the H–R diagram. In this case, a uniform distribution is a good and efficient approximation for Δt_{extra} .

The double-WD merger delay Δt_{merger} originates from the binary evolution before the merger. We refer to binary population synthesis results [e.g., 167] and approximate the delay by

$$p(\Delta t_{\text{merger}}) \propto \Delta t_{\text{merger}}^{-0.7} \quad (\text{V.28})$$

for $\Delta t_{\text{merger}} > 0.5$ Gyr and zero for smaller Δt_{merger} . Unlike the extra cooling delay, we do not set any free parameter for this merger-delay distribution.

Parameterization of the AVR

We define the U , V , W axes as pointing toward $(l = 0^\circ, b = 0^\circ)$, $(l = 90^\circ, b = 0^\circ)$, and $(b = 90^\circ)$, respectively, and assume that the main axes of the Gaussian velocity distribution are aligned with these directions with dispersion $\sigma_U(\tau)$, $\sigma_V(\tau)$, and $\sigma_W(\tau)$. Observations show that the AVR in each direction can be fit by a shifted power law. The power index of the in-disk components are around 0.35 and that of the W component is around 0.5 [e.g., 204, 205]. For old stars including thick-disk members, the AVR is still a matter of debate [e.g., 206, 207]. So, in each direction, we use a shifted power law to parameterize the AVR of the younger, thin-disk stars:

$$\sigma(\tau) = \sigma^{\tau=0} + \Delta\sigma^{0 \rightarrow 4} \times \left(\frac{\tau}{4}\right)^\beta, \quad \tau < 7 \text{ Gyr}, \quad (\text{V.29})$$

and we use a constant value σ^{thick} to represent the velocity dispersion of stars older than 10 Gyr (thick-disk stars); in between 7 and 10 Gyr, we linearly interpolate the values from the two ends to reflect the increasing fraction of thick-disk stars. The shape of the AVR with our parameterization is shown in Figure V-12. The ratio of the two in-disk components σ_V and σ_U should be a constant for a local sample [e.g., 157], so we set $\sigma_V(\tau) = k \sigma_U(\tau)$. As the assumption for the velocity distribution to be Gaussian gradually breaks down when σ_U increases, we allow the ratio k to be different for the thin and thick disks. Thus, we use in total 10 parameters to model the anisotropic AVR: two initial velocity dispersions $\sigma_{U,W}^{\tau=0}$, two dispersion increases $\Delta\sigma_{U,W}^{0 \rightarrow 4}$ between 0 and 4 Gyr, two power indices $\beta_{U,W}$, two thick-disk dispersions $\sigma_{U,W}^{\text{thick}}$, and two in-disk dispersion ratios k^{thin} and k^{thick} . The best-fitting values of these parameters can be checked against existing estimates presented in the literature.

The Bayesian framework and elimination of selection effects

We follow a Bayesian approach to build our model. This means that we can first build a forward model outputting the likelihood probability density function (PDF) of observables \mathbf{y} given the model parameters $\boldsymbol{\theta}$:

$$\mathcal{L} \equiv p(\mathbf{y}|\boldsymbol{\theta}), \quad (\text{V.30})$$

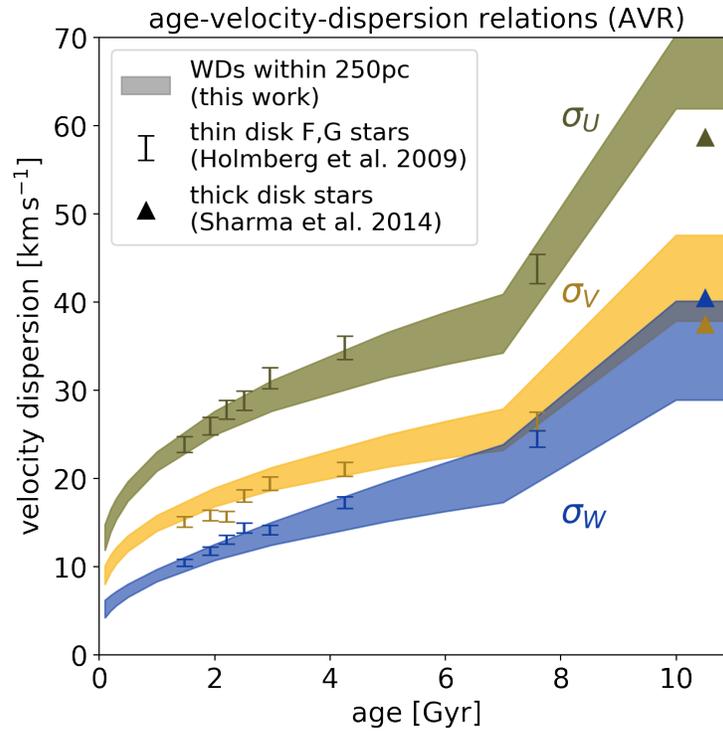


Figure V-12. The comparison of AVRs constrained by this work and in the literature. The shaded regions show the 16th and 84th percentiles of the AVR posterior constrained by our high-mass WD sample. Symbols with error bars show the AVR measured for main-sequence stars by GCS and RAVE [204, 205]. In our model, the $\tau < 3.5$ Gyr part of the AVR is mainly constrained by the normal-cooling WDs (population 's' and 'm'), and the older part by the extra-delayed WDs. Note that the turnings at 7 and 10 Gyr reflect our parameterization of the AVR.

and then we obtain the posterior PDF of model parameters $p(\boldsymbol{\theta}|\mathbf{y})$ from the observed value of \mathbf{y} through the Bayes' theorem:

$$p(\boldsymbol{\theta}|\mathbf{y}) \propto \mathcal{L} \cdot p(\boldsymbol{\theta}), \quad (\text{V.31})$$

where $p(\boldsymbol{\theta})$ is the prior PDF of the parameters. Finally, we use the MCMC method to sample the posterior distribution and estimate the parameters of interest after marginalizing nuisance parameters. Among these three steps, the key part is to construct the likelihood.

As each white dwarf provides an independent observation, the likelihood \mathcal{L} in our model can be written as the product of the likelihoods of each individual white dwarf:

$$\mathcal{L} = \prod_i p_i(\mathbf{y}_i|\boldsymbol{\theta}). \quad (\text{V.32})$$

To avoid a direct dependence on selection effects, we use the conditional likelihood to let the constraining power originate only from velocity distributions: we define the individual likelihood p_i as the probability density for the i th white dwarf to have transverse the velocity \mathbf{v}_T given all other observables of this white dwarf:

$$p_i \equiv p(\mathbf{v}_T|\{\tau_{\text{phot}}, m_{\text{WD}}, l, b\}_i, \boldsymbol{\theta}). \quad (\text{V.33})$$

We condition on τ_{phot} and m_{WD} because their distributions are influenced by the detection completeness, quality cuts, and white dwarf spatial distribution. Moreover, the mass m_{WD} in equation V.33 model is only used to identify whether a white dwarf is on the Q branch. In order to decompose the different populations, we derive

$$\begin{aligned} p(\mathbf{v}_T|\tau_{\text{phot}}, m_{\text{WD}}) &= \frac{p(\mathbf{v}_T, \tau_{\text{phot}}, m_{\text{WD}})}{p(\tau_{\text{phot}}, m_{\text{WD}})} \\ &= \frac{\sum_x f_x p_x(\mathbf{v}_T, \tau_{\text{phot}}, m_{\text{WD}})}{\sum_x f_x p_x(\tau_{\text{phot}}, m_{\text{WD}})}, \end{aligned} \quad (\text{V.34})$$

where the sums are taken over x with possible values 's', 'm', and 'extra' representing different populations.

To express these observable distributions by the AVR and star formation history, we employ the probability identity:

$$\begin{aligned} p(\mathbf{v}_T, \tau_{\text{phot}}) &= \int p(\mathbf{v}_T|\tau_{\text{phot}}, \tau) p(\tau_{\text{phot}}, \tau) d\tau \\ &= \int p(\mathbf{v}_T|\tau) p(\tau_{\text{phot}}, \tau) d\tau \end{aligned} \quad (\text{V.35})$$

(where the second step is valid because the velocity is only a function of the true age τ) and another identity:

$$p(\tau_{\text{phot}}) = \int p(\tau_{\text{phot}}, \tau) d\tau. \quad (\text{V.36})$$

We also assume that the age distribution is uniform, $\tau \sim U[0, 11 \text{ Gyr}]$. In this way, the likelihood PDF in Equation V.32 and V.33 can be expressed through the delay distributions $p(\Delta t)$ and a velocity model $p(\mathbf{v}_T|\tau)$.

The Gaussian velocity model

Here, we describe the PDF of transverse velocities \mathbf{v}_T and their true stellar ages τ using the AVR. The velocity distribution of disk stars in the solar neighborhood with respect to the local standard of rest can be approximated as superposition of 3D Gaussian distributions [e.g., 157]:

$$p(\mathbf{v}|\tau) = \frac{\exp[-\frac{1}{2}(\mathbf{v} - \mathbf{v}_0)^T \boldsymbol{\Sigma}(\tau)^{-1}(\mathbf{v} - \mathbf{v}_0)]}{\sqrt{8\pi^3 |\boldsymbol{\Sigma}(\tau)|}}, \quad (\text{V.37})$$

whose mean and covariance matrix are determined by stellar age τ . The mean velocity $\mathbf{v}_0(\tau)$ is determined by two effects: the solar reflex motion $(-U_\odot, -V_\odot, -W_\odot)$ with respect to the local standard of rest, and the asymmetric drift in V direction by $-\sigma_U^2/80 \text{ km s}^{-1}$ [e.g., 157]. We set the solar motion as free parameters and use them to check the validity of our model.

To obtain the distribution of the observable transverse velocity $\mathbf{v}_T = (v_l, v_b)^T$, we project the 3D Gaussian distribution $p(\mathbf{v}|\tau)$ onto the tangential plane for each white dwarf and marginalize the radial component v_R . Because the resulting distribution is still a Gaussian distribution for a given age τ , the only task is to find its covariance matrix and mean vector. Let $\mathbf{v}_{XYZ} = (U, V, W)$ and $\mathbf{v}_{LBR} = (v_L, v_B, v_R)$ be the expressions of the same vector \mathbf{v} in the XYZ and LBR coordinate systems, respectively, and matrix M the rotation transformation matrix between the two systems:

$$(\mathbf{v} - \mathbf{v}_0)_{XYZ} = M \cdot (\mathbf{v} - \mathbf{v}_0)_{LBR}, \quad (\text{V.38})$$

where

$$M = \begin{bmatrix} \sin l & -\sin b \cos l & \cos b \cos l \\ \cos l & -\sin b \sin l & \cos b \sin l \\ 0 & \cos b & \sin b \end{bmatrix}. \quad (\text{V.39})$$

We ignore the small in-disk rotation between the Cartesian coordinate XYZ and the galactic polar coordinate. Then, we write the 3D Gaussian distribution in both coordinate systems (note that the Jacobian determinant of rotation transform is unity) and obtain the following relation:

$$(\mathbf{v} - \mathbf{v}_0)_{XYZ}^T \boldsymbol{\Sigma}_{XYZ}^{-1} (\mathbf{v} - \mathbf{v}_0)_{XYZ} = (\mathbf{v} - \mathbf{v}_0)_{LBR}^T \boldsymbol{\Sigma}_{LBR}^{-1} (\mathbf{v} - \mathbf{v}_0)_{LBR}. \quad (\text{V.40})$$

Substituting equation V.38, we obtain:

$$\boldsymbol{\Sigma}_{LBR} = (M^T \boldsymbol{\Sigma}_{XYZ}^{-1} M)^{-1} = M^T \boldsymbol{\Sigma}_{XYZ} M, \quad (\text{V.41})$$

where we assume

$$\boldsymbol{\Sigma}_{XYZ} = \begin{bmatrix} \sigma_U^2 & 0 & 0 \\ 0 & \sigma_V^2 & 0 \\ 0 & 0 & \sigma_W^2 \end{bmatrix}. \quad (\text{V.42})$$

The covariance matrix $\boldsymbol{\Sigma}_{LB}$ for the 2D Gaussian distribution marginalized along the R direction is the top left 2×2 sub-matrix of $\boldsymbol{\Sigma}_{LBR}$. The mean vector $\mathbf{v}_{T0} = (\mathbf{v}_0)_{LB}$ can be derived directly from vector rotation and projection. Then, the conditional PDF of the two transverse components $\mathbf{v}_T = (v_l, v_b)$ of \mathbf{v} can be written as

$$p(\mathbf{v}_T|\tau, l, b) = \frac{\exp[-\frac{1}{2}(\mathbf{v}_T - \mathbf{v}_{T0})^T \boldsymbol{\Sigma}_{LB}^{-1}(\mathbf{v}_T - \mathbf{v}_{T0})]}{\sqrt{4\pi^2 |\boldsymbol{\Sigma}_{LB}|}}, \quad (\text{V.43})$$

where the covariance matrix $\boldsymbol{\Sigma}_{LB}(\tau, l, b)$ and mean $\mathbf{v}_{T0}(\tau, l, b)$ depends on (τ, l, b) . We use the condition on (l, b) for each white dwarf to avoid modeling the spatial selection effects and reduce unnecessary parameters and biases.

MCMC Settings

In our Bayesian model, we assume uniform distributions for parameter priors. We feed the affine invariant MCMC sampler `emcee` [88] with the natural logarithm of the likelihood function defined in equation V.32. We use 500 walkers to explore the parameter space. After 200 steps of burn-in, the chains are checked to converge by comparing the percentile values of the parameters in each chain. Then, we run another 400 steps and use this sampling to represent the posterior distribution of each parameter. Figure V-13 shows the marginal posteriors of the 10 parameters of the AVR and their correlations, under the setup 1. Setup 2 leads to similar constraints of the AVR.

Results

To constrain the extra cooling delay properties and merger fraction, we feed our Bayesian model with the 1070 white dwarfs selected in Section 5. We use the Markov chain Monte Carlo (MCMC) sampler `emcee` [88] to obtain the posterior distribution of the parameters. Details of the settings are described in 5.

