



HAL
open science

Regularity and localized representations of random-phase textures

Samuel Ronsin

► **To cite this version:**

Samuel Ronsin. Regularity and localized representations of random-phase textures. General Mathematics [math.GM]. Université René Descartes - Paris V, 2014. English. NNT : 2014PA05S024 . tel-01160321

HAL Id: tel-01160321

<https://theses.hal.science/tel-01160321>

Submitted on 5 Jun 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Université Paris Descartes
Laboratoire MAP5 UMR CNRS 8145
École doctorale 386 : Sciences Mathématiques de Paris Centre

THÈSE

présentée pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ PARIS DESCARTES

Spécialité : Mathématiques

Présentée par

Samuel Ronsin

**Régularité et Représentations Localisées de Textures
à Phases Aléatoires**

**Regularity and Localized Representations
of Random-Phase Textures**

Soutenue le 15 Décembre 2014 devant le jury composé de :

Directeurs de Thèse :

Hermine Biermé Université de Poitiers
Lionel Moisan Université Paris Descartes

Rapporteurs :

Yann Gousseau Télécom ParisTech
Stéphane Jaffard Université Paris-Est

Examineurs :

Bruno Galerne Université Paris Descartes
Céline Lacaux École des Mines de Nancy
Frédéric Richard Université Aix-Marseille

Remerciements

Merci tout d’abord à mes deux directeurs de thèse Hermine Biermé et Lionel Moisan, pour m’avoir accordé leur confiance et offert l’opportunité de passer quelques années d’une grande richesse intellectuelle. Toute trace d’intelligibilité dans mes démonstrations de probabilité est largement imputable à Hermine. Les heures que j’y ai consacrées doivent à sa rigueur leur sauvetage des méandres du brouillon. Si quelque intuition a jamais traversé mon travail en traitement d’images, il faut la rendre à Lionel. Sans elle je serais, pixelliquement, aveugle. Leur supervision et l’attention constante qu’ils se sont épuisés à porter à mon travail a fait échouer tous mes plans de ne jamais terminer cette thèse. Directement et indirectement, techniquement et humainement, j’ai beaucoup appris grâce à eux – je leur en suis profondément reconnaissant.

Merci à Yann Gousseau – dont les recherches sur le modèle de textures à phases aléatoires ont largement inspiré mon travail – et à Stéphane Jaffard pour m’avoir fait l’honneur d’accepter la mission de rapporter cette thèse. La lecture précise et éclairée qu’ils en ont faite a grandement bénéficié à la qualité de ce manuscrit.

Merci à Bruno Galerne, Céline Lacaux, et Frédérique Richard pour avoir accepté de faire partie de mon jury. Leurs travaux, leurs questions, leurs commentaires et les discussions qui s’en sont suivies – notamment à l’occasion des réunions du groupe de travail MATAIM – ont été à l’origine de nombreux développements présentés ici.

Merci aussi à Agnès Desolneux auprès de qui j’ai eu la chance d’apprendre beaucoup, et dont l’apport pour les principaux résultats de notre *texton* est inestimable. Merci à tous mes admirables professeurs, au premier rang desquels je salue Yves Meyer dont les cours à l’ENS de Cachan m’ont passionné – sans ses encouragements, cette thèse n’aurait simplement pas eu lieu – ainsi qu’Emmanuel Vesin et Pierre Sautier – qui ont tant contribué à éveiller mon intérêt pour les sciences et les mathématiques.

Merci à tous les membres du laboratoire MAP5, tout particulièrement Annie Raoult pour en avoir fait un lieu de travail aussi accueillant et stimulant, et Mireille Chaleyat-Maurel pour l’organisation du Colloquium qui fut pour moi un rendez-vous incontournable. Merci à l’équipe administrative, notamment Marie-Hélène Gbaguidi, pour son efficacité et sa disponibilité.

Merci aux membres de l’UFR mathématiques et informatique: Christine Graffigne pour m’avoir offert de participer à l’équipe d’enseignement, François Patte pour avoir supervisé le cours de Mathématiques et Calcul, Thierry Raedersdorff pour sa grande expertise informatique qui m’a épargné de nombreuses heures de privation de connexion au réseau, Michel Guillemot, Isabelle Valero et Christophe Castellani pour leur extrême diligence.

Merci à tous les “permanents” de l’équipe de recherche du MAP5, particulièrement Avner Bar-Hen, Florent Benaych, Anne Estrade et Joan Glaunès pour leurs encouragements bienveillants et leur gentillesse.

Merci ~~à mes compagnons de galère~~ aux doctorants du MAP5 (sans oublier Charlotte Laclau du LIPADE voisin) pour m’avoir épaulé au cours des oscillations cérébrales qui ont ponctué ces dernières années. Merci à Kevin Kuoch pour m’avoir doublé dans l’agenda des soutenances. Merci à Raphaël Lachièze-Rey et Mario Micheli pour avoir pris les devants

dans la poursuite du rêve américain. Merci à tous les autres “non-permanents” pour avoir tant contribué par leur expertise et leur sympathie à rendre ces journées utiles et agréables: Rémy Abergel, Rebecca Bauer, Mélina Bec, Gaëlle Chagny, Baptiste Coulange, Christophe Denis, Mariela Dimiccoli, Fanny Doré, Charlotte Dion, Diarra Fall, Arthur Leclair, The-Mihn Luong, Anne-Sophie Macé, Alkeos Michail, Nicolas Muller, Maël Primet, Jean Rochet.

Merci à tous mes amis pour m’ avoir empêché de passer l’ essentiel de mes soirées à travailler – votre amitié est un privilège. Merci en particulier à Adrien et François pour m’ avoir guidé sur le chemin de la procrastination. Merci à Ferdinand et Edward pour avoir partagé avec moi leur connaissance encyclopédique des lieux de la convivialité parisienne. Merci aux parisiennes et aux parisiens, installés ou de passage: Akpéli, Anna, Arnaud, Aujain, Camila, Chloé, Cosimo, Erwan, Florence, Juliana, Lætitia, Laura, Lilianne, Luc, Marie, Mélanie, Paul, Sabine...

Merci à ceux qui m’ ont donné le courage de reprendre mes études de mathématiques, en particulier Alexis, Armand, Augustin et Jérôme. Merci aussi à mon regretté ami Bertrand.

Merci à mes amis de Bruxelles et de partout en Europe – en particulier Rinchen – qui m’ ont si généreusement offert les bols d’ air que Paris rend parfois bénéfiques. Merci à mes amis Américains (au sens large), particulièrement Adam, Ben, Dirk, Guillaume, Jérémie, Paul et Will pour m’ avoir tant appris de New York, San Francisco et la Nouvelle Orléans.

Merci à mes amis entrepreneurs, en particulier Oussama pour avoir accueilli mon projet Cyclop à TheFamily, Paul pour son expertise du financement public, et l’ équipe du Thousand Network pour leur dynamisme et leur enthousiasme contagieux.

Merci à mes amis de Saint-Julien en Genevois – Bruno, Damien, Julien, Mark, Vincent et *tutti quanti* – dont non seulement la majorité a eu le culot d’ obtenir avant moi le titre de docteur, mais qui ne cessent de moquer, du haut des parcours de golf que leur rapport privilégié au marché du travail suisse leur permet de fréquenter assidûment, le dépouillement de la vie loin de nos montagnes.

Merci à toi Marly pour ton soutien sans faille et le bonheur que toi seule sais insuffler aux jours, aux semaines, aux mois et aux années que nous passons ensemble.

Merci enfin à mes parents et à ma sœur. Votre amour est un pilier dans tout ce que j’ entreprends.

Contents

Remerciements	6
1 Introduction	9
1.1 Textures, algorithms and models	9
1.2 Micro-textures	15
1.3 Sparse representation of micro-textures	18
1.4 Phase and projection of signals	24
1.5 Asymptotics of the random phase algorithm	26
1.6 Outline	32
2 A Texton for Random Phase and Gaussian Textures	35
2.1 Introduction	35
2.2 Definition and properties of the texton	40
2.3 Sparse representations	54
2.4 Textons for color images	61
2.5 Appendix	70
3 Large Error Approximations of Images	79
3.1 Introduction	79
3.2 The importance of phase in signals as a LEA problem	81
3.3 Four large error projections with constraints on the phase	85
3.4 The importance of the Fourier transform	92
3.5 Concluding remarks	95
4 The Billard Theorem for Multiple Random Fourier Series	97
4.1 Introduction	97
4.2 Notations and Hypotheses	98
4.3 Methods of summation in \mathbb{Z}^d	100
4.4 Billard's theorem in arbitrary finite dimension	103
4.5 Properties of Random Phase Noise Processes	109
4.6 Appendix: Proofs	112
5 On the Regularity of some Multiple Random Fourier Series	117
5.1 Introduction	117
5.2 Notations and Hypotheses	118

5.3	Conditions for continuity	119
5.4	Almost sure unboundedness	126
5.5	Hölder regularity	130
5.6	Extensions	135
6	Conclusion	139
	Bibliography	141

Chapter 1

Introduction

1.1 Textures, algorithms and models

Textures are one of the few concepts to apply to all the human senses. This synesthetic essence makes the texture category easy to grasp intuitively, albeit quite challenging to define properly. As opposed to *intensity* or *persistence*, generic categories which address quantitative aspects of perception, *textures* seem to be intrinsically qualitative. In other words qualities such as *intensity* and *persistence* are straightforwardly measurable – respectively by *energy* and *scale* – but the natural scientific framework for *textures* is much more complex and certainly not to be characterized by a single quantity.

A commonly mentioned characteristic of textures is overall “homogeneity” or “uniformity”, as opposed to shapes whose purpose is to encapsulate and distinguish separate entities. Indeed, textures and shapes are very complementary concepts, and one could even argue that textures describe precisely everything that shapes don’t. One way to apprehend this complementarity comes from the modeling perspective. Mathematically, a shape might be best described as a “deterministic” object (a curve, a surface etc.), whereas a probabilistic model (random functions or fields) is often best suited to describe a texture.

The modern study of the visual perception of textures goes back to Julesz, who made the first hypotheses on their fundamental probabilistic properties [63]. He famously claimed that textures that shared first and second order statistics were somehow indistinguishable, a hypothesis that was later proven incorrect [68]. Julesz also defended the decomposition of textures into atomic elements – *textons* – and spatially homogeneous probabilistic distributions of these elements [65]. Although mostly driven by the perception of textures, Julesz built the foundations of texture analysis which has been drawing a lot of interest, with many applications *e.g.* in medical, industrial, satellite or astronomical imagery. For instance, mammographic density – a quantity that can be approached through texture analysis of mammographs – happens to be a very relevant parameter with respect to the risk of breast cancer, as discussed in [19]. The measure of anisotropy is also relevant in the evaluation of the risk of bone fracture induced by osteoporosis, as argued in [18]. Generally speaking, modeling textures involves two parts:

- A set of parameters (*e.g.* a mean and an autocovariance function).

- A stochastic field on a domain D , *e.g.* \mathbb{R}^2 , that is specified by the aforementioned set of parameters (*e.g.* a stationary Gaussian field).

A model often conveys – at least indirectly – statistical estimators of the parameters that specify the model (*e.g.* the average and the auto-correlation function of input images).

Texture analysis can also be used to feed processed data to a texture synthesis algorithm, a very active area in computer graphics. Texture synthesis refers to a clearly defined task: outputting images that conform to the chosen description of a texture, and a set of desired characteristics, typically resemblance to an input image. Hence, texture synthesis is performed by an algorithm, fed by data. More precisely:

- A dataset, *e.g.* an image (or a collection of images), or a set of numerical parameters captures the “local” information of the texture at different relevant scales.
- The algorithm may or may not be stochastic. It should take the dataset as a first argument, and arbitrary dimensions as a second argument. It should output an image with the designed dimensions.

Two requirements on the output of the texture algorithm should impose two fundamental invariance principles that textures require, namely invariance with respect to dimensions (a texture should not be limited in space) and translations (a texture should be the same everywhere). To guarantee this spatial homogeneity, the law of any local subpart of the output image should only vary with the first argument (the data input) and never with the second argument (the dimensions input). Moreover, such a law should not depend on the location of the local section.

Our work is mainly concerned with texture analysis in a broad sense, and to a lesser extent synthesis. Most of the work presented here can be traced back to a very specific type of textures, namely Gaussian and Random Phase Noise textures as introduced in [83], [127] and more recently discussed in [51] where they are referred to as “micro-textures”. As we shall discuss, these texture models are only based on first and second order statistics, such as the mean and the covariance function. They intensively involve the Fourier transform and the fast Fourier transform algorithm in their implementation. As a consequence textures are implicitly assumed to be periodical, an assumption that we shall discuss below. A strength of these models is the convenient framework for many aspects of texture analysis they offer. In particular, as we shall discuss in Chapter 2, this allows for a strongly localized representation which paves the way for sparse and localized representations of Gaussian and Random Phase Noise textures. Sparse representations of signals have recently been a very active topic in applied mathematics, mostly thanks to the results of compressed sensing (see *e.g.* [21], [22] and [38]). However, to our knowledge, sparsity within representations of textures has been mainly left out of focus, with the exception of [79] and [114] which follow approaches very different from ours. The importance of the Fourier phase in our work (especially in Chapter 2) relates to a widely documented phenomenon, namely “the importance of phase in signals” [106], that we further discuss and somewhat challenge in Chapter 3. The “micro-textures” framework also allows for asymptotical analysis of texture models, both while scaling up the

dimensions of the texture (Chapter 2) and zooming in within a continuous model (Chapters 4 and 5). The continuous model is connected to the study of random Fourier series, which has been first developed by Zygmund along with Paley [108], [109], [110] and Salem [119], and widely discussed still in the one-dimensional case by Kahane [69] and in more generality by Marcus and Pisier [96]. In Chapter 4 and 5, we focus on extending to the general finite-dimensional case some results which are well known in the one-dimensional case with respect to convergence ([15] and [69]), continuity ([15], [69], [45] and [32]) and regularity ([69] and [32]). Recent papers have been working on a similar path of generalizations from the one-dimensional to the finite-dimensional case with a focus on continuity, let us cite *e.g.* [27].

In the remaining of this section, we propose a non-exhaustive tour of texture algorithms and models. As discussed *supra*, randomness plays a very important role within texture algorithms/models. As we shall see, they can indeed be characterized by a varying degree of randomness. Indeed, a continuum can be drawn from deterministic tilings to filtered white noise algorithms, which can be translated in terms of information theory by an increasing entropy from deterministic textures (strongly structured) to noises (strongly unstructured). In between these two extremes lie two important families of texture algorithms: exemplar-based synthesis, which can be characterized as “weakly structured”, and noise synthesis based on statistical analysis qualified as “weakly unstructured”. The former roughly consists in adding randomness to deterministic algorithms supported by rigid structures, and the latter in adding control to noisy images through constraints based on statistical analysis of some inputs.

1.1.1 Structured textures

Tiling. Tiling is one of the simplest and most structured texture synthesis algorithm. It consists in repeating tiles from an image dataset over the plane according some given tiling rule, or algorithm. There are many different ways of tiling the plane, periodical or even non-periodical – see *e.g.* Penrose [111] – which may be applied to render strongly structured textures such as brick walls, floor tilings etc. However, one requirement for visual relevance is that the tiles need to match seamlessly, which is of both great interest and difficulty for photo realistic input images. Moisan proposed an original solution to the seamless single tile problem [99]. However, a major drawback of tiling algorithms is that, except for small dimensions (with respect to the input image), the display of the super-structure (the tiling rule) tends to be visually overwhelming for the observer – see *e.g.* Rao [116] for empirical evidence of the importance of periodicity in texture perception. Hence, the repetition due to the underlying structure may discard outputs of tiling algorithms as realistic textures for many applications.

Weakly structured textures. Exemplar-based texture synthesis algorithms can be considered as stochastic adaptations of tiling algorithms. Indeed, exemplar-based texture algorithms aim at relaxing the tiling rule by injecting a dose of randomness into it, hence the qualification “weakly structured”. The purpose is to avoid the strict repetition and the trivial

predictability of patterns and the display of super-structure. Another shortcoming of tilings is overcome as weakly structured models avoid exact matches between independent outputs.

These algorithms perform a stochastic mapping from the output domain D_{out} to the initial domain D_{in} over which the input image u_{in} is defined. More precisely, each coordinate (x_{out}, y_{out}) from D_{out} is stochastically associated with a coordinate $P_{(x_{out}, y_{out})}^{in}$ in D_{in} , and the final texture u_{out} is defined by applying u_{in} to the mapped coordinate: $u_{out}(x_{out}, y_{out}) = u_{in}(P_{(x_{out}, y_{out})}^{in})$ (see Figure Figure 1.1). Thus, the main problem becomes to build algorithms that provide relevant stochastic mapping functions. Efros and Leung breakthrough in their seminal paper [42], started a rich literature of computer vision, which was thoroughly surveyed in [129]. They proposed a construction of P^{in} based on an approximate Markov random field model. The texture is grown sequentially, pixel by pixel (typically in raster or print order), by matching the output neighborhood $\mathcal{N}_{(x_{out}, y_{out})}$ of the last synthesized pixel to its closest equivalents from the input sample u_{in} and sampling randomly among these neighborhoods. Many improvement were invented: Wei and Levoy [130] introduced a data structure to accelerate the repeated exhaustive nearest-neighbor searches ; Efros and Freeman in [41] replaced the “one pixel at a time” synthesis by a “one patch at a time” approach, improving the overall quality and speed of the synthesis ; Kwatra *et al.* [75] introduced graphcut techniques to find the optimal cut path over the overlapping patches ; parallel implementations of the patch-based synthesis, to take advantage of multi-core processing units that have been recently flourishing, see [81] for a pyramidal synthesis where all parts of the output is simultaneously synthesized, the synthesis being refined along a decreasing sequence of scales.

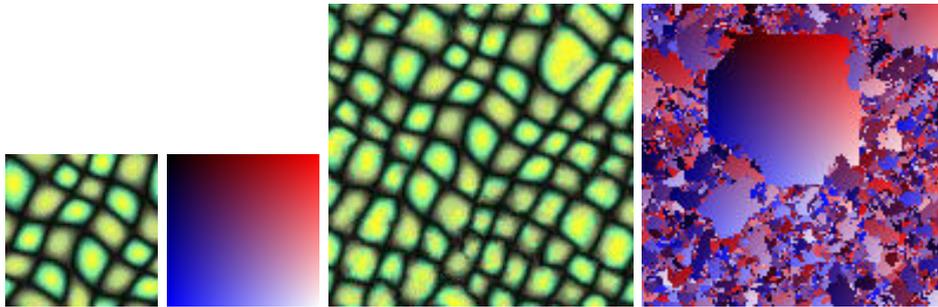


Figure 1.1: Left: input image “cell” and its canonical mapping. Right: output image obtained with the Efros-Leung algorithm [42] (implementation [2]) and its stochastic mapping.

From many perspectives, these algorithms are very well suited for texture synthesis, and their output are often of spectacular quality. However, exemplar-based synthesis does not require any statistical processing of the input, and has thus been of very limited interest from a texture analysis perspective so far.

1.1.2 Noisy textures

Procedural noise. At the upper extreme on the entropy spectrum lie the strongly unstructured “procedural noise textures”. Perlin invented what is now known as Perlin noise –

see his seminal paper “An Image Synthesizer” [112] – after working on the graphics of the Disney movie “Tron” [87].

Since then the topic has attracted considerable attention, and is still – as of today – an indispensable source of texture synthesis in computer graphics. The literature on the subject encompasses hundreds of articles for which we refer to the recent survey in [76], Section 3. Basically, the setup of these algorithms is as follow:

- The dataset consists in a few explicit parameters (*e.g.* an exponent that characterizes a frequency spectrum).
- The algorithm synthesizes noise images of any size, parametered by the input parameters.

Most of the time, these algorithms output approximations of a filtered Gaussian white noise, which are in turn assembled to render visually compelling textures. Indeed, many functions can be applied (such as scalar products, norms, functions defined over the domain etc.) to any raw output field, which can in turn be superposed in order to emulate different classes of textures such as water, fire, clouds, marble etc. The main advantage of such algorithms is that they are local: the value at one point only depend on the value of a few neighbors, which can be very useful for computer graphics performance such as linear computational cost and high memory efficiency. Several improvements and variants have been proposed. For instance, Cook and DeRose [29] proposed the “wavelet noise” algorithm in order to gain precision over the band-limits of the spectrum. Indeed, lack of precision imposes a tradeoff between aliasing and blur at high frequencies for Perlin noise.

A slightly less computationally efficient noise synthesis consists in Gaussian white noise filtering in the Fourier domain as developped in [3] (see also [120] and [128] for fractal models). Direct filtering allows perfect control over the frequency spectrum and thus directional partitions of the spectrum which is of particular interest for anisotropy [55]. These techniques were later used by Van Wijk for spot noise synthesis [127], where the frequency spectrum is specified by the Fourier transform of an input “spot”. Although less computationally efficient, and less memory efficient (non-local), Fourier synthesis performance became competitive thanks to the fast Fourier transform (FFT) implementations [31] [30]. For 2-dimensional signals, the FFT runs with $\mathcal{O}(MN \log(MN))$ operations instead of $\mathcal{O}((MN)(M + N))$ with more naive approaches.

These techniques provide an efficient way to create building blocks that can in turn be assembled in order to render more elaborate textures. However, a major drawback is that this part of the process can hardly be automated. It thus still requires a lot of creative work, and some artistic talent.

Noise and statistical analysis. The purpose of statistical-based noise textures is to automate the “artisanal” process of turning the building blocks provided by procedural synthesis into a realistic texture. The idea is to estimate parameters (that may or may not be selected through a learning strategy) from an input texture and use an appropriate synthesis algorithm with these estimated parameters. Hence, two main tasks arise with such algorithms: first,

find the relevant statistics (analysis), and second, implement efficient synthesis algorithms based on these statistics. Relevance of the statistics is hard to define objectively. This task is often rooted in the study of human perception of textures – *e.g.* Julesz [63], [68], [65], [67], Bergen [9], Malik and Perona [90], see Landy and Graham [78] for a quite recent survey – and neuroscientific study of the visual cortex, especially V1 – *e.g.* Knierim and Van Essen [73], Olshausen [103], along with Field [104], [105], and with Simoncelli [124]. The purpose is to provide relevant mathematical tools to describe the vision of textures. An important feature of these models is the sparsity of the set of relevant parameters. Hence, these algorithms successively perform texture analysis and synthesis as follows.

- Analysis: estimate the relevant parameters of an input image. The relevance of parameters can be either predetermined or learned.
- Synthesis: according to a given model that can be specified via the estimated parameters, synthesize random samples of a field compliant with the stochastic law of the model.

Let us now run through a few fundamental examples. Among these, we can extract two main families, according to the processing methods used for texture analysis. The wavelet-based and the Fourier-based algorithms, following two prominent representations of signals. We refer to Section 4 in [76] for a more complete overview.

Wavelet analysis The wavelet-based methods can be connected to studies suggesting that a process similar to wavelet pyramidal decomposition is likely performed by the human visual cortex for texture discrimination [9], [90], [103]. Let us mention a few examples here.

In their pioneering work, Heeger and Bergen [59] draw an analogy with the use of trichromatic color synthesis – based in the knowledge of color perception in the eye as shown in [125] – and suggest that a similar strategy should be adopted in order to provide a robust illusion of texture fidelity. They proposed to match histograms profiles of both plain RGB values and learned coefficients in a multiscale (pyramidal) wavelet decomposition of images, starting from Gaussian white noise. The convergence is then “visually” validated.

Zhu, Wu and Mumford developed a very general framework in [134] and [135], and proposed to derive the maximum entropy law, given learned constraints based on wavelet coefficients. These are incrementally chosen to maximize a distance between the observed distribution and the theoretical distribution derived by entropy maximization. Major advantages of this method are that no prior assumption over the distributions is needed, and that the marginal laws for each selected statistic are independent.

Portilla and Simoncelli [115] developed a framework combining a Markov chains Monte Carlo approach (as initiated by [133]) and an analysis performed over the wavelet coefficients of a collection of input image. As in Heeger and Berger’s initial method, the variables of interest are learned and picked semi-automatically. As the number of such relevant statistics increases, more and more textures from the database end up being adequately synthesized.

A major drawback of these algorithms is that the analysis is often very computationally

costly or needs to be performed semi-automatically, due to the large set of coefficients from which the selection is drawn. A design to select efficiently a critical set of relevant statistics in a multi-scale wavelet decomposition is still an active area of research. Moreover, the synthesis is also very costly as multiple steps are often needed in order to provide a compelling result (multiple white noise projections in [59] and [115], multiple entropy maximisations in [134]).

Fourier analysis The discrete Fourier transform (DFT) of images is very largely used in signal processing. The Fourier representation of signals is not localized, so each coefficient catches a “global” feature of the input. More precisely, the action of a translation over the Fourier transform of an image is simply a pointwise multiplication: the translations are diagonal operators in the Fourier basis. This is a strong theoretical argument to consider the Fourier representation as a fundamental tool for textures, since translations are of particular importance in the field.

Fourier synthesis has been largely used since the introduction of the fast Fourier transform. One reason is that in the Fourier domain, periodic convolutions are simple pointwise multiplications. Hence, the cost of a periodic convolution to output an image of dimensions $M \times N$ can be lowered down to $\mathcal{O}(MN \log(MN))$. This compares favorably to direct convolution, except maybe for extremely sparse convolution kernels S , namely when $|S| = \mathcal{O}(\log(MN))$ where $|S|$ denotes the cardinality of S .

As discussed above, Fourier synthesis is deeply connected with Gaussian textures.

Lewis [83] was one of the first to suggest a synthesis algorithm based on the spectrum of a sub-sample taken from an image texture, both through FFT and sparse convolution.

Van Wijk [127] who proposed a texture design based on the Fourier transform of spots. Moreover, he developed an idea that Lewis had briefly suggested, which consists in taking the phases at random without modifying about the law of the modulus. However, he surprisingly rejected the idea of taking a texture image as input “spot” because of the “tautological character of this solution”.

The recent paper [51] bridged that gap. In a nutshell, they proposed to take the average and the periodic auto-correlation of an input texture image – obtained thanks to the FFT – as estimators of the mean and the covariance function of an underlying periodic Gaussian field. They clearly established that the “random phase noise” model was not equivalent to the periodic Gaussian model, although they both produce extremely similar outcomes.

1.2 Micro-textures

We now discuss in more detail the class of “micro-textures” as defined in [51], upon which most of the work of this thesis is built. In a nutshell, these textures have the property to display no structure at all, but can still be parametrized by input images and render compelling noisy textures such as clouds, water, grass (from a distance) etc. In [51], Galerne *et al.* introduced “micro-textures” as the class of input images that could be faithfully synthesized through either asymptotic spot noise or Fourier phase randomization, which are detailed *infra*. They further noticed that the two algorithm mostly output visually indistinguishable results. Rectifying an unproven claim by Van Wijk, they noticed that the output of the phase

randomisation synthesis algorithm would be a non-Gaussian field. They renamed this field “Random Phase Noise”, and we shall refer to it as RPN in the following.



Figure 1.2: Left: input image “wood”. Right: output image obtained with the RPN algorithm [51] (implementation [50]).