Constraints on the main parameters

In Figure V-14 we present the constraints we obtain for the parameters of interest: f_{extra} , t_{extra} , and f_{m} . We find that the extra-delayed population fraction is

$$\begin{aligned} f_{\text{extra}} &= 6.4_{-1.5}^{+2.3} \% \text{ (setup 1)} \\ &= 9.2_{-2.7}^{+4.4} \% \text{ (setup 2)}, \end{aligned} \tag{V.44}$$

and the length of the extra cooling delay is

$$\begin{aligned} t_{\text{extra}} &\geq 8 \text{ Gyr (setup 1)} \\ &\geq 10 \text{ Gyr (setup 2)}. \end{aligned} \tag{V.45}$$

These constraints confirm our qualitative conclusion that f_{extra} is small and t_{extra} is long in Section 5. We point out that the difference of t_{extra} in the two setups is exactly where the peak of the merger-delay distribution is located (2 Gyr), which is expected. The lower limit for t_{extra} slightly depends on the parameterization of the AVR: if we adopt a younger thick-disk age [7–11 Gyr instead of 9–11 Gyr, 207], this lower limit will also decrease by about 1–2 Gyr. t_{extra} may be overestimated for the fact that at a high level of dispersion, the velocity distribution is not exactly Gaussian. But it is unlikely that we overestimate t_{extra} too much, because we set the thick-disk dispersion as a free parameter. We have also checked that reasonable variations of the input delay-time distribution of the mergers [e.g., 161] do not change these two constraints significantly.

The fraction of merger products *without* the extra cooling delay is found to be $f_{\text{m}} = 15_{-5}^{+6} \%$ and $20_{-5}^{+6} \%$ (setups 1 and 2). Therefore, the total fraction of double-WD merger products is

$$\begin{aligned} f_{\text{extra}} + f_{\text{m}} &= 22_{-5}^{+7} \% \text{ (setup 1)} \\ f_{\text{m}} &= 20_{-5}^{+6} \% \text{ (setup 2)} \end{aligned} \tag{V.46}$$

among 1.08–1.23 M_{\odot} white dwarfs. This total fraction is mainly constrained by the fast white dwarfs in the early segment (where the two setups do not differ from each other), so the constraints

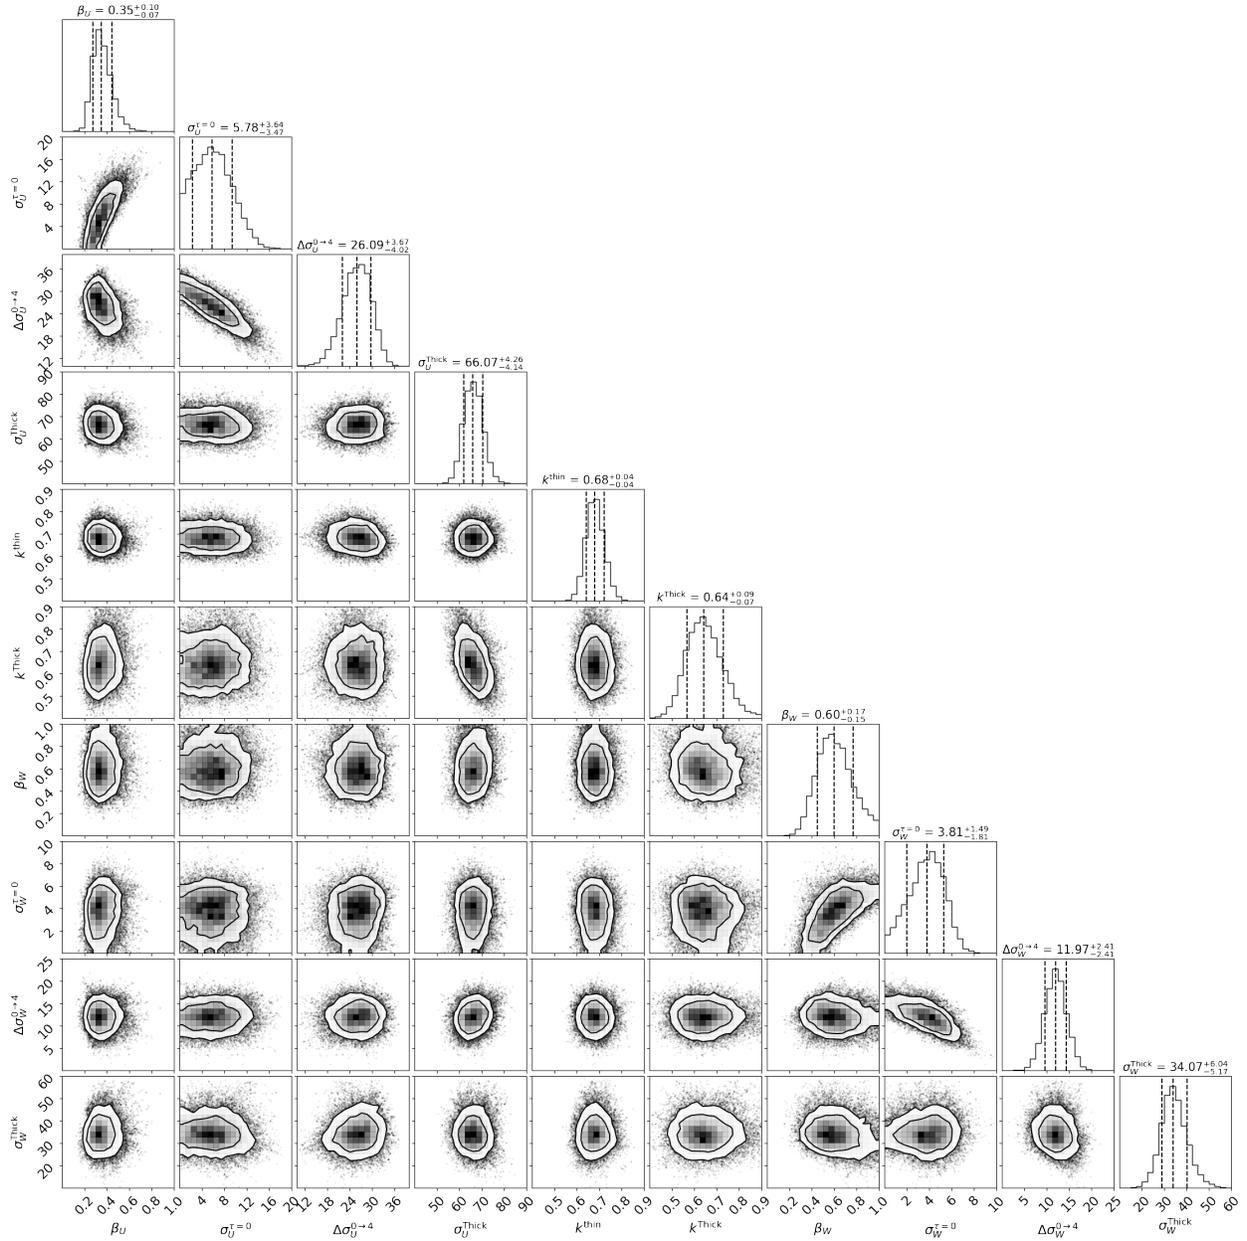


Figure V-13. The corner plot of the posteriors of the AVR parameters. We use flat priors for these parameters within the ranges shown on this figure. We have checked that there are no correlations between these parameters and the three main parameters f_{extra} , t_{extra} , and f_m , and the three components of solar motion in our model.

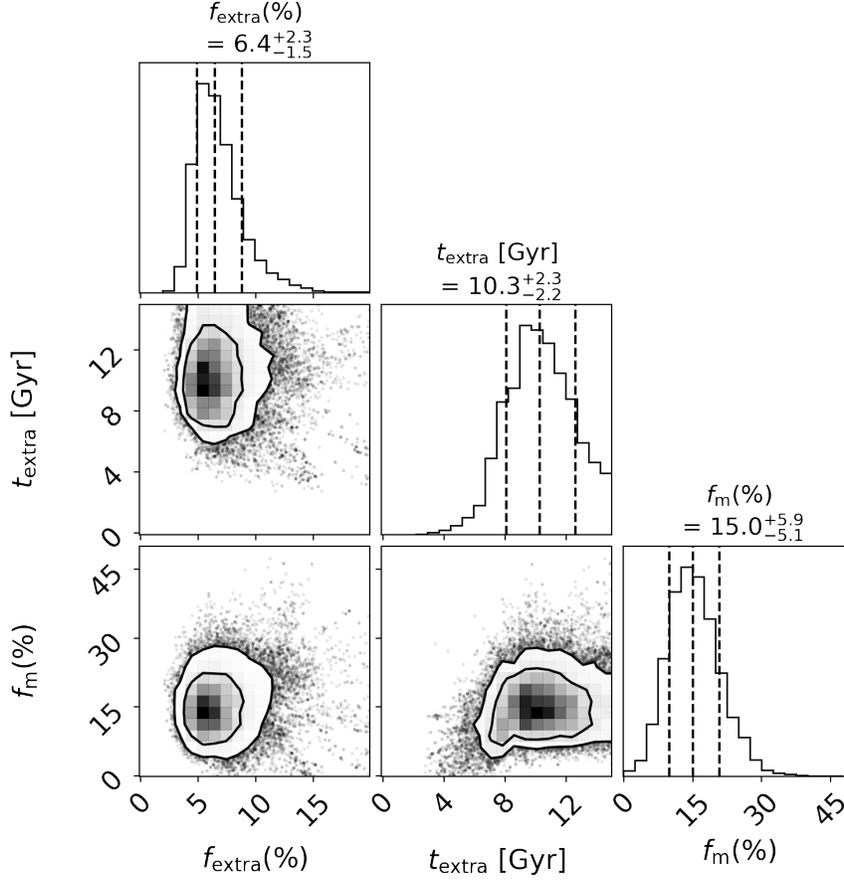


Figure V-14. The posterior distribution of the main parameters for setup 1. f_{extra} is the fraction of extra-delayed population, t_{extra} is the length of the extra cooling delay, and f_m is the fraction of normal-cooling double-WD merger products. Note that in setup 1, it is $f_{\text{extra}} + f_m$ rather than f_m that is the total fraction of merger products.

on this fraction under setups 1 and 2 are similar. A more detailed analysis of the merger products among high-mass white dwarfs is presented in Cheng *et al.* [118].

Finally, we calculate the contribution of the extra-delayed population in the Q-branch segment according to the above fractions:

$$\begin{aligned}
 F_{\text{extra}} &= \frac{f_{\text{extra}}(t_{\text{extra}}/\Delta t_{\text{branch}} + 1)}{1 + f_{\text{extra}}t_{\text{extra}}/\Delta t_{\text{branch}}} \\
 &= (47 \pm 8) \% ,
 \end{aligned}
 \tag{V.47}$$

where Δt_{branch} is the width of the Q branch (see section 5). This fraction, derived purely from the velocity information, is consistent with the factor ~ 2 estimate obtained directly from the number-density enhancement in Figure V-8, which confirms that our extra-cooling-delay scenario is a good phenomenological model for the Q branch.

Validation of the model

To further validate our model and results, we first check our constraints on the nuisance parameters. For the solar motion we obtain

$$(U, V, W)_{\odot} = (10.3 \pm 1.0, 7.3 \pm 1.0, 6.7 \pm 0.5) \text{ km s}^{-1}, \quad (\text{V.48})$$

which is consistent with the measurement of Rowell & Kilic [208] based on mainly standard-mass white dwarfs. Our values of U_{\odot} and W_{\odot} are also consistent with the results of Schönrich *et al.* [203]. The discrepancy of the V_{\odot} measurement comes from the different treatments of the asymmetric drift and is beyond the scope of this paper.

Figure V-12 shows our constraint on the white-dwarf AVR, which is consistent with the AVR of thin- and thick-disk main-sequence stars [204, 205]. Removing either the extra-delayed population or the merger population leads to unreasonably higher AVRs. Before *Gaia* DR2 came out, Anguiano *et al.* [209] reported an unexpectedly high AVR for young white dwarfs (their figures 21 and 22), without considering the extra cooling delay or merger delay in their age estimate. This unreasonably high AVR is exactly what we see when we remove the extra-delayed and/or merger population from our model. Thus, we verify that both the extra cooling delay and the merger delay are necessary.

To check the goodness of fit, we compare the observed and modeled velocity distributions in Figure V-15. Our best-fitting models (in both setups 1 and 2) provide good characterizations of the observed velocity distribution in all the early, Q-branch, and late segments. Adopting a different star formation history introduces no significant changes to our results. We test both a linearly decreasing star formation rate with a five-time higher star formation rate in the past, and a star formation history with a bump at 2.5 Gyr ago [e.g., 210], and find that the changes in best-fitting values are smaller than their uncertainties. This insensitivity to the assumed star formation history is expected because our model mainly uses the velocity information.

To further argue for our extra-delayed scenario against other explanations of the velocity anomaly, such as an accretion event of the Milky Way, we run a simple test where the velocities of fast white dwarfs on the Q branch are parameterized by only one Gaussian distribution. We find that the mean of the U and W components are consistent with zero, and the mean of V is $-50 \pm 6 \text{ km s}^{-1}$. Moreover, the U component has a dispersion of $60 \pm 6 \text{ km s}^{-1}$ and the ratio between the U and V dispersion is 0.60 ± 0.08 . All of these values satisfy the relations for a disk in equilibrium [e.g., 157]: asymmetric drift $\bar{V} = -\sigma_U^2 / (80 \text{ km s}^{-1})$ and dispersion ratio $\sigma_V / \sigma_U = 0.67$. It is unlikely for accreted stars to exactly reproduce the disk kinematics.

The physics behind the extra cooling delay: ²²Ne settling?

In previous sections, we showed that a previously unreported cooling delay is required to explain the velocity distribution of white dwarfs on the Q branch. Physically, this extra cooling delay requires an energy source satisfying the following conditions:

1. It has a highly peaked effect on the Q branch;
2. It is selective and applies to only $f_{\text{extra}} \sim 6\%$ of high-mass white dwarfs;
3. It is powerful enough to create a $t_{\text{extra}} \sim 8 \text{ Gyr}$ delay (in addition to crystallization delay and merger delay).

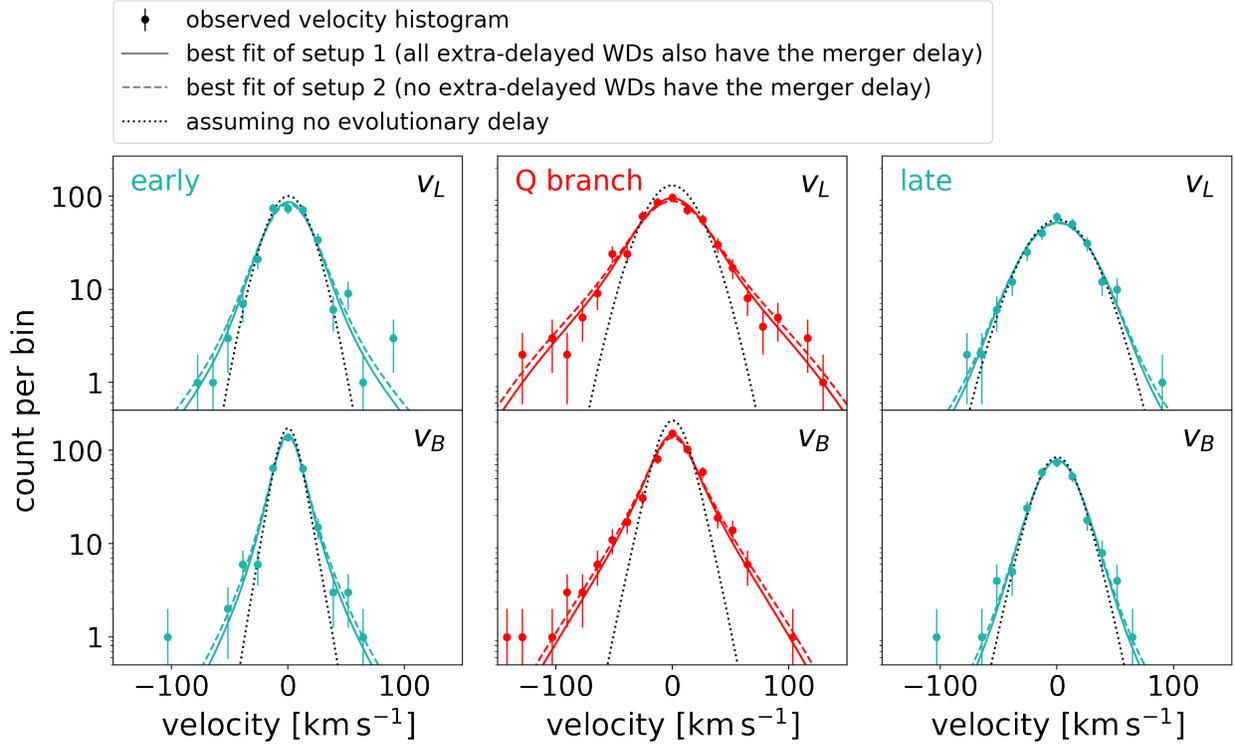


Figure V-15. The observed and modeled velocity distributions. v_l and v_b are the Galactic longitude and latitude components of the transverse velocity. For the observed distribution, we present the histogram between -150 and 150 km s^{-1} with 23 bins and the Poisson error of each bin. Note that the y-axis is in logarithmic scale. The solid and dashed curves, which are not very different, are the velocity distributions of the best-fitting models under setup 1 and 2, respectively. Both models fit the observations quite well. The dotted curves are the velocity distributions when no white dwarf has the extra cooling delay or merger delay. Its discrepancy to observed histograms shows the necessity of the two delays.

These requirements are very demanding. For example, a higher energy release from latent heat or phase separation is ruled out because their effects are not peaked enough and they are not selective. Besides crystallization, another possible energy source in a white dwarf is the settling of ^{22}Ne [211, 212]. Below, we show that ^{22}Ne settling could account for the extra cooling delay.

Different from the large amount of ^{20}Ne in O/Ne-core white dwarfs, the neutron-rich ^{22}Ne is produced from C, N, and O in the core of the progenitor stars. At the hydrogen burning stage, the CNO cycle builds up the slowest reactant ^{14}N , and at the helium burning stage, all ^{14}N is converted into ^{22}Ne . This leads to an abundance $X_{^{22}\text{Ne}}^{\text{WD}} \approx X_{\text{CNO}}^{\text{star}} \approx 0.014$ for solar metallicity stars. Due to the additional two neutrons, ^{22}Ne nuclei feel more downward force from gravity than the upward force from the electron-pressure gradient. So, they gradually settle down to the white dwarf center and release gravitational energy [212].

Now, let us check if ^{22}Ne settling satisfies the three requirements. We first emphasize that the delay *effect* only depends on the *fractional* contribution of the extra energy source to the white dwarf luminosity ($L_{\text{extra}}/L_{\text{surf}}$, see 5 for details). Therefore, to create a peaked effect, L_{extra} need not be also peaked.

The luminosity of ^{22}Ne settling ($L_{\text{extra}}^{\text{Ne}}$) relies on the ^{22}Ne abundance, mass, and core composition

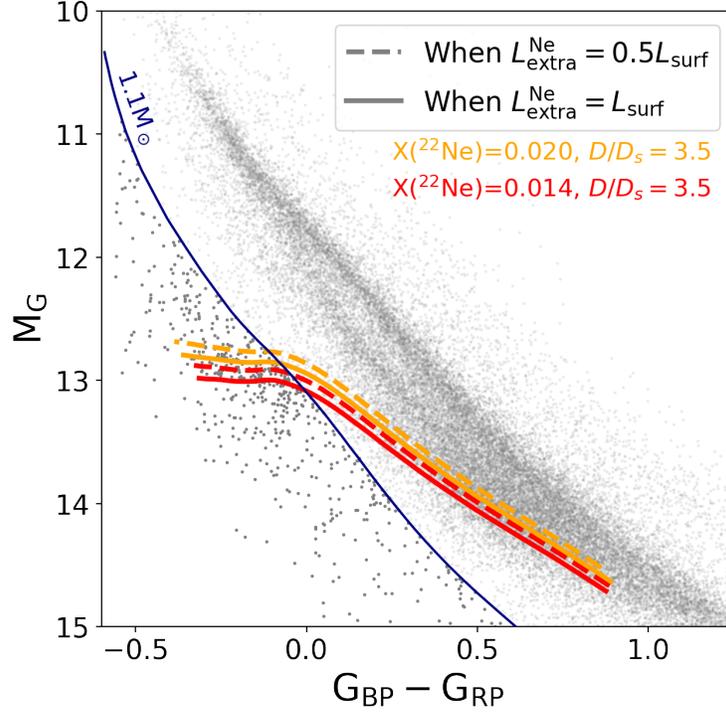


Figure V-16. The effective zone of ^{22}Ne settling for C/O-core DA WDs with $D/D_s = 3.5$, assuming no suppression from crystallization. ^{22}Ne settling significantly delays the WD cooling between the solid and dashed curves, when $L_{\text{extra}}^{\text{Ne}}$ is close to L_{surf} . The colors represent two ^{22}Ne abundances of WDs, corresponding to $[M/H] = 0$ and 0.15 in their progenitor stars. We observe that the position, trend, and narrowness of the effective zone of ^{22}Ne settling match the Q branch quite well.

of the white dwarf, and the inter-to-self-diffusion factor D/D_s , which is of order unity but not well-determined. As a white dwarf cools down, $L_{\text{extra}}^{\text{Ne}}$ does not change much, whereas L_{surf} drops quickly with temperature [e.g., Figure 2 of 212]. So, if no suppression of ^{22}Ne settling, the two luminosities will meet at some temperature. Around this meeting point is the effective zone of ^{22}Ne settling, where the white dwarf cooling rate is influenced significantly. On the other hand, the meeting temperature is a function of white dwarf mass. We derive this temperature–mass relation in 5 and translate it into H–R diagram coordinates. In Figure V-16 we show the results for $X(^{22}\text{Ne}) = 0.014$ and 0.020 ($[M/H] = 0$ and 0.15 in the progenitor stars), $D/D_s = 3.5$, C/O-core white dwarfs. The effective zone of ^{22}Ne settling is indeed highly peaked, and it matches the position and shape of the Q branch well.

Crystallization is a mechanism that may suppress the ^{22}Ne settling by reducing its mobility in the plasma [e.g., 212, 213]. Therefore, in order to see a strong effect of ^{22}Ne settling, $L_{\text{extra}}^{\text{Ne}}$ must be high enough to let the meeting point precede crystallization (see 5). Because ^{22}Ne settling favors C/O-core and previously metal-rich white dwarfs versus O/Ne-core and/or previously metal-poor white dwarfs, and crystallization sets in earlier in O/Ne-core white dwarfs than C/O-core white dwarfs, the delay effect of ^{22}Ne settling is indeed selective. It is worth noting that high-mass C/O-core white dwarfs are believed not to be singly evolved [e.g., 200, 201], which means that if the extra cooling delay is really caused by ^{22}Ne settling, then the extra-delayed white dwarfs should originate from double-WD mergers, i.e., our setup 1 is correct.

The gravitational energy of ^{22}Ne stored in 1.0 and 1.2 M_\odot white dwarfs ($Z = 0.02$) are 6.8×10^{47} and 1.5×10^{48} ergs [212]; the surface luminosity of white dwarfs on the Q branch is $10^{-3.2}$ and $10^{-2.7} L_\odot$ for the two masses. If crystallization sets in later than this luminosity, ^{22}Ne settling can stop their cooling for around 8.9 and 6.2 Gyr, respectively, close to our observational constraint for the extra cooling delay. Existing numerical simulations [213–216] give shorter delays (0.2–4.1 Gyr) for white dwarfs with even the highest possible $L_{\text{extra}}^{\text{Ne}}$. However, the delay time is sensitive to the choice of D/D_s and temperature of crystallization, but existing models have only sparsely sampled the parameter space. Moreover, for the two-component C/O plasma, the updated phase diagram [217, 218] suggests a much lower melting temperature than the widely used phase diagram of Segretain & Chabrier [219] and the naive prescription of using the same condition as in one-component plasma ($\Gamma = 178$). This low melting temperature means a later crystallization, which can lengthen the delay of ^{22}Ne settling.

In summary, we propose ^{22}Ne settling as a promising candidate for the physical origin of the extra cooling delay. ^{22}Ne settling has a more significant effect in C/O-core white dwarfs, which suggests that the extra-delayed white dwarfs are also merger products. To test our idea, detailed cooling models of high-mass C/O white dwarfs are needed.

The peak of ^{22}Ne -settling effect

The delay effect depends on the fractional contribution $L_{\text{extra}}/L_{\text{surf}}$ of the extra source luminosity to the surface luminosity of the white dwarf because the more this extra energy contributes, the less does the white dwarf need to consume its thermal energy and to cool down. The pile-up factor of this effect can be expressed as

$$A = \frac{\zeta^{-1}}{\zeta_0^{-1}} - 1 = \frac{1}{L_{\text{surf}}/L_{\text{extra}} - 1}, \quad (\text{V.49})$$

where A is the same as defined in Equation V.23, ζ and ζ_0 are the cooling rates with and without the extra energy. Assuming no crystallization suppression, $L_{\text{extra}}^{\text{Ne}}$ will meet L_{surf} at some surface temperature T_{eff} . If this occurs, ^{22}Ne settling will stop the white dwarf cooling at this temperature until ^{22}Ne is exhausted, creating a peaked delay effect. Here, we estimate the dependence of this meeting temperature as a function of white dwarf mass, which can be translated into a curve on the H–R diagram.

According to Bildsten & Hall [212], the energy release of ^{22}Ne settling can be expressed by

$$L_{\text{extra}}^{\text{Ne}} = \int_0^R FVn(^{22}\text{Ne})4\pi r^2 dr, \quad (\text{V.50})$$

where $F = 2m_p g_r$ is the net force felt by each nucleus of ^{22}Ne , g_r is the gravity at radius r ; $V = (D/D_s)18m_p g_r / (Ze\Gamma^{1/3}\sqrt{4\pi\rho})$ is the drift velocity of the settling, $\Gamma \equiv (Ze)^2 / (akT) \propto \rho^{1/3} Z^2 / (TA^{1/3})$ is the Coulomb coupling parameter, D is the inter-diffusion coefficient of ^{22}Ne , and D_s is the one-component self-diffusion coefficient, which can be used as a reference value. Substituting these quantities in Equation V.50, we obtain:

$$L_{\text{extra}}^{\text{Ne}} \propto \int_0^R g_r \cdot \frac{D}{D_s} X(^{22}\text{Ne}) \frac{A^{0.11}}{Z^{1.67}} g_r \rho^{-0.61} T_c^{0.33} \cdot \rho r^2 dr, \quad (\text{V.51})$$

where X is the element abundance in mass, T_c is the core temperature. For the main composition of a white dwarf, the charge-to-mass ratio $Z/A = 0.5$ is a constant. Assuming the following

approximations: $g_r \sim g \sim M/R^2$, $\rho \sim g/R$, $\int_0^R dr \sim R$, we obtain

$$L_{\text{extra}}^{\text{Ne}} \propto \frac{D}{D_s} X(^{22}\text{Ne}) Z^{-1.56} T_c^{0.33} g^{2.39} R^{0.61} \cdot R^2 \quad (\text{V.52})$$

Before the convective coupling between the core and atmosphere, the core temperature T_c scales with the surface temperature T_{eff} and gravity g as

$$T_c \propto (T_{\text{eff}}^4/g)^{0.41}, \quad (\text{V.53})$$

which is obtained empirically from existing white dwarf models [145]. This scaling relation is more realistic than the one used by Mestel [220]. Substituting Equation V.53 for T_c , we obtain

$$L_{\text{extra}}^{\text{Ne}} \propto \frac{D}{D_s} X(^{22}\text{Ne}) Z^{-1.56} T_{\text{eff}}^{0.55} g^{2.25} R^{0.61} \cdot R^2. \quad (\text{V.54})$$

We checked the simulation results from Figures 7 and 8 of García-Berro *et al.* [214] and found that our scaling relation is accurate to within 10%, which is sufficient for our purpose.