Gaussian textures. Let us consider a stationary periodic real Gaussian field U_g defined over the discrete rectangular domain D . U_g is entirely characterized by its expectation $\mathbb{E}[U(\mathbf{x})]$ which does not depend on \mathbf{x} thanks to stationarity, and its covariance function γ_{U_g} defined thanks to stationarity by

$$\gamma_{U_g}(x - y) := \mathbb{E}[U_g(x)U_g(y)] \quad (1.1)$$

for x and y in D . U_g is equivalently characterized but its expectation and its power spectrum $(\mathbb{E}[|\widehat{U}_g(\xi)|^2])_{\xi \in D}$. Moreover, the discrete Fourier transform of such a stationary Gaussian field U_g has the following properties:

- for each ξ in $D \setminus \{\mathbf{0}\}$, $\widehat{U}_g(\xi)$ is a complex Gaussian random variable
- for each $D' \subset D \setminus \mathbf{0}$ such that $-D' \cap D' = \emptyset$, $(\widehat{U}_g(\xi))_{\xi \in D'}$ is a complex Gaussian vector with independent entries.

A property of any complex Gaussian variable Z is to write as $Z = Re^{i\Phi}$ with R real non-negative random variable and Φ variable in $\mathbb{R}/2\pi\mathbb{Z}$. Moreover, R and Φ have the following properties: Φ is uniformly distributed over $\mathbb{R}/2\pi\mathbb{Z}$; R follows a Rayleigh law; R and Φ are independent. Thus, the Fourier transform of a Gaussian texture synthesis based on image u can be written as

$$\widehat{U}_g(\xi) = \widehat{u}(\xi)R(\xi)e^{i\Phi(\xi)} \quad (1.2)$$

where $(R(\xi)e^{i\Phi(\xi)})_{\xi \in D}$ is the Fourier transformation of a real Gaussian white noise over D . Two main algorithms: white noise convolved with an input images or a sparse sample of input image [83]; asymptotic discrete spot noise (ADSN) obtained by throwing an image (or a spot) u with uniform Poisson density over a periodic field, summing and renormalizing [127], [51]. The main shortcoming of the ADSN is its low convergence speed, since convergence is typically of order $\mathcal{O}(\frac{1}{\sqrt{k}})$ with k shots.

Random phase textures. In [127] Van Wijk proposed an alternative synthesis algorithm, named the Random Phase Noise (RPN) algorithm in [51]. In a nutshell, this algorithm is directly based on the Fourier transform, and only the phase part of an image u is randomized (the Rayleigh noise R is removed). In this model, a texture based on image u can be described in the Fourier domain as

$$\widehat{U_{RPN}}(\xi) = \widehat{u}(\xi)e^{i\Phi(\xi)} \quad (1.3)$$

where $(\Phi(\xi))_{\xi \in D}$ is a random phase, that is

- $\Phi(-\xi) = -\Phi(\xi)$ (modulo 2π) for each ξ in D
- $\Phi(\xi)$ is uniformly distributed over $\mathbb{R}/2\pi\mathbb{Z}$ for $\xi \neq -\xi$.

As noticed in [51], the RPN algorithm outputs textures that are indistinguishable from their Gaussian counterparts with a very large majority of inputs – actually only synthetic (and somewhat pathological) inputs yield visually distinguishable outputs. As we shall discuss in Chapter 2 (see also [36]), this can be explained by the fact that a Gaussian texture can be deduced from a RPN texture by the convolution of a signal that is very close to a Dirac mass at the spatial origin. To fully understand the connection between RPN and Gaussian textures, let us assume that we observe a unique sample of a Gaussian texture U_0 . We want to simulate a texture U_1

- that is different of the observation $U_0(\omega)$
- that has the same law as U_0 .

Of course, one possibility would be to convolve U_0 with a Gaussian white noise W assumed to be independent from U_0 . However, this solution has many drawbacks: the resulting texture would fail to be Gaussian (it would only be Gaussian conditionally on U_0) ; iterations of this procedure converge almost surely to a constant field. Synthesizing U_1 based on U_0 with the RPN algorithm preserves the original Gaussian law and shows an interesting stability property: successive samples based on the last iteration still follow the same (Gaussian) law. Moreover

$$\mathbb{E}[\|U_0 - U_1\|_2^2 | U_0] = \sum_{\xi \in D} \mathbb{E}[|\widehat{U}_0(\xi)|^2 |e^{i\Phi_{U_0}(\xi)} - e^{i\Phi(\xi)}|^2 | U_0] = 2 \|U_0\|_2^2 \quad (1.4)$$

so U_0 and U_1 differ significantly. Hence, the RPN algorithm is well suited to resynthesize a Gaussian random field from a single Gaussian sample.

Going one step further in the discussion, we prove in Chapter 2 (see also [36]) that a suitably normalized RPN synthesis based on a deterministic spectrum converges in finite dimensional law, as its dimensions tend to infinity, to a discrete Gaussian field over \mathbb{Z}^2 . Another interpretation of this similarity between RPN and Gaussian textures consists in noticing that the two algorithms essentially differ from the convolution with a random signal that is very close to a weighted Dirac mass located at the spatial origin.

Fourier transform texture analysis. Both micro-texture synthesis algorithms can be decomposed in an analysis and a synthesis part. Importantly, there are only two parameters that matter for both of them: the mean and the covariance function.

- Analysis: estimate the mean and the covariance function, typically by computing the auto-correlation of a given input image. This step can be performed in the frequency domain for optimized computational cost through FFT.
- Synthesis: a field corresponding to the estimated covariance function is synthesized. This step can be performed by computing the inverse Fourier transform of the Hermitian-symmetric complex Gaussian field associated with the frequency-spectrum estimated thanks to texture analysis.

Both RPN and ADSN algorithms estimate the mean by the average and the covariance by the periodic auto-correlation. However, other estimation strategies of the power-spectrum/covariance have been proposed. Indeed, since the auto-correlation function and the power spectrum of a texture image are often quite noisy, a solution is to look for smooth (in the Fourier domain) estimations of the power spectrum. Gabor kernels have been proposed as a projection basis by Gilet in [54], and further developed in [52], in order to bring an interactive texture design tool to the computer graphics community. Sparse approximations of the *texton* of an image, as introduced in [117], are also smooth spectrum estimations as discussed in Chapter 2.

1.3 Sparse representation of micro-textures

Noise samples are somewhat incompressible from the information theory point of view [28] since they inherently have a high Shannon entropy [122]. However, their probabilistic law can often be accurately described with simple stochastic rules and a small set of parameters. Their pseudo-random implementations might thus be very concise – in other words, they might have a low Kolmogorov “algorithmic” complexity [84].

As discussed *supra*, textures can be characterized by a high level of redundancy of patterns, and the overall homogeneity of their aspect. In other words, under the homogeneity hypothesis, zooming in at two distant neighborhoods of a given domain shall yield statistically indistinguishable results (at least over several independent samples). Thus, at a given level of details, the dimensions or size of a texture image appears to be an information that is irrelevant to its probabilistic nature. Hence, for each texture, both a critical (minimal) size and an amount of information bounded by this size should intuitively exist and characterize the texture.

An important motivation to focus on sparse texture representations is their direct contribution to the dimensionality reduction literature, a very dynamic research area involving both the statistics, machine learning and signal processing communities [48]. Low dimensionality is of particular interest for predictive models, since the training time of learning algorithms is often a non-decreasing function of the number of parameters. Moreover, lowering the dimension of the input might help avoid overfitting, a well known problem in machine learning and statistical estimation.

Another argument in favor of the study of sparsity in textures is that it has long been hypothesized that human vision itself performs high level sparse representation of scenes (see *e.g.* [5] and [6]), and in particular textures. In celebrated articles (see *e.g.* [67]), Julesz developed a theory that preattentive texture discrimination involved two main characteristics: the density (particularly its first and second order statistics [63]) and the shape of elementary structures that he named *textons*.

Let us state the general “texture sparsity problem” as follows: “given a texture model, determine a minimal set of information that characterizes, exactly or approximatively, this texture model”. Of course, this so-called “minimal set” may be specific to the underlying model.

Research on sparse texture representations has been surprisingly scarce to this day. Let us mention however the works of Lazebnik *et al.* [79] and Peyré [114]. In [79] the authors propose the extraction of a sparse set of local descriptors that are intended for texture classification. They depart from the homogeneity paradigm and impose for these descriptors to be robust under the transformations of a 3-dimensional scene such as rotations. In [114], for each texture, a redundant dictionary is learned, which provides a basis for both sparse synthesis and analysis.

Our approach to sparsity can be traced back to the work of Lewis [83] and van Wijk [127]. In [83] the author proposes to draw directly a sparse sample of an input texture to perform a synthesis through convolution with a white noise. In [127] a similar design based on a synthetic spot is introduced, along with the explicit random phase algorithm, further studied in [51]. More precisely, in the asymptotic shot noise and random phase models à la van Wijk, a stronger version of the “texture sparsity problem” can be stated as follows: based on one single sample, is it possible to find, exactly or approximately, the sparse spot that yielded a given texture? The answer of this question is still quite open, but the deep relationship between synthesis and analysis of micro-textures makes them particularly promising subjects for the “texture sparsity problem”.

The research presented in Chapter 2 (see also [117], [35] and [36]) stands for a first step towards a solution to these problems. Let us now introduce our approach and present some of our contributions here. It has been known for decades that the Fourier phase is an important part of the signal, as stated in detail in [106]. Precisely, the phase has been accounted for encoding the shapes of images, which we shall discuss in more detail below. On the other hand, discrete periodic real Gaussian fields have, modulo Hermitian symmetry, independent Fourier coefficients, and in particular independent phase [83] [127] [51]. As a consequence, these texture synthesis models are invariant by the multiplication of any Hermitian symmetric phase field. Let us moreover mention that invariance of the distribution of probabilistic images with respect to periodic translations also implies a weak form of phase independence, namely pairwise independence [117]. Thus, given an input image u defined over a domain D , we introduced the set of images $\mathcal{M}_u = \{s \in \mathbb{R}^D; \widehat{s}(\mathbf{0}) = \widehat{u}(\mathbf{0}) \text{ and } |\widehat{s}(\xi)| = |\widehat{u}(\xi)| \forall \xi \neq \mathbf{0}\}$ that is the class of images that define the same ADSN and RPN models as u and experimented deformations of u obtained by manipulating its phase.

Interestingly, the null-phase representative element of the class \mathcal{M}_u displays particular concentration around the spatial origin, and we chose to name this special representant the *texton* of its micro-texture class, after the naming by Julesz in [65].

Definition (Texton of an image). *The texton of an image $u : D \rightarrow \mathbb{R}$ is the image $T(u) : D \rightarrow \mathbb{R}$ that has the same mean value as u , the same Fourier amplitude as u , and identically null phases except maybe at the spatial origin $\mathbf{0}$. In others words, $T(u)$ is characterized in Fourier domain by*

$$\widehat{T(u)}(\mathbf{0}) = \widehat{u}(\mathbf{0}) \quad \text{and} \quad \forall \xi \in \Omega \setminus \{\mathbf{0}\}, \quad \widehat{T(u)}(\xi) = |\widehat{u}(\xi)| \quad (1.5)$$

or, equivalently, in the spatial domain by

$$\forall \mathbf{x} \in D, \quad T(u)(\mathbf{x}) = \frac{1}{|D|} \widehat{u}(\mathbf{0}) + \frac{1}{|D|} \sum_{\xi \in D, \xi \neq \mathbf{0}} |\widehat{u}(\mathbf{x})| e^{2i\pi \langle x, \xi \rangle}. \quad (1.6)$$

Interestingly, the *texton* can be characterized as a solution of some variational problems. In [117] we introduced

$$(\mathbf{P1}) : \max_{v \in \mathcal{M}_u} v(\mathbf{0})$$

and proved the following result.

Proposition (Property of spatial concentration). *For any image $u : \Omega \rightarrow \mathbb{R}$, $T(u)$ is the unique solution of **(P1)**.*

Another variational formulation of the concentration problem, first introduced in [34], can be stated as follows.

$$(\mathbf{P2}) : \min_{v \in \mathcal{M}_u} \sum_{\mathbf{x} \in \Omega} A(\mathbf{x}) v(\mathbf{x})^2, \quad (1.7)$$

where $A(x)$ is a penalty function that should increase as $|x|$ increases. We proved a range of results about optimal concentration problems at the spatial origin **(P2)** satisfied by the *texton*.

Proposition (Property of spatial concentration). *Let $u : D \rightarrow \mathbb{R}$ be an image with a non-negative mean value. Let $A : D \rightarrow \mathbb{R}$ be a symmetric weight image such that $A(\mathbf{0}) = 0$ and*

$$\forall \xi \in D \setminus \{\mathbf{0}\}, \quad \hat{A}(\xi) \leq 0. \quad (1.8)$$

*Then, A is non-negative and $T(u)$ is a solution of the optimization problem **(P2)** associated to A (Equation (2.15)).*

Starting with this concentrated representation of micro-textures, we studied several strategies to design sparse and faithful approximations. Such approximations are best judged depending on the task that they are used to performed. For texture synthesis, the criterion is visual and based on the distinguishability of textures generated by the *texton* and its sparse approximation. In this case, hard thresholding appears to be the best suited procedure to

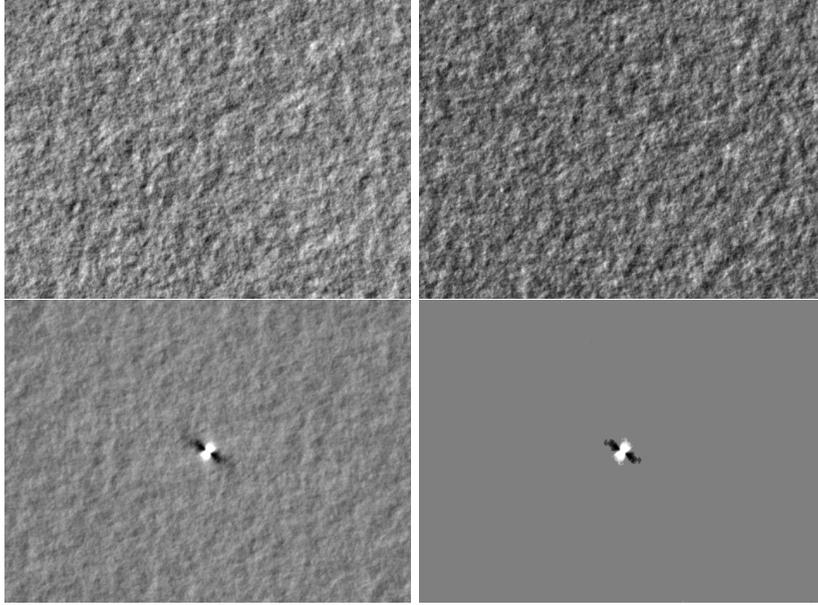


Figure 1.3: Top: the micro-texture samples based on the texton (left) and its sparse approximation (right) are indistinguishable up to an affine contrast change. Bottom: the texton (left) is very concentrated around the spatial origin and a sparse approximation (right) can be obtained through hard thresholding.

sparsely approximate textons. In this procedure, $T(u)$ is approximated by $v = f_{\alpha,\beta}^{\text{hard}}(T(u))$, where $f_{\alpha,\beta}^{\text{hard}}$ is a hard-thresholding function

$$f_{\alpha,\beta}^{\text{hard}}(t) = \begin{cases} t & \text{if } t \leq -\alpha \\ 0 & \text{if } -\alpha < t < \beta \\ t & \text{if } \beta \leq t \end{cases}$$

and the choice of α and β is to be optimized. Our approach thus provides a new tool, which could be of useful for texture analysis tasks *e.g.* texture classification and segmentation. To illustrate this point, let us cite a result about the Fréchet distance between textures found in [39] and [132].

Theorem ([39], [132]). *Let u_0 and u_1 be two images with zero mean. Then, the Fréchet distance between the two Gaussian distributions $\text{ADSN}(u_0)$ and $\text{ADSN}(u_1)$ is given by the L^2 distance between their respective textons:*

$$d(\text{ADSN}(u_0), \text{ADSN}(u_1))^2 = \min_{\substack{U_0 \sim \text{ADSN}(u_0) \\ U_1 \sim \text{ADSN}(u_1)}} \mathbb{E} (\|U_0 - U_1\|_2^2) = \|T(u_0) - T(u_1)\|_2^2.$$

The same result holds also for the Fréchet distance between the two RPN distributions $\text{RPN}(u_0)$ and $\text{RPN}(u_1)$:

$$d(\text{RPN}(u_0), \text{RPN}(u_1))^2 = \min_{\substack{U_0 \sim \text{RPN}(u_0) \\ U_1 \sim \text{RPN}(u_1)}} \mathbb{E} (\|U_0 - U_1\|_2^2) = \|T(u_0) - T(u_1)\|_2^2.$$

Sparse color textures. The extension of our approach to color images raised many difficulties and proved to be a quite challenging step. A strategy consisted to keep working directly on the phases of each RGB channel. In the case of color images, the link between the empirical periodic covariance $\Gamma_{\mathbf{u}}$ and the Fourier transform is given by

$$\forall \xi \in D, \widehat{\Gamma}_{\mathbf{u}}(\xi) = \frac{1}{|D|} \widehat{\mathbf{u}}(\xi) \widehat{\mathbf{u}}(\xi)^*,$$

where $\widehat{\mathbf{u}}(\xi)$ is considered as a column matrix in \mathbb{C}^3 and the notation $*$ denotes the conjugate transpose of a complex matrix. A natural extension of the class $\mathcal{M}_{\mathbf{u}}$ to color images requires that both the color mean and auto-correlation are preserved

Proposition. *A necessary and sufficient condition for \mathbf{u} and \mathbf{v} to have the same color mean and auto-correlation is that there exists a phase field $\varphi : D \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ such that $\varphi(-\xi) = -\varphi(\xi)$ for all ξ in D and*

$$\forall \xi \in D, \widehat{T(\mathbf{u})}(\xi) = e^{i\varphi(\xi)} \widehat{\mathbf{u}}(\xi).$$

We thus define, for each color image \mathbf{u} , the class

$$\mathcal{M}_{\mathbf{u}} = \{\mathbf{s} \in \mathbb{R}^{3D}; \exists \varphi : D \rightarrow 2\pi\mathbb{R}/\mathbb{Z}, (\forall \xi \in D), (\varphi(-\xi) = -\varphi(\xi)) \text{ and } (\widehat{\mathbf{s}}(\xi) = e^{i\varphi(\xi)} \widehat{\mathbf{u}}(\xi))\}$$

and a color texton that relies on the phase of a grey-scale projection of the input color image \mathbf{u} as follows.

Definition (α -color texton). *For a phase field $\varphi : D \rightarrow \mathbb{R}$ satisfying the Hermitian-symmetry condition $\varphi(-\xi) = -\varphi(\xi)$ for all $\xi \in D$, let us denote by S_{φ} the operator defined on color images \mathbf{u} by shifting the phases of all the channels with the phase field φ . That is, if \mathbf{u} is a color image, then $S_{\varphi}\mathbf{u}$ is also a color image, given by*

$$\forall \xi \in \Omega, \widehat{S_{\varphi}\mathbf{u}}(\xi) = e^{-i\varphi(\xi)} \widehat{\mathbf{u}}(\xi).$$

For $\alpha \in \mathbb{R}^3$, let us define the α -color texton by the operator

$$\mathbf{u} \mapsto T_{\alpha}(\mathbf{u}) := S_{\varphi_{\alpha \cdot \mathbf{u}}}\mathbf{u}, \text{ where } \varphi_{\alpha \cdot \mathbf{u}} \text{ is the phase field of } \alpha \cdot \widehat{\mathbf{u}}.$$

This approach was arguably motivated by the following result on the projections of such α -color textons onto real (grey-scale) images.

Proposition. *For any color image \mathbf{u} and any α in \mathbb{R}^3 , we have the identity*

$$\alpha \cdot T_{\alpha}(\mathbf{u}) = T(\alpha \cdot \mathbf{u}),$$

where T is the texton operator on grey-level images defined in Equation (2.16). As a consequence, the α -color texton is solution of the two following optimization problems:

$$\alpha \cdot T_{\alpha}(\mathbf{u})(\mathbf{0}) = \max_{\mathbf{v} \in \mathcal{M}_{\mathbf{u}}} \alpha \cdot \mathbf{v}(\mathbf{0}) \text{ and } T_{\alpha}(\mathbf{u}) = \operatorname{Argmin}_{\mathbf{v} \in \Theta_{\mathbf{u}}} \sum_{\mathbf{x} \in D} A(\mathbf{x}) (\alpha \cdot \mathbf{v}(\mathbf{x}))^2,$$

where A is a real-valued image with non-positive Fourier transform, as defined in Proposition 2.5.

The choice of α can be optimized for each input \mathbf{u} through principal component analysis. Indeed, the principal component $\tilde{\alpha}$ seems to be empirically always the best choice in terms of spatial concentration. This can be interpreted in light of the following proposition which shows that the color direction $\tilde{\alpha}$ is the one that captures, in expectation, most of the “energy” of the phase-shifted images.

Proposition. *Let \mathbf{u} be a color image and let $\alpha \in \mathbb{R}^3$ be any color direction. Let A be a real-valued image with non-positive Fourier transform, as defined in Proposition 2.5, and let the weighted energy of a phase-shift $S_\varphi \mathbf{u}$ in the α direction be defined by*

$$E_{A,\alpha}(S_\varphi \mathbf{u}) = \sum_{\mathbf{x} \in D} A(\mathbf{x}) |\alpha \cdot S_\varphi \mathbf{u}(\mathbf{x})|^2.$$

Then, taking the expectation of this energy when the $\varphi(\xi)$ are i.i.d. (up to the Hermitian symmetry condition) uniform on $[0, 2\pi)$, we get

$$\mathbb{E}[E_{A,\alpha}(S_\varphi \mathbf{u})] = 2|D| \left(\sum_{\xi \in D_+ \setminus \{\mathbf{0}\}} \lambda_\xi \right) {}^t \alpha \Gamma_{\mathbf{u}}(\mathbf{0}) \alpha,$$

where $\lambda_\xi = -\frac{1}{|D|} \widehat{A}(\xi)$ are positive coefficients. Therefore, whatever A , $\mathbb{E}[E_{A,\alpha}(S_\varphi \mathbf{u})]$ as a function of $\alpha \in S^2$ (S^2 is here the unit sphere of \mathbb{R}^3) is maximal when $\alpha = \tilde{\alpha}$.

Computing a sparse approximation to the $\tilde{\alpha}$ -texton turns out to be more challenging than its grey-scale counterpart. Indeed, not only intra-channel variance is lost by the process, which was already observed for grey-scale images, but also cross-channel correlations are modified, which can radically alter the color spectrum of textures. One simple way to compensate for this color loss follows. For any given sparsely supported function χ (a support either imposed or derived *e.g.* by hard thresholding):

- compute m_0 the empirical mean and Γ_0 the color-covariance matrix of the original color image u_0
- compute m_χ the empirical mean and Γ_χ the color-covariance matrix of the texton approximation $\chi T_\alpha(u)$
- perform orthogonal diagonalizations $\Gamma_0 = O_0 D_0 O_0^*$ and $\Gamma_\chi = O_\chi D_\chi O_\chi^*$
- compute the equalization matrix $M_{eq} = O_0 D_0^{1/2} D_c^{-1/2} O_c^*$
- compute the adjusted sparse-approximate color texton as $M_{eq}(\chi(\mathbf{x})T(u)(\mathbf{x}) - m_\chi) + m_0$ for each \mathbf{x} in D .

The results for color micro-texture synthesis after this equalization step are visually compelling (see figures in Chapter 2).

1.4 Phase and projection of signals

In a celebrated article [106], Oppenheim and Lim discussed how the information contained in the phase (denoted φ_u where $\widehat{u} = |\widehat{u}|e^{i\varphi_u}$) is of particular importance as regards the shapes contained in image u . Interestingly, the definition of the “texton” for micro-textures is yet another example of a very general phenomenon in signal processing, namely “the importance of phase”.

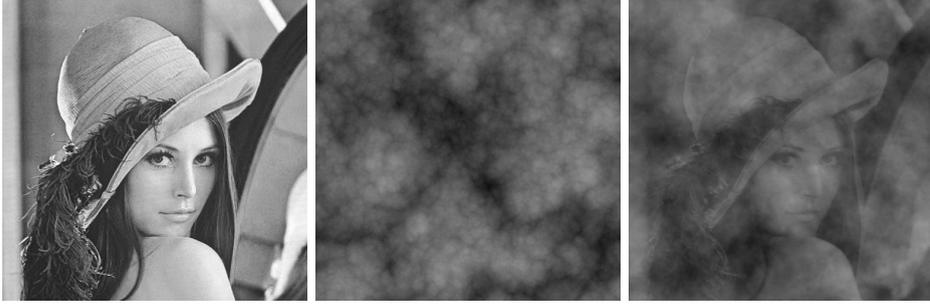


Figure 1.4: The importance of phase in images. From left to right: “Lena”, “clouds” and the image obtained by taking the phase of “Lena” and the modulus of “clouds” as defined in (3.4).

In the grey-scale setting, the texton is indeed a particular case of this phenomenon. More precisely, it is straightforward to see that the texton is the image with the phase of $\delta_{\mathbf{0}}$ – the Dirac mass at the origin – and the modulus of the input texture image. In other words, the remarkable concentration of the texton at the spatial origin (Propositions 2.3 and 2.5) illustrates what could be referred to as a strong “plasticity” property of the set of images with a given Fourier modulus such as \mathcal{M}_u (recall that $\mathcal{M}_u = \{s \in \mathbb{C}^D; |\widehat{s}(\mathbf{0})| = \widehat{u}(\mathbf{0}) \text{ and } |\widehat{s}(\xi)| = |\widehat{u}(\xi)| (\forall \xi \in D \setminus \{\mathbf{0}\})\}$).

Thus, our research on the phase lead us to investigate beyond their commonly accepted role. Indeed, given two images u and v , we realized that the phase/modulus swap defined by

$$\widehat{w}(\xi) = \mathbf{1}_{\{\widehat{u}(\xi) \neq 0\}} \frac{\widehat{u}(\xi)}{|\widehat{u}(\xi)|} |\widehat{v}(\xi)| \quad (1.9)$$

could be interpreted as an orthogonal projection onto some (non-convex) set \mathcal{M}_v , as

$$w \in \arg \min_{s \in \mathcal{M}_v} \|s - u\|_2. \quad (1.10)$$

Hence, we adopted this more general framework and performed a few experiments of orthogonal projection of an image u onto different sets of images, that were designed to both impose the phase of the projection and leave enough degrees of freedom (plasticity) to find visually compelling approximations of the projected image. Most notably, we introduced the sets $\mathcal{D}_v^{(\pi)} = \mathbb{R}^D \cap \{s \in \mathbb{C}^D; \widehat{s}(\xi) \in \mathbb{R} \cdot \widehat{v}(\xi) (\forall \xi \in D)\}$ and $\mathcal{D}_v^{(2\pi)} = \mathbb{R}^D \cap \{s \in \mathbb{C}^D; \widehat{s}(\xi) \in \mathbb{R}_+ \cdot \widehat{v}(\xi) (\forall \xi \in D)\}$ which have the property of imposing respectively modulo π or modulo 2π the phase of images. We then performed orthogonal projections of images u onto these constraint sets.

Proposition. *The set $\mathcal{D}_v^{(\pi)}$ is a linear subspace of \mathbb{R}^D and $\mathcal{D}_v^{(2\pi)}$ is a convex cone. Moreover*

$$\arg \min_{s \in \mathcal{D}_v^{(\pi)}} \|s - u\|_2 = \{P_{\mathcal{D}_v^{(\pi)}}(u)\} \quad (1.11)$$

and

$$\arg \min_{s \in \mathcal{D}_v^{(2\pi)}} \|s - u\|_2 = \{P_{\mathcal{D}_v^{(2\pi)}}(u)\} \quad (1.12)$$

where $P_{\mathcal{D}_v^{(\pi)}}(u)$ (resp. $P_{\mathcal{D}_v^{(2\pi)}}(u)$) is the orthogonal projections of u onto $\mathcal{D}_v^{(\pi)}$ (resp. $\mathcal{D}_v^{(2\pi)}$) which can be defined through their discrete Fourier transform by

$$P_{\mathcal{D}_v^{(\pi)}}(u)(\xi) = \mathbb{1}_{\{\widehat{v}(\xi) \neq 0\}} \operatorname{Re} \left(\widehat{u}(\xi) \overline{\widehat{v}(\xi)} \right) \frac{\widehat{v}(\xi)}{|\widehat{v}(\xi)|^2} \quad (1.13)$$

and

$$P_{\mathcal{D}_v^{(2\pi)}}(u)(\xi) = \mathbb{1}_{\{\operatorname{Re}(\widehat{u}(\xi) \overline{\widehat{v}(\xi)}) > 0\}} \operatorname{Re} \left(\widehat{u}(\xi) \overline{\widehat{v}(\xi)} \right) \frac{\widehat{v}(\xi)}{|\widehat{v}(\xi)|^2}. \quad (1.14)$$

The interest of these constructions lies both in the fact that they can be easily implemented, thanks to their algebraic simplicity, and in the constraints that they impose. Since the phase of both $P_{\mathcal{D}_v^{(2\pi)}}(u)$ and $P_{\mathcal{D}_v^{(\pi)}}(u)$ are constrained to be radically different from the phase of u , a common interpretation of Oppenheim and Lim suggests that the “shapes” of u should be absent of these images. Surprisingly, we observe that some shapes of u can easily



Figure 1.5: Images of the projections of “Lena” (as the target image u) onto two constraint sets defined by “clouds” (as v). In particular, all the images shown here have the phase of “clouds” either modulo π (left) or modulo 2π (right), but “Lena” is somewhat recognizable in each projection. Left: $P_{\mathcal{D}_v^{(\pi)}}(u)$ (phase of “clouds” modulo π). Right: $P_{\mathcal{D}_v^{(2\pi)}}(u)$ (phase of “clouds” modulo 2π).

be recognized by human vision in both $P_{\mathcal{D}_v^{(2\pi)}}(u)$ and $P_{\mathcal{D}_v^{(\pi)}}(u)$. Moreover, we can show that these projections allow for exact reconstruction, and direct formulas can even be provided, under easily satisfied hypotheses. These reconstructions arguably compete with the exact reconstruction scheme proposed by Oppenheim and Lim in their seminal paper [106], as up to

50% of the signal can be reconstructed based solely on projections (see discussion in Chapter 3).

The perceptual phenomenon of recognition of the original image also occurs in more restricted sets of images, where not only the phase, but also the modulus can be somehow constrained, see the experiments in Chapter 3. Interestingly, this recognition phenomenon also occurs – albeit to a noticeably lesser extent – when the discrete Fourier transform is replaced by a unitary operator drawn at random. This further supports the hypothesis that the importance of the phase in signal is actually due to the fact that the phase characterizes the projection onto a large enough set (see (3.4)). Such a replacement, however, allows for recognition results of a lesser quality than with the original Fourier transform. A tentative explanation lies in the particular interaction between the Fourier transform and some edge detector operators, like gradients or smoothed gradients.

For instance we show that for the constraint set \mathcal{S} being either \mathcal{M}_v , $\mathcal{D}_v^{(\pi)}$ or $\mathcal{D}_v^{(2\pi)}$, the projections onto these sets satisfy

$$P_{\mathcal{S}}(u) \in \arg \min_{s \in \mathcal{S}} \|\nabla s - \nabla u\|_2 \quad (1.15)$$

where the gradient operator is classically defined by $\nabla u = (\partial_1 u, \partial_2 u) : D \rightarrow \mathbb{R}^2$ for any image $u : D \rightarrow \mathbb{R}$.

1.5 Asymptotics of the random phase algorithm

At this stage of the discussion, images have always been considered to be defined over a finite grid, typically $\mathbb{Z}/M\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z}$. Shannon-Nyquist theory shows that this setting is well suited to actually sample “real” images (defined over the infinite domain \mathbb{R}^2) perfectly, provided that they can satisfy two hypotheses, namely

- compactness: an image assumed to be periodical or compactly supported
- smoothness: the Fourier transform is itself assumed to have a compact support.

These two assumptions are ubiquitous and often implicit in image processing. They are also generally considered to form a framework adapted to computer vision. Indeed, physical limitations of optical devices – most notably diffraction inherent to the measurement of local light intensities with photoreceptors – legitimate the smoothness assumption. The compactness assumption seems to stand for an obvious consequence of spatial restrictions of any measurement system.

However, as regards textures, these hypotheses might appear somehow limiting: both infinite size and arbitrary levels of non-trivial detail might be of particular interest. As we shall see, these two directions out of the discrete-finite paradigm amounts to studying limits of the random phase model by growing the domain $\mathbb{Z}/M\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z}$ either to \mathbb{Z}^2 (non-local textures) or to $\mathbb{R}^2/\mathbb{Z}^2$ (non-discrete or continuous textures).

1.5.1 Non-local textures

As argued *supra*, a texture is an object of arbitrary dimensions, and thus a non-compact framework seems quite natural. In Chapter 2 (also [36]) we investigated the asymptotics of the random phase algorithm when adapted to output textures over \mathbb{Z}^2 . In other words, assuming a constant level of details, the domain is expanded to infinity and the periodicity assumption is thus waived. Interestingly, we found the limit to be a Gaussian field. More precisely, let us define a spectrum S as a real non-negative function defined over \mathbb{R}^2 , periodic (with a group of periods containing $2\pi\mathbb{Z}^2$) with the following assumptions:

- S is symmetric: $S(-\xi) = S(\xi)$;
- $S(\mathbf{0}) = 0$;
- S is bounded and piecewise continuous on \mathbb{R}^2 .

For any integers M and N , we can define a RPN model U_{MN} based on the spectrum function S which dimensions are $M \times N$ (that is, which is $M\mathbb{Z} \times N\mathbb{Z}$ periodic). In D_{MN}

$$\widehat{U_{MN}}(\xi) = S(\xi)e^{i\varphi(\xi)}$$

for any ξ in $\frac{2\pi}{M}\mathbb{Z} \times \frac{2\pi}{N}\mathbb{Z}$ with $\varphi(\xi)$ i.i.d. modulo Hermitian-symmetry over D_{MN} . In the spatial domain \mathbb{Z}^2 , DFT translates this to

$$U_{MN}(x) = \frac{1}{MN} \sum_{\xi \in D_{MN}} S(\xi)e^{i\varphi(\xi)} e^{2i\pi x \cdot \xi}$$

for any x in \mathbb{Z}^2 . Interestingly, the finite-dimensional distributions of a proper renormalisation of U_{MN} converge to a Gaussian field over \mathbb{Z}^2 .

Proposition. *Under the set of hypotheses upon S above, $\tilde{U}_{MN} = \sqrt{MN}U_{MN}$ converges towards a Gaussian field U defined over \mathbb{Z}^2 in finite-dimensional distributions, as (M, N) tends towards infinity. The covariance of the limit random field U is given by*

$$\text{Cov}(U(x), U(x')) = \frac{1}{4\pi^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} S(\xi_1, \xi_2)^2 \cos((x - x') \cdot (\xi_1, \xi_2)) d\xi_1 d\xi_2 \quad (1.16)$$

for any x and x' in \mathbb{Z}^2 .

This further explains the similarities of the outputs of the two algorithms ADSN and RPN as mentioned in [51].

1.5.2 Continuous textures

Textures are ubiquitous in images. For an image with infinite resolution, one could argue that locally “almost every pixel” (in the sense of measure theory) is part of a texture sub-image. Thus an important aspect of texture analysis is to understand the local properties when gathering an arbitrary level of detail.

In particular, we are interested in the study of sample paths regularity from fields generated by the RPN algorithm from a given spectrum defined over \mathbb{Z}^2 . Regularity of random fields is an important area of modern mathematical research. For instance, multifractal analysis is mainly concerned with regularity analysis of random processes and fields, and has drawn considerable attention since the seminal work of Mandelbrot [92]. These developments have been of great interest for the image processing research community [113], [131].

The study of the limit of the RPN model over $\mathbb{R}^2/\mathbb{Z}^2$ turns out to be, in many respects, less straightforward than the \mathbb{Z}^2 Gaussian limit presented above. Reaching an infinite resolution amounts to grow the initial discrete domain $\mathbb{Z}/M\mathbb{Z} \times \mathbb{Z}/N\mathbb{Z} \equiv (\frac{1}{M}\mathbb{Z})/\mathbb{Z} \times (\frac{1}{N}\mathbb{Z})/\mathbb{Z}$ to the continuous torus $\mathbb{R}^2/\mathbb{Z}^2$, and the Fourier domain to infinite discrete plane \mathbb{Z}^2 .

We chose to extend the random phase algorithm by considering stochastic sums over \mathbb{Z}^d

$$\sum_{n \in \mathbb{Z}^d} a_n e^{i(n \cdot t + \Phi_n)} \quad (1.17)$$

with the following properties:

- \mathcal{H}_1 : “ $(a_n)_{n \in \mathbb{Z}^d}$ is a deterministic, real, non-negative, even ($a_{-n} = a_n$ for all n in \mathbb{Z}^d) square summable family with $a_0 = 0$ ”
- \mathcal{H}_2 : “ $(\Phi_n)_{n \in \mathbb{Z}^d}$ is a pure phase noise field, that is for all n in \mathbb{Z}^d , $\Phi_{-n} = -\Phi_n$ (modulo 2π) almost surely, Φ_n has uniform distribution over $\mathbb{R}/2\pi\mathbb{Z}$ and $(\Phi_n)_{n \in \mathcal{A}}$ are independent for all $\mathcal{A} \subset \mathbb{Z}^d$ such that \mathcal{A} and $-\mathcal{A}$ do not intersect.”

Properties of such fields are the object of Chapters 4 and 5 (and articles [118] and [14]), whose results are summarized below.

Marginal law. To highlight the difference of nature between this work and the asymptotics presented above, we notice that for any square summable family $(b_n)_{n \in \mathbb{N}}$ the sum $\sum_{n \in \mathbb{N}} b_n e^{i\Phi_n}$ is non-Gaussian – actually non-infinitely divisible – as its characteristic function vanishes. In Chapter 4 (also [118]) we prove that the density function of the marginal law is uniformly continuous and bounded over \mathbb{R} as long as at least three (non-symmetrical) components do not vanish.

Sums over \mathbb{Z}^d . In this work, we extended results already known for random series of functions $\sum_{n \in \mathbb{Z}} a_n e^{i(nt + \Phi_n)}$ ($d = 1$) for sums over \mathbb{Z}^d with arbitrary d , as defined in (1.17). In particular, we have been interested in sample paths continuity without modification. The celebrated Kolmogorov continuity Theorem (see [80] Chapter 11) was thus of little help. The study of both a pointwise and a uniform limit was needed. Interestingly, as we shall discuss below (see Theorem 4.1), these limits happen to be equivalent under some reasonable hypotheses.

One difficulty arising when dealing with Fourier sums over \mathbb{Z}^d is that there is no canonical way to define infinite summation, which makes the definition of both pointwise and uniform limits ambiguous.

We thus defined “methods of summation”, inspired by Kahane’s summation matrices [69] in the one-dimensional configuration, as follows.

Definition (Method of summation). $(\mathcal{A}_k)_{k \in \mathbb{N}}$ is said to be a (symmetrical) method of summation over \mathbb{Z}^d if

1. for each k , \mathcal{A}_k is a finite subset of \mathbb{Z}^d (such that $-\mathcal{A}_k = \mathcal{A}_k$);
2. for each k , $\mathcal{A}_k \subset \mathcal{A}_{k+1}$;
3. $\bigcup_{k \in \mathbb{N}} \mathcal{A}_k = \mathbb{Z}^d$.

Given a Banach space B , a family $(x_n)_{n \in \mathbb{Z}^d}$ of elements in B is said to be summable according to $(\mathcal{A}_k)_{k \in \mathbb{N}}$ if $\sum_{n \in \mathcal{A}_k} x_n$ converges in B as $k \rightarrow \infty$.

Interestingly, in our case as defined by Equation (1.17), independence with respect to the methods of summation can be obtained, just as with summation matrices in the one-dimensional case ([69] Chapter 2).

Proposition. Let $(\mathcal{A}_k)_{k \in \mathbb{N}}$ be a symmetrical method of summation in \mathbb{Z}^d . Assume that, almost surely, the sequence of functions

$$S_{\mathcal{A}_k} : t \mapsto \sum_{n \in \mathcal{A}_k} a_n e^{i(n \cdot t + \Phi_n)} \quad (1.18)$$

converges uniformly (resp. is uniformly bounded) on \mathbb{T}^d as $k \rightarrow \infty$ and call $S_{\mathcal{A}}$ its limit. Then, under the hypotheses \mathcal{H}_1 and \mathcal{H}_2 , for any other method of summation $(\mathcal{B}_k)_{k \in \mathbb{N}}$, the sequence of functions

$$S_{\mathcal{B}_k} : t \mapsto \sum_{n \in \mathcal{B}_k} a_n e^{i(n \cdot t + \Phi_n)} \quad (1.19)$$

converges a. s. uniformly to $S_{\mathcal{A}}$ (resp. is a. s. uniformly bounded) on \mathbb{T}^d as $k \rightarrow \infty$.

Continuity of samples. This independence result allows an extension of a theorem by Billard [15] and Kahane [69] to the more general case where $d \geq 2$, which is the main result of Chapter 3. Part of this result had already been proved in the general case where the circle is replaced by any compact abelian or non-abelian group (see [47]). Our purpose was to extend the result of Billard and Kahane to our modeling of micro-textures (random Fourier sums defined over $\mathbb{R}^2/\mathbb{Z}^2$) and the step from \mathbb{R}/\mathbb{Z} to $\mathbb{R}^2/\mathbb{Z}^2$ paved the way to the generalization to $\mathbb{R}^d/\mathbb{Z}^d$.

Theorem (Billard-Kahane, extended in dimension d). Under the hypotheses \mathcal{H}_1 and \mathcal{H}_2 , the following conditions are equivalent:

- (i) ω -almost surely, there exists a continuous function $X(\omega, \cdot)$, such that $(a_n e^{i\Phi_n(\omega)})_{n \in \mathbb{Z}^d}$ are the Fourier coefficients of $X(\omega, \cdot)$;
- (ii) there exists a method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$ such that, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ converges uniformly;
- (iii) for all methods of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ converges uniformly;

- (iv) there exists a method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$ such that, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is bounded;
- (v) for all methods of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is bounded;
- (vi) for all methods of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, almost surely, for all t in \mathbb{T}^d , $(S_{\mathcal{A}_k}(t))_{k \in \mathbb{N}}$ converges.

This result turns out to be quite unpractical, and the results proved in Chapter 4 aim at complementing its claims with a more tractable approach. We are indeed looking for tractable conditions for the equivalence chain – denoted by (\star) – in Theorem 1.5.2 to hold true. This area of research has been also widely studied in the one-dimensional case [69], [32]. Let us also mention the breakthrough by Fernique [45], first introducing metric entropy methods to the field, which paved the way for necessary and sufficient conditions on random Fourier series defined over \mathbb{R} . Unfortunately, these latter conditions have many drawbacks: they are practically intractable and they are hard to extend to dimensionality greater than one. Hence, we developed a framework of hypotheses to provide continuity results for random sums (1.17).

Assume that Hypothesis \mathcal{H}_1 is satisfied and assume that $(N_k)_{k \in \mathbb{N}}$ is an increasing sequence of integers and that $(\mathcal{A}_k)_{k \in \mathbb{N}}$ is a method of summation such that, for each k , $\mathcal{A}_k \subset B_\infty(N_k)$ the l^∞ ball. Let us state the two hypothesis

- $\mathcal{H}_3: \sum_{k \in \mathbb{N}} \left(\log(N_{k+1}) \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} a_n^2 \right)^{1/2} < \infty$
- $\mathcal{H}_4: \sum_{k \in \mathbb{N}} \frac{1}{N_k} < \infty$.

Theorem. *Assume that hypotheses \mathcal{H}_1 , \mathcal{H}_2 , \mathcal{H}_3 and \mathcal{H}_4 are satisfied. Then (\star) holds.*

Interestingly, an almost converse results can be stated. For all k in \mathbb{N} , let us define $\sigma_k \geq 0$ by

$$\sigma_k^2 := \sum_{2^k < |n|_\infty \leq 2^{k+1}} a_n^2. \quad (1.20)$$

Theorem. *Assume hypotheses \mathcal{H}_1 , \mathcal{H}_2 and*

$$\mathcal{H}_7: \sum_{k \in \mathbb{N}} \sigma_k = \infty$$

are satisfied, with σ_k^2 defined by (5.40). Then (\star) does not hold, and $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is almost surely unbounded for every method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$.

Hölder regularity. In order to quantify the degree of regularity of the samples of random phase fields, we introduced the Hölder regularity and extended again results already known in the case $d = 1$ to the general case. In order to gain some insight over anisotropic properties of these fields, we introduced coefficients and summation sets defined by quasi-norms [13]. Let $E = \text{diag}(\alpha_1, \dots, \alpha_d)$ be a diagonal matrix with positive eigenvalues $\alpha_1, \dots, \alpha_d \in (0, +\infty)$ and τ_E be a quasi-norm associated with E [13] and let us denote $\underline{\alpha} = \min_{1 \leq j \leq d} \alpha_j$. Assume that Hypothesis \mathcal{H}_1 is satisfied and, for all k in \mathbb{N} , define $\sigma_{\tau_E, k} \geq 0$ by

$$\sigma_{\tau_E, k}^2 := \sum_{2^k < \tau_E(n) \leq 2^{k+1}} \alpha_n^2. \quad (1.21)$$

Theorem. Assume that hypotheses \mathcal{H}_1 , \mathcal{H}_2 and

$$\mathcal{H}_8(\tau_E) : \exists \beta \in (0, \underline{\alpha}) \text{ and } C > 0; \forall k \in \mathbb{N}, \sigma_{\tau_E, k} \leq C 2^{-\beta k}$$

are satisfied, with $\sigma_{\tau_E, k}^2$ defined by (5.69). Then, almost surely, (\star) holds and, for any method of summation, the limit function X satisfies that almost surely, there exists a constant $C > 0$ such that

$$\forall t, s \in \mathbb{T}^d, |X(t) - X(s)| \leq C \tau_E(t - s)^\beta \log(1 + \tau_E(t - s)^{-1})^{1/2}.$$

We may obtain a partial converse in the isotropic case for $E = I$, considering $\tau_E = |\cdot|_\infty$. Actually, in this setting, we extend the classic condition for the case $d = 1$ (see [69], Chapter 7, Theorem 3).

Proposition. Let $f : \mathbb{T}^d \mapsto \mathbb{R}$ be a α -Hölder function. Then for all k in \mathbb{N}

$$\sum_{2^k < |n|_\infty \leq 2^{k+1}} |\hat{f}_n|^2 \leq C_f 2^{-\alpha k} \quad (1.22)$$

holds for some constant C_f , where \hat{f}_n denotes the n -th Fourier coefficient given by

$$\hat{f}_n = \frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} f(t) e^{-in \cdot t} dt.$$

These conditions allow to derive many tractable examples of spectrum profiles and to compute their Hölder regularity with a fair precision. For instance, we derive the critical Hölder exponent with isotropic power spectrum, where $a_n = |n|_2^{-\alpha}$ for some α in $(0, 1)$. Indeed, the theorem stated in the last paragraph (Theorem 5.1 in Chapter 5) yields that if $\alpha > d/2$, the condition (\star) holds. Moreover sum-integral comparison allows to derive two constants C_1 and C_2 such that

$$C_1 2^{k(d-2\alpha)/2} \leq \sigma_k \leq C_2 2^{k(d-2\alpha)/2}. \quad (1.23)$$

Thus, for any norm ν the random phase field associated with $a_n = \nu(n)^{-\alpha}$ is almost surely

- β -Hölder for all $\beta < \alpha - \frac{d}{2}$ thanks to the theorem in this paragraph (Theorem 5.3 in Chapter 5)
- not β -Hölder for $\beta > \alpha - \frac{d}{2}$ thanks to the proposition in this paragraph (Proposition 5.3 in Chapter 5).

1.6 Outline

The rest of this thesis is organized as follows.

Chapter 2 gathers our work on what we defined as *texton* and a sparse and localized representation of textures, along with the asymptotical analysis of the RPN model defined over \mathbb{Z}^2 . In this Chapter, we first present the basic properties of the gray-scale *texton*, and provide a few results with respect to its spatial concentration properties. After discussing scaling properties of Gaussian periodic textures, we establish the convergence of Random Phase Noise textures towards a Gaussian field with appropriate renormalization. We then move on to build sparsely supported approximations of *textons*. The last section of Chapter 2 is devoted to studying the extension of the gray-scale case to color images, and two approaches are proposed: first, a color *texton* is constructed via the phase field of a projection onto real (gray-scale) images; second, a matricial color *texton* is introduced as a matricial square root of the covariance function. Finally, sparse approximation strategies for color *textons* are discussed.

Part of this work has already been published as conference proceedings, see [34] and [35]. A more thorough version has been submitted for publication in a journal.

Chapter 3 focuses on our experiments and tentative explanations on the phase-constrained projections. In this Chapter, we develop the following argument: the “importance of phase in signals”, as detailed by Oppenheim and Lim in [106], can be considered as a case of a more general human vision recognition phenomenon. Namely, it appears that recognition of large error approximations (LEA) of images occurs for a much broader range of sets than the sets defined by constraining the Fourier modulus from an image. We illustrate our purpose by choosing sets where the phase itself is constrained, and show a few surprising results. We extend our range of experiments by replacing the Fourier transform by more generic unitary transforms and discuss how the fact that the Fourier transform diagonalizes translation operators could explain why it is so well suited to LEA.

Chapter 4 dives into the continuous texture asymptotics. We provide basic results on the continuous limit of the random phase noise model. We introduce our “method of summation” framework to deal with infinite summations over multi-dimensional domains (without total order). We propose a generalization of the Billard-Kahane Theorem on the equivalence between many different almost sure convergences, most notably uniform convergence, pointwise convergence, and boundedness. This work has been accepted for publication in the Journal of Fourier Analysis and Applications.

Chapter 5 studies continuous texture asymptotics further, by focusing on conditions for almost sure regularity. Both continuity and Hölder regularity are investigated and in each case, both sufficient and necessary conditions are stated. Taking advantage of the multi-dimensional framework, we emphasize the anisotropic aspect of textures. This work has been submitted for publication.

Chapter 6 concludes by summing up the main findings and the some research perspective that can be drawn from our work.

Chapter 2

A Texton for Random Phase and Gaussian Textures

2.1 Introduction

2.1.1 The notion of texton

It is a very general problem to find a good descriptor of a texture that can be used at the same time for the analysis of the texture (or its discrimination) and for its synthesis. In his seminal work, Julesz ([63], [64], [65], [66]) introduced the notion of *texton* to describe “the putative units of pre-attentive human texture perception”. This notion, first introduced for artificial texture patterns, has then been generalized in the literature for different types of textures. Let us mainly mention the paper of Zhu et al. “What are Textons” [136] and the ones of Leung and Malik [82] and Malik et al. [89]. In these papers, the authors aim at answering the problem raised by Zhu et al.: “Unfortunately, the word “texton” remains a vague concept in the literature for lack of a good mathematical model.” This will be one of the aim of this chapter: to give a good mathematical model of the texton. In Zhu et al., the authors learn the texton dictionary of a natural texture image by fitting a generative model to the observed images. They study the geometric (for static texture images), dynamic (for video sequences, and moving textons are called “motons”) and photometric (for images representing a 3D surface under varying illuminations, and the texton is then called a “lighton” in that framework) structures of the texton representation. In Leung and Malik, the textons are local image patches obtained as the K centers (obtained by K -means) of the filter responses to a filter bank (made of Gaussian and difference of Gaussian filters). This definition is also extended to relief textures (images of a material under different illuminations), and in that framework they obtained a so-called 3D texton model.

Here, in this chapter, we will consider two particular models of texture that are either Gaussian stationary random fields or random phase noises. These are two mathematical models of textures that are precisely defined, and this will allow us to give a precise mathematical definition of their texton. We will also be able to give theoretical properties of this texton, in particular about its sparsity. Gaussian and Random Phase textures are models of

textures that are widely used in Computer Graphics (because they are limits of spot noise type models [127], see also *e.g.* [77]) but also for medical images. In particular, Gaussian textures encompass all $1/f^\alpha$ noise models (sometimes also called power-law noise or pink noise). All these are micro-textures models, as opposed to macro-textures, according to the discussion in [51]. Macro-textures can be described as images containing spatially organized visual elements, like a periodic brick of wall, or a cheetah skin image for instance. On the contrary, micro-textures do not contain well-identified “objects” and are characterized by the fact that they are perceptually invariant under randomizing the phases in the Fourier transform. This property is not true for macro-textures, because changing the phases then completely destroys the “objects” of the image. Among all possible realizations for the phase field, we will here focus on the one that has identically null phases. This is a very simple choice, but it leads to very interesting properties for both texture analysis and texture synthesis. We will call this particular image the *texton* of the micro-texture, and will we show that it can be related to the two fundamental aspects of textures according to Julesz: the second-order statistics and the texton as a notion of an elementary shape that characterizes the texture.

In that framework, the texton we define can be used for the synthesis of the texture on domains of arbitrary size but also for the discrimination of textures. Indeed, we have a natural distance between textures (given by the optimal transportation distance between probabilistic distributions) that we will prove to be simply equal, in the case of gray-level images, to the L^2 distance between their respective texton. Being able to discriminate different textures via their texton can then be used for image segmentation as done for instance by Malik et. al. in [89]. The notion of texton we develop here has been first described in a preliminary version of this work [35]. It has been used in [132] for texture synthesis and mixing (using the optimal transportation distance), and also extended to dynamic textures.

This chapter is organized as follows. We first start by defining the framework and the notations. We also recall the precise definitions of Gaussian Textures (GT) and of Random Phase Noise (RPN). In Section 2.2 we define the texton of a gray-level image (as the image that has the same Fourier amplitudes and null phases) and we give many properties of this texton, in particular in terms of “concentration”. Then in Section 2.3, we discuss the way a small support texton can be obtained from the original texton, and the way a texture can be synthesized on a domain of arbitrary size. In Section 2.4, the definition of texton is extended to color texture images. The extension is not straightforward, and there are two different definitions: the texton as a color image or the texton as a matrix-valued image (this second definition is more general and contains the first one). Finally, we end the chapter with some conclusions and perspectives.