When the surface luminosity $L_{\text{surf}} \propto T_{\text{eff}}^4 \cdot R^2$ is set equal to $L_{\text{extra}}^{\text{Ne}}$, we obtain

$$T_{\text{eff}} = K^{0.29} g^{0.65} R^{0.18}, \quad (\text{V.55})$$

where $K \propto \frac{D}{D_s} X(^{22}\text{Ne}) Z^{-1.56}$ is a constant that has no relevance to the white dwarf mass M , and both g and R are determined by M . The proportional factor within K can be evaluated from an existing simulation of ^{22}Ne settling.

Discussion

In this section, we discuss two other observational features of the Q branch: the concentration of DQ white dwarfs, and the lack of wide-binary systems. Both of them support the idea that the extra-delayed white dwarfs may also be double-WD merger products, which has been suggested from the ^{22}Ne -settling explanation.

Concentration of DQ white dwarfs on the Q branch

The Q branch is named after the presence of DQ-type white dwarfs [137]. To explore this dimension, we cross-match our white dwarf sample with the Montreal white dwarf database, MWDD [221]¹⁰. We note that most high-mass DQs are concentrated on the branch (Figure V-17) and the fraction of fast DQs on the branch is very high (Table V-III). Therefore, all of these Q-branch DQs are likely to belong to the extra-delayed population. However, *not all extra-delayed white dwarfs are DQs*. We estimate the fraction of DQs in the extra-delayed population to be

$$F_{\text{DQ}} = 19/(76F_{\text{extra}}) = (53 \pm 16) \%, \quad (\text{V.56})$$

based on the total number of DQs and DAs in Table V-III. Changing the distance limit of the sample does not influence this result much. It remains unclear but is of further interest to investigate the reason why half of the extra-delayed white dwarfs are DQs while the other half are DAs.

¹⁰<http://www.montrealwhitedwarfdatabase.org>

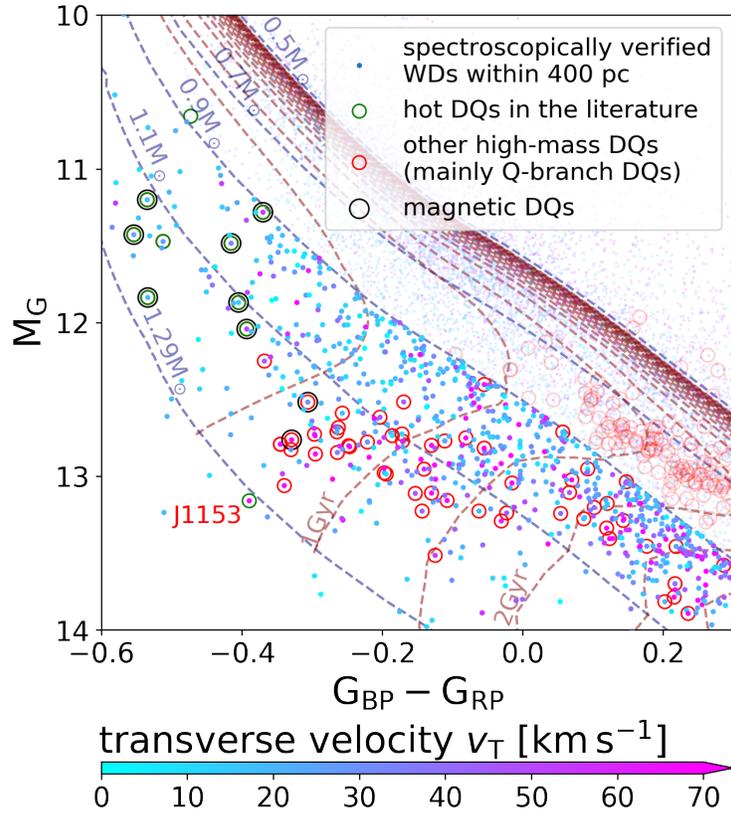


Figure V-17. A part of the H–R diagram showing the spectroscopically verified WDs, with Q-branch DQ and hot-DQ white dwarfs marked by red and green open circles. The dots without circles are mostly DA white dwarfs. We estimate that half of the extra-delayed white dwarfs are DAs, because half of the fast-moving white dwarfs on the Q branch are DAs. We also mark the known magnetic DQs with larger black circles. Note that the mass range here ($> 0.9 M_{\odot}$) is larger than for the sample in our main analysis (1.08–1.23 M_{\odot}).

Table V-III. The statistics of velocity and spectral type of white dwarfs on the Q branch. The fraction of fast DQs is consistent with it belonging purely to the extra-delayed population.

250 pc spectroscopic sample	all	DQ	DA
all v_T	76	19	53
$v_T > 50 \text{ km s}^{-1}$	23	8	14
	$30 \pm 6\%$	$42 \pm 15\%$	$26 \pm 7\%$
$v_T > 60 \text{ km s}^{-1}$	16	7	8
	$21 \pm 5\%$	$37 \pm 14\%$	$15 \pm 6\%$
$v_T > 70 \text{ km s}^{-1}$	9	2	6
	$12 \pm 4\%$	$11 \pm 7\%$	$11 \pm 5\%$

The DQs on the Q-branch are anomalous because the convection zone in a normal white dwarf with similar temperature is not deep enough to dredge up carbon [222]. In a similar way, the hot-DQ white dwarfs discovered by Dufour *et al.* [223] are also abnormal. In Figure V-17 we show the distributions of the Q-branch DQs and hot-DQs on the H–R diagram. Below, we argue that although these two groups of DQs are observationally different, they may be related through an evolutionary relation.

The hot-DQs and Q-branch DQs appear to be different in some aspects. Hot-DQ white dwarfs are characterized by the high temperature ($>18,000$ K), highly carbon-dominant atmosphere [224], high rate of having a magnetic field [225, 226], high rate of being variable [e.g., 227–230], and rarity [e.g., 231]. In contrast, the Q-branch DQ white dwarfs are concentrated on the Q branch, have helium-dominant atmospheres with 10^{-4} – 10^{-1} carbon [232–234], and have undetectable or no magnetic field (see Figure V-17; a caveat for the magnetic field is that most hot-DQs have been examined with high-resolution spectroscopy, so their magnetic fields are more likely to be found). As for kinematics, hot-DQs are mildly faster than normal white dwarfs, which is an indication of being merger products [133], whereas Q-branch DQs are much faster, which needs the long extra cooling delay to explain. Dunlap & Clemens [133] discussed one strange hot-DQ (SDSS J115305.47+005645.8 or J1153) with a very high proper motion. We note that J1153 has not been reported to have magnetic field or variability and lies on the Q branch (Figure V-17), which means that J1153 can be classified as a Q-branch DQ.

We now turn to the similarities between Q-branch DQs and hot-DQs: both of them have high masses, and both of them might have a merger origin. These two similarities raise a serious question: *are the Q-branch DQs evolved from the hot-DQs?* We use number counts to explore this possibility. Hot-DQs are rare; based on a spectroscopically verified white dwarfs sample (as shown in Figure V-17), we find that the fraction of hot-DQs is $8/203 = 4.0 \pm 1.4\%$ in the region earlier than the Q branch ($\tau_{\text{phot}} < 0.5$ Gyr, $m_{\text{WD}} > 0.9 M_{\odot}$). As a comparison, our estimate of the extra-delayed population fraction (f_{extra}) in this region is 6.4% (Equation V.44) and about half of them are DQs (Equation V.56). So, these number counts are consistent with the scenario that hot-DQs are the evolutionary counterparts of Q-branch DQs.

In summary, based on the velocity distribution, we argue that all of Q-branch DQs belong to the extra-delayed population, and they account for $53 \pm 16\%$ of this population. In terms of observational properties, Q-branch DQs form a new class of DQ white dwarfs in addition to hot-DQs and the well-understood standard-mass DQs. However, number counts show that hot-DQs may

evolve into Q-branch DQs, and both of them are likely to originate from double-WD mergers.

Lack of wide binaries on the Q branch

One additional way to test if the extra-delayed white dwarfs are also merger products is to check the wide binary fraction. The kick velocity of a few km s^{-1} [estimated from the results of 152] from a merger may destroy many wide-separation binaries, making the wide-binary fraction lower. Because the extra-delayed population is significantly enhanced on the Q branch, if the extra-delayed white dwarfs are double-WD merger products, one would expect to see a lower wide-binary fraction on the Q branch.

We cross-match the wide binaries in *Gaia* DR2 [158, 235] with our high-mass white dwarf sample. In the early, Q-branch, and late segments, we find 5, 4, and 7 white dwarfs with wide-separation companions out of 309, 510, and 251 white dwarfs, respectively. So, the wide-binary fraction *on* the Q branch is $0.8 \pm 0.4\%$, 2σ lower than the value *off* the branch ($2.2 \pm 0.5\%$). If we assume that the extra-delayed population contributes no wide-binary system, then the wide-binary fraction of the normal-cooling populations *on* the Q branch becomes $4/[510 \times (1 - F_{\text{extra}})] = (1.7 \pm 0.8)\%$, consistent with the *off*-branch value $2.2 \pm 0.5\%$ within 1σ . Therefore, the fraction of wide binaries provides additional support for the idea that the extra-delayed white dwarfs are double-WD merger products.

Conclusion

The white dwarf H–R diagram derived from *Gaia* data has revealed a number-density enhancement of high-mass white dwarfs, called the Q branch. This branch coincides with the crystallization branch, but it is more peaked than what crystallization can create. Adding transverse-velocity information to the H–R diagram, we find a clear excess of fast white dwarfs on the Q branch (Figure V-7). According to the age–velocity–dispersion relation (AVR) of Milky Way disk stars, these fast white dwarfs are much older than their photometric isochrone ages. Therefore, both the number count and velocity distribution suggest an extra cooling delay on the Q branch.

Motivated by these simple observations, we build a Bayesian model to quantitatively investigate this extra cooling delay. Because double-WD merger products also contribute to high-mass white dwarfs, we consider in our model both the extra cooling delay and the double-WD merger delay. Our model includes three white dwarf populations: one with no evolutionary delay, one with only the merger delay, and one with the extra cooling delay. We explore both situations in which all (setup 1) and none (setup 2) of the extra-delayed white dwarfs also have the merger delay. Our statistical model uses the discrepancy between the dynamical age inferred from transverse velocity and the photometric age to constraint the fraction of each white dwarf population and the length of the extra cooling delay. To eliminate selection effects, we model the conditional probability distribution of the transverse velocity of each white dwarf given its H–R diagram coordinate and spatial position. To avoid systematic errors of the model, we fit the solar motion and the anisotropic AVR together with the main parameters of interest.

We feed the model with 1070 high-mass white dwarfs ($1.08\text{--}1.23 M_{\odot}$, $0.1 \text{ Gyr} < \tau_{\text{phot}} < 3.5 \text{ Gyr}$, and $d < 250 \text{ pc}$) selected from *Gaia* DR2. Having checked that the AVR and solar motion parameters are all in agreement with standard values from the literature and that our best-fitting model provides a good fit to the observed velocity distribution, we find

1. about 6% of the high-mass white dwarfs experience an extra cooling delay that significantly slows down their cooling and makes them stay on the Q branch for about 8 Gyr;
2. in the Q-branch region, an enhanced fraction (about a half) of the white dwarfs are extra-delayed due to the pile-up effect;
3. half of the extra-delayed white dwarfs are DQs;
4. as a byproduct of our analysis, the double-WD merger fraction is estimated to be about 20% in our mass range.

The results for the two setups are similar.

This previously unreported extra cooling delay on the Q branch is a challenge to the white dwarf cooling model and our understanding of white dwarf physics. We propose that ^{22}Ne settling [212] could be the physical origin of this extra cooling delay. ^{22}Ne settling favors C/O-core versus O/Ne-core white dwarfs, suggesting that the extra-delayed white dwarfs are also double-WD merger products, i.e., our setup 1 is correct. This idea is also supported by the concentration of DQ white dwarfs and lack of wide-separation binaries on the Q branch also support this idea. To further investigate the nature of this extra cooling delay, detailed cooling models for $m_{\text{WD}} > 1.1 M_{\odot}$ C/O white dwarfs with the ^{22}Ne settling will be needed.

High-mass white dwarfs have been used to explore the AVR, star formation history, and white dwarf mass distribution. Given the existence of the extra cooling delay, the relevant results of those functions in the literature should be reconsidered. In future analyses of these functions, one could reduce the influence of the extra cooling delay by using only the high-mass white dwarfs above the Q branch or modeling the extra cooling delay.