2.1.2 Framework and notations

We first describe here the framework and the notations that will be used in the following. For an integer $n \geq 1$, we will denote by I_n the discrete interval defined by $\{-\frac{n-1}{2}, \dots, -1, 0, 1, \dots, \frac{n-1}{2}\}$ when n is odd and by $\{-\frac{n}{2}, \dots, -1, 0, 1, \dots, \frac{n}{2} - 1\}$ when n is even, also associated with the cyclic group $\mathbb{Z}/n\mathbb{Z}$. In all cases the cardinality of I_n is n . Moreover we will denote by I_n^0 the

points of I_n that do not have a distinct symmetric point in I_n , that is, $I_n^0 = \{0\}$ when n is odd and $I_n^0 = \{0, -\frac{n}{2}\}$ when n is even.

Let $u : D \rightarrow \mathbb{R}$ be a discrete $M \times N$ gray-level image, defined on the domain $D = I_M \times I_N$. For $\mathbf{x} \in D$, $u(\mathbf{x})$ represents the intensity of the pixel \mathbf{x} . Notice that it is not very usual to have negative spatial coordinates for the pixels, but because of the properties of the texton (symmetry and concentration around the point $\mathbf{0} = (0, 0)$), it is more convenient to consider that the point $\mathbf{0}$ is located at the center of the image.

The (non-unitary) Discrete Fourier Transform (DFT) of u is the complex-valued function defined on D by

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{u}(\boldsymbol{\xi}) = \sum_{\mathbf{x} \in D} u(\mathbf{x}) e^{-2i\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle}, \quad (2.1)$$

where the inner product between $\mathbf{x} = (x_1, x_2)$ and $\boldsymbol{\xi} = (\xi_1, \xi_2)$ is defined by $\langle \mathbf{x}, \boldsymbol{\xi} \rangle = \frac{1}{M} x_1 \xi_1 + \frac{1}{N} x_2 \xi_2$ to have simpler formulas. As usual, the image u can be recovered from its Fourier transform \widehat{u} by the (non-unitary) Inverse Discrete Fourier Transform:

$$\forall \mathbf{x} \in D, \quad u(\mathbf{x}) = \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D} \widehat{u}(\boldsymbol{\xi}) e^{2i\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle}, \quad (2.2)$$

where $|D| = MN$ is the size of the domain D . In particular, the mean value of u , defined as

$$m_u = \frac{1}{|D|} \sum_{\mathbf{x} \in D} u(\mathbf{x}), \quad (2.3)$$

satisfies $|D|m_u = \widehat{u}(\mathbf{0})$.

A function $\varphi : D \rightarrow S^1$ is a phase of u if it satisfies $\widehat{u}(\boldsymbol{\xi}) = |\widehat{u}(\boldsymbol{\xi})| e^{i\varphi(\boldsymbol{\xi})}$ for all $\boldsymbol{\xi} \in D$ (note that $\varphi(\boldsymbol{\xi})$ may take any value when $\widehat{u}(\boldsymbol{\xi}) = 0$). Since the image u is real-valued, we necessarily have $\varphi(-\boldsymbol{\xi}) = -\varphi(\boldsymbol{\xi})$ for all $\boldsymbol{\xi} \in D$ such that $\widehat{u}(\boldsymbol{\xi}) \neq 0$. More precisely, if \widehat{u} does not vanish on D , the image u is real-valued if and only if the following constraints are satisfied:

$$\begin{cases} |\widehat{u}(-\boldsymbol{\xi})| = |\widehat{u}(\boldsymbol{\xi})| \text{ and } \varphi(-\boldsymbol{\xi}) = -\varphi(\boldsymbol{\xi}) & \text{for } \boldsymbol{\xi} \in D \setminus I_M^0 \times I_N^0, \\ \varphi(\boldsymbol{\xi}) \in \{0, \pi\} & \text{for } \boldsymbol{\xi} \in I_M^0 \times I_N^0. \end{cases} \quad (2.4)$$

In the following, we will impose this set of constraints (2.4) even when the Fourier transform of u may vanish. To have simpler notations, unless specified otherwise we will assume in the following that M and N are odd, and define

$$D_+ = \left(\{0\} \times \left\{ 1, 2, \dots, \frac{N-1}{2} \right\} \right) \cup \left(\left\{ 1, 2, \dots, \frac{M-1}{2} \right\} \times I_N \right), \quad (2.5)$$

so that the domain D can be partitioned into $D = D_+ \cup (-D_+) \cup \{\mathbf{0}\}$ (disjoint union). Then, to define a real-valued image on D , thanks to (2.4) it is enough to specify the values of its Fourier phases and amplitudes on $D_+ \cup \{\mathbf{0}\}$. Let us insist on the fact that we assume M and N odd in the theoretical results only to simplify the proofs but in practice, from the numerical and experimental point of view, M and N can be of any parity.

Given an image $u : D \rightarrow \mathbb{R}$, we may consider its periodic extension to \mathbb{Z}^2 as the image $\dot{u} : \mathbb{Z}^2 \rightarrow \mathbb{R}$ defined by

$$\forall \mathbf{x} = (x_1, x_2) \in \mathbb{Z}^2, \quad \dot{u}(\mathbf{x}) = u(x_1 \pmod{M}, x_2 \pmod{N}). \quad (2.6)$$

This allows us to define the periodic convolution between two images $u : D \rightarrow \mathbb{R}$ and $v : D \rightarrow \mathbb{R}$ as the image $u \star v : D \rightarrow \mathbb{R}$ given by

$$\forall \mathbf{x} \in D, \quad (u \star v)(\mathbf{x}) = \sum_{\mathbf{y} \in D} u(\mathbf{y}) v(\mathbf{x} - \mathbf{y}).$$

This convolution is described more simply in Fourier domain, since the convolution/product exchange property, which will be extensively used throughout this chapter, states that

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{u \star v}(\boldsymbol{\xi}) = \widehat{u}(\boldsymbol{\xi}) \widehat{v}(\boldsymbol{\xi}). \quad (2.7)$$

We shall also make use of Parseval's Theorem: if the L^2 norm of a (real-valued or complex-valued) image $u : D \rightarrow \mathbb{C}$ is defined by $\|u\|_2 = (\sum_{\mathbf{x} \in D} |u(\mathbf{x})|^2)^{1/2}$, then one simply has $\|\widehat{u}\|_2^2 = |D| \cdot \|u\|_2^2$.

2.1.3 Gaussian random fields and Fourier transform

Stationary Gaussian textures form a widely used model of textures. For instance, the ADSN model that we introduce *infra* from [51] – the limit of a renormalized shot noise field with infinitely many shots – are known to be Gaussian. Such models are very well characterized in the Fourier domain (in the periodic case). Indeed, we have the following result.

Theorem 2.1. *Let $(U(\mathbf{x}))_{\mathbf{x} \in D}$ be a real-valued random field on D . Then $(U(\mathbf{x}))_{\mathbf{x} \in D}$ is a zero-mean Gaussian periodic stationary random field if and only if the random variables*

$$\left\{ \widehat{U}(\mathbf{0}), \operatorname{Re} \widehat{U}(\boldsymbol{\xi}), \operatorname{Im} \widehat{U}(\boldsymbol{\xi}); \boldsymbol{\xi} \in D_+ \right\}$$

are independent zero-mean Gaussian variables. Moreover, in this case, if Γ denotes the covariance of U defined by $\Gamma(\mathbf{x}) = \operatorname{Cov}(U(\mathbf{x}), U(\mathbf{0}))$ for all $\mathbf{x} \in D$, then

$$\operatorname{Var}(\widehat{U}(\mathbf{0})) = |D| \cdot \widehat{\Gamma}(\mathbf{0}) \quad \text{and} \quad \forall \boldsymbol{\xi} \in D_+, \quad \operatorname{Var}(\operatorname{Re} \widehat{U}(\boldsymbol{\xi})) = \operatorname{Var}(\operatorname{Im} \widehat{U}(\boldsymbol{\xi})) = \frac{1}{2} |D| \cdot \widehat{\Gamma}(\boldsymbol{\xi}). \quad (2.8)$$

The proof is given in the Appendix. This theorem gives a characterization of periodic stationary Gaussian textures: the phases $(\varphi(\boldsymbol{\xi}))_{\boldsymbol{\xi} \in D_+}$ are independent identically distributed uniformly on $[0, 2\pi)$ and independent from the Fourier amplitudes $(R(\boldsymbol{\xi}))_{\boldsymbol{\xi} \in D_+}$ which are independent random variables following each a Rayleigh distribution of parameter $\sigma = \sqrt{\frac{1}{2} |D| \cdot \widehat{\Gamma}(\boldsymbol{\xi})}$; except at $\boldsymbol{\xi} = \mathbf{0}$ where $\widehat{U}(\mathbf{0})$ follows a zero-mean normal distribution with variance $|D| \cdot \widehat{\Gamma}(\mathbf{0})$. The simplest case of a Gaussian stationary random field is the white noise: the $U(\mathbf{x})$ are i.i.d. following the $\mathcal{N}(0, 1)$ distribution. Then, in this case $\Gamma(\mathbf{x}) = \delta_{\mathbf{0}}$ (the indicator function of $\{\mathbf{0}\}$), and thus $\widehat{\Gamma}(\boldsymbol{\xi}) = 1$ for all $\boldsymbol{\xi} \in D$.

In the following, we will denote by $\operatorname{GT}(\Gamma)$ (GT stands for Gaussian Texture) the law of the zero-mean Gaussian periodic stationary random field with covariance function Γ .

2.1.4 Two mathematical models of micro-textures

The mathematical models of micro-textures that we will consider are the two following models of discrete random fields, called (following the terminology of Galerne et al. [51]) respectively RPN (for Random Phase Noise) and ADSN (for Asymptotic Discrete Spot Noise). These two models define, from an original image u , a stationary random field (a random image) whose first and second order moments are given by the corresponding empirical moments of u . Since the first moment (the mean value of the field) can be treated separately by adding a constant value to the random image, we will systematically drop the DC term and focus on zero-mean random fields. This convention slightly changes the definitions of the RPN and ADSN models, but yields a simplification that will be particularly convenient in Section 2.3.

Definition 2.1. *The ADSN associated to a kernel $k : D \rightarrow \mathbb{R}$ is the random image*

$$U = k \star W, \quad (2.9)$$

where W is a white noise with variance 1, that is, the random variables $W(\mathbf{x})$, $\mathbf{x} \in D$ are i.i.d. following the normal distribution $\mathcal{N}(0, 1)$. We will denote by $\text{ADSN}(k)$ the law of U .

Definition 2.2. *The (extended) RPN associated to a kernel $k : D \rightarrow \mathbb{R}$ is the random image $U : D \rightarrow \mathbb{R}$ defined (in Fourier domain) by*

$$\widehat{U}(0) = \sqrt{|D|} \widehat{k}(0)W \quad \text{and} \quad \forall \boldsymbol{\xi} \in D \setminus \{\mathbf{0}\}, \quad \widehat{U}(\boldsymbol{\xi}) = \sqrt{|D|} |\widehat{k}(\boldsymbol{\xi})| e^{i\Phi(\boldsymbol{\xi})}, \quad (2.10)$$

where the $\Phi(\boldsymbol{\xi})$, $\boldsymbol{\xi} \in D_+$ are uniformly distributed on $[0, 2\pi)$ and independent modulo the constraint (2.4), and independent of $W \sim \mathcal{N}(0, 1)$. We will denote by $\text{RPN}(k)$ the law of U .

We first recall the first and second-order statistics of these two models.

Proposition 2.1 (statistics of the ADSN and extended RPN models). *If $U \sim \text{ADSN}(k)$ or $U \sim \text{RPN}(k)$, then $\mathbb{E}(U(\mathbf{x})) = 0$ for all $\mathbf{x} \in D$ and*

$$\forall \mathbf{x}, \mathbf{y} \in D, \quad \mathbb{E}(U(\mathbf{x})U(\mathbf{y})) = \Gamma_k(\mathbf{y} - \mathbf{x}),$$

where Γ_k is the D -periodic function defined by

$$\forall \mathbf{z} \in \mathbb{Z}^2, \quad \Gamma_k(\mathbf{z}) = \frac{1}{|D|} \sum_{\mathbf{x} \in D} k(\mathbf{x}) \widehat{k}(\mathbf{x} - \mathbf{z}). \quad (2.11)$$

Proof. First, notice that $\mathbb{E}[\widehat{U}(\boldsymbol{\xi})] = 0$ for all $\boldsymbol{\xi}$ so thanks to linearity of the inverse Fourier transform and expectation operators $\mathbb{E}(U(\mathbf{x})) = 0$ for all $\mathbf{x} \in D$.

Moreover for all \mathbf{x} and \mathbf{y} in D , $\mathbb{E}(U(\mathbf{x})U(\mathbf{y})) = \mathbb{E}[\frac{1}{|D|} \sum_{\mathbf{z} \in D} U(\mathbf{x} + \mathbf{z}) \widehat{U}(\mathbf{y} + \mathbf{z})] = \mathbb{E}[U \star U(\mathbf{x} - \mathbf{y})]$ thanks to the stationarity of U . Now simply notice that $\mathbb{E}[\widehat{U \star U}(\boldsymbol{\xi})] = \mathbb{E}[|\widehat{U}|^2(\boldsymbol{\xi})] = |\widehat{u}|^2(\boldsymbol{\xi})$ for all $\boldsymbol{\xi}$ to conclude. \square

The function Γ_{u-m_u} that naturally appears for the RPN model is nothing but the empirical covariance of u , also called auto-correlation. When $m_u = 0$, the two models have the same statistics up to second order. Contrary to the RPN model, the ADSN model is Gaussian: the random field $U \sim \text{ADSN}(u)$ is a stationary zero-mean Gaussian random field whose covariance function is Γ_u (that is, $\text{ADSN}(u) = \text{GT}(\Gamma_u)$). Conversely, for any D -periodic covariance function $\Gamma : \mathbb{Z}^2 \rightarrow \mathbb{R}$, $\text{GT}(\Gamma) = \text{ADSN}(u)$ whenever u satisfies $|\widehat{u}(\boldsymbol{\xi})|^2 = |D| \cdot \widehat{\Gamma}(\boldsymbol{\xi})$ for all $\boldsymbol{\xi} \in D$. This means that the two models GT and ADSN are the same for gray-level textures. We will see however that this property is not true anymore for color textures.

In practice, as we explained above, the synthesis of a RPN or ADSN texture that looks like an given exemplar texture image u is obtained with $m_u + U$ where $U \sim \text{RPN}(u)$ or $U \sim \text{ADSN}(u - m_u)$. Notice that for any $c \in \mathbb{R}$, one has $\text{RPN}(u + c) = \text{RPN}(u)$ (and in particular $\text{RPN}(u - m_u) = \text{RPN}(u)$), but $\text{ADSN}(u + c) \neq \text{ADSN}(u)$ if $c \neq 0$.

The difference between the models RPN and ADSN can be simply understood in Fourier domain. Indeed, if $U \sim \text{ADSN}(u)$, then

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{U}(\boldsymbol{\xi}) = \widehat{u}(\boldsymbol{\xi}) \widehat{W}(\boldsymbol{\xi}). \quad (2.12)$$

The random variables $\widehat{U}(\boldsymbol{\xi})$, $\boldsymbol{\xi} \in D_+$, are i.i.d. and can be written $\widehat{U}(\boldsymbol{\xi}) = R(\boldsymbol{\xi})e^{i\Phi(\boldsymbol{\xi})}$, where each $R(\boldsymbol{\xi})$ follows a Rayleigh distribution with parameter $\sigma = |\widehat{u}(\boldsymbol{\xi})|/\sqrt{2}$ and is independent from $\Phi(\boldsymbol{\xi})$, which is uniformly distributed on $[0, 2\pi)$. Hence, the ADSN model can be viewed as a RPN process (uniform phase randomization) followed by a random perturbation (a Rayleigh-distributed multiplicative noise) of the modulus of the Fourier transform. This perturbation has no perceptual impact (see [51]), but it makes the model Gaussian. It is also interesting to notice that contrary to the RPN model for which $|\widehat{U}(\boldsymbol{\xi})| = |\widehat{u}(\boldsymbol{\xi})|$ is a deterministic equality for any $\boldsymbol{\xi} \neq \mathbf{0}$, one has for the ADSN model

$$\forall \boldsymbol{\xi} \in D, \quad \mathbb{E} \left(|\widehat{U}(\boldsymbol{\xi})|^2 \right) = |\widehat{u}(\boldsymbol{\xi})|^2 \quad \text{and} \quad \mathbb{E} \left(|\widehat{U}(\boldsymbol{\xi})| \right) = \begin{cases} \frac{\sqrt{\pi}}{2} |\widehat{u}(\boldsymbol{\xi})| & \text{if } \boldsymbol{\xi} \neq \mathbf{0}, \\ \sqrt{\frac{2}{\pi}} |\widehat{u}(\mathbf{0})| & \text{if } \boldsymbol{\xi} = \mathbf{0}. \end{cases} \quad (2.13)$$

2.2 Definition and properties of the texton

2.2.1 A variational characterization of the texton

As we already explained in the introduction, a micro-texture image is characterized by the fact that it is perceptually invariant under phase randomization. This can be formalized by associating, to an image $u : D \rightarrow \mathbb{R}$, the set \mathcal{M}_u of images that have the same mean value as u and only differ from u by their phase function:

$$\mathcal{M}_u = \{v : D \rightarrow \mathbb{R}; \widehat{v}(\mathbf{0}) = \widehat{u}(\mathbf{0}) \text{ and } \forall \boldsymbol{\xi} \neq \mathbf{0}, |\widehat{v}(\boldsymbol{\xi})| = |\widehat{u}(\boldsymbol{\xi})|\}. \quad (2.14)$$

Any element of \mathcal{M}_u defines the same RPN or ADSN model as u . Now, we would like to find, among all equivalent members of \mathcal{M}_u , one representative that will be as most ‘‘concentrated’’ or ‘‘sparse’’ as possible. There are many ways to measure the ‘‘concentration’’ of an image. Since applying a periodic translation of an image only changes its phase function, we can

choose to measure the concentration around a specific point, chosen to be $\mathbf{0}$, the center of the image. We can then formulate the problem of finding the “more concentrated” image in a variational way by looking for instance to the solution of

$$(\mathbf{P1}) : \max_{v \in \mathcal{M}_u} v(\mathbf{0}).$$

Another possible variational formulation of the concentration problem is

$$(\mathbf{P2}) : \min_{v \in \mathcal{M}_u} \sum_{\mathbf{x} \in D} A(\mathbf{x})v(\mathbf{x})^2, \quad (2.15)$$

where $A(\mathbf{x})$ is a penalty function that should increase as $|\mathbf{x}|$ increases. As we will see in details in Section 2.2.3, it turns out that for some particular A , the solution of $(\mathbf{P2})$ is simple and is the same as the one of $(\mathbf{P1})$: just take $\widehat{v}(\boldsymbol{\xi}) = |\widehat{u}(\boldsymbol{\xi})|$ for all $\boldsymbol{\xi} \neq \mathbf{0}$. In other terms, among all possible phase functions that define the different elements of \mathcal{M}_u , the null function is, as we shall see, particularly interesting. This is why we introduce the following definition for the texton.

Definition 2.3 (Texton of an image). *The texton of an image $u : D \rightarrow \mathbb{R}$ is the image $T(u) : D \rightarrow \mathbb{R}$ that has the same mean value as u , the same Fourier amplitude as u , and identically null phases outside $\mathbf{0}$. In others words, $T(u)$ is characterized in Fourier domain by*

$$\widehat{T(u)}(\mathbf{0}) = \widehat{u}(\mathbf{0}) \quad \text{and} \quad \forall \boldsymbol{\xi} \in D \setminus \{\mathbf{0}\}, \quad \widehat{T(u)}(\boldsymbol{\xi}) = |\widehat{u}(\boldsymbol{\xi})| \quad (2.16)$$

or, equivalently, in the spatial domain by

$$\forall \mathbf{x} \in D, \quad T(u)(\mathbf{x}) = \frac{1}{|D|} \widehat{u}(\mathbf{0}) + \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D, \boldsymbol{\xi} \neq \mathbf{0}} |\widehat{u}(\boldsymbol{\xi})| e^{2i\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle}. \quad (2.17)$$

The reason for which the coefficient $\widehat{T(u)}(\mathbf{0})$ has not exactly the same definition as the others is that we want the texton $T(u)$ to have the same mean value as u (which is equivalent to have the same Fourier coefficient in $\boldsymbol{\xi} = \mathbf{0}$). Thus, an equivalent definition would be to define first T by $\widehat{T(v)} = |\widehat{v}|$ for any zero-mean image v , and then extend it to any $u : D \rightarrow \mathbb{R}$ by $T(u) = T(u - m_u) + m_u$. Hence, since it is not a restriction, we shall consider later in several occasions only images that have a null mean value.

2.2.2 Elementary properties

Proposition 2.2 (Elementary properties). *The texton operator T has the following elementary properties:*

1. For any image u , $T(u)$ is a symmetric (and real-valued) image, that is, $T(u)(-\mathbf{x}) = T(u)(\mathbf{x})$ for all $\mathbf{x} \in D$.
2. $T(T(u)) = T(u)$, which means that $T(u)$ is its own texton.
3. The operator T is 1-Lipschitz for the L^2 -norm, that is: if u and v are images on D , then $\|T(u) - T(v)\|_2 \leq \|u - v\|_2$.

4. The texton is translation invariant: given any $\mathbf{y} \in \mathbb{Z}^2$, if $u_{\mathbf{y}}$ is the image defined on D by $u_{\mathbf{y}}(\mathbf{x}) = \hat{u}(\mathbf{x} - \mathbf{y})$, then $T(u_{\mathbf{y}}) = T(u)$.
5. The set of images with mean value 0 is stable under the operator T .
6. For any image u , and for any real numbers α and β , we have $T(u + \beta) = T(u) + \beta$ and $T(\alpha u) = |\alpha|T(u) + (\alpha - |\alpha|)m_u$. In particular when $\alpha \geq 0$, one has $T(\alpha u) = \alpha T(u)$.
7. For any images u and v , $T(u \star v) = T(u) \star T(v)$.

Proof. Since u is a real-valued image, we have $\widehat{u}(\boldsymbol{\xi}) = \widehat{u}(-\boldsymbol{\xi})^*$ for all $\boldsymbol{\xi} \in D$ (here the star denotes the complex conjugate). Thus $\widehat{T(u)}(\boldsymbol{\xi}) = |\widehat{u}(\boldsymbol{\xi})| = \widehat{T(u)}(-\boldsymbol{\xi}) = \widehat{T(u)}(\boldsymbol{\xi})^*$, which implies that $T(u)$ is real-valued and symmetric. The second property comes from the definition of the texton by Equation (2.16), because $\widehat{T(u)}(\boldsymbol{\xi}) = |\widehat{T(u)}(\boldsymbol{\xi})|$ for $\boldsymbol{\xi} \neq \mathbf{0}$ and $\widehat{T(u)}(\mathbf{0}) = \widehat{u}(\mathbf{0})$. Property 3 is a consequence of Parseval's Theorem and the second triangular inequality. Indeed,

$$\|T(u) - T(v)\|_2^2 = \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D} (|\widehat{u}(\boldsymbol{\xi})| - |\widehat{v}(\boldsymbol{\xi})|)^2 \leq \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D} |\widehat{u}(\boldsymbol{\xi}) - \widehat{v}(\boldsymbol{\xi})|^2 = \|u - v\|_2^2.$$

Property 4 is a consequence of the fact the translating the image is equivalent to shifting the phases of its Fourier transform while keeping the amplitudes unchanged (more precisely, we have $\widehat{u_{\mathbf{y}}}(\boldsymbol{\xi}) = e^{2i\pi\langle \mathbf{y}, \boldsymbol{\xi} \rangle} \widehat{u}(\boldsymbol{\xi})$ for any $\boldsymbol{\xi} \in D$). As for Property 5, it is a direct consequence of the equality $\widehat{u}(\mathbf{0}) = |D|m_u$. Property 6 is obtained by computing the Fourier coefficients of $v = \alpha u + \beta$. Indeed, since $\widehat{v} = \alpha \widehat{u} + |D|\beta \delta_0$, one has $\widehat{T(v)}(\boldsymbol{\xi}) = |\alpha| |\widehat{u}(\boldsymbol{\xi})|$ for $\boldsymbol{\xi} \neq \mathbf{0}$, and $\widehat{T(v)}(\mathbf{0}) = \alpha \widehat{u}(\mathbf{0}) + |D|\beta = |\alpha| \widehat{u}(\mathbf{0}) + |D|(\alpha - |\alpha|)m_u + |D|\beta$. Then, taking the inverse Fourier transform, we have the result. Property 7 is a consequence of the product/convolution property of the Fourier transform: for all $\boldsymbol{\xi} \in D$, we have $\widehat{u \star v}(\boldsymbol{\xi}) = \widehat{u}(\boldsymbol{\xi}) \widehat{v}(\boldsymbol{\xi})$, so that $\widehat{T(u \star v)}(\boldsymbol{\xi}) = \widehat{T(u)}(\boldsymbol{\xi}) \widehat{T(v)}(\boldsymbol{\xi})$ and thus $T(u \star v) = T(u) \star T(v)$. \square

2.2.3 Spatial concentration properties

In this section, we are interested in the spatial concentration properties of the texton. In particular, we prove that the texton solves the concentration properties formulated by **(P1)** and **(P2)**.

Proposition 2.3 (Property of spatial concentration). *For any image $u : D \rightarrow \mathbb{R}$, $T(u)$ is the unique solution of **(P1)**.*

Proof. As mentioned earlier, it is enough to prove the property for u with zero-mean ($m_u = 0$). Then, for any image $v \in \mathcal{M}_u$ we have

$$v(\mathbf{0}) = \text{Re}(v(\mathbf{0})) = \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D} \text{Re}(\widehat{v}(\boldsymbol{\xi})) \leq \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D} |\widehat{v}(\boldsymbol{\xi})| = \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D} |\widehat{u}(\boldsymbol{\xi})| = T(u)(\mathbf{0}).$$

Moreover, the equality holds if and only if $\text{Re}(\widehat{v}(\boldsymbol{\xi})) = |\widehat{v}(\boldsymbol{\xi})| = |\widehat{u}(\boldsymbol{\xi})|$ for all $\boldsymbol{\xi} \in D$, which means that $v = T(u)$. \square

Notice also that the texton $T(u)$ achieves its maximum value at the spatial origin, since

$$\forall \mathbf{x} \in D, \quad T(u)(\mathbf{x}) = \frac{1}{|D|} \hat{u}(\mathbf{0}) + \frac{2}{|D|} \sum_{\boldsymbol{\xi} \in D_+} |\hat{u}(\boldsymbol{\xi})| \cos(2\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle) \leq T(u)(\mathbf{0}).$$

Moreover, it can also be shown that the texton $T(u)$ is, among all images that have same Fourier modulus as u , the one that optimizes the continuous even-order partial derivatives at 0. To be more precise, we have that if we define, for any m, n integers,

$$\frac{\partial^{2m+2n} v}{\partial^{2m} x_1 \partial^{2n} x_2}(\mathbf{0}) := \left(\frac{2i\pi}{M}\right)^{2m} \left(\frac{2i\pi}{N}\right)^{2n} \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D} \xi_1^{2m} \xi_2^{2n} \hat{v}(\boldsymbol{\xi})$$

then

$$\max_{v \in \mathcal{M}_u} (-1)^{m+n} \frac{\partial^{2m+2n} v}{\partial^{2m} x_1 \partial^{2n} x_2}(\mathbf{0}) \text{ is achieved when } v = T(u).$$

Notice that all odd-order derivatives are equal to 0 at the spatial origin $\mathbf{0}$ of D because of the symmetry property of the texton.