References

1. Mallat, S. Group Invariant Scattering. *Communications on Pure and Applied Mathematics* **65**, 1331–1398. arXiv: [1101.2286](https://arxiv.org/abs/1101.2286) [[math.FA](#)]. <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.21413> (2012).
2. Andén, J. & Mallat, S. Multiscale Scattering for Audio Classification. *International Society for Music Information Retrieval Conference*, 657–662 (Oct. 2011).
3. Bruna, J. & Mallat, S. Invariant Scattering Convolution Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**, 1872–1886. arXiv: [1203.1513](https://arxiv.org/abs/1203.1513) [[cs.CV](#)] (2013).
4. Sifre, L. & Mallat, S. *the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1233–1240 (2013).
5. Andén, J. & Mallat, S. Deep Scattering Spectrum. *IEEE Transactions on Signal Processing* **62**, 4114–4128 (2014).
6. Bruna, J., Mallat, S., Bacry, E. & J.-F., M. *The Annals of Statistics* **43**, 323. arXiv: [1311.4104](https://arxiv.org/abs/1311.4104) [[stat.ME](#)] (2015).
7. Hirn, M., Mallat, S. & Poilvert, N. Wavelet Scattering Regression of Quantum Chemical Energies. *Multiscale Modeling & Simulation* **15**, 827–863 (2017).
8. Eickenberg, M., Exarchakis, G., Hirn, M., Mallat, S. & Thiry, L. Solid harmonic wavelet scattering for predictions of molecule properties. **148**, 241732. arXiv: [1805.00571](https://arxiv.org/abs/1805.00571) [[physics.chem-ph](#)] (June 2018).
9. Sinz, P. *et al.* Wavelet Scattering Networks for Atomistic Systems with Extrapolation of Material Properties. *arXiv e-prints*, arXiv:2006.01247. arXiv: [2006.01247](https://arxiv.org/abs/2006.01247) [[physics.comp-ph](#)] (June 2020).
10. Gama, F., Ribeiro, A. & Bruna, J. Diffusion Scattering Transforms on Graphs. *arXiv e-prints*, arXiv:1806.08829. arXiv: [1806.08829](https://arxiv.org/abs/1806.08829) [[cs.LG](#)] (June 2018).
11. Kavalerov, I., Li, W., Czaja, W. & Chellappa, R. Three-Dimensional Fourier Scattering Transform and Classification of Hyperspectral Images. *arXiv e-prints*, arXiv:1906.06804. arXiv: [1906.06804](https://arxiv.org/abs/1906.06804) [[cs.CV](#)] (June 2019).
12. Allys, E. *et al.* The RWST, a comprehensive statistical description of the non-Gaussian structures in the ISM. **629**, A115. arXiv: [1905.01372](https://arxiv.org/abs/1905.01372) [[astro-ph.CO](#)] (Sept. 2019).
13. Saydjari, A. K. *et al.* Classification of Magnetohydrodynamic Simulations using Wavelet Scattering Transforms. *arXiv e-prints*, arXiv:2010.11963. arXiv: [2010.11963](https://arxiv.org/abs/2010.11963) [[astro-ph.GA](#)] (Oct. 2020).
14. Regaldo-Saint Blancard, B., Levrier, F., Allys, E., Bellomi, E. & Boulanger, F. Statistical description of dust polarized emission from the diffuse interstellar medium. A RWST approach. **642**, A217. arXiv: [2007.08242](https://arxiv.org/abs/2007.08242) [[astro-ph.CO](#)] (Oct. 2020).

15. Cheng, S., Ting, Y.-S., Ménard, B. & Bruna, J. A new approach to observational cosmology using the scattering transform. arXiv: [2006.08561](https://arxiv.org/abs/2006.08561) [[astro-ph.CO](#)] (Oct. 2020).
16. Cheng, S. & Ménard, B. Weak lensing scattering transform: dark energy and neutrino mass sensitivity. *arXiv e-prints*, arXiv:2103.09247. arXiv: [2103.09247](https://arxiv.org/abs/2103.09247) [[astro-ph.CO](#)] (Mar. 2021).
17. Bruna, J. & Mallat, S. Multiscale sparse microcanonical models. *Mathematical Statistics and Learning* **1**, 257–315 (2019).
18. Waldspurger, I. Exponential decay of scattering coefficients. *arXiv e-prints*, arXiv:1605.07464. arXiv: [1605.07464](https://arxiv.org/abs/1605.07464) [[math.FA](#)] (May 2016).
19. Hubel, D. & Wiesel, T. Receptive fields and functional architecture of monkey striate cortex. *The Journal of physiology* **195**, 215–243 (1968).
20. Krizhevsky, A., Sutskever, I. & Hinton, G. E. *ImageNet Classification with Deep Convolutional Neural Networks* in *Advances in Neural Information Processing Systems* (eds Pereira, F., Burges, C. J. C., Bottou, L. & Weinberger, K. Q.) **25** (Curran Associates, Inc., 2012). <https://proceedings.neurips.cc/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf>.
21. Olshausen, B. & Field, D. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–609 (June 1996).
22. Carron, J. On the Incompleteness of the Moment and Correlation Function Hierarchy as Probes of the Lognormal Field. **738**, 86. arXiv: [1105.4467](https://arxiv.org/abs/1105.4467) [[astro-ph.CO](#)] (Sept. 2011).
23. Mallat, S. *Recursive Interferometric Representations* in *Proceedings of 18th European Signal Processing Conference, Denmark* (2010), 716–720.
24. Mallat, S. *A wavelet tour of signal processing. The Sparse Way*. Third (Academic Press, Boston, 2009).
25. Mallat, S., Zhang, S. & Rochette, G. Phase Harmonic Correlations and Convolutional Neural Networks. *arXiv e-prints*, arXiv:1810.12136. arXiv: [1810.12136](https://arxiv.org/abs/1810.12136) [[eess.SP](#)] (Oct. 2018).
26. Portilla, J. & Simoncelli, E. A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients. *International Journal of Computer Vision* **40**, 49–70 (Oct. 2000).
27. Peebles, P. J. E. *The large-scale structure of the universe* (1980).
28. Spergel, D. N. *et al.* First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Determination of Cosmological Parameters. **148**, 175–194. arXiv: [astro-ph/0302209](https://arxiv.org/abs/astro-ph/0302209) [[astro-ph](#)] (Sept. 2003).
29. Eisenstein, D. J. *et al.* Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies. **633**, 560–574. arXiv: [astro-ph/0501171](https://arxiv.org/abs/astro-ph/0501171) [[astro-ph](#)] (Nov. 2005).
30. Rimes, C. D. & Hamilton, A. J. S. Information content of the non-linear matter power spectrum. **360**, L82–L86. arXiv: [astro-ph/0502081](https://arxiv.org/abs/astro-ph/0502081) [[astro-ph](#)] (June 2005).
31. Rimes, C. D. & Hamilton, A. J. S. Information content of the non-linear power spectrum: the effect of beat-coupling to large scales. **371**, 1205–1215. arXiv: [astro-ph/0511418](https://arxiv.org/abs/astro-ph/0511418) [[astro-ph](#)] (Sept. 2006).
32. Neyrinck, M. C., Szapudi, I. & Rimes, C. D. Information content in the halo-model dark-matter power spectrum. **370**, L66–L70. arXiv: [astro-ph/0604282](https://arxiv.org/abs/astro-ph/0604282) [[astro-ph](#)] (July 2006).

33. Neyrinck, M. C. & Szapudi, I. Information content in the halo-model dark-matter power spectrum - II. Multiple cosmological parameters. **375**, L51–L55. arXiv: [astro-ph/0610211](#) [[astro-ph](#)] (Feb. 2007).
34. Springel, V. *et al.* Simulations of the formation, evolution and clustering of galaxies and quasars. **435**, 629–636. arXiv: [astro-ph/0504097](#) [[astro-ph](#)] (June 2005).
35. Bernardeau, F., Mellier, Y. & van Waerbeke, L. Detection of non-Gaussian signatures in the VIRMOS-DESCART lensing survey. **389**, L28–L32. arXiv: [astro-ph/0201032](#) [[astro-ph](#)] (July 2002).
36. Takada, M. & Jain, B. The Three-Point Correlation Function for Spin-2 Fields. **583**, L49–L52. arXiv: [astro-ph/0210261](#) [[astro-ph](#)] (Jan. 2003).
37. Semboloni, E. *et al.* Weak lensing from space: first cosmological constraints from three-point shear statistics. **410**, 143–160. arXiv: [1005.4941](#) [[astro-ph.CO](#)] (Jan. 2011).
38. Fu, L. *et al.* CFHTLenS: cosmological constraints from a combination of cosmic shear two-point and three-point correlations. **441**, 2725–2743. arXiv: [1404.5469](#) [[astro-ph.CO](#)] (July 2014).
39. Sefusatti, E., Crocce, M., Pueblas, S. & Scoccimarro, R. Cosmology and the bispectrum. **74**, 023522. arXiv: [astro-ph/0604505](#) [[astro-ph](#)] (July 2006).
40. Welling, M. Robust Higher Order Statistics. *AISTATS*, 7 (2005).
41. Neyrinck, M. C., Szapudi, I. & Szalay, A. S. Rejuvenating Power Spectra. II. The Gaussianized Galaxy Density Field. **731**, 116. arXiv: [1009.5680](#) [[astro-ph.CO](#)] (Apr. 2011).
42. Simpson, F., James, J. B., Heavens, A. F. & Heymans, C. Clipping the Cosmos: The Bias and Bispectrum of Large Scale Structure. **107**, 271301. arXiv: [1107.5169](#) [[astro-ph.CO](#)] (Dec. 2011).
43. Carron, J. On the assumption of Gaussianity for cosmological two-point statistics and parameter dependent covariance matrices. **551**, A88. arXiv: [1204.4724](#) [[astro-ph.CO](#)] (Mar. 2013).
44. Giblin, B. *et al.* KiDS-450: enhancing cosmic shear with clipping transformations. **480**, 5529–5549. arXiv: [1805.12084](#) [[astro-ph.CO](#)] (Nov. 2018).
45. Mecke, K. R., Buchert, T. & Wagner, H. Robust morphological measures for large-scale structure in the Universe. **288**, 697–704. arXiv: [astro-ph/9312028](#) [[astro-ph](#)] (Aug. 1994).
46. Hikage, C. *et al.* Minkowski Functionals of SDSS Galaxies I : Analysis of Excursion Sets. **55**, 911–931. arXiv: [astro-ph/0304455](#) [[astro-ph](#)] (Oct. 2003).
47. Shirasaki, M. & Yoshida, N. Statistical and Systematic Errors in the Measurement of Weak-Lensing Minkowski Functionals: Application to the Canada-France-Hawaii Lensing Survey. **786**, 43. arXiv: [1312.5032](#) [[astro-ph.CO](#)] (May 2014).
48. Kratochvil, J. M. *et al.* Probing cosmology with weak lensing Minkowski functionals. **85**, 103513. arXiv: [1109.6334](#) [[astro-ph.CO](#)] (May 2012).
49. Jain, B. & Van Waerbeke, L. Statistics of Dark Matter Halos from Gravitational Lensing. **530**, L1–L4. arXiv: [astro-ph/9910459](#) [[astro-ph](#)] (Feb. 2000).
50. Marian, L., Smith, R. E. & Bernstein, G. M. The Cosmology Dependence of Weak Lensing Cluster Counts. **698**, L33–L36. arXiv: [0811.1991](#) [[astro-ph](#)] (June 2009).

51. Kratochvil, J. M., Haiman, Z. & May, M. Probing cosmology with weak lensing peak counts. **81**, 043519. arXiv: [0907.0486 \[astro-ph.CO\]](#) (Feb. 2010).
52. Liu, J. *et al.* Cosmology constraints from the weak lensing peak counts and the power spectrum in CFHTLenS data. **91**, 063507. arXiv: [1412.0757 \[astro-ph.CO\]](#) (Mar. 2015).
53. Liu, X. *et al.* Cosmological constraints from weak lensing peak statistics with Canada-France-Hawaii Telescope Stripe 82 Survey. **450**, 2888–2902. arXiv: [1412.3683 \[astro-ph.CO\]](#) (July 2015).
54. Pisani, A. *et al.* Cosmic voids: a novel probe to shed light on our Universe. **51**, 40. arXiv: [1903.05161 \[astro-ph.CO\]](#) (May 2019).
55. Lecun, Y., Bottou, L., Bengio, Y. & Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**, 2278–2324 (1998).
56. Gupta, A., Matilla, J. M. Z., Hsu, D. & Haiman, Z. Non-Gaussian information from weak lensing data via deep learning. **97**, 103515. arXiv: [1802.01212 \[astro-ph.CO\]](#) (May 2018).
57. Ribli, D., Pataki, B. Á. & Csabai, I. An improved cosmological parameter inference scheme motivated by deep learning. *Nature Astronomy* **3**, 93–98. arXiv: [1806.05995 \[astro-ph.CO\]](#) (Jan. 2019).
58. Ribli, D. *et al.* Weak lensing cosmology with convolutional neural networks on noisy data. **490**, 1843–1860. arXiv: [1902.03663 \[astro-ph.CO\]](#) (Dec. 2019).
59. Bartelmann, M. & Schneider, P. Weak gravitational lensing. **340**, 291–472. arXiv: [astro-ph/9912508 \[astro-ph\]](#) (Jan. 2001).
60. Dodelson, S. *Gravitational Lensing* (2017).
61. Kilbinger, M. Cosmology with cosmic shear observations: a review. *Reports on Progress in Physics* **78**, 086901. arXiv: [1411.0115 \[astro-ph.CO\]](#) (July 2015).
62. Mandelbaum, R. Weak Lensing for Precision Cosmology. **56**, 393–433. arXiv: [1710.03235 \[astro-ph.CO\]](#) (Sept. 2018).
63. Zorrilla Matilla, J. M., Haiman, Z., Hsu, D., Gupta, A. & Petri, A. Do dark matter halos explain lensing peaks? **94**, 083506. arXiv: [1609.03973 \[astro-ph.CO\]](#) (Oct. 2016).
64. Allys, E. *et al.* New Interpretable Statistics for Large Scale Structure Analysis and Generation. *arXiv e-prints*, arXiv:2006.06298. arXiv: [2006.06298 \[astro-ph.CO\]](#) (June 2020).
65. van Waerbeke, L. Noise properties of gravitational lens mass reconstruction. **313**, 524–532. arXiv: [astro-ph/9909160 \[astro-ph\]](#) (Apr. 2000).
66. Kilbinger, M. *et al.* CFHTLenS: combined probe cosmological model comparison using 2D weak gravitational lensing. **430**, 2200–2220. arXiv: [1212.3338 \[astro-ph.CO\]](#) (Apr. 2013).
67. Abbott, T. M. C. *et al.* The Dark Energy Survey: Data Release 1. **239**, 18. arXiv: [1801.03181 \[astro-ph.IM\]](#) (Dec. 2018).
68. Mandelbaum, R. *et al.* The first-year shear catalog of the Subaru Hyper Suprime-Cam Subaru Strategic Program Survey. **70**, S25. arXiv: [1705.06745 \[astro-ph.CO\]](#) (Jan. 2018).
69. Andreux, M. *et al.* Kymatio: Scattering transforms in python. *Journal of Machine Learning Research* **21**, 1–6. <http://jmlr.org/papers/v21/19-047.html> (2020).
70. Petri, A. Mocking the weak lensing universe: The LensTools Python computing package. *Astronomy and Computing* **17**, 73–79. arXiv: [1606.01903 \[astro-ph.CO\]](#) (Oct. 2016).