We will be interested in investigating further decay properties of the texton. Now, it is well-known that the decay property of a function can be seen on the regularity of its Fourier transform. This is why we will here be interested in the regularity of the Fourier transform of the texton of an image. We will then, in a second step, analyze the consequences of this result on the decay of the texton itself.

We first show a general property of the texton in Fourier domain: it is optimally smooth in the sense of generalized gradient L^p norms.

Proposition 2.4 (Property of regularity in the Fourier domain). *Let $u : D \rightarrow \mathbb{R}$ be an image with a non-negative mean value. Let $\{\lambda_\omega\}_{\omega \in D_+}$ be a family of non-negative real numbers, and let $p > 0$. Then $T(u)$ is a solution of the following optimization problem:*

$$(\mathbf{P3}) : \min_{v \in \mathcal{M}_u} \sum_{\boldsymbol{\xi} \in D} \sum_{\boldsymbol{\omega} \in D_+} \lambda_\omega |\hat{v}(\boldsymbol{\xi} + \boldsymbol{\omega}) - \hat{v}(\boldsymbol{\xi})|^p. \quad (2.18)$$

An interesting particular case corresponds to the family $\{\lambda_\omega\}$ that has only two non-zero components $\lambda_{(0,1)} = \lambda_{(1,0)} = 1$. In this case, **(P3)** corresponds to the minimization of the L^p norm (for $p \geq 1$) of the usual discrete gradient of \hat{v} .

Proof. Given two non-negative numbers r_1 and r_2 , one clearly has

$$\min_{\theta_1, \theta_2 \in S^1} \left| r_1 e^{i\theta_1} - r_2 e^{i\theta_2} \right|^p = |r_1 - r_2|^p,$$

the minimum being attained if and only if $\theta_1 = \theta_2$. Consequently, for any $\boldsymbol{\xi} \in D$ and $\boldsymbol{\omega} \in D_+$, the minimum value of $|\hat{v}(\boldsymbol{\xi} + \boldsymbol{\omega}) - \hat{v}(\boldsymbol{\xi})|^p$ (over $v \in \mathcal{M}_u$) is achieved if and only if the complex numbers $\hat{v}(\boldsymbol{\xi} + \boldsymbol{\omega})$ and $\hat{v}(\boldsymbol{\xi})$ have the same argument. When $v = T(u)$, all values of \hat{v} are non-negative real numbers (argument 0), so this condition is fulfilled for all $\boldsymbol{\xi}, \boldsymbol{\omega}$. Since all λ_ω are non-negative, this proves that $v = T(u)$ is a solution of **(P3)**. \square

A remark on uniqueness. The question of the uniqueness of the solution of **(P3)** can be addressed by considering the graph associated to the family $\lambda = \{\lambda_\omega\}$ and to u as follows. Let $G(\hat{u}, \lambda)$ be the non-oriented graph whose vertices are the points $\boldsymbol{\xi} \in D$ such that $\hat{u}(\boldsymbol{\xi}) \neq 0$ (that is the support of \hat{u} , denoted by $\text{Supp}(\hat{u})$) and whose edges are the $(\boldsymbol{\xi}, \boldsymbol{\xi}') \in \text{Supp}(\hat{u})$ such that $\lambda_{\boldsymbol{\xi}-\boldsymbol{\xi}'} \neq 0$ or $\lambda_{\boldsymbol{\xi}'-\boldsymbol{\xi}} \neq 0$. Then the solution of **(P3)** is unique if and only if the graph $G(\hat{u}, \lambda)$ is connected. Otherwise, one can choose a constant phase equal to 0 or π (these are the only possible values to have real-valued images) independently on each connected component of the graph that do not contain $\mathbf{0}$. An important characterization of the texton is that for any image u , there exists a family λ such that the graph $G(\hat{u}, \lambda)$ is connected. Indeed, any family λ with support the whole set D works.

Proposition 2.4 can be transposed in the spatial domain to get another result of spatial concentration for the texton.

Proposition 2.5 (Property of spatial concentration). *Let $u : D \rightarrow \mathbb{R}$ be an image with a non-negative mean value. Let $A : D \rightarrow \mathbb{R}$ be a symmetric weight image such that $A(\mathbf{0}) = 0$ and*

$$\forall \boldsymbol{\xi} \in D \setminus \{\mathbf{0}\}, \quad \hat{A}(\boldsymbol{\xi}) \leq 0. \quad (2.19)$$

*Then, A is non-negative and $T(u)$ is a solution of the optimization problem **(P2)** associated to A (Equation (2.15)).*

Proof. Since A is real and symmetric, so is \hat{A} and the Fourier Inversion Theorem yields

$$\forall \mathbf{x} \in D, \quad A(\mathbf{x}) = \frac{1}{|D|} \hat{A}(\mathbf{0}) + \frac{2}{|D|} \sum_{\boldsymbol{\omega} \in D_+} \hat{A}(\boldsymbol{\omega}) \cos(2\pi \langle \mathbf{x}, \boldsymbol{\omega} \rangle). \quad (2.20)$$

Subtracting the same equality applied to $\mathbf{x} = \mathbf{0}$, we get, since $A(\mathbf{0}) = 0$,

$$\forall \mathbf{x} \in D, \quad A(\mathbf{x}) = \frac{2}{|D|} \sum_{\boldsymbol{\omega} \in D_+} (-\hat{A}(\boldsymbol{\omega})) \cdot (1 - \cos(2\pi \langle \mathbf{x}, \boldsymbol{\omega} \rangle)), \quad (2.21)$$

so that $A(\mathbf{x}) \geq 0$ as all the terms of the sum are non-negative. Moreover, if we consider the family of non-negative real numbers $\{\lambda_\omega\}_{\omega \in D_+}$ defined by $\lambda_\omega = -\hat{A}(\omega)/|D|^2$, we can rewrite (2.21) as

$$\forall \mathbf{x} \in D, \quad A(\mathbf{x}) = |D| \sum_{\boldsymbol{\omega} \in D_+} \lambda_\omega |e^{-2i\pi \langle \mathbf{x}, \boldsymbol{\omega} \rangle} - 1|^2. \quad (2.22)$$

Now Parseval's Theorem yields

$$\sum_{\mathbf{x} \in D} A(\mathbf{x}) v(\mathbf{x})^2 = \sum_{\boldsymbol{\omega} \in D_+} \lambda_\omega |D| \sum_{\mathbf{x} \in D} |(e^{-2i\pi \langle \mathbf{x}, \boldsymbol{\omega} \rangle} - 1)v(\mathbf{x})|^2 = \sum_{\boldsymbol{\omega} \in D_+} \lambda_\omega \sum_{\boldsymbol{\xi} \in D} |\hat{v}(\boldsymbol{\xi} + \boldsymbol{\omega}) - \hat{v}(\boldsymbol{\xi})|^2, \quad (2.23)$$

which proves that the energies to be minimized in **(P2)** and **(P3)** (in the case $p = 2$) are the same. We conclude, thanks to Proposition 2.4, that $v = T(u)$ is a solution of **(P2)**. \square

For example the function

$$A(x_1, x_2) = \sin^2\left(\pi \frac{x_1}{M}\right) + \sin^2\left(\pi \frac{x_2}{N}\right), \quad (2.24)$$

which corresponds to $\lambda_{(1,0)} = \lambda_{(0,1)} = (4|D|)^{-1}$ and $\lambda_{\omega} = 0$ otherwise, is an admissible function. Since it satisfies $A(\mathbf{0}) = 0$ and increases as x_1 and x_2 move away from 0, the variational formulation **(P2)** is a good characterization of the concentration (or decay) of the texton.

We can also ask about the regularity of the texton. Now, since the texton has no special property of decay in the Fourier domain (the amplitudes are exactly the same as the original image), one cannot hope for more spatial regularity. And indeed, we have the following result.

Proposition 2.6. *For any image $v \in \mathcal{M}_u$, one has $\|v\|_2 = \|u\|_2$ and $\|\nabla v\|_2 = \|\nabla u\|_2$. In particular, $\|T(u)\|_2 = \|u\|_2$ and $\|\nabla T(u)\|_2 = \|\nabla u\|_2$.*

Proof. Let $v \in \mathcal{M}_u$. Because of Parseval's Theorem, it is straightforward that $\|v\|_2 = \|u\|_2$. For the gradient, thanks again to Parseval's theorem, we can write

$$\|\nabla v\|_2^2 = \sum_{\boldsymbol{\xi}=(\xi_1,\xi_2)\in D} |\widehat{v}(\boldsymbol{\xi})|^2 \left(\sin^2\left(\frac{\pi\xi_1}{M}\right) + \sin^2\left(\frac{\pi\xi_2}{N}\right) \right). \quad (2.25)$$

And since $|\widehat{v}(\boldsymbol{\xi})| = |\widehat{u}(\boldsymbol{\xi})|$ for all $\boldsymbol{\xi}$, we have the result. \square

Notice that the proof is valid for any (generalized) gradient form. Hence, with respect to this particular measure of the regularity of an image, the texton $T(u)$ is not smoother than the original image u , or than any phase-shifted version of u . This is of particular interest as regards the RPN algorithm: all realization of the RPN model associated to a given image u have the same gradient norm as u , which means that these RPN textures are just as smooth as the original image (with respect of this particular regularity measure). This can be extended to the ADSN algorithm, provided that we consider the expectation of the norm of the gradient.

Proposition 2.7. *Let $u : D \rightarrow \mathbb{R}$ be an image. If U is a random image distributed according to the law $\text{RPN}(u)$, then*

$$\|\nabla U\|_2^2 = \|\nabla u\|_2^2. \quad (2.26)$$

If U is a random image distributed according to the law $\text{ADSN}(u)$, then $\|\nabla U\|_2^2$ is a (non-constant) random variable that satisfies

$$\mathbb{E}(\|\nabla U\|_2^2) = \|\nabla u\|_2^2. \quad (2.27)$$

Proof. Equation (2.26) directly follows from Proposition 2.6 above. In the case of the ADSN model, Equation (2.27) is a direct consequence of (2.13) and (2.25) applied to $v = U$. \square

Another way to see the spatial concentration properties of texton images is to make the link between the texton and the empirical covariance of the image. This is the purpose of the following section.

2.2.4 Texton and empirical covariance

A simple texture analysis task consists in computing its second-order statistics, namely its empirical variance and its empirical covariance function. Since these quantities are unchanged when adding a constant to the image, we will, in the following, consider images that have a null mean value ($m_u = 0$). Let us start by establishing the link between the texton and the empirical covariance of an image.

Proposition 2.8. *Let $u : D \rightarrow \mathbb{R}$ be an image with mean value 0, and let the empirical covariance of u , denoted C_u , be defined by*

$$\forall \mathbf{y} \in D, \quad C_u(\mathbf{y}) = \frac{1}{|D|} \sum_{\mathbf{x} \in D} u(\mathbf{x})u(\mathbf{x} - \mathbf{y}). \quad (2.28)$$

Then $\widehat{C}_u = \frac{1}{|D|}|\widehat{u}|^2$, or, equivalently,

$$C_u = \frac{1}{|D|}T(u) \star T(u), \quad (2.29)$$

which means that C_u is, up to a constant, the auto-convolution of the texton $T(u)$.

Proof. The proof is a direct consequence of the convolution/product property of the Fourier transform. Indeed, if u^- denotes the symmetric image of u given by $u^-(\mathbf{y}) = u(-\mathbf{y})$ then $C_u = \frac{1}{|D|}u \star u^-$. Consequently, for all $\boldsymbol{\xi} \in D$, we have

$$\widehat{C}_u(\boldsymbol{\xi}) = \frac{1}{|D|}\widehat{u}(\boldsymbol{\xi})\widehat{u}^-(\boldsymbol{\xi}) = \frac{1}{|D|}\widehat{u}(\boldsymbol{\xi})\widehat{u}(\boldsymbol{\xi})^* = \frac{1}{|D|}|\widehat{u}(\boldsymbol{\xi})|^2.$$

Then, taking the inverse Fourier transform and using the symmetry property of the texton, we obtain $C_u = \frac{1}{|D|}T(u) \star T(u)$ as announced. \square

A consequence of (2.29) is that the texton is generally less blurry than the covariance, which, as we shall see later, can be easily noticed on experiments. It is also generally more “concentrated” (intuitively, the area of the “support” of the covariance is four times the one of the texton). Another advantage of the texton (compared to the empirical covariance function) is that it belongs to the space of images: its values are in the same unit as the grey levels of the image, whereas the covariance is rather seen in the units of an “energy”. A second consequence of Proposition 2.8 is that the empirical covariance is exactly preserved by the realizations of the RPN model, and preserved in expectation for the ADSN model. Let us display an example of a texton based on a natural texture image taken from the website http://www.lemog.fr/lemog_textures/.

We can see in Figure 2.1 the original texture image, a RPN sample from it, its empirical covariance, its texton and we measure the concentration of the texton by plotting the proportion of energy outside a disc of radius r centered at the origin as a function of the percentage of pixels inside this disc. More precisely, for r between 0 and $\max(M, N)$, let $Disc_r$ denote the discrete disc of radius r centered at $\mathbf{0}$. Then we plot

$$\frac{\sum_{\mathbf{x} \in D \setminus Disc_r} (u(\mathbf{x}) - m_u)^2}{\sum_{\mathbf{x} \in D} (u(\mathbf{x}) - m_u)^2} \text{ as a function of } 100 \times \frac{|Disc_r|}{|D|}, \quad (2.30)$$

for both the texture image and its texton. For the texture image, this is an almost linear function, since the energy in a domain is, up to statistical fluctuations, proportional to the size of the domain. Now, for the texton, this curve decreases very fast in around 0 because of the concentration of the texton. See also Figure 2.4 for more examples.

Notice that in order to avoid boundary effects (all the definitions and properties assume the periodicity of the texture image), we first process any sample image u by replacing it by its periodic component obtained from its periodic+smooth decomposition [99]. As remarked in [51], this does not affect the geometrical and statistical properties of the texture image and prevents undesirable artefacts that could be caused by the non-periodic nature of the sample image used to define the RPN or ADSN texture models.

2.2.5 Distance between textures

To compare two texture models, one can use the Fréchet distance [49], as was done in [132]. Let us recall this definition.

Definition 2.4. *Let \mathcal{P} and \mathcal{Q} be two probability laws with values in the same space V . Assume that both \mathcal{P} and \mathcal{Q} are the of laws of L^2 random variables. The Fréchet distance between \mathcal{P} and \mathcal{Q} is defined by*

$$d(\mathcal{P}, \mathcal{Q})^2 = \min_{\substack{(X,Y) \\ X \sim \mathcal{P} \\ Y \sim \mathcal{Q}}} \mathbb{E} (\|X - Y\|_2^2).$$

The distance between two Gaussian models can be very simply expressed in function of their respective covariances [39], but we shall see now that an even simpler expression can be obtained by using their respective textons. The result is also extended to the RPN model.

Theorem 2.2. *Let u_0 and u_1 be two images with null mean value. Then, the Fréchet distance, between the two Gaussian distributions $\text{ADSN}(u_0)$ and $\text{ADSN}(u_1)$ is given by the L^2 distance between their respective textons:*

$$d(\text{ADSN}(u_0), \text{ADSN}(u_1))^2 = \min_{\substack{U_0 \sim \text{ADSN}(u_0) \\ U_1 \sim \text{ADSN}(u_1)}} \mathbb{E} (\|U_0 - U_1\|_2^2) = \|T(u_0) - T(u_1)\|_2^2.$$

The same result holds also for the Fréchet distance between the two RPN distributions $\text{RPN}(u_0)$ and $\text{RPN}(u_1)$:

$$d(\text{RPN}(u_0), \text{RPN}(u_1))^2 = \min_{\substack{U_0 \sim \text{RPN}(u_0) \\ U_1 \sim \text{RPN}(u_1)}} \mathbb{E} (\|U_0 - U_1\|_2^2) = \|T(u_0) - T(u_1)\|_2^2.$$

Proof. The distribution $\text{ADSN}(u_j)$ (for $j = 0$ or 1) is a centered multivariate normal distribution on \mathbb{R}^D with covariance matrix Γ_j given by $\Gamma_j(\mathbf{x}, \mathbf{y}) = C_{u_j}(\mathbf{x} - \mathbf{y})$. Then according to [132] and [39], the Fréchet distance between $\text{ADSN}(u_0)$ and $\text{ADSN}(u_1)$ is given by

$$d(\text{ADSN}(u_0), \text{ADSN}(u_1))^2 = \text{tr} \left(\Gamma_0 + \Gamma_1 - 2(\Gamma_1^{1/2} \Gamma_0 \Gamma_1^{1/2})^{1/2} \right), \quad (2.31)$$

where $A^{1/2}$ is the unique symmetric positive semi-definite square root of a symmetric positive semi-definite matrix A .

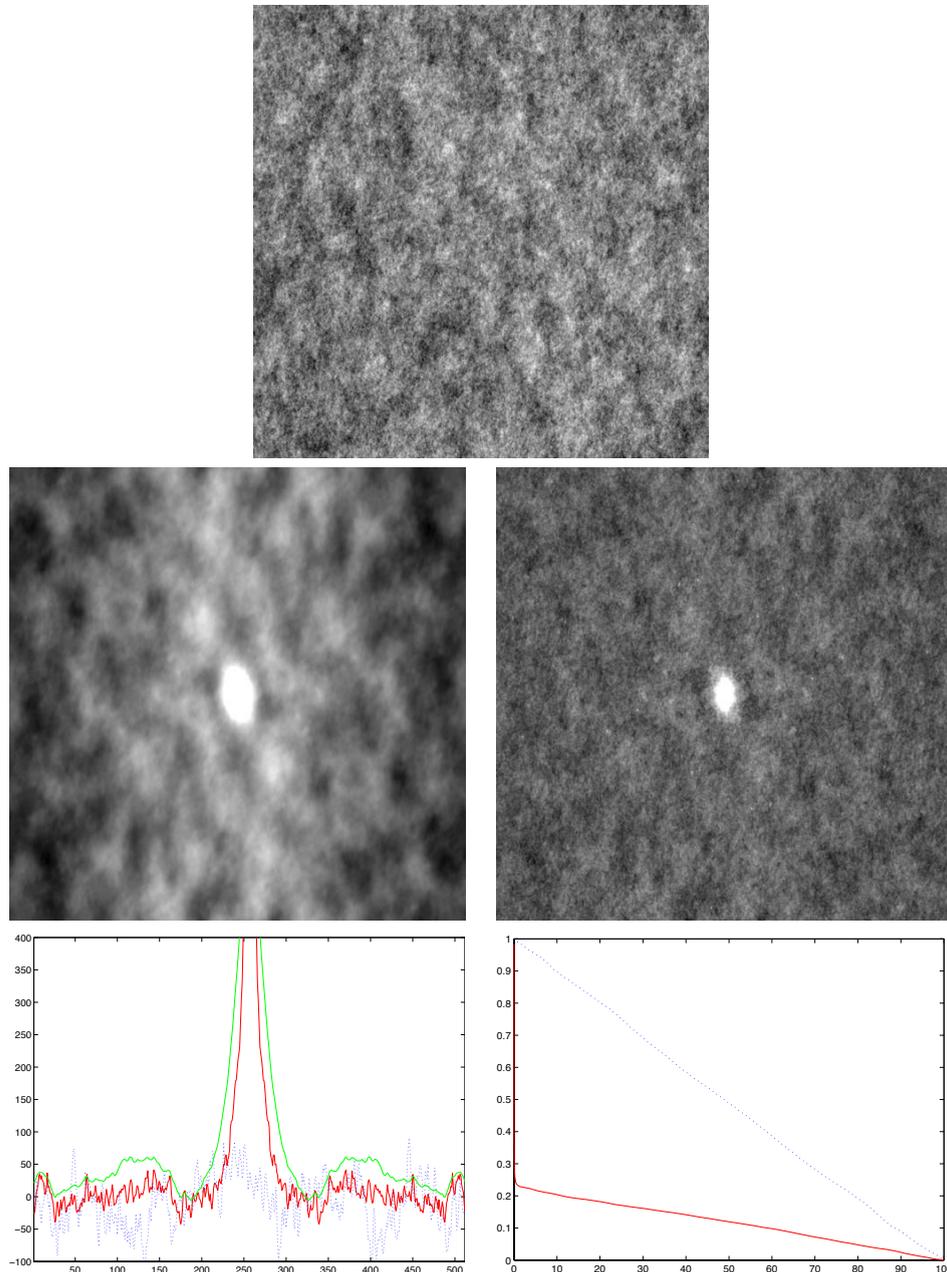


Figure 2.1: Top: the original stone wall texture image u . Second line: on the left, the empirical covariance of u ; on the right, the texton $T(u)$. The visualization of the texton is not easy since its values around the spatial origin $\mathbf{0}$ are generally much larger than the usual grey level values between 0 and 255. To visualize it here, we have simply saturated the values above 255. One can clearly observe that the texton is less blurry and more concentrated than the covariance. This is confirmed by the next graphic. Third line: on the left we show a 1D slice through the origin of the texton (in red), the empirical covariance (in green) and of the original texture u (dotted line). On the right, we give a measure of the concentration of the texton by plotting the proportion of energy outside a disc of radius r centered at the origin as a function of the percentage of pixels inside this disc. For the original texture this gives the linear dotted curve (the energy in the disc is proportional to its area), whereas for the texton (red curve), we see that a disc of radius approximately 50 (that thus contains 10% of the pixels) already captures 80% of the total energy.

Now since the covariance matrices Γ_0 and Γ_1 are circulant, they are diagonalized in the Fourier basis. More precisely, let F denote the $|D| \times |D|$ normalized Fourier Transform matrix, whose coefficients are $F(\boldsymbol{\xi}, \mathbf{x}) = \frac{1}{\sqrt{|D|}} e^{-2i\pi(\mathbf{x}, \boldsymbol{\xi})}$, then F is a unitary matrix and one has $\hat{u}(\boldsymbol{\xi}) = \sqrt{|D|}(Fu)(\boldsymbol{\xi})$ for any $u \in \mathbb{R}^D$. Now, a simple computation shows that for all $\boldsymbol{\xi}, \boldsymbol{\xi}' \in D$,

$$(F\Gamma_j F^*)(\boldsymbol{\xi}, \boldsymbol{\xi}') = \sum_{\mathbf{z}, \mathbf{t} \in D} F(\boldsymbol{\xi}, \mathbf{z}) C_{u_j}(\mathbf{z} - \mathbf{t}) F^*(\mathbf{t}, \boldsymbol{\xi}') = \widehat{C}_{u_j}(\boldsymbol{\xi}) \delta_{\boldsymbol{\xi}=\boldsymbol{\xi}'} = |\widehat{u}_j(\boldsymbol{\xi})|^2 \delta_{\boldsymbol{\xi}=\boldsymbol{\xi}'}$$

Then, using the fact that $\text{tr}(A) = \text{tr}(FAF^*)$ and $\text{tr}(A^{1/2}) = \text{tr}((FAF^*)^{1/2})$ for any symmetric positive semi-definite matrix A , we obtain

$$\begin{aligned} \text{tr}\left(\Gamma_0 + \Gamma_1 - 2(\Gamma_1^{1/2}\Gamma_0\Gamma_1^{1/2})^{1/2}\right) &= \text{tr}\left(F\Gamma_0F^* + F\Gamma_1F^* - 2F(\Gamma_1^{1/2}\Gamma_0\Gamma_1^{1/2})^{1/2}F^*\right) \\ &= \sum_{\boldsymbol{\xi} \in D} (|\widehat{u}_0(\boldsymbol{\xi})|^2 + |\widehat{u}_1(\boldsymbol{\xi})|^2 - 2|\widehat{u}_0(\boldsymbol{\xi})||\widehat{u}_1(\boldsymbol{\xi})|) \\ &= \sum_{\boldsymbol{\xi} \in D} (|\widehat{u}_0(\boldsymbol{\xi})| - |\widehat{u}_1(\boldsymbol{\xi})|)^2 = \|T(u_0) - T(u_1)\|_2^2. \end{aligned}$$

For the Fréchet distance between two RPN distributions, we first notice that if $U_0 \sim \text{RPN}(u_0)$ and $U_1 \sim \text{RPN}(u_1)$, then if Φ_0 and Φ_1 are uniform distributions on $[0, 2\pi)$ then

$$\|U_0 - U_1\|_2^2 = \sum_{\boldsymbol{\xi}} \left| |\widehat{u}_0(\boldsymbol{\xi})| e^{i\Phi_0(\boldsymbol{\xi})} - |\widehat{u}_1(\boldsymbol{\xi})| e^{i\Phi_1(\boldsymbol{\xi})} \right|^2 \geq \sum_{\boldsymbol{\xi}} (|\widehat{u}_0(\boldsymbol{\xi})| - |\widehat{u}_1(\boldsymbol{\xi})|)^2$$

and the equality holds when $\Phi_0 = \Phi_1$, which ends the proof. \square

2.2.6 Textons of some micro-textures

In this section we give several examples of textons. These examples are either theoretical examples, synthetic examples or examples from natural texture images. Let us start with theoretical examples.