71. Lin, C.-A. & Kilbinger, M. A new model to predict weak-lensing peak counts. II. Parameter constraint strategies. **583**, A70. arXiv: [1506.01076](https://arxiv.org/abs/1506.01076) [[astro-ph.CO](#)] (Nov. 2015).
72. Fisher, R. A. The Logic of Inductive Inference. *Journal of the Royal Statistical Society* **98**, 39–82. <http://www.jstor.org/stable/2342435> (1935).
73. Tegmark, M., Taylor, A. N. & Heavens, A. F. Karhunen-Loève Eigenvalue Problems in Cosmology: How Should We Tackle Large Data Sets? **480**, 22–35. arXiv: [astro-ph/9603021](https://arxiv.org/abs/astro-ph/9603021) [[astro-ph](#)] (May 1997).
74. Hartlap, J., Simon, P. & Schneider, P. Why your model parameter confidences might be too optimistic. Unbiased estimation of the inverse covariance matrix. **464**, 399–404. arXiv: [astro-ph/0608064](https://arxiv.org/abs/astro-ph/0608064) [[astro-ph](#)] (Mar. 2007).
75. Planck Collaboration *et al.* Planck 2015 results. XIII. Cosmological parameters. **594**, A13. arXiv: [1502.01589](https://arxiv.org/abs/1502.01589) [[astro-ph.CO](#)] (Sept. 2016).
76. Li, P., Dodelson, S. & Croft, R. A. C. Large scale structure reconstruction with short-wavelength modes. **101**, 083510. arXiv: [2001.02780](https://arxiv.org/abs/2001.02780) [[astro-ph.CO](#)] (Apr. 2020).
77. Bradač, M., Lombardi, M. & Schneider, P. Mass-sheet degeneracy: Fundamental limit on the cluster mass reconstruction from statistical (weak) lensing. **424**, 13–22. arXiv: [astro-ph/0405357](https://arxiv.org/abs/astro-ph/0405357) [[astro-ph](#)] (Sept. 2004).
78. Szegedy, C. *et al.* Intriguing properties of neural networks. *arXiv e-prints*, arXiv:1312.6199. arXiv: [1312.6199](https://arxiv.org/abs/1312.6199) [[cs.CV](#)] (Dec. 2013).
79. Jaynes, E. T. Information Theory and Statistical Mechanics. *Physical Review* **106**, 620–630 (May 1957).
80. Vogeley, M. S. & Szalay, A. S. Eigenmode Analysis of Galaxy Redshift Surveys. I. Theory and Methods. **465**, 34. arXiv: [astro-ph/9601185](https://arxiv.org/abs/astro-ph/9601185) [[astro-ph](#)] (July 1996).
81. Liu, J. *et al.* CMB lensing beyond the power spectrum: Cosmological constraints from the one-point probability distribution function and peak counts. **94**, 103501. arXiv: [1608.03169](https://arxiv.org/abs/1608.03169) [[astro-ph.CO](#)] (Nov. 2016).
82. Liu, J. *et al.* MassiveNuS: cosmological massive neutrino simulations. **2018**, 049. arXiv: [1711.10524](https://arxiv.org/abs/1711.10524) [[astro-ph.CO](#)] (Mar. 2018).
83. Chevallier, M. & Polarski, D. Accelerating Universes with Scaling Dark Matter. *International Journal of Modern Physics D* **10**, 213–223. arXiv: [gr-qc/0009008](https://arxiv.org/abs/gr-qc/0009008) [[gr-qc](#)] (Jan. 2001).
84. LSST Science Collaboration *et al.* LSST Science Book, Version 2.0. *arXiv e-prints*, arXiv:0912.0201. arXiv: [0912.0201](https://arxiv.org/abs/0912.0201) [[astro-ph.IM](#)] (Dec. 2009).
85. Kaiser, N. & Squires, G. Mapping the Dark Matter with Weak Gravitational Lensing. **404**, 441 (Feb. 1993).
86. Pires, S. *et al.* Euclid: Reconstruction of weak-lensing mass maps for non-Gaussianity studies. **638**, A141. arXiv: [1910.03106](https://arxiv.org/abs/1910.03106) [[astro-ph.CO](#)] (June 2020).
87. Jeffrey, N., Lanusse, F., Lahav, O. & Starck, J.-L. Deep learning dark matter map reconstructions from DES SV weak lensing data. **492**, 5023–5029. arXiv: [1908.00543](https://arxiv.org/abs/1908.00543) [[astro-ph.CO](#)] (Mar. 2020).
88. Foreman-Mackey, D., Hogg, D. W., Lang, D. & Goodman, J. emcee: The MCMC Hammer. **125**, 306. arXiv: [1202.3665](https://arxiv.org/abs/1202.3665) [[astro-ph.IM](#)] (Mar. 2013).

89. Sellentin, E. & Heavens, A. F. On the insufficiency of arbitrarily precise covariance matrices: non-Gaussian weak-lensing likelihoods. **473**, 2355–2363. arXiv: [1707.04488 \[astro-ph.CO\]](#) (Jan. 2018).
90. Hahn, C. *et al.* Likelihood non-Gaussianity in large-scale structure analyses. **485**, 2956–2969. arXiv: [1803.06348 \[astro-ph.CO\]](#) (May 2019).
91. Diaz Rivero, A. & Dvorkin, C. Flow-based likelihoods for non-Gaussian inference. **102**, 103507. arXiv: [2007.05535 \[astro-ph.CO\]](#) (Nov. 2020).
92. Jeffrey, N., Alsing, J. & Lanusse, F. Likelihood-free inference with neural compression of DES SV weak lensing map statistics. **501**, 954–969. arXiv: [2009.08459 \[astro-ph.CO\]](#) (Feb. 2021).
93. Liu, J. & Madhavacheril, M. S. Constraining neutrino mass with the tomographic weak lensing one-point probability distribution function and power spectrum. **99**, 083508. arXiv: [1809.10747 \[astro-ph.CO\]](#) (Apr. 2019).
94. Coulton, W. R., Liu, J., Madhavacheril, M. S., Böhm, V. & Spergel, D. N. Constraining neutrino mass with the tomographic weak lensing bispectrum. **2019**, 043. arXiv: [1810.02374 \[astro-ph.CO\]](#) (May 2019).
95. Li, Z., Liu, J., Matilla, J. M. Z. & Coulton, W. R. Constraining neutrino mass with tomographic weak lensing peak counts. **99**, 063527. arXiv: [1810.01781 \[astro-ph.CO\]](#) (Mar. 2019).
96. Ajani, V. *et al.* Constraining neutrino masses with weak-lensing multiscale peak counts. **102**, 103531. arXiv: [2001.10993 \[astro-ph.CO\]](#) (Nov. 2020).
97. Coulton, W. R., Liu, J., McCarthy, I. G. & Osato, K. Weak lensing minima and peaks: Cosmological constraints and the impact of baryons. **495**, 2531–2542. arXiv: [1910.04171 \[astro-ph.CO\]](#) (Apr. 2020).
98. Marques, G. A. *et al.* Constraining neutrino mass with weak lensing Minkowski Functionals. **2019**, 019. arXiv: [1812.08206 \[astro-ph.CO\]](#) (June 2019).
99. Ajani, V., Starck, J.-L. & Pettorino, V. Starlet l1-norm for weak lensing cosmology. *arXiv e-prints*, arXiv:2101.01542. arXiv: [2101.01542 \[astro-ph.CO\]](#) (Jan. 2021).
100. Aihara, H. *et al.* The Hyper Suprime-Cam SSP Survey: Overview and survey design. **70**, S4. arXiv: [1704.05858 \[astro-ph.IM\]](#) (Jan. 2018).
101. Hikage, C. *et al.* Cosmology from cosmic shear power spectra with Subaru Hyper Suprime-Cam first-year data. **71**, 43. arXiv: [1809.09148 \[astro-ph.CO\]](#) (Apr. 2019).
102. Tanaka, M. *et al.* Photometric redshifts for Hyper Suprime-Cam Subaru Strategic Program Data Release 1. **70**, S9. arXiv: [1704.05988 \[astro-ph.GA\]](#) (Jan. 2018).
103. Hirata, C. & Seljak, U. Shear calibration biases in weak-lensing surveys. **343**, 459–480. arXiv: [astro-ph/0301054 \[astro-ph\]](#) (Aug. 2003).
104. Kaiser, N., Squires, G. & Broadhurst, T. A Method for Weak Lensing Observations. **449**, 460. arXiv: [astro-ph/9411005 \[astro-ph\]](#) (Aug. 1995).
105. Bernstein, G. M. & Jarvis, M. Shapes and Shears, Stars and Smears: Optimal Measurements for Weak Lensing. **123**, 583–618. arXiv: [astro-ph/0107431 \[astro-ph\]](#) (Feb. 2002).
106. Shirasaki, M., Hamana, T., Takada, M., Takahashi, R. & Miyatake, H. Mock galaxy shape catalogues in the Subaru Hyper Suprime-Cam Survey. **486**, 52–69. arXiv: [1901.09488 \[astro-ph.CO\]](#) (June 2019).

107. Takahashi, R. *et al.* Full-sky Gravitational Lensing Simulation for Large-area Galaxy Surveys and Cosmic Microwave Background Experiments. **850**, 24. arXiv: [1706.01472](https://arxiv.org/abs/1706.01472) [[astro-ph.CO](#)] (Nov. 2017).
108. Hinshaw, G. *et al.* Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results. **208**, 19. arXiv: [1212.5226](https://arxiv.org/abs/1212.5226) [[astro-ph.CO](#)] (Oct. 2013).
109. Prochaska, J. X., Cornillon, P. C. & Reiman, D. M. Deep Learning of Sea Surface Temperature Patterns to Identify Ocean Extremes. *Remote Sensing* **13**. <https://www.mdpi.com/2072-4292/13/4/744> (2021).
110. Ma, L. *et al.* Deep learning in remote sensing applications: A meta-analysis and review. *ISPRS journal of photogrammetry and remote sensing* **152**, 166–177 (2019).
111. Zhang, Z. *et al.* Monthly and quarterly sea surface temperature prediction based on gated recurrent unit neural network. *Journal of Marine Science and Engineering* **8**, 249 (2020).
112. Ratnam, J., Dijkstra, H. & Behera, S. K. A machine learning based prediction system for the Indian Ocean Dipole. *Scientific reports* **10**, 1–11 (2020).
113. Yu, X. *et al.* A novel method for sea surface temperature prediction based on deep learning. *Mathematical Problems in Engineering* **2020** (2020).
114. Hubble, E. P. Extragalactic nebulae. **64**, 321–369 (Dec. 1926).
115. Dieleman, S., Willett, K. W. & Dambre, J. Rotation-invariant convolutional neural networks for galaxy morphology prediction. **450**, 1441–1459. arXiv: [1503.07077](https://arxiv.org/abs/1503.07077) [[astro-ph.IM](#)] (June 2015).
116. Zhu, X.-P. *et al.* Galaxy morphology classification with deep convolutional neural networks. **364**, 55. arXiv: [1807.10406](https://arxiv.org/abs/1807.10406) [[astro-ph.GA](#)] (Apr. 2019).
117. McInnes, L., Healy, J. & Melville, J. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv e-prints*, arXiv:1802.03426. arXiv: [1802.03426](https://arxiv.org/abs/1802.03426) [[stat.ML](#)] (Feb. 2018).
118. Cheng, S., Cummings, J. D., Ménard, B. & Toonen, S. Double White Dwarf Merger Products among High-mass White Dwarfs. **891**, 160. arXiv: [1910.09558](https://arxiv.org/abs/1910.09558) [[astro-ph.SR](#)] (Mar. 2020).
119. Marsh, T. R. The discovery of a short-period double-degenerate binary star. **275**, L1–L5. arXiv: [astro-ph/9504088](https://arxiv.org/abs/astro-ph/9504088) [[astro-ph](#)] (July 1995).
120. Marsh, T. R., Dhillon, V. S. & Duck, S. R. Low-Mass White Dwarfs Need Friends - Five New Double-Degenerate Close Binary Stars. **275**, 828 (Aug. 1995).
121. Iben Icko, J., Tutukov, A. V. & Yungelson, L. R. On the Origin of Hydrogen-deficient Supergiants and Their Relation to R Coronae Borealis Stars and Non-DA White Dwarfs. **456**, 750 (Jan. 1996).
122. Han, Z. The formation of double degenerates and related objects. **296**, 1019–1040 (June 1998).
123. Badenes, C. & Maoz, D. The Merger Rate of Binary White Dwarfs in the Galactic Disk. **749**, L11. arXiv: [1202.5472](https://arxiv.org/abs/1202.5472) [[astro-ph.SR](#)] (Apr. 2012).
124. Maoz, D., Hallakoun, N. & Badenes, C. The separation distribution and merger rate of double white dwarfs: improved constraints. **476**, 2584–2590. arXiv: [1801.04275](https://arxiv.org/abs/1801.04275) [[astro-ph.SR](#)] (May 2018).

125. Brown, W. R. *et al.* The ELM Survey. VIII. Ninety-eight Double White Dwarf Binaries. **889**, 49. arXiv: [2002.00064 \[astro-ph.SR\]](#) (Jan. 2020).
126. Lorén-Aguilar, P., Isern, J. & García-Berro, E. High-resolution smoothed particle hydrodynamics simulations of the merger of binary white dwarfs. **500**, 1193–1205 (June 2009).
127. Toonen, S., Hollands, M., Gänsicke, B. T. & Boekholt, T. The binarity of the local white dwarf population. **602**, A16. arXiv: [1703.06893 \[astro-ph.SR\]](#) (June 2017).
128. Temmink, K. D., Toonen, S., Zapartas, E., Justham, S. & Gänsicke, B. T. The fault in our stars: underestimating the age of single white dwarfs due to binary mergers. *arXiv e-prints*, arXiv:1910.05335. arXiv: [1910.05335 \[astro-ph.SR\]](#) (Oct. 2019).
129. Giammichele, N., Bergeron, P. & Dufour, P. Know Your Neighborhood: A Detailed Model Atmosphere Analysis of Nearby White Dwarfs. **199**, 29. arXiv: [1202.5581 \[astro-ph.SR\]](#) (Apr. 2012).
130. Rebassa-Mansergas, A., Rybicka, M., Liu, X.-W., Han, Z. & García-Berro, E. The mass function of hydrogen-rich white dwarfs: robust observational evidence for a distinctive high-mass excess near 1 Msun. **452**, 1637–1642. arXiv: [1506.06802 \[astro-ph.SR\]](#) (Sept. 2015).
131. Tremblay, P.-E. *et al.* The field white dwarf mass distribution. **461**, 2100–2114. arXiv: [1606.05292 \[astro-ph.SR\]](#) (Sept. 2016).
132. Wegg, C. & Phinney, E. S. White dwarf kinematics versus mass. **426**, 427–439. arXiv: [1206.1056](#) (Oct. 2012).
133. Dunlap, B. H. & Clemens, J. C. *Hot DQ White Dwarf Stars as Failed Type Ia Supernovae in 19th European Workshop on White Dwarfs* (eds Dufour, P., Bergeron, P. & Fontaine, G.) **493** (June 2015), 547.
134. Nordström, B. *et al.* The Geneva-Copenhagen survey of the Solar neighbourhood. Ages, metallicities, and kinematic properties of 14 000 F and G dwarfs. **418**, 989–1019. arXiv: [astro-ph/0405198 \[astro-ph\]](#) (May 2004).
135. Gaia Collaboration *et al.* The Gaia mission. **595**, A1. arXiv: [1609.04153 \[astro-ph.IM\]](#) (Nov. 2016).
136. Cheng, S., Cummings, J. D. & Ménard, B. A Cooling Anomaly of High-mass White Dwarfs. **886**, 100. arXiv: [1905.12710 \[astro-ph.SR\]](#) (Dec. 2019).
137. Gaia Collaboration *et al.* Gaia Data Release 2. Observational Hertzsprung-Russell diagrams. **616**, A10. arXiv: [1804.09378 \[astro-ph.SR\]](#) (Aug. 2018).
138. Iben Jr., I. & Tutukov, A. V. Supernovae of type I as end products of the evolution of binaries with components of moderate initial mass (M not greater than about 9 solar masses). **54**, 335–372 (Feb. 1984).
139. Webbink, R. F. Double white dwarfs as progenitors of R Coronae Borealis stars and Type I supernovae. **277**, 355–360 (Feb. 1984).
140. Lindgren, L. *et al.* Gaia Data Release 2. The astrometric solution. **616**, A2. arXiv: [1804.09366 \[astro-ph.IM\]](#) (Aug. 2018).
141. Evans, D. W. *et al.* Gaia Data Release 2. Photometric content and validation. **616**, A4. arXiv: [1804.09368 \[astro-ph.IM\]](#) (Aug. 2018).
142. Riello, M. *et al.* Gaia Data Release 2. Processing of the photometric data. **616**, A3. arXiv: [1804.09367 \[astro-ph.IM\]](#) (Aug. 2018).

143. Gentile Fusillo, N. P. *et al.* A Gaia Data Release 2 catalogue of white dwarfs and a comparison with SDSS. **482**, 4570–4591. arXiv: [1807.03315 \[astro-ph.SR\]](#) (Feb. 2019).
144. Gaia Collaboration *et al.* Gaia Data Release 2. Summary of the contents and survey properties. **616**, A1. arXiv: [1804.09365](#) (Aug. 2018).
145. Fontaine, G., Brassard, P. & Bergeron, P. The Potential of White Dwarf Cosmochronology. **113**, 409–435 (Apr. 2001).
146. Holberg, J. B. & Bergeron, P. Calibration of Synthetic Photometry Using DA White Dwarfs. **132**, 1221–1233 (Sept. 2006).
147. Kowalski, P. M. & Saumon, D. Found: The Missing Blue Opacity in Atmosphere Models of Cool Hydrogen White Dwarfs. **651**, L137–L140. eprint: [astro-ph/0609720](#) (Nov. 2006).
148. Tremblay, P.-E., Bergeron, P. & Gianninas, A. An Improved Spectroscopic Analysis of DA White Dwarfs from the Sloan Digital Sky Survey Data Release 4. **730**, 128. arXiv: [1102.0056 \[astro-ph.SR\]](#) (Apr. 2011).
149. Camisassa, M. E. *et al.* The evolution of ultra-massive white dwarfs. **625**, A87. arXiv: [1807.03894 \[astro-ph.SR\]](#) (May 2019).
150. Cummings, J. D., Kalirai, J. S., Tremblay, P.-E., Ramirez-Ruiz, E. & Choi, J. The White Dwarf Initial-Final Mass Relation for Progenitor Stars from 0.85 to 7.5 Msun. **866**, 21. arXiv: [1809.01673 \[astro-ph.SR\]](#) (Oct. 2018).
151. Choi, J. *et al.* Mesa Isochrones and Stellar Tracks (MIST). I. Solar-scaled Models. **823**, 102. arXiv: [1604.08592 \[astro-ph.SR\]](#) (June 2016).
152. Dan, M., Rosswog, S., Brüggen, M. & Podsiadlowski, P. The structure and fate of white dwarf merger remnants. **438**, 14–34. arXiv: [1308.1667 \[astro-ph.HE\]](#) (Feb. 2014).
153. Sollima, A. The stellar initial mass function of the solar neighbourhood revealed by Gaia. **489**, 2377–2394. arXiv: [1904.05646 \[astro-ph.GA\]](#) (Oct. 2019).
154. Arenou, F. *et al.* Gaia Data Release 2. Catalogue validation. **616**, A17. arXiv: [1804.09375 \[astro-ph.GA\]](#) (Aug. 2018).
155. Bovy, J. Galactic rotation in Gaia DR1. **468**, L63–L67. arXiv: [1610.07610](#) (June 2017).
156. Maoz, D., Sharon, K. & Gal-Yam, A. The Supernova Delay Time Distribution in Galaxy Clusters and Implications for Type-Ia Progenitors and Metal Enrichment. **722**, 1879–1894. arXiv: [1006.3576](#) (Oct. 2010).
157. Binney, J. & Tremaine, S. *Galactic Dynamics: Second Edition* (Princeton University Press, 2008).
158. El-Badry, K. & Rix, H.-W. Imprints of white dwarf recoil in the separation distribution of Gaia wide binaries. **480**, 4884–4902. arXiv: [1807.06011 \[astro-ph.SR\]](#) (Nov. 2018).
159. Li, W. *et al.* Nearby supernova rates from the Lick Observatory Supernova Search - III. The rate-size relation, and the rates as a function of galaxy Hubble type and colour. **412**, 1473–1507. arXiv: [1006.4613 \[astro-ph.SR\]](#) (Apr. 2011).
160. Portegies Zwart, S. F. & Verbunt, F. Population synthesis of high-mass binaries. **309**, 179–196 (May 1996).
161. Toonen, S., Nelemans, G. & Portegies Zwart, S. Supernova Type Ia progenitors from merging double white dwarfs. Using a new population synthesis model. **546**, A70. arXiv: [1208.6446 \[astro-ph.HE\]](#) (Oct. 2012).

162. Kroupa, P., Tout, C. A. & Gilmore, G. The distribution of low-mass stars in the Galactic disc. **262**, 545–587 (June 1993).
163. Raghavan, D. *et al.* A Survey of Stellar Families: Multiplicity of Solar-type Stars. **190**, 1–42. arXiv: [1007.0414 \[astro-ph.SR\]](#) (Sept. 2010).
164. Duchêne, G. & Kraus, A. Stellar Multiplicity. **51**, 269–310. arXiv: [1303.3028 \[astro-ph.SR\]](#) (Aug. 2013).
165. Abt, H. A. Normal and abnormal binary frequencies. **21**, 343–372 (1983).
166. Heggie, D. C. Binary evolution in stellar dynamics. **173**, 729–787 (Dec. 1975).
167. Toonen, S., Claeys, J. S. W., Mennekens, N. & Ruiter, A. J. PopCORN: Hunting down the differences between binary population synthesis codes. **562**, A14. arXiv: [1311.6503 \[astro-ph.SR\]](#) (Feb. 2014).
168. Ivanova, N. *et al.* Common envelope evolution: where we stand and how we can move forward. **21**, 59. arXiv: [1209.4302 \[astro-ph.HE\]](#) (Feb. 2013).
169. Nelemans, G., Verbunt, F., Yungelson, L. R. & Portegies Zwart, S. F. Reconstructing the evolution of double helium white dwarfs: envelope loss without spiral-in. **360**, 1011–1018 (Aug. 2000).
170. Moe, M. & Di Stefano, R. Mind Your Ps and Qs: The Interrelation between Period (P) and Mass-ratio (Q) Distributions of Binary Stars. **230**, 15. arXiv: [1606.05347 \[astro-ph.SR\]](#) (June 2017).
171. Just, A. & Jahreiß, H. Towards a fully consistent Milky Way disc model - I. The local model based on kinematic and photometric data. **402**, 461–478. arXiv: [0910.3481 \[astro-ph.GA\]](#) (Feb. 2010).
172. McMillan, P. J. Mass models of the Milky Way. **414**, 2446–2457. arXiv: [1102.4340](#) (July 2011).
173. Schwab, J. Evolutionary Models for R Coronae Borealis Stars. **885**, 27. arXiv: [1909.02569 \[astro-ph.SR\]](#) (Nov. 2019).
174. Bhattacharya, M., Mukhopadhyay, B. & Mukerjee, S. Luminosity and cooling of highly magnetized white dwarfs: suppression of luminosity by strong magnetic fields. **477**, 2705–2715. arXiv: [1509.00936 \[astro-ph.SR\]](#) (June 2018).
175. Maoz, D., Badenes, C. & Bickerton, S. J. Characterizing the Galactic White Dwarf Binary Population with Sparsely Sampled Radial Velocity Data. **751**, 143. arXiv: [1202.5467 \[astro-ph.SR\]](#) (June 2012).
176. Maoz, D. & Hallakoun, N. The binary fraction, separation distribution, and merger rate of white dwarfs from SPY. **467**, 1414–1425. arXiv: [1609.02156 \[astro-ph.SR\]](#) (May 2017).
177. Brown, W. R., Kilic, M., Kenyon, S. J. & Gianninas, A. Most Double Degenerate Low-mass White Dwarf Binaries Merge. **824**, 46. arXiv: [1604.04269 \[astro-ph.SR\]](#) (June 2016).
178. Brown, W. R., Gianninas, A., Kilic, M., Kenyon, S. J. & Allende Prieto, C. The ELM Survey. VII. Orbital Properties of Low-Mass White Dwarf Binaries. **818**, 155. arXiv: [1604.04268 \[astro-ph.SR\]](#) (Feb. 2016).
179. Maoz, D., Mannucci, F. & Nelemans, G. Observational Clues to the Progenitors of Type Ia Supernovae. **52**, 107–170. arXiv: [1312.0628](#) (Aug. 2014).

180. Tutukov, A. V., Yungelson, L. R. & Iben Jr., I. The frequencies of supernovae in binaries. **386**, 197–205 (Feb. 1992).
181. Mennekens, N., Vanbeveren, D., De Greve, J. P. & De Donder, E. The delay-time distribution of Type Ia supernovae: a comparison between theory and observation. **515**, A89. arXiv: [1003.2491 \[astro-ph.SR\]](#) (June 2010).
182. Sato, Y. *et al.* A Systematic Study of Carbon-Oxygen White Dwarf Mergers: Mass Combinations for Type Ia Supernovae. **807**, 105. arXiv: [1505.01646 \[astro-ph.HE\]](#) (July 2015).
183. Liu, D., Wang, B., Wu, C. & Han, Z. Merging of a CO WD and a He-rich WD to produce a type Ia supernovae. **606**, A136. arXiv: [1707.07387 \[astro-ph.SR\]](#) (Oct. 2017).
184. Shen, K. J., Kasen, D., Miles, B. J. & Townsley, D. M. Sub-Chandrasekhar-mass White Dwarf Detonations Revisited. **854**, 52. arXiv: [1706.01898 \[astro-ph.HE\]](#) (Feb. 2018).
185. Shen, K. J. *et al.* Three Hypervelocity White Dwarfs in Gaia DR2: Evidence for Dynamically Driven Double-degenerate Double-detonation Type Ia Supernovae. **865**, 15. arXiv: [1804.11163 \[astro-ph.SR\]](#) (Sept. 2018).
186. Perets, H. B., Zenati, Y., Toonen, S. & Bobrick, A. Normal type Ia supernovae from disruptions of hybrid He-CO white-dwarfs by CO white-dwarfs. *arXiv e-prints*, arXiv:1910.07532. arXiv: [1910.07532 \[astro-ph.HE\]](#) (Oct. 2019).
187. Ruiter, A. J., Belczynski, K. & Fryer, C. Rates and Delay Times of Type Ia Supernovae. **699**, 2026–2036. arXiv: [0904.3108 \[astro-ph.SR\]](#) (July 2009).
188. Yungelson, L. R. & Kuranov, A. G. Merging white dwarfs and Type Ia supernovae. **464**, 1607–1632 (Jan. 2017).
189. Shen, K. J., Toonen, S. & Graur, O. The Evolution of the Type Ia Supernova Luminosity Function. **851**, L50. arXiv: [1710.09384 \[astro-ph.HE\]](#) (Dec. 2017).
190. Maoz, D. & Graur, O. Star Formation, Supernovae, Iron, and α : Consistent Cosmic and Galactic Histories. **848**, 25. arXiv: [1703.04540 \[astro-ph.HE\]](#) (Oct. 2017).
191. Kollmeier, J. A. *et al.* SDSS-V: Pioneering Panoptic Spectroscopy. *arXiv e-prints*, arXiv:1711.03234. arXiv: [1711.03234 \[astro-ph.GA\]](#) (Nov. 2017).
192. Bergeron, P. *et al.* On the Measurement of Fundamental Parameters of White Dwarfs in the Gaia Era. **876**, 67. arXiv: [1904.02022 \[astro-ph.SR\]](#) (May 2019).
193. Tremblay, P.-E. *et al.* Core crystallization and pile-up in the cooling sequence of evolving white dwarfs. **565**, 202–205 (Jan. 2019).
194. van Horn, H. M. Crystallization of White Dwarfs. **151**, 227 (Jan. 1968).
195. Garcia-Berro, E., Hernanz, M., Mochkovitch, R. & Isern, J. Theoretical white-dwarf luminosity functions for two phase diagrams of the carbon-oxygen dense plasma. **193**, 141–147 (Mar. 1988).
196. Segretain, L. *et al.* Cooling theory of crystallized white dwarfs. **434**, 641–651 (Oct. 1994).
197. Isern, J., Mochkovitch, R., Garcia-Berro, E. & Hernanz, M. The Physics of Crystallizing White Dwarfs. **485**, 308–312. eprint: [astro-ph/9703028](#) (Aug. 1997).
198. Astropy Collaboration *et al.* Astropy: A community Python package for astronomy. **558**, A33. arXiv: [1307.6212 \[astro-ph.IM\]](#) (Oct. 2013).
199. Astropy Collaboration *et al.* The Astropy Project: Building an Open-science Project and Status of the v2.0 Core Package. **156**, 123. arXiv: [1801.02634 \[astro-ph.IM\]](#) (Sept. 2018).