Theoretical examples

Let u be a zero-mean image, we are interested in the textons of realizations of the RPN and ADSN models associated to u . Are they equal or at least “close” to the texton of u ? For the RPN model ($U \sim \text{RPN}(u)$), since Fourier amplitudes are exactly preserved, we have $T(U) = T(u)$, that is, the texton remains unchanged. For the ADSN model, $U \sim \text{ADSN}(u)$ means that $U = u \star W$, where W is a Gaussian white noise, and as a consequence (see last property of Proposition 2.2) we have $T(U) = T(u) \star T(W)$. We thus need to compute the law of $T(W)$, which is the aim of the following

Proposition 2.9. *Let W be a white noise image defined on D , with variance $\frac{1}{|D|}$ (that is, the $W(\mathbf{x})$ are i.i.d. and $\mathcal{N}(0, \frac{1}{|D|})$ distributed). Then $T(W)$ is a random image and its moments are*

$$\mathbb{E}(T(W)(\mathbf{0})) = \frac{\sqrt{\pi}}{2} \left(1 - \frac{1}{|D|}\right) \quad \text{and} \quad \text{Var}(T(W)(\mathbf{0})) = \frac{1}{|D|^2} + \frac{4 - \pi}{2|D|} \left(1 - \frac{1}{|D|}\right), \quad (2.32)$$

$$\forall \mathbf{x} \in D \setminus \{\mathbf{0}\}, \quad \mathbb{E}(T(W)(\mathbf{x})) = -\frac{\sqrt{\pi}}{2|D|} \quad \text{and} \quad \text{Var}(T(W)(\mathbf{x})) = \frac{1}{|D|^2} + \frac{4-\pi}{4|D|} \left(1 - \frac{2}{|D|}\right), \quad (2.33)$$

$$\forall \mathbf{x}, \mathbf{y} \in D, \quad \mathbf{x} \neq \mathbf{y}, \mathbf{x} \neq -\mathbf{y}, \quad \text{Cov}(T(W)(\mathbf{x}), T(W)(\mathbf{y})) = \frac{\pi-2}{2|D|^2}. \quad (2.34)$$

Proof. We can write, thanks to the inverse Fourier transform,

$$\forall \mathbf{x} \in D, \quad T(W)(\mathbf{x}) = \frac{\widehat{W}(\mathbf{0})}{|D|} + \frac{2}{|D|} \sum_{\boldsymbol{\xi} \in D_+} |\widehat{W}(\boldsymbol{\xi})| \cos(2\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle). \quad (2.35)$$

Now the random variables $\{\widehat{W}(\mathbf{0}); |\widehat{W}(\boldsymbol{\xi})|, \boldsymbol{\xi} \in D_+\}$ are independent, and $\widehat{W}(\mathbf{0}) \sim \mathcal{N}(0, 1)$. In addition, for any $\boldsymbol{\xi} \in D_+$, $|\widehat{W}(\boldsymbol{\xi})|$ follows a Rayleigh distribution with parameter $\sigma = \frac{1}{\sqrt{2}}$, so that

$$\mathbb{E}(|\widehat{W}(\boldsymbol{\xi})|) = \frac{\sqrt{\pi}}{2} \quad \text{and} \quad \text{Var}(|\widehat{W}(\boldsymbol{\xi})|) = \frac{4-\pi}{4}. \quad (2.36)$$

Recalling that $|D_+| = (|D| - 1)/2$, we easily obtain from (2.35) the two formulas of (2.32). Besides, for any $\mathbf{x} \in D \setminus \{\mathbf{0}\}$, one has $\sum_{\boldsymbol{\xi} \in D_+} \cos(2\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle) = -1/2$ (just write the Fourier representation of the Dirac image $\delta_{\mathbf{0}}$), while

$$\sum_{\boldsymbol{\xi} \in D_+} \cos^2(2\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle) = \frac{|D_+|}{2} + \frac{1}{2} \sum_{\boldsymbol{\xi} \in D_+} \cos(2\pi \langle 2\mathbf{x}, \boldsymbol{\xi} \rangle) = \frac{|D| - 2}{4}. \quad (2.37)$$

By using these relations in conjunction with (2.36) in (2.35), we easily obtain the two formulas of (2.32). Now let us consider $\mathbf{x}, \mathbf{y} \in D$ with $\mathbf{x} \neq \mathbf{y}$ and $\mathbf{x} \neq -\mathbf{y}$. The bilinearity of the covariance implies, thanks to (2.35) and the fact that the random variables $\{|\widehat{W}(\boldsymbol{\xi})|\}_{\boldsymbol{\xi} \in D_+ \cup \{\mathbf{0}\}}$ are independent, that

$$\begin{aligned} \text{Cov}(T(W)(\mathbf{x}), T(W)(\mathbf{y})) &= \frac{1}{|D|^2} \text{Var}(\widehat{W}(\mathbf{0})) + \frac{4}{|D|^2} \sum_{\boldsymbol{\xi} \in D_+} \text{Var}(|\widehat{W}(\boldsymbol{\xi})|) \cos(2\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle) \cos(2\pi \langle \mathbf{y}, \boldsymbol{\xi} \rangle) \\ &= \frac{1}{|D|^2} + \frac{4-\pi}{|D|^2} \cdot \frac{1}{2} \sum_{\boldsymbol{\xi} \in D_+} (\cos(2\pi \langle \mathbf{x} + \mathbf{y}, \boldsymbol{\xi} \rangle) + \cos(2\pi \langle \mathbf{x} - \mathbf{y}, \boldsymbol{\xi} \rangle)) \\ &= \frac{1}{|D|^2} - \frac{4-\pi}{2|D|^2} = \frac{\pi-2}{2|D|^2}. \end{aligned}$$

□

Proposition 2.9 shows that as the size of D increases, the first moments of $T(W)$ (up to order 2) converge to those of a weighted Dirac mass located in $\mathbf{0}$, that is, $\frac{\sqrt{\pi}}{2} \delta_{\mathbf{0}}$. As a consequence, when $U \sim \text{ADSN}(u)$ and $|D|$ is large, one expects to observe that

$$T(U) \approx \frac{\sqrt{\pi}}{2} T(u).$$

Remark: As precised in the introduction, we here assumed that the dimensions M and N of the rectangular domain D were odd numbers. When M and/or N are even, the formulas that appear in Proposition 2.9 for the moments of $T(W)$ are slightly changed, but the asymptotic behavior of the first moments of $T(W)$ remains the same.

Synthetic examples

A first simple way to produce a synthetic texton is to specify some symmetric positive Fourier coefficients and take the inverse Fourier transform. For instance, one can define the texton T_α of a pink noise by specifying its Fourier coefficients:

$$\forall \boldsymbol{\xi} \in D \setminus \{\mathbf{0}\}, \widehat{T}_\alpha(\boldsymbol{\xi}) = \frac{1}{\|\boldsymbol{\xi}\|^\alpha} \text{ and } \widehat{T}_\alpha(\mathbf{0}) = 0.$$

An example of the texton T_α thus defined is shown on the first column of Figure 2.2, with $\alpha = 1.8$ and $M = N = 256$.

A second way to synthesize textons is to take the auto-convolution of an image. This will indeed define a texton since the Fourier transform of such an image is the squared modulus of the Fourier transform of the original image. For instance, if we start with a centered binary disc we will obtain what we will call a cone texton. Or if we start with a centered binary indicator function of a snake-like shape, we will obtain what we will call a snake texton. These two examples are illustrated in Figure 2.2.

A third way to build a texton is to start from an image u , and then take the inverse Fourier transform of its Fourier amplitude (this is exactly the definition of the texton). Some examples of this are given in Figure 2.3. For each example, we show the original image u – which is here the indicator function of respectively a square, two discs, a T-shaped polygon – the texton $T(u)$ (which, in these cases, is not equal to u , because the Fourier transform of u is not positive), and some samples of $\text{RPN}(u)$. The geometry of the texton of a simple indicator function as the ones of Figure 2.3 is already quite complex.

Remark: The examples above raise a natural question: Can an indicator function be its own texton? The answer is yes, but all solutions are very peculiar. More precisely, a solution has to be the indicator function of a sublattice (that is, a discrete additive subgroup) of D , and its Fourier transform has to be the indicator function of the dual lattice. In other words, it can be proved that if $u = \mathbf{1}_R$ with $R \subset D$ is such that $T(u) = u$ (which is equivalent to say that $\widehat{u}(\boldsymbol{\xi})$ is real and positive for all $\boldsymbol{\xi} \in D$), then $\widehat{u} = |R| \mathbf{1}_{R'}$ with $|R| \times |R'| = |D|$. Moreover: (a) for all \mathbf{x} and \mathbf{y} in R , one has $\mathbf{x} - \mathbf{y} \in R$ (which means that R is a lattice in D); (b) R' is also a lattice of D ; and (c) they are dual, in the sense that $\forall \mathbf{x} \in R, \forall \boldsymbol{\xi} \in R', \langle \mathbf{x}, \boldsymbol{\xi} \rangle = 0$. The simplest example is obtained when u is the indicator function of D or of $\{\mathbf{0}\}$. But there are other non degenerate lattices, for instance $R = I_M \times \{0\}$, which leads to $R' = \{0\} \times I_N$ – or, assuming that m divides M and n divides N , $R = \{km\}_{k \in I_{M/m}} \times \{kn\}_{k \in I_{N/n}}$, which leads to $R' = \{kM/m\}_{k \in I_m} \times \{kN/n\}_{k \in I_n}$.

Textons of some natural textures

We end this section by showing in Figure 2.4 the texton of several natural texture images taken from the website http://www.lemog.fr/lemog_textures/.

For each texture, we show the texton and we illustrate the concentration of the texton by plotting the relative energy in a disc centered at the origin as a function of the percentage of pixels in it (as defined in (2.30)). As expected, the texton shows a high amount of concentration, but important variations can be observed depending on the considered texture. For

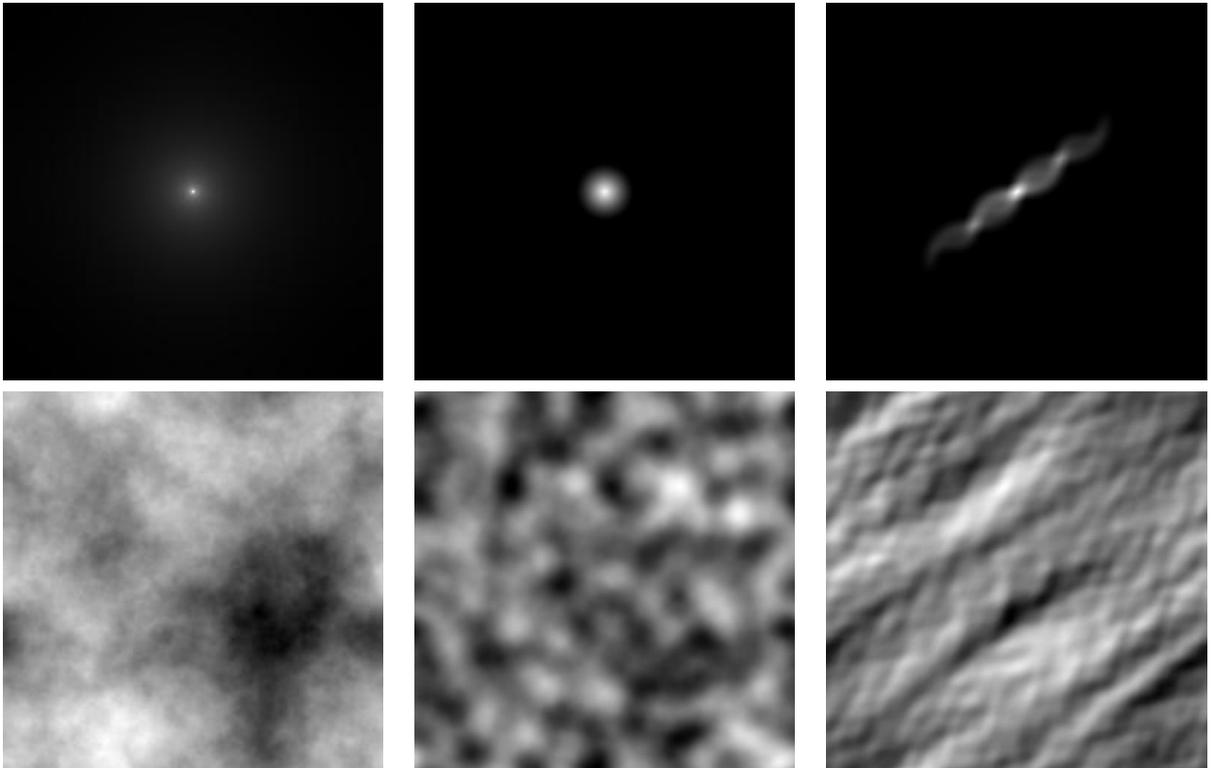


Figure 2.2: Synthetic examples of textons. On the first line we show three textons, and on the second line we show a RPN realization for each of them. From left to right, the three textons are: a pink noise (power law with exponent $\alpha = 1.8$) texton given by $\widehat{T}_\alpha(\boldsymbol{\xi}) = 1/\|\boldsymbol{\xi}\|^\alpha$ for $\boldsymbol{\xi} \neq \mathbf{0}$ and with $M = N = 256$; a 512×512 “cone” texton (obtained as the auto-convolution of the indicator function of a centered disc of radius 20); a 512×512 “snake” texton (obtained as the auto-convolution of the indicator function of a snake-like shape).

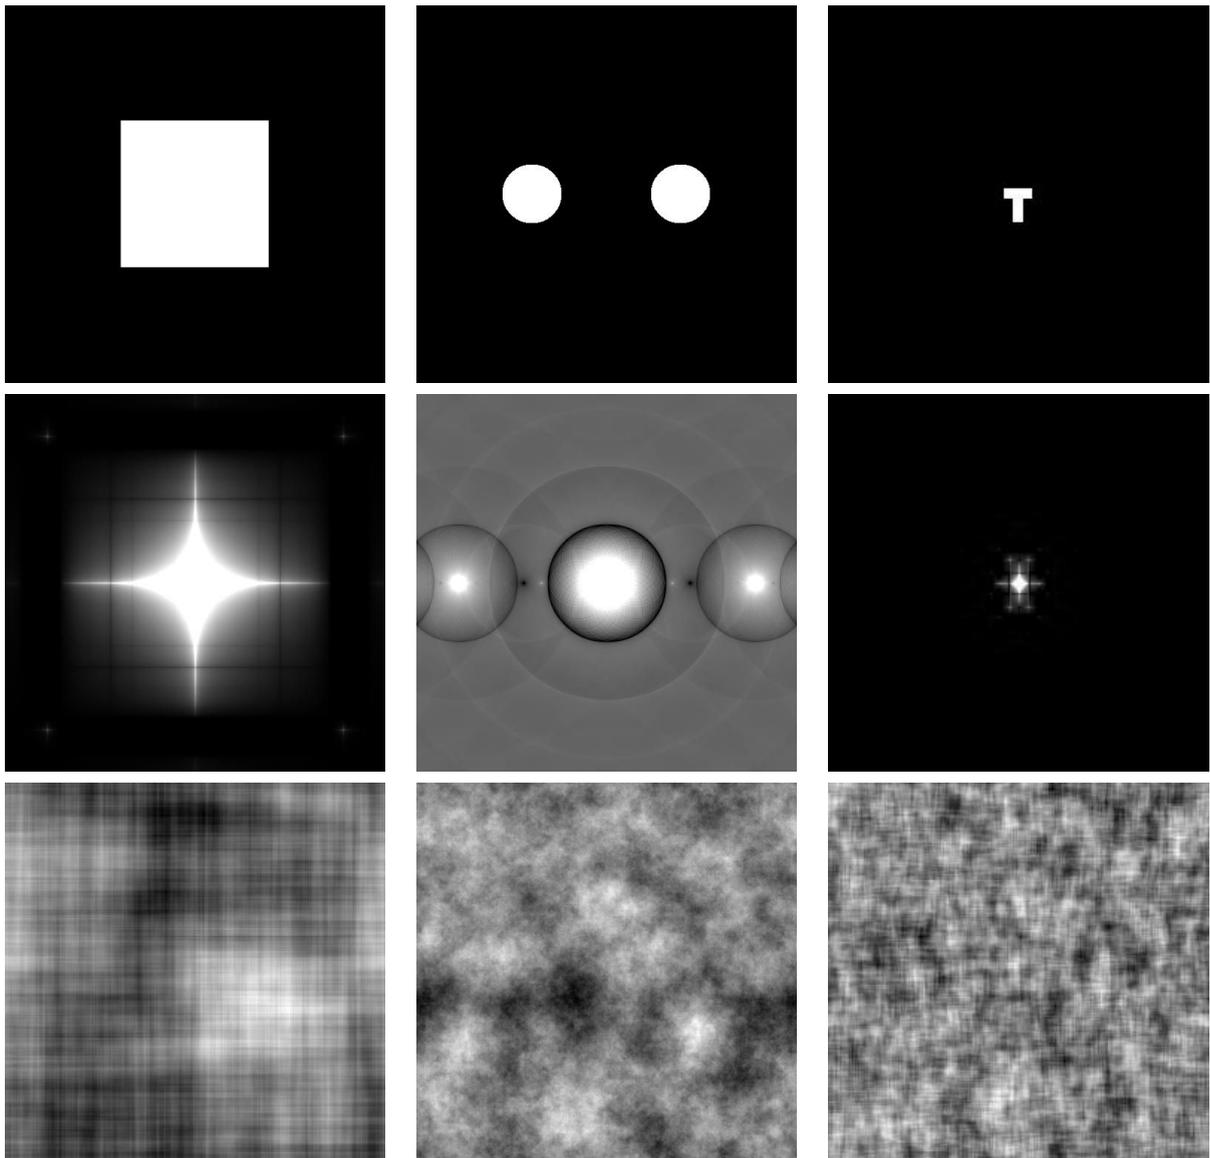


Figure 2.3: First row: we show the original image u which is here the indicator function of respectively a square (of side size 200 in a 512×512 domain, two discs of radius 40 in a 512×512 domain and a T shape in a 256×256 domain. Second row: the texton $T(u)$. For a better display, the grey level values that exceed 255 are saturated. Third row: samples from $RPN(u)$.

instance, the texture displayed on the right of Figure 2.4 contains a periodic pattern, and its concentration curve is significantly above the one of the fur texture (displayed on the left).

2.3 Sparse representations

The notion of texton we have introduced lets us represent a RPN or a Gaussian texture model by an image which is concentrated around the origin. However, we would like in this section to go a step further and consider representations that are not only concentrated, but have a “small” support. There are several motivations for such a kind of sparse representation. One is the possibility of representing a texture with a small number of coefficients, which may be useful both for texture synthesis and texture analysis. Another motivation is, as we shall see now, the possibility of conveniently defining textures on arbitrary large domains.

2.3.1 Extending Gaussian textures to \mathbb{Z}^2

Given a zero-mean Gaussian texture model defined by an image $v : D \rightarrow \mathbb{R}$ (which has not necessarily mean 0), how to produce Gaussian texture samples defined on a larger domain? In other terms, how to extend to a larger domain the Gaussian texture model? The relation between the Gaussian model and the shot noise model (see *e.g.* [51]) is useful to address this issue. As we saw in (2.9), the Gaussian model of D -periodic textures can be written

$$\forall \mathbf{y} \in \mathbb{Z}^2, \quad U(\mathbf{y}) = v \star W_D(\mathbf{y}) = \sum_{\mathbf{x} \in D} v(\mathbf{x}) W_D(\mathbf{y} - \mathbf{x}), \quad (2.38)$$

where $(W_D(\mathbf{x}))_{\mathbf{x} \in D}$ is a white Gaussian noise with marginal distribution $\mathcal{N}(0, |D|^{-1})$. Note that the symbol \star in (2.38) corresponds to a periodic convolution, which means that the right-hand term (here W_D) is extended to \mathbb{Z}^2 as a D -periodic image. This periodicity convention yields a nice relationship with the Discrete Fourier Transform and with the RPN model. Its consequence is that the random image U is also periodic, which means that it can be nicely tiled on the plane without creating spurious boundaries. This property is not necessarily desirable, since real-world textures are not periodic, and tiling might not be a satisfactory way to produce large texture samples. Hence, it is natural to consider the non-periodic variant of (2.38), which defines a stationary Gaussian texture on the whole plane \mathbb{Z}^2 by

$$\forall \mathbf{y} \in \mathbb{Z}^2, \quad U(\mathbf{y}) = v \star W_{\mathbb{Z}^2}(\mathbf{y}) = \sum_{\mathbf{x} \in D} v(\mathbf{x}) W_{\mathbb{Z}^2}(\mathbf{y} - \mathbf{x}), \quad (2.39)$$

where $(W_{\mathbb{Z}^2}(\mathbf{x}))_{\mathbf{x} \in \mathbb{Z}^2}$ is a set of i.i.d. random variables distributed as $\mathcal{N}(0, |D|^{-1})$. The difference between models (2.38) and (2.39) can be analyzed through the covariance function of both Gaussian fields. In the periodic case corresponding to (2.38), the covariance is the function Γ_v defined in (2.11). In the non-periodic case (0-padding of v) corresponding to (2.39), the covariance of U writes

$$\forall \mathbf{z} \in \mathbb{Z}^2, \quad \Gamma_v^0(\mathbf{z}) = \frac{1}{|D|} \sum_{\substack{\mathbf{x} \in D, \mathbf{y} \in D \\ \mathbf{y} - \mathbf{x} = \mathbf{z}}} v(\mathbf{x})v(\mathbf{y}). \quad (2.40)$$

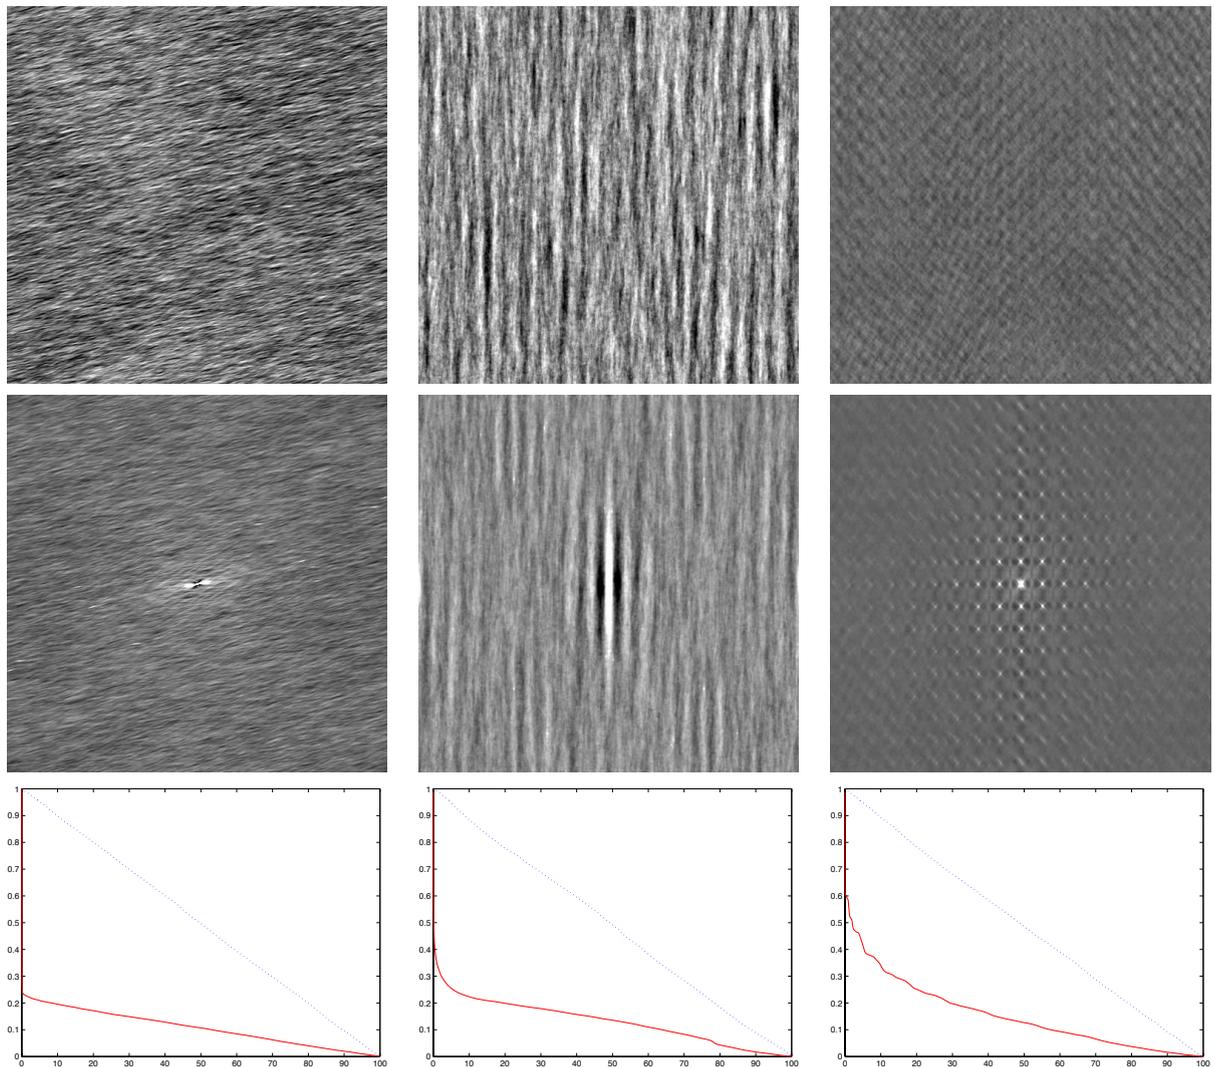


Figure 2.4: First row: the texture image u . Second row: the texton $T(u)$. Third row: the proportion of energy in a disc centered at the origin as a function of the percentage of pixels it contains (see Equation (2.30)), for both the texton (red curve), and the original texture (dotted curve).

Writing $S = \text{Supp}(v)$ the support of v , we can notice that Γ_v^0 has a finite support, contained in $S + (-S) = \{\mathbf{y} - \mathbf{x}; \mathbf{x} \in D, \mathbf{y} \in D\}$, whereas Γ_v is D -periodic. More precisely, if $D = I_M \times I_N$, then

$$\forall (z_1, z_2) \in \mathbb{Z}^2, \quad \Gamma_v(z_1, z_2) = \sum_{(k,l) \in \mathbb{Z}^2} \Gamma_v^0(z_1 + kM, z_2 + lN),$$

which is nothing but the D -periodization of Γ_v^0 . If $S + (-S) \subset D$, then Γ_v^0 and Γ_v are identical on the subset D of \mathbb{Z}^2 , which means that the model defined by (2.38) is simply obtained by considering the restriction of (2.39) to D , and thus shows that (2.39) is a natural way to extend to \mathbb{Z}^2 the periodic Gaussian model (2.38). Of course, if we want the condition $S + (-S) \subset D$ to be satisfied, we shall not take for v the (centered) original texture sample $u - m_u$, which has no reason to vanish outside a neighborhood of $\mathbf{0}$ since it is supposed to look like the realization of a stationary random field. The texton $v = T(u)$ is an interesting candidate, since it defines through (2.38) the same periodic Gaussian field as u . Its support $S = \text{Supp}(v)$ does not fulfill the relation $S + (-S) \subset D$ either in general, but since it is concentrated around $\mathbf{0}$ we can expect to find a good approximation of it that satisfies this condition. Several ways to achieve this approximation (and, more generally, to approximate a texton by an image supported by a neighborhood of $\mathbf{0}$) will be discussed in Section 2.3.4.

2.3.2 Extending Gaussian textures to a finite domain

The extension of a Gaussian texture to \mathbb{Z}^2 that we discussed above yields a natural extension to a finite domain $D_1 \supset D$. Indeed, there are essentially two possibilities.

Periodic extension. First, we can consider the extension as a D_1 -periodic texture, obtained by

$$\forall \mathbf{y}_1, \quad U(\mathbf{y}) = \sum_{\mathbf{x} \in D} v(\mathbf{x}) W_{D_1}(\mathbf{y} - \mathbf{x}), \quad (2.41)$$

where $(W_{D_1}(\mathbf{x}))_{\mathbf{x}_1}$ is a set of i.i.d. random variables distributed as $\mathcal{N}(0, |D|^{-1})$. In other terms, U is the Gaussian model on D_1 associated to the image αv_{D_1} , where v_{D_1} is the extension of v to D_1 obtained by zero-padding (that is, setting $v_{D_1}(\mathbf{x}) = 0$ for all \mathbf{x}_1) and $\alpha = (|D_1|/|D|)^{1/2}$ is a normalization constant that compensates for the fact that the marginal variance of W_{D_1} is $|D_1|^{-1}$ (and not $|D|^{-1}$) in the Gaussian model. Ideally, D_1 should be large enough to ensure that $D_1 \supset D + (-D)$ (or, at least, $D_1 \supset S + (-S)$ where $S = \text{Supp}(v)$), so that the covariance is preserved around 0. Taking the Discrete Fourier Transform of (2.41) in the domain D_1 , we obtain

$$\forall \xi_1, \quad \widehat{U}(\xi) = \widehat{v_{D_1}}(\xi) \widehat{W_{D_1}}(\xi), \quad (2.42)$$

where v_{D_1} is the extension of v to D_1 obtained by setting $v_{D_1}(\mathbf{x}) = 0$ for all \mathbf{x}_1 . Thanks to (2.42), the simulation of U is thus direct in Fourier Domain.