200. Siess, L. Evolution of massive AGB stars. II. model properties at non-solar metallicity and the fate of Super-AGB stars. **476**, 893–909 (Dec. 2007).
201. Lauffer, G. R., Romero, A. D. & Kepler, S. O. New full evolutionary sequences of H- and He-atmosphere massive white dwarf stars using MESA. **480**, 1547–1562. arXiv: [1807.04774 \[astro-ph.SR\]](#) (Oct. 2018).
202. Bergeron, P. *et al.* A Comprehensive Spectroscopic Analysis of DB White Dwarfs. **737**, 28. arXiv: [1105.5433 \[astro-ph.SR\]](#) (Aug. 2011).
203. Schönrich, R., Binney, J. & Dehnen, W. Local kinematics and the local standard of rest. **403**, 1829–1833. arXiv: [0912.3693](#) (Apr. 2010).
204. Holmberg, J., Nordström, B. & Andersen, J. The Geneva-Copenhagen survey of the solar neighbourhood. III. Improved distances, ages, and kinematics. **501**, 941–947. arXiv: [0811.3982](#) (July 2009).
205. Sharma, S. *et al.* Kinematic Modeling of the Milky Way Using the RAVE and GCS Stellar Surveys. **793**, 51. arXiv: [1405.7435](#) (Sept. 2014).
206. Yu, J. & Liu, C. The age-velocity dispersion relation of the Galactic discs from LAMOST-Gaia data. **475**, 1093–1103. arXiv: [1712.03965](#) (Mar. 2018).
207. Mackereth, J. T. *et al.* Dynamical heating across the Milky Way disc using APOGEE and Gaia. **489**, 176–195. arXiv: [1901.04502 \[astro-ph.GA\]](#) (Oct. 2019).
208. Rowell, N. & Kilic, M. The kinematics of Galactic disc white dwarfs in Gaia DR2. **484**, 3544–3551. arXiv: [1901.04948](#) (Apr. 2019).
209. Anguiano, B. *et al.* The kinematics of the white dwarf population from the SDSS DR12. **469**, 2102–2120. arXiv: [1703.09152 \[astro-ph.GA\]](#) (Aug. 2017).
210. Mor, R., Robin, A. C., Figueras, F., Roca-Fàbrega, S. & Luri, X. Gaia DR2 reveals a star formation burst in the disc 2-3 Gyr ago. **624**, L1. arXiv: [1901.07564 \[astro-ph.GA\]](#) (Apr. 2019).
211. Isern, J., Hernanz, M., Mochkovitch, R. & Garcia-Berro, E. The role of the minor chemical species in the cooling of white dwarfs. **241**, L29–L32 (Jan. 1991).
212. Bildsten, L. & Hall, D. M. Gravitational Settling of ^{22}Ne in Liquid White Dwarf Interiors. **549**, L219–L223. eprint: [astro-ph/0101365](#) (Mar. 2001).
213. Deloye, C. J. & Bildsten, L. Gravitational Settling of ^{22}Ne in Liquid White Dwarf Interiors: Cooling and Seismological Effects. **580**, 1077–1090. eprint: [astro-ph/0207623](#) (Dec. 2002).
214. García-Berro, E., Althaus, L. G., Córscico, A. H. & Isern, J. Gravitational Settling of ^{22}Ne and White Dwarf Evolution. **677**, 473–482. arXiv: [0712.1212](#) (Apr. 2008).
215. Althaus, L. G. *et al.* Evolution of White Dwarf Stars with High-metallicity Progenitors: The Role of ^{22}Ne Diffusion. **719**, 612–621. arXiv: [1006.4170 \[astro-ph.SR\]](#) (Aug. 2010).
216. Camisassa, M. E. *et al.* The Effect of ^{22}Ne Diffusion in the Evolution and Pulsational Properties of White Dwarfs with Solar Metallicity Progenitors. **823**, 158. arXiv: [1604.01744 \[astro-ph.SR\]](#) (June 2016).
217. Horowitz, C. J., Schneider, A. S. & Berry, D. K. Crystallization of Carbon-Oxygen Mixtures in White Dwarf Stars. *Physical Review Letters* **104**, 231101. arXiv: [1005.2441 \[astro-ph.SR\]](#) (June 2010).

218. Hughto, J. *et al.* Direct molecular dynamics simulation of liquid-solid phase equilibria for a three-component plasma. *Physical Review E* **86**, 066413. arXiv: [1211.0891 \[astro-ph.SR\]](#) (Dec. 2012).
219. Segretain, L. & Chabrier, G. Crystallization of binary ionic mixtures in dense stellar plasmas. *Astronomy and Astrophysics* **271**, L13 (Apr. 1993).
220. Mestel, L. On the theory of white dwarf stars. I. The energy sources of white dwarfs. **112**, 583 (Jan. 1952).
221. Dufour, P. *et al.* *The Montreal White Dwarf Database: A Tool for the Community in 20th European White Dwarf Workshop* (eds Tremblay, P.-E., Gaensicke, B. & Marsh, T.) **509** (Mar. 2017), 3. arXiv: [1610.00986 \[astro-ph.SR\]](#).
222. Dufour, P., Bergeron, P. & Fontaine, G. Detailed Spectroscopic and Photometric Analysis of DQ White Dwarfs. **627**, 404–417. eprint: [astro-ph/0503112](#) (July 2005).
223. Dufour, P., Liebert, J., Fontaine, G. & Behara, N. White dwarf stars with carbon atmospheres. **450**, 522–524. arXiv: [0711.3227](#) (Nov. 2007).
224. Williams, K. A. *et al.* Photometric Variability in a Warm, Strongly Magnetic DQ White Dwarf, SDSS J103655.39+652252.2. **769**, 123. arXiv: [1304.3165 \[astro-ph.SR\]](#) (June 2013).
225. Dufour, P. *et al.* *HST Observations of Hot DQ White Dwarfs in American Institute of Physics Conference Series* (eds Werner, K. & Rauch, T.) **1273** (Nov. 2010), 64–69.
226. Dufour, P., Vornanen, T., Bergeron, P. & Fontaine A., B. *White Dwarfs with Carbon Dominated Atmosphere: New Observations and Analysis in 18th European White Dwarf Workshop*. **469** (Jan. 2013), 167.
227. Dufour, P. *et al.* Follow-Up Observations of the Second and Third Known Pulsating Hot DQ White Dwarfs. **703**, 240–251. arXiv: [0907.4293 \[astro-ph.SR\]](#) (Sept. 2009).
228. Dunlap, B. H., Barlow, B. N. & Clemens, J. C. A New Small-amplitude Variable Hot DQ White Dwarf. **720**, L159–L163. arXiv: [1007.5293 \[astro-ph.SR\]](#) (Sept. 2010).
229. Dufour, P., Béland, S., Fontaine, G., Chayer, P. & Bergeron, P. Taking Advantage of the COS Time-tag Capability: Observations of Pulsating Hot DQ White Dwarfs and Discovery of a New One. **733**, L19. arXiv: [1104.2543 \[astro-ph.SR\]](#) (June 2011).
230. Williams, K. A., Montgomery, M. H., Winget, D. E., Falcon, R. E. & Bierwagen, M. Variability in Hot Carbon-dominated Atmosphere (Hot DQ) White Dwarfs: Rapid Rotation? **817**, 27. arXiv: [1511.08834 \[astro-ph.SR\]](#) (Jan. 2016).
231. Dufour, P., Fontaine, G., Liebert, J., Schmidt, G. D. & Behara, N. Hot DQ White Dwarfs: Something Different. **683**, 978–989. arXiv: [0805.0331](#) (Aug. 2008).
232. Kepler, S. O. *et al.* New white dwarf stars in the Sloan Digital Sky Survey Data Release 10. **446**, 4078–4087. arXiv: [1411.4149 \[astro-ph.SR\]](#) (Feb. 2015).
233. Kepler, S. O. *et al.* New white dwarf and subdwarf stars in the Sloan Digital Sky Survey Data Release 12. **455**, 3413–3423. arXiv: [1510.08409 \[astro-ph.SR\]](#) (Feb. 2016).
234. Coutu, S. *et al.* Analysis of Helium-Rich White Dwarfs Polluted by Heavy Elements in the Gaia Era. *arXiv e-prints*, arXiv:1907.05932. arXiv: [1907.05932 \[astro-ph.SR\]](#) (July 2019).
235. El-Badry, K. & Rix, H.-W. The wide binary fraction of solar-type stars: emergence of metallicity dependence at a < 200 au. **482**, L139–L144. arXiv: [1809.06860 \[astro-ph.SR\]](#) (Jan. 2019).

3400 North Charles Street
 Baltimore, Maryland 21218, USA
 (+1) 443.207.1532

<https://sihaocheng.github.io>
s.cheng@jhu.edu

EDUCATION AND DEGREES

2017–2021 Graduate student, Department of Physics & Astronomy, Johns Hopkins University

2012–2016 Undergraduate student, Department of Astronomy, Peking University (China)

RESEARCH INTEREST

I am interested in using statistical analysis of large data set to obtain more understanding of our Universe, including topics in observational cosmology, Milky Way, stellar physics, and extrasolar planets. I am mainly working on two topics now: 1) a new statistic borrowing ideas from convolutional neural nets and its cosmological applications; 2) a cooling anomaly (which I discovered) of white dwarfs and the physics behind it.

AWARDS

2020 Outstanding Publication in Astrostatistics Award	Jan 2021
IAU travel grant for Symposium No.357	Oct 2019
Lin-Qiao Prize for Undergraduate Research at Peking University	Sept 2015
Wu-Si Scholarship	May 2015, May 2014

PUBLICATIONS

[Weak lensing scattering transform: dark energy and neutrino mass sensitivity](#)

Visualised what the scattering transform sees from a lensing mass map, and emphasised its advantage of having Gaussian likelihood

Sihao Cheng and Brice Ménard

2021, arXiv:2103.09247, accepted to *MNRAS*

[Forever young white dwarfs: when stellar ageing stops](#)

Camisassa, M. **et al.**

2021, *A&A Letters*, 649, 7

[A new approach to observational cosmology using the scattering transform](#)

Introduced a new statistic inspired by convolutional neural nets to observational cosmology, and demonstrated that it outperforms classic estimators

Sihao Cheng, Yuan-Sen Ting, Brice Ménard, and Joan Bruna

2020, *MNRAS*, 499, 5902

[An Increase in Small-planet Occurrence with Metallicity for Late-type Dwarf Stars in the Kepler Field and Its Implications for Planet Formation](#)

Cicero X. Lu, Kevin C. Schlaufman, and **Sihao Cheng**

2020, *AJ*, 160, 253

[Multi-Gigayear White Dwarf Cooling Delays from Clustering-Enhanced Gravitational Sedimentation](#)

Evan B. Bauer, Josiah Schwab, Lars Bildsten, and **Sihao Cheng**

2020, *ApJ*, 902, 93

[A Gravitational Redshift Measurement of the White Dwarf Mass–Radius Relation](#)

Used populational gravitational redshift to probe the white dwarf mass–radius relation over a wide mass range

Vedant Chandra, Hsiang-Chih Hwang, Nadia L. Zakamska, and **Sihao Cheng**

2020, *ApJ*, 899, 146

[Carbon star formation as seen through the non-monotonic initial–final mass relation](#)

Marigo, P. **et al.**

2020, *Nature Astronomy*

[Double White Dwarf Merger Products among High-mass White Dwarfs](#)

Measured the white dwarf merger rate with unprecedented high precision using a novel kinematic method

Sihao Cheng, Jeffrey D. Cummings, Brice Ménard, and Silvia Toonen
2020, *ApJ*, 891, 160

[Two delays in white dwarf evolution revealed by *Gaia*](#)

Sihao Cheng

2019, *Proceedings of IAU*, 15 (S357), 175

[A Cooling Anomaly of High-mass White Dwarfs](#)

Discovered an unexpected, extremely long cooling delay in a population of white dwarfs using *Gaia* data

Sihao Cheng, Jeffrey D. Cummings, and Brice Ménard

2019, *ApJ*, 886, 100

[Meteor spectral observation with DSLR, normal lens and prism](#)

Sihao Cheng and Simiao Cheng

2011, *JIMO*, 39, 39