Non-periodic extension. Another possibility is to consider the extension to D_1 as a restriction of the extension to \mathbb{Z}^2 . Given $v : D \rightarrow \mathbb{R}$, we follow (2.39) to define the non-periodic Gaussian field $U : D_1 \rightarrow \mathbb{R}$ by

$$\forall \mathbf{y}_1, \quad U(\mathbf{y}) = \sum_{\mathbf{x} \in D} v(\mathbf{x}) W_{D_2}(\mathbf{y} - \mathbf{x}), \quad (2.43)$$

where $D_2 = D_1 + (-D)$ and $(W_{D_2}(\mathbf{x}))_{\mathbf{x}_2}$ is a set of i.i.d. random variables distributed as $\mathcal{N}(0, |D|^{-1})$. Interestingly enough, the synthesis of this Gaussian field can also be realized in Fourier domain. Indeed, since no term of (2.43) involves a periodic extension of W_{D_2} , we can see that the right term of (2.43) is the periodic convolution of v_{D_2} (v extended to D_2 by zero-padding) with W_{D_2} , evaluated in a given point of D_1 . Therefore, the simulation of U can be realized through the simulation of the D_2 -periodic field $v_{D_2} \star W_{D_2}$ (which is, as we remarked before, direct in Fourier domain), followed by a simple restriction to D_1 .

2.3.3 Extending RPN textures

Taking the Discrete Fourier Transform of (2.38) leads to

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{U}(\boldsymbol{\xi}) = \widehat{v}(\boldsymbol{\xi}) \cdot \widehat{W}_D(\boldsymbol{\xi}), \quad (2.44)$$

and the structure of the field $(\widehat{W}_D(\boldsymbol{\xi}))_{\boldsymbol{\xi} \in D}$ can be made completely explicit thanks to Theorem 2.1. This establishes a natural link with the RPN model that can be written, like (2.44), under the form

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{U}(\boldsymbol{\xi}) = \widehat{v}(\boldsymbol{\xi}) \cdot Z(\boldsymbol{\xi}), \quad (2.45)$$

where $Z(\mathbf{0}) = 0$ and for all $\boldsymbol{\xi} \neq \mathbf{0}$, $Z(\boldsymbol{\xi})$ is equal to $\widehat{W}(\boldsymbol{\xi})$ conditionally to $|\widehat{W}(\boldsymbol{\xi})| = 1$. In other terms, the complex-valued random variables $(Z(\boldsymbol{\xi}))_{\boldsymbol{\xi} \in D_+}$ are uniformly distributed on the complex unit circle $|z| = 1$, and one has $Z(-\boldsymbol{\xi}) = Z(\boldsymbol{\xi})^*$ for all $\boldsymbol{\xi} \in D_+$. Considering the discussion about the Gaussian model above, the extension of the RPN model to a larger finite domain D_1 is quite straightforward. Indeed, given an image $v : D \rightarrow \mathbb{R}^2$, the associated RPN model on $D_1 \supset D$ is simply the classical RPN model on D_1 associated to αv_{D_1} , where, like in the Gaussian model, v_{D_1} is the extension of v to D_1 obtained by zero-padding and $\alpha = (|D_1|/|D|)^{1/2}$. Note that this coefficient α ensures that the covariance functions of the models $\text{RPN}(v)$ and $\text{RPN}(\alpha v_{D_1})$ are identical on D .

Now a natural question arises: is it possible to extend the RPN model to non-periodic textures defined on \mathbb{Z}^2 , as we did for the Gaussian model? Given an image $v : \mathbb{Z}^2 \rightarrow \mathbb{R}$ supported by D , we can define its Fourier Transform by

$$\forall \boldsymbol{\xi} \in \mathbb{R}^2, \quad \widehat{v}(\boldsymbol{\xi}) = \sum_{\mathbf{x} \in \mathbb{Z}^2} v(\mathbf{x}) e^{-i\mathbf{x} \cdot \boldsymbol{\xi}} = \sum_{\mathbf{x} \in D} v(\mathbf{x}) e^{-i\mathbf{x} \cdot \boldsymbol{\xi}},$$

where $(x_1, x_2) \cdot (\xi_1, \xi_2) = x_1 \xi_1 + x_2 \xi_2$. This function \widehat{v} is 2π -periodic along each direction, and we have the reconstruction formula

$$\forall \mathbf{x} \in \mathbb{Z}^2, \quad v(\mathbf{x}) = \frac{1}{4\pi^2} \int_{[-\pi, \pi]^2} \widehat{v}(\boldsymbol{\xi}) e^{i\mathbf{x} \cdot \boldsymbol{\xi}} d\boldsymbol{\xi}.$$

Hence, given a random field $\zeta : \mathbb{R}^2 \rightarrow \mathbb{R}$ that is 2π -periodic along both directions, we could consider, when it makes sense, the random field U defined by

$$\forall \mathbf{x} \in \mathbb{Z}^2, \quad U(\mathbf{x}) = \frac{1}{4\pi^2} \int_{[-\pi, \pi]^2} \widehat{v}(\boldsymbol{\xi}) e^{i\mathbf{x} \cdot \boldsymbol{\xi}} \zeta(\boldsymbol{\xi}) d\boldsymbol{\xi}. \quad (2.46)$$

When ζ is a Gaussian random field, a sense can be given to this integral (a Wiener integral) and we can obtain a spectral formulation of the random field specified by (2.39). However,

to the best of our knowledge, there is not way to give a sense to (2.46) when ζ satisfy the constraints of the RPN model, that is, when the random variables $(\zeta(\boldsymbol{\xi}))_{\boldsymbol{\xi} \in [-\pi, \pi]^2}$ are independent (modulo the symmetry constraints that ensure that U is real-valued) and uniformly distributed on the complex circle $|z| = 1$.

Moreover, the following result indicates that the “natural” extension of the RPN model to \mathbb{Z}^2 is in fact the Gaussian model (2.39). Indeed, we show that asymptotically, a given subpart of a RPN model with growing support tends to a Gaussian field.

Proposition 2.10. *Let $S : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a non-negative, $2\pi\mathbb{Z}^2$ -periodic, symmetric, bounded, and piecewise continuous function. For any integers M and N , we consider the RPN random field U_{MN} defined on $D_{MN} = I_M \times I_N$ by*

$$U_{MN}(x_1, x_2) = \frac{1}{\sqrt{MN}} \sum_{(\xi_1, \xi_2) \in D_{MN} \setminus \{\mathbf{0}\}} S(\xi_1, \xi_2) e^{i\Phi(\xi_1, \xi_2)} e^{2i\pi(\frac{x_1\xi_1}{M} + \frac{x_2\xi_2}{N})},$$

where $(\Phi(\boldsymbol{\xi}))_{\boldsymbol{\xi} \in (D_{MN})_+}$ are independent and uniformly distributed on $[0, 2\pi)$, and $\Phi(-\boldsymbol{\xi}) = -\Phi(\boldsymbol{\xi})$. Then, U_{MN} converges towards a Gaussian random field $U : \mathbb{Z}^2 \rightarrow \mathbb{R}$ in the sense of finite-dimensional distributions as (M, N) tends towards infinity, and the covariance of the limit random field U is given by

$$\forall \mathbf{x}, \mathbf{y} \in \mathbb{Z}^2, \quad \text{Cov}(U(\mathbf{x}), U(\mathbf{y})) = \frac{1}{4\pi^2} \int_{[-\pi, \pi]^2} S(\boldsymbol{\xi})^2 \cos((\mathbf{x} - \mathbf{y}) \cdot \boldsymbol{\xi}) d\boldsymbol{\xi}. \quad (2.47)$$

The proof is given in Appendix.

2.3.4 Compact approximations of a texton

We saw in the previous sections the interest of a compact representation of a texture u , that is, an image v with small support that defines the same Gaussian and RPN model as u . In practice, of course, the exact representation of u with a compact image is quite hard, but interesting approximations can be found based on its texton. In this section, we discuss several possibilities to approximate the texton $T(u)$ associated to a texture u with an image v that has a small support.

Notice that looking for an approximation of $T(u)$ with a small support is akin to looking for a smooth approximation of $\widehat{T(u)} = |\widehat{u}|$. Moreover, extending such an approximation to a larger support through zero padding is exactly equivalent to zooming in its Fourier spectrum with the sin C interpolation.

Remarks on texton approximation

In the following, we define and compare several simple strategies to approximate the texture model associated to u (or, equivalently, to $T(u)$) and the corresponding model for another image v with small support. A natural way to discuss the quality of an approximation comes through the transport distance between the Gaussian textures defined on D by u and v ,

$$\left(|\widehat{u}(\mathbf{0}) - \widehat{v}(\mathbf{0})|^2 + \sum_{\boldsymbol{\xi} \in D, \boldsymbol{\xi} \neq \mathbf{0}} (|\widehat{u}(\boldsymbol{\xi})| - |\widehat{v}(\boldsymbol{\xi})|)^2 \right)^{1/2}. \quad (2.48)$$

As we said above, we consider zero-mean texture models (since the constant (DC) term can be added separately to the synthesized image), so that the original image u satisfies $\hat{u}(\mathbf{0}) = 0$ (that is, $m_u = 0$). Hence, the Gaussian model we shall consider with the texton approximation v will not be $\text{ADSN}(v)$, but $\text{ADSN}(v - m_v)$, so that the term $|\hat{u}(\mathbf{0}) - \hat{v}(\mathbf{0})|^2$ is zero. With that convention, the distance we obtain corresponds to the Fréchet distance between the RPN models, that is,

$$d(u, v) = \|T(u) - T(v)\|_2 = \left(\sum_{\xi \in D, \xi \neq 0} (|\hat{u}(\xi)| - |\hat{v}(\xi)|)^2 \right)^{1/2}. \quad (2.49)$$

Let us remark that if the support S of v is a strict subset of D , then v and $v + \alpha \mathbf{1}_S$ do not define the same RPN and ADSN models. More precisely, there exists a value of α that minimizes the distance $d(u, v + \alpha \mathbf{1}_S)$, but this value may not be unique since the convex function $\alpha \rightarrow d(u, v + \alpha \mathbf{1}_S)$ may not be strictly convex.

Let us now discuss the issue of the normalization of v . More precisely, let us consider the rescaled versions αv (with $\alpha > 0$) of v . On the one hand, we want to impose that v is optimally scaled for the distance measure (2.49), which means that the minimum of $\alpha \rightarrow d(u, \alpha v)$ is obtained for $\alpha = 1$. Under this constraint, we easily derive that $\|\hat{v}\|_2^2 = \langle |\hat{u}|, |\hat{v}| \rangle$, which means that in general we will observe that for such a v , $\|\hat{v}\|_2 < \|\hat{u}\|_2$. On the other hand, we might rather want to impose the normalization $\|\hat{v}\|_2 = \|\hat{u}\|_2$, since this ensures that, in the Gaussian model, the marginal distribution of each pixel are the same in the exact (u) and approximate (v) models. These two constraints (optimality and energy conservation) are generally incompatible, and we may want to adapt the rescaling factor depending on the final objective. Thus, in the following we shall consider a modified texton “pseudo-distance”, which does not depend on a rescaling factor

$$d'(u, v) = d\left(u, \frac{\|u\|_2}{\|v\|_2} \cdot v\right). \quad (2.50)$$

Specified support

A first simple strategy to obtain an approximation of a texton $T(u) : D \rightarrow \mathbb{R}$ by an image v with small support consists in restricting $T(u)$ to a given support, that is,

$$v = T(u) \cdot \mathbf{1}_S, \quad (2.51)$$

where S is a (generally symmetric) subdomain of D . It is often interesting to choose, instead of a fixed domain S , a parametric collection of domains $(S_\theta)_{\theta \in \Theta}$ and to select θ according to the respective values of $|S_\theta|$ (the size of the domain S_θ) and $d'(u, v_\theta)$ (where $v_\theta = T(u) \cdot \mathbf{1}_{S_\theta}$). Among classical cases are

$$S_r = \{(k, l) \in D; k^2 + l^2 < r^2\} \quad (\text{discs}),$$

$$S_{a,b} = \{(k, l) \in D; |k| < a, |l| < b\} \quad (\text{rectangles}),$$

$$\text{and } S_{a,b,\alpha} = \left\{ (k, l) \in D; \frac{(k \cos \alpha + l \sin \alpha)^2}{a^2} + \frac{(-k \sin \alpha + l \cos \alpha)^2}{b^2} < 1 \right\} \quad (\text{ellipses}).$$

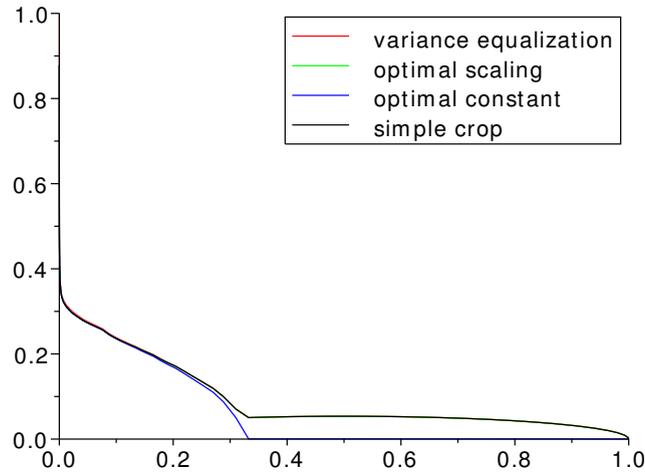


Figure 2.5: **Texture approximation with texton cropping (discs).** The texton of *grass* image (left) is multiplied by the indicator function of the discrete disc $Disc_R = \{\mathbf{x} \in \mathbb{Z}^2, |\mathbf{x}| \leq R\}$, and the resulting texture model V_R is compared to the original one, U (obtained with the full texton). Each plot displays the relative distance between U and V_R (see (2.50)) as a function of $|Disc_R|$ (the area of $Disc_R$). The different graphs correspond to different approximation strategies: simple crop (black curve), variance equalization (red), optimal scaling (blue), and optimal shift and scaling (green). We can observe that in terms of distance, the four strategies are very similar: only the variance equalization process is slightly worse than the simple crop. Another remark is that the first half of the distance is attained for a very small support, but the gain brought by increasing further the support size is quite slow.

We show in Figure 2.5 an example of such a texton approximation and the corresponding distance plots

$$A \mapsto \inf\{d'(u, v_\theta); |S_\theta| \leq A\},$$

that show, for each $A \in \mathbb{N}$, the minimum error that can be achieved by approximating the texton $T(u)$ with an image whose support has an area less than A .

Thresholding

Another strategy consists in applying a thresholding procedure to approximate $T(u)$ by $v = f(T(u))$, where f can be the hard-thresholding function

$$f_{\alpha,\beta}^{\text{hard}}(t) = \begin{cases} t & \text{if } t \leq -\alpha, \\ 0 & \text{if } -\alpha < t < \beta, \\ t & \text{if } \beta \leq t, \end{cases}$$

or the soft-thresholding function

$$f_{\alpha,\beta}^{\text{soft}}(t) = \begin{cases} t + \alpha & \text{if } t \leq -\alpha, \\ 0 & \text{if } -\alpha < t < \beta, \\ t - \beta & \text{if } \beta \leq t \end{cases}$$

(here, α and β are two positive parameters that can be chosen equal or not). Notice that in both cases, the approximations can be decomposed in two steps:

1. the selection of the support

$$S_{\alpha,\beta} = \{\mathbf{x} \in D, T(u) \notin (-\alpha, \beta)\};$$

2. the approximation of $T(u)$ in that support, according to the chosen f .

Thus, additional geometric constraints can be easily incorporated in this approximation strategy, in particular :

- the constraint that the support S of v is connected (take for S the connected component of $S_{\alpha,\beta}$ that contains 0);
- the constraint $S \subset D(0, R)$, simply enforced by setting $S = S_{\alpha,\beta} \cap D(0, R)$.

As before, the distance plot

$$A \mapsto \inf\{d'(u, v_{\alpha,\beta}); |S_{\alpha,\beta}| \leq A\},$$

can be used to select the thresholding procedure and the parameters α, β , as illustrated in Fig. Figure 2.5.

2.4 Textons for color images

In the following, we investigate the generalization of the notion of texton to color images. Let us consider a color image $\mathbf{u} : D \rightarrow \mathbb{R}^3$ (we assume as usual that color is represented by a 3-D vector corresponding to red, green and blue channels, but all what follows can be easily transposed in the more general case of a vector image $\mathbf{u} : D \rightarrow \mathbb{R}^d$). The color RPN model associated to \mathbf{u} is the random color image $\mathbf{U} : D \rightarrow \mathbb{R}^3$ defined by

$$\forall \mathbf{x} \in D, \quad \mathbf{U}(\mathbf{x}) = \frac{1}{|D|} R_0 \hat{\mathbf{u}}(\mathbf{0}) + \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D, \boldsymbol{\xi} \neq \mathbf{0}} e^{-i\varphi(\boldsymbol{\xi})} e^{2i\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle} \hat{\mathbf{u}}(\boldsymbol{\xi}), \quad (2.52)$$

where $\hat{\mathbf{u}}$ is the DFT of the color image \mathbf{u} (obtained by processing each channel independently), R_0 is a $\mathcal{N}(0, 1)$ random variable and $\varphi(\boldsymbol{\xi})$ a random set of uniform phases, independent modulo the Hermitian symmetry constraint. It is important to notice that in the color RPN model, the phase randomization is realized simultaneously on all channels, with the same phase translation function φ . An independent phase randomization of each channel yields poor results in general [51]. As in the monochromatic case, the ADSN model associated to \mathbf{u}

can be defined in several equivalent ways, either by adding in (2.52) a multiplicative Rayleigh noise term $R(\boldsymbol{\xi})$ of parameter $\frac{1}{2}$, with $R(\boldsymbol{\xi}) = R(-\boldsymbol{\xi})$, leading to

$$\forall \mathbf{x} \in D, \quad \mathbf{U}_g(\mathbf{x}) = \frac{1}{|D|} R_0 \hat{\mathbf{u}}(\mathbf{0}) + \frac{1}{|D|} \sum_{\boldsymbol{\xi} \in D, \boldsymbol{\xi} \neq \mathbf{0}} R(\boldsymbol{\xi}) e^{-i\varphi(\boldsymbol{\xi})} e^{2i\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle} \hat{\mathbf{u}}(\boldsymbol{\xi}) \quad (2.53)$$

(note that $R(\boldsymbol{\xi})e^{-i\varphi(\boldsymbol{\xi})}$ is a Gaussian complex random variable), or directly with

$$\forall \mathbf{x} \in D, \quad \mathbf{U}_g(\mathbf{x}) = (W \star \mathbf{u})(\mathbf{x}) = \sum_{\mathbf{y} \in D} W(\mathbf{x} - \mathbf{y}) \mathbf{u}(\mathbf{y}), \quad (2.54)$$

where $W : D \rightarrow \mathbb{R}$ is a Gaussian white noise (extended to \mathbb{Z}^2 by D -periodicity) such that each $W(\mathbf{x})$ has variance 1.

As in the grey-level case, we will assume that all images have a zero mean on D , that is

$$\mathbf{m}_{\mathbf{u}} = \frac{1}{|D|} \sum_{\mathbf{x} \in D} \mathbf{u}(\mathbf{x}) = \mathbf{0},$$

which is equivalent to say that each channel has zero mean. We will again use the notation $\mathcal{M}_0(\Omega)$ to denote the set of \mathbb{R}^3 -valued images with zero mean.

For a color image $\mathbf{u} = (u_1, u_2, u_3) \in \mathcal{M}_0(\Omega)$, we define its empirical covariance as

$$\forall \mathbf{y} \in D, \quad \Gamma_{\mathbf{u}}(\mathbf{y}) = (\Gamma_{\mathbf{u}}(\mathbf{y})_{k,l})_{1 \leq k, l \leq 3} \quad \text{with} \quad \Gamma_{\mathbf{u}}(\mathbf{y})_{k,l} = \frac{1}{|D|} \sum_{\mathbf{x} \in D} u_k(\mathbf{x}) u_l(\mathbf{x} - \mathbf{y}).$$

In the case of color images, the link between the covariance and the Fourier transform is given by

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{\Gamma}_{\mathbf{u}}(\boldsymbol{\xi}) = \frac{1}{|D|} \hat{\mathbf{u}}(\boldsymbol{\xi}) \hat{\mathbf{u}}(\boldsymbol{\xi})^*,$$

where $\hat{\mathbf{u}}(\boldsymbol{\xi})$ is considered as a column matrix in \mathbb{C}^3 and the notation A^* denotes the conjugate transpose of a (complex) matrix A .

As in the grey-level case, we have that the color RPN model (Equation (2.52)) exactly preserves the empirical covariance of the image, whereas the color Gaussian model (defined by Equation (2.53) or equivalently by Equation (2.54)) preserves it only in expectation. That is

$$\Gamma_{\mathbf{U}} = \Gamma_{\mathbf{u}} \quad \text{when} \quad \mathbf{U} \sim \text{RPN}(\mathbf{u}) \quad \text{and} \quad \mathbb{E}(\Gamma_{\mathbf{U}_g}) = \Gamma_{\mathbf{u}} \quad \text{when} \quad \mathbf{U}_g \sim \text{ADSN}(\mathbf{u}).$$

Proposition 2.11. *The optimal transport (OT) distance, also called Fréchet distance, between two color Gaussian textures $\text{ADSN}(\mathbf{u})$ and $\text{ADSN}(\mathbf{v})$ is equal to the distance between their RPN models, and it is given by:*

$$d(\text{ADSN}(\mathbf{u}), \text{ADSN}(\mathbf{v}))^2 = d(\text{RPN}(\mathbf{u}), \text{RPN}(\mathbf{v}))^2 = \sum_{\boldsymbol{\xi} \in D} (\|\hat{\mathbf{u}}(\boldsymbol{\xi})\|^2 + \|\hat{\mathbf{v}}(\boldsymbol{\xi})\|^2 - 2|\hat{\mathbf{u}}(\boldsymbol{\xi})^* \hat{\mathbf{v}}(\boldsymbol{\xi})|).$$

Proof. As in the monochromatic case, the proof uses Parseval theorem and the following result: if \mathbf{a} and \mathbf{b} are in \mathbb{C}^3 and seen as complex column matrices, then

$$\min_{\theta \in [0, 2\pi)} \|\mathbf{a} - e^{i\theta} \mathbf{b}\|^2 = \|\mathbf{a}\|^2 + \|\mathbf{b}\|^2 - 2|\mathbf{a}^* \mathbf{b}|.$$

□

Notice that here the term $|\widehat{\mathbf{u}}(\boldsymbol{\xi})^* \widehat{\mathbf{v}}(\boldsymbol{\xi})|$ is in general, contrarily to the monochromatic case, not equal to $\|\widehat{\mathbf{u}}(\boldsymbol{\xi})\| \|\widehat{\mathbf{v}}(\boldsymbol{\xi})\|$, and therefore the optimal transport distance will not be as simple as in the case of grey-level textures.

2.4.1 The α -color texton

As already mentioned, the important point about the color RPN model and the color Gaussian model is that the same phase field φ is used simultaneously on all channels. The poor results obtained when performing an independant phase randomization on each channel is explained by the fact that such a randomization does not preserve the values $\widehat{u}_k(\boldsymbol{\xi}) \widehat{u}_l(\boldsymbol{\xi})^*$ for $k \neq l$ (only the covariances of each channel are preserved but not the cross-covariances between different channels). Therefore, using the same phase field on all channels is the only possible choice to preserve the covariance of the image.

Definition and properties

To define the texton of a color image, we have different possibilities. A first possibility is to find a phase field φ such that when adding it to the phases of each channel, we obtain a color image that has some optimal concentration properties. In the grey-level case, with such a point of view, we saw that the optimal choice was to take for φ the opposite of the phase of \widehat{u} , in such a way that the resulting texton was a grey-level image with null phases, that is real positive Fourier transform. In the case of a color image \mathbf{u} , one can project it as a grey-level image by considering for any $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \alpha_3) \in \mathbb{R}^3$, the real-valued image

$$\boldsymbol{\alpha} \cdot \mathbf{u} = \sum_{k=1}^3 \alpha_k u_k.$$

For instance, the intensity of the color image corresponds to $\boldsymbol{\alpha} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, while its luminance (perceived brightness) is obtained for $\boldsymbol{\alpha} = (0.299, 0.587, 0.114)$. Taking the phases of $\boldsymbol{\alpha} \cdot \mathbf{u}$ to shift the phases of the three channels of the image is a first possibility to define a color texton, and we will see that its has some of the desired concentration properties. To be more precise, here is the definition.

Definition 2.5 (α -color texton). *For a phase field $\varphi : D \rightarrow \mathbb{R}$ satisfying the anti-symmetry condition $\varphi(-\boldsymbol{\xi}) = -\varphi(\boldsymbol{\xi})$ for all $\boldsymbol{\xi} \in D$, let us denote by S_φ the operator defined on color images \mathbf{u} by shifting the phases of all the channels with the phase field φ . That is, if \mathbf{u} is a color image, then $S_\varphi \mathbf{u}$ is also a color image, given by*

$$\forall \boldsymbol{\xi} \in D, \widehat{S_\varphi \mathbf{u}}(\boldsymbol{\xi}) = \widehat{\mathbf{u}}(\boldsymbol{\xi}) e^{-i\varphi(\boldsymbol{\xi})}.$$

For $\boldsymbol{\alpha} \in \mathbb{R}^3$, let us define the α -color texton by the operator

$$\mathbf{u} \mapsto T_\alpha(\mathbf{u}) := S_{\varphi_{\boldsymbol{\alpha} \cdot \mathbf{u}}} \mathbf{u}, \text{ where } \varphi_{\boldsymbol{\alpha} \cdot \mathbf{u}} \text{ is the phase field of } \boldsymbol{\alpha} \cdot \widehat{\mathbf{u}}.$$

As in the case of grey-level textures, we also introduce the set of all images obtained from \mathbf{u} by a phase shift, i.e. we define

$$\Theta_{\mathbf{u}} = \{ \mathbf{v} = S_\varphi \mathbf{u} \text{ for } \varphi \text{ anti-symmetric phase field } \}.$$

Notice that for any λ real, $\boldsymbol{\alpha}$ and $\lambda\boldsymbol{\alpha}$ define the same color texton. The link between the $\boldsymbol{\alpha}$ -color texton and the texton of grey-level images that we have defined and studied in the first sections of this chapter is given by the following proposition.

Proposition 2.12. *For any color image \mathbf{u} and any $\boldsymbol{\alpha}$ in \mathbb{R}^3 , we have the identity*

$$\boldsymbol{\alpha} \cdot T_{\boldsymbol{\alpha}}(\mathbf{u}) = T(\boldsymbol{\alpha} \cdot \mathbf{u}),$$

where T is the texton operator on grey-level images defined in Equation (2.16). As a consequence, the $\boldsymbol{\alpha}$ -color texton is solution of the two following optimization problems:

$$\boldsymbol{\alpha} \cdot T_{\boldsymbol{\alpha}}(\mathbf{u})(\mathbf{0}) = \max_{\mathbf{v} \in \Theta_{\mathbf{u}}} \boldsymbol{\alpha} \cdot \mathbf{v}(\mathbf{0}) \quad \text{and} \quad T_{\boldsymbol{\alpha}}(\mathbf{u}) = \underset{\mathbf{v} \in \Theta_{\mathbf{u}}}{\text{Argmin}} \sum_{\mathbf{x} \in D} A(\mathbf{x}) (\boldsymbol{\alpha} \cdot \mathbf{v}(\mathbf{x}))^2,$$

where A is a real-valued image with non-positive Fourier transform, as defined in Proposition 2.5.

Proof. Let $\boldsymbol{\alpha} \in \mathbb{R}^3$ and let \mathbf{u} be a color image. In the Fourier domain, we have for all $\boldsymbol{\xi} \in D$,

$$\begin{aligned} \boldsymbol{\alpha} \cdot \widehat{T_{\boldsymbol{\alpha}}(\mathbf{u})}(\boldsymbol{\xi}) &= \boldsymbol{\alpha} \cdot \widehat{T(\boldsymbol{\alpha} \cdot \mathbf{u})}(\boldsymbol{\xi}) = \boldsymbol{\alpha} \cdot \widehat{\mathbf{u}}(\boldsymbol{\xi}) e^{-i\varphi_{\boldsymbol{\alpha} \cdot \mathbf{u}}(\boldsymbol{\xi})} \\ &= |\boldsymbol{\alpha} \cdot \widehat{\mathbf{u}}(\boldsymbol{\xi})| = \widehat{T(\boldsymbol{\alpha} \cdot \mathbf{u})}(\boldsymbol{\xi}). \end{aligned}$$

The second part of the proposition is a direct consequence of the optimality properties of the grey-level texton given in Propositions 2.4 and 2.5. \square

The synthesis of textures from an $\boldsymbol{\alpha}$ -color texton is straightforward: we simply use either the color RPN model of Equation (2.52) or the color Gaussian model of Equations (2.53) or (2.54) with the color image $T_{\boldsymbol{\alpha}}(\mathbf{u})$. Since $T_{\boldsymbol{\alpha}}(\mathbf{u})$ is obtained from \mathbf{u} by a simple phase shift, both models have the same law as with the original image \mathbf{u} .

On the choice of $\boldsymbol{\alpha}$

Of course, one of the main questions is the choice of $\boldsymbol{\alpha}$. Many choices are possible, for instance: one of the channel (that is $\boldsymbol{\alpha} = (1, 0, 0)$, or $(0, 1, 0)$ or $(0, 0, 1)$), the intensity ($\boldsymbol{\alpha}_{int} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$), the usual luminance ($\boldsymbol{\alpha}_{lum} = (0.299, 0.587, 0.114)$), or eigenvectors of the covariance matrix $\Gamma_{\mathbf{u}}(\mathbf{0})$. Indeed, the eigenvector corresponding to the largest eigenvalue of this matrix defines a color direction $\tilde{\boldsymbol{\alpha}}$ that captures most the information between the different channels of the color image, and is therefore a good choice for defining a $\boldsymbol{\alpha}$ -color texton. We will illustrate this on some examples (see bottom left graphics of Figure 2.7, Figure 2.8, Figure 2.9 and Figure 2.10). On the contrary, the two other eigenvectors are far from being good choices. To make this statement more precise, we compare the concentration properties of the different $\boldsymbol{\alpha}$ -color texton by plotting, in a way similar to the grey-level case, the energy outside a disk of radius r centered at the origin as a function of the percentage of pixels inside this disk. The lower the obtained curve is, the better the concentration of the color texton is. More precisely, we plot

$$\frac{\sum_{\mathbf{x} \in D \setminus D_r} \|T_{\boldsymbol{\alpha}}(\mathbf{u})(\mathbf{x})\|^2}{\sum_{\mathbf{x} \in D} \|T_{\boldsymbol{\alpha}}(\mathbf{u})(\mathbf{x})\|^2} \text{ as a function of } 100 \times \frac{|D_r|}{|D|}, \quad (2.55)$$

where $\|\cdot\|$ is the euclidean norm on \mathbb{R}^3 . The plots obtained for different possible α 's and for different textures are shown on Figure 2.7, Figure 2.8, Figure 2.9 and Figure 2.10.

These curves are all bounded from below by the one that corresponds to the color image made of the three textons of the three channels, i.e. the image given by $(T(u_1), T(u_2), T(u_3))$. Recall that this is not an admissible choice since, in general, it is not the result of a single phase shift. However, the energy curve it gives (denoted "3 textons" in the plots) is the best bound that can be achieved, or not, by a phase shift S_φ .

From these examples, and from other experiments we have made, we can draw the following conclusions: the eigenvector $\tilde{\alpha}$ corresponding to the highest eigenvalue of the covariance matrix $\Gamma_{\mathbf{u}}(\mathbf{0})$ is always the best possible α , while the two other eigenvectors are the worst choices. For natural textures, α_{int} and α_{lum} have concentration performances that are very close to the ones of $\tilde{\alpha}$, and very close also to best possible ones given by the 3 textons. For synthetic examples (example 3 that is the RPN of a natural but not micro-texture, and example 4 that is a color image made of three different grey-level textures in the three channels), $\tilde{\alpha}$ is still the best choice, but it does not reach the optimal bound.

The fact that $\tilde{\alpha}$ seems to be always the best choice is a consequence of the following proposition. Indeed, this proposition shows that the color direction $\tilde{\alpha}$ is the one that captures, in expectation, most of the "energy" of the phase-shifted images.

Proposition 2.13. *Let \mathbf{u} be a color image and let $\alpha \in \mathbb{R}^3$ be any color direction. Let A be a real-valued image with non-positive Fourier transform, as defined in Proposition 2.5, and let the weighted energy of a phase-shift $S_\varphi \mathbf{u}$ in the α direction be defined by*

$$E_{A,\alpha}(S_\varphi \mathbf{u}) = \sum_{\mathbf{x} \in D} A(\mathbf{x}) |\alpha \cdot S_\varphi \mathbf{u}(\mathbf{x})|^2.$$

Then, taking the expectation of this energy when the $\varphi(\xi)$ are i.i.d. (up to the Hermitian symmetry condition) uniform on $[0, 2\pi)$, we get

$$\mathbb{E}[E_{A,\alpha}(S_\varphi \mathbf{u})] = 2|D| \left(\sum_{\xi \in D_+ \setminus \{0\}} \lambda_\xi \right) {}^t \alpha \Gamma_{\mathbf{u}}(\mathbf{0}) \alpha,$$

where $\lambda_\xi = -\frac{1}{|D|} \widehat{A}(\xi)$ are positive coefficients. Therefore, whatever A , $\mathbb{E}[E_{A,\alpha}(S_\varphi \mathbf{u})]$ as a function of $\alpha \in S^2$ (S^2 is here the unit sphere of \mathbb{R}^3) is maximal when $\alpha = \tilde{\alpha}$.

Proof. By linearity of the expectation we have

$$\mathbb{E}(E_{A,\alpha}(S_\varphi \mathbf{u})) = \sum_{\mathbf{x} \in D} A(\mathbf{x}) {}^t \alpha \mathbb{E}(S_\varphi \mathbf{u}(\mathbf{x}) {}^t S_\varphi \mathbf{u}(\mathbf{x})) \alpha.$$

Now, because of the stationarity of the model RPN(\mathbf{u}) we have that $\mathbb{E}(S_\varphi \mathbf{u}(\mathbf{x}) {}^t S_\varphi \mathbf{u}(\mathbf{x}))$ is independent of \mathbf{x} and it is equal to $\Gamma_{\mathbf{u}}(\mathbf{0})$. Therefore

$$\mathbb{E}(E_{A,\alpha}(S_\varphi \mathbf{u})) = \left(\sum_{\mathbf{x} \in D} A(\mathbf{x}) \right) {}^t \alpha \Gamma_{\mathbf{u}}(\mathbf{0}) \alpha.$$

To end the proof, we simply notice that, according to Proposition 2.5,

$$\sum_{\mathbf{x} \in D} A(\mathbf{x}) = \widehat{A}(\mathbf{0}) = 2|D| \sum_{\boldsymbol{\xi} \in D_+ \setminus \{\mathbf{0}\}} \lambda_{\boldsymbol{\xi}}.$$

□

The conclusion of this section is that we can define the color texton of a color image by choosing a color direction $\boldsymbol{\alpha}$. Many possible choices are available for $\boldsymbol{\alpha}$, and in terms of concentration property, the best choice seems to be $\boldsymbol{\alpha} = \tilde{\boldsymbol{\alpha}}$, the eigenvector associated with the largest eigenvalue of the covariance matrix $\Gamma_{\mathbf{u}}(\mathbf{0})$. However, for natural microtexture images, other reasonable choices of $\boldsymbol{\alpha}$ (intensity, or luminance for instance), lead to similar concentration performances. To analyze further the concentration performances of the different possible $\boldsymbol{\alpha}$'s, we will consider their ability to re-synthesize the original texture after being cropped, as a function of the spatial size of the remaining support. This will be analyzed in Section 2.4.3.

2.4.2 The matricial color texton

Recall that the texton of a grey-scale image \mathbf{u} is a square root – for the convolution operation – of the periodic auto-correlation. For color images, one may want to follow that idea, and define alternative textons by $\mathbf{x} \mapsto M(\mathbf{x})$ where for each $\mathbf{x} \in D$, $M(\mathbf{x})$ is matrix of rank at most one and such that $M \star M = \Gamma_{\mathbf{u}}$, thus of dimensions either 3×1 , 3×2 or 3×3 . In the Fourier domain, this is equivalent to $\widehat{M}(\boldsymbol{\xi})\widehat{M}(\boldsymbol{\xi})^* = \widehat{\Gamma}_{\mathbf{u}}(\boldsymbol{\xi})$ for all $\boldsymbol{\xi} \in D$. We propose an analysis where $M(\mathbf{x})$ is assumed to be a 3×3 matrix with rank less than one for all \mathbf{x} in the following.

As in the case of the grey-level texton, among all possible choices for M , we will choose the one that is the most concentrated at $\mathbf{0}$, and we get the following definition.

Definition 2.6. *The matricial texton of a color image \mathbf{u} is the matrix-valued function $\mathbf{x} \mapsto M_T(\mathbf{u})(\mathbf{x}) \in \mathcal{M}_3(\mathbb{R})$ such that its Fourier transform is given by*

$$\forall \boldsymbol{\xi} \in D, \widehat{M_T(\mathbf{u})}(\boldsymbol{\xi}) = \frac{1}{\|\widehat{\mathbf{u}}(\boldsymbol{\xi})\|} \widehat{\mathbf{u}}(\boldsymbol{\xi})\widehat{\mathbf{u}}(\boldsymbol{\xi})^*.$$

Moreover we have the property that, among all matrices M such that $\widehat{M}(\boldsymbol{\xi})\widehat{M}(\boldsymbol{\xi})^* = \widehat{\Gamma}_{\mathbf{u}}(\boldsymbol{\xi})$, then

$$\text{tr}(M(\mathbf{0})) \text{ is maximal when } M = M_T(\mathbf{u}).$$

The maximal property of $\text{tr}(M(\mathbf{0}))$ comes from the following computation: assume that M is such that $\widehat{M}(\boldsymbol{\xi})\widehat{M}(\boldsymbol{\xi})^* = \widehat{\Gamma}_{\mathbf{u}}(\boldsymbol{\xi})$, then since $\widehat{\Gamma}_{\mathbf{u}}(\boldsymbol{\xi})$ is a matrix with rank at most one with all columns proportional to $\widehat{\mathbf{u}}(\boldsymbol{\xi})$, $\widehat{M}(\boldsymbol{\xi})$ is also necessarily a matrix with rank at most one. Moreover, $\widehat{\Gamma}_{\mathbf{u}}(\boldsymbol{\xi})$ must be of the form $\widehat{M}(\boldsymbol{\xi}) = \widehat{\mathbf{u}}(\boldsymbol{\xi})\boldsymbol{\gamma}(\boldsymbol{\xi})^*$, where $\boldsymbol{\gamma}(\boldsymbol{\xi}) \in \mathbb{C}^3$ satisfies $\boldsymbol{\gamma}(\boldsymbol{\xi})^*\boldsymbol{\gamma}(\boldsymbol{\xi}) = 1$. Then

$$\text{tr}(M(\mathbf{0})) = \sum_{\boldsymbol{\xi}} \text{tr}(\widehat{M}(\boldsymbol{\xi})) = \sum_{\boldsymbol{\xi}} \text{tr}(\widehat{\mathbf{u}}(\boldsymbol{\xi})\boldsymbol{\gamma}(\boldsymbol{\xi})^*),$$

and it is maximal when $\gamma(\boldsymbol{\xi}) = \frac{1}{\|\widehat{\mathbf{u}}(\boldsymbol{\xi})\|} \widehat{\mathbf{u}}(\boldsymbol{\xi})$.

Notice furthermore that the matricial texton of a color image \mathbf{u} is given, in the Fourier domain, by Hermitian matrices. Indeed, we have

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{M}_T(\mathbf{u})(\boldsymbol{\xi}) = \widehat{M}_T(\mathbf{u})(\boldsymbol{\xi})^*.$$

Thus, as defined above, the matricial color texton is the unique 3×3 -matrix-valued function, whose values are all symmetrical and whose convolution square is $\Gamma_{\mathbf{u}}$, the auto-correlation function of \mathbf{u} .

We have just seen that the matricial texton $M_T(\mathbf{u})$ of a color image \mathbf{u} satisfies a concentration property in $\mathbf{0}$ that is analogous to the property **(P1)** in the grey-level case. In a straightforward way, we also have that it satisfies another concentration property analogous to **(P2)**. More precisely, we have that $M_T(\mathbf{u}) = \text{Argmin}_M \sum_{\mathbf{x}} A(\mathbf{x}) \|M(\mathbf{x})\|_F^2$, where A is a real-valued weight image with $\widehat{A}(\boldsymbol{\xi}) \leq 0$ for all $\boldsymbol{\xi} \in D \setminus \{\mathbf{0}\}$ (as in Proposition 2.5), $\|\cdot\|_F$ denotes the Frobenius norm, and where the matrices M on which the minimum is taken satisfy $\widehat{M}(\boldsymbol{\xi})\widehat{M}(\boldsymbol{\xi})^* = \widehat{\Gamma}_{\mathbf{u}}(\boldsymbol{\xi})$ for all $\boldsymbol{\xi}$, $\widehat{M}(\boldsymbol{\xi})$ is Hermitian for all $\boldsymbol{\xi}$ and $\widehat{M}(\mathbf{0})$ is moreover positive.

To explore the link between the matricial color texton and the $\boldsymbol{\alpha}$ -color textons, we claim that the matricial texton is more general in the sense that we can always recover the $\boldsymbol{\alpha}$ -color texton for any color direction $\boldsymbol{\alpha} \in \mathbb{R}^3$ with the formula

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{T}_{\boldsymbol{\alpha}}(\mathbf{u})(\boldsymbol{\xi}) = \widehat{\mathbf{u}}(\boldsymbol{\xi}) \frac{\widehat{\mathbf{u}}(\boldsymbol{\xi})^* \boldsymbol{\alpha}}{|\widehat{\mathbf{u}}(\boldsymbol{\xi})^* \boldsymbol{\alpha}|} = \frac{\text{tr}(\widehat{M}_T(\mathbf{u})(\boldsymbol{\xi}))}{\sqrt{\boldsymbol{\alpha}^* \widehat{M}_T(\mathbf{u})(\boldsymbol{\xi})^* \widehat{M}_T(\mathbf{u})(\boldsymbol{\xi}) \boldsymbol{\alpha}}} \widehat{M}_T(\mathbf{u})(\boldsymbol{\xi}) \boldsymbol{\alpha}.$$

The synthesis of the texture models $\text{ADSN}(\mathbf{u})$ or $\text{RPN}(\mathbf{u})$ from the matricial texton are obtained the following way. Let $\mathbf{W} : D \mapsto \mathbb{R}^3$ be a vector-valued Gaussian white noise with variance $1/|D|$, which means that $\mathbf{W} = (W_1, W_2, W_3)$ with W_k independent real-valued Gaussian white noises with variance $1/|D|$. Then, the $\text{ADSN}(\mathbf{u})$ and $\text{RPN}(\mathbf{u})$ models are respectively synthesized by defining in the Fourier domain

$$\widehat{\mathbf{U}}_g(\boldsymbol{\xi}) = \widehat{M}_T(\mathbf{u})(\boldsymbol{\xi}) \widehat{\mathbf{W}}(\boldsymbol{\xi}).$$

and

$$\widehat{\mathbf{U}}(\boldsymbol{\xi}) = \widehat{M}_T(\mathbf{u})(\boldsymbol{\xi}) \widehat{\mathbf{Z}}(\boldsymbol{\xi}), \quad \text{with } \widehat{\mathbf{Z}}(\boldsymbol{\xi}) = \frac{e^{i\varphi(\boldsymbol{\xi})}}{\|\widehat{\mathbf{u}}(\boldsymbol{\xi})\|} \widehat{\mathbf{u}}(\boldsymbol{\xi}),$$

where the $\varphi(\boldsymbol{\xi})$ are random uniform phases, independent up to the Hermitian symmetry constraint.

As we already mentioned, the matricial texton $M_T(\mathbf{u})$ of a color image \mathbf{u} has the property of being such that all $\widehat{M}_T(\mathbf{u})(\boldsymbol{\xi})$ are matrices of rank at most one. But it is theoretically possible to define matricial “textons” that yield matrices in the Fourier domain of any rank up to 3. Although such textons cannot be directly based on any exemplar image, the case might arise from compression (cropping around $\mathbf{0}$).

Let us first define general Gaussian models for color textures. In the case of grey-level images, we have the equivalence: U is a real-valued Gaussian periodic stationary random field on D if and only if there exists a real-valued image u such that $U \sim \text{ADSN}(u)$. Now,

for color images, the situation is different, all Gaussian periodic stationary color textures are not of the form ADSN(\mathbf{u}). They are, in the general case, given in the Fourier domain by

$$\widehat{\mathbf{U}}(\boldsymbol{\xi}) = \widehat{M}(\boldsymbol{\xi})\widehat{\mathbf{W}}(\boldsymbol{\xi}),$$

where $\widehat{\mathbf{W}}$ is a vector-valued complex white noise and where $\widehat{M}(\boldsymbol{\xi})$ is a complex matrix, not necessarily of rank one. In the most general case, the covariance of the Gaussian model is given by

$$\widehat{\Gamma}(\boldsymbol{\xi}) = \widehat{M}(\boldsymbol{\xi})\widehat{M}(\boldsymbol{\xi})^* = \mathbb{E}(\widehat{\Gamma}_{\mathbf{U}}(\boldsymbol{\xi})).$$

This covariance is in general a Hermitian matrix of any rank up to 3, whereas all Fourier transform coefficients of an empirical covariance have a rank less than one. In the following we will denote by GT(Γ) such a model of color Gaussian texture. For any color image \mathbf{u} , then ADSN(\mathbf{u}) and GT($\Gamma_{\mathbf{u}}$) define the same model. But on the inverse, a model GT(Γ) is equal to a model ADSN(\mathbf{u}) if and only if all $\widehat{\Gamma}(\boldsymbol{\xi})$, for $\boldsymbol{\xi} \in D$ are rank one matrices.

The texton of a general GT(Γ) model can be defined the following way:

$$\forall \boldsymbol{\xi} \in D, \quad \widehat{M}_T(\boldsymbol{\xi}) = \widehat{\Gamma}(\boldsymbol{\xi})^{1/2},$$

where the notation $\widehat{\Gamma}(\boldsymbol{\xi})^{1/2}$ denotes the unique positive square root of the positive matrix $\widehat{\Gamma}(\boldsymbol{\xi})$. This definition implies in particular that here again the Fourier transform $\widehat{M}_T(\boldsymbol{\xi})$ of the matricial texton are Hermitian matrices. The synthesis from the general matricial texton is then a consequence of the fact $\mathbf{U} \sim \text{GT}(\Gamma)$ is equivalent in the Fourier domain to $\widehat{\mathbf{U}}(\boldsymbol{\xi}) = \widehat{M}_T(\boldsymbol{\xi})\widehat{\mathbf{W}}(\boldsymbol{\xi})$ or in the space domain to $\mathbf{U} = M_T \star \mathbf{W}$, where \mathbf{W} is vector-valued white noise.

Although the pertinence of such general Gaussian (or RPN) synthesis is debatable to synthesize realistic visual textures, it might be of interest in order to synthesize new kinds of textures

On the different examples shown on Figure 2.7, Figure 2.8, Figure 2.9 and Figure 2.10, we show for each texture its $\tilde{\alpha}$ -color texton, its matricial texton (shown as two color images: one corresponds to the diagonal elements of the matricial texton, and the other one to the off-diagonal elements of the matricial texton). We also compare the concentration performances of the different textons on the bottom left graphic of these figures.

2.4.3 Sparse color representation

Since the α -color textons and the matricial texton of a color image are concentrated around $\mathbf{x} = \mathbf{0}$, we can obtain, as in the grey-level case, a sparse representation of a color texture \mathbf{u} by a simple crop of one of these textons. More precisely, let χ denote a crop function (for instance it is the indicator function of a disk of radius r centered at $\mathbf{0}$). Then we can consider the cropped color texton $T_c : \mathbf{x} \mapsto \chi(\mathbf{x})T_{\alpha}(\mathbf{u})(\mathbf{x})$, or the cropped matricial texton $M_c : \mathbf{x} \mapsto \chi(\mathbf{x})M_T(\mathbf{u})(\mathbf{x})$. And we can synthesize again color textures from these two models (either with RPN or ADSN for T_c , and with a GT for M_c).

Optimal constant before the crop

As in the case of grey-level images, instead of performing directly the crop of the color texton $T_{\mathbf{u}}$, we can look for an optimal constant $\beta \in \mathbb{R}^3$ such that the OT distance between the original texture \mathbf{u} and the texture obtained from the crop of $T_{\mathbf{u}} + \beta$ is as small as possible. From Proposition 2.11, we have that this distance is given by

$$d(\text{ADSN}(\mathbf{u}), \text{ADSN}(\mathbf{v}))^2 = d(\text{RPN}(\mathbf{u}), \text{RPN}(\mathbf{v}))^2 = \sum_{\xi \in D} (\|\widehat{\mathbf{u}}(\xi)\|^2 + \|\widehat{\mathbf{v}}(\xi)\|^2 - 2|\widehat{\mathbf{u}}(\xi)^* \widehat{\mathbf{v}}(\xi)|),$$

with $\mathbf{v} = \chi(T_{\mathbf{u}} + \beta)$. There is no closed formula to get the optimal β , but as in the grey-level case, we can easily compute the gradient of the above OT distance seen as a function of β .

Analogous computations can be performed with the matricial color texton, where in that case the constant β becomes a 3×3 matrix.

Covariance equalization after the crop

In the case of grey-level images, it is straightforward to see that our way of cropping textons induces a loss of variance for the texture and that therefore a variance correction had to be performed. Here, for color textures, the same phenomenon appears, in a more striking way because not only the variance of each channel is reduced, but the co-variances between the different channels is also modified, resulting sometimes in the “lost” of some colors. This is illustrated on Figure 2.6 where we show why a covariance equalization is needed after the crop. More precisely, the sparse color representation of a color texture \mathbf{u} is performed the following way:

1. Choose a crop function χ and compute either the cropped color texton T_c or the cropped matricial texton M_c .
2. Perform the covariance equalization, that is find a matrix $B \in \mathcal{M}_3(\mathbb{R})$ such that

$$B \sum_{\xi} \widehat{T}_c(\xi) \widehat{T}_c(\xi)^* B^* = \sum_{\xi} \widehat{\mathbf{u}}(\xi) \widehat{\mathbf{u}}(\xi)^* = \Gamma_{\mathbf{u}}(\mathbf{0}), \text{ or } B \sum_{\xi} \widehat{M}_c(\xi) \widehat{M}_c(\xi)^* B^* = \Gamma_{\mathbf{u}}(\mathbf{0}).$$

3. Synthesize new textures from BT_c with $\text{ADSN}(BT_c)$ or $\text{RPN}(BT_c)$, or from BM_c with $\text{GT}(BM_c M_c^* B^*)$.

On the choice of B . Notice that there are a priori infinitely many possible choices for the matrix B used to equalize the covariances. Indeed, let us denote Γ_c the matrix $\sum_{\xi} \widehat{T}_c(\xi) \widehat{T}_c(\xi)^*$ (or $\sum_{\xi} \widehat{M}_c(\xi) \widehat{M}_c(\xi)^*$) and $\Gamma_0 = \Gamma_{\mathbf{u}}(\mathbf{0})$, then these are two positive definite symmetric matrices and we can choose for B any matrix that satisfies $B\Gamma_c B^* = \Gamma_0$. If we diagonalize Γ_0 and Γ_c in orthogonal basis we can write

$$\Gamma_0 = O_0 D_0 O_0^* \text{ and } \Gamma_c = O_c D_c O_c^*,$$

where O_0 and O_c are orthogonal matrices (with determinant +1) and D_0 and D_c are diagonal matrices with positive coefficients arranged in decreasing order on the diagonal. Then,

$$\forall O \in O_3(\mathbb{R}), \quad B = O_0 D_0^{1/2} O D_c^{-1/2} O_c^* \text{ satisfies } B\Gamma_c B^* = \Gamma_0. \quad (2.56)$$

Now, let us address the problem of the choice for B among all these possibilities. A first possible criterion is to choose B such that it changes the image the less, which means for instance that we choose B such that $\|B - I\|_F^2 = \text{tr}((B - I)(B^* - I))$ is as small as possible (I denotes here the identity matrix in \mathbb{R}^3 and $\|\cdot\|_F$ stands for the Frobenius norm). Another possibility is to take B such that the OT distance between the original model and the cropped and equalized model $\text{ADSN}(BT_c)$ (or $\text{GT}(BM_cM_c^*B^*)$) is the smallest. Unfortunately, unlike the monochromatic case, there is no closed formula to solve this optimization problem. We propose to choose $O = I_3$ (Identity of \mathbb{R}^3) in (2.56), leading to the equalization matrix

$$B = O_0 D_0^{1/2} D_c^{-1/2} O_c^*$$

which yield in practice very good results for both performance criteria.

To compare the synthesis performances of the different α -color textons and of the matricial texton, we have plotted for different textures the OT-distance between the original model and the model after the crop and the covariance equalization. That is, we plot

$$d(\text{ADSN}(BT_c), \text{ADSN}(\mathbf{u})) \text{ as a function of } 100 \times \frac{\sum_{\mathbf{x}} \chi(\mathbf{x})}{D},$$

where we take for χ the indicator function of a ball of radius r .

We do it the same way for the cropped matricial texton. But in that case, after the crop and the equalization, the matrices $\widehat{BM}_c(\boldsymbol{\xi})$ are not necessarily of rank one, and therefore the Gaussian model given by BM_c is not necessarily an ADSN model. In that case, to compute the OT-distance the two stationary Gaussian models, we need to use the full formula:

$$d(\text{GT}(\Gamma_1), \text{GT}(\Gamma_2)) = \sum_{\boldsymbol{\xi}} \text{tr}(\widehat{\Gamma}_1(\boldsymbol{\xi}) + \widehat{\Gamma}_2(\boldsymbol{\xi}) - 2(\widehat{\Gamma}_1(\boldsymbol{\xi})^{1/2} \widehat{\Gamma}_2(\boldsymbol{\xi}) \widehat{\Gamma}_1(\boldsymbol{\xi})^{1/2})^{1/2}).$$

Remark: To compute numerically this distance in an efficient way we first notice that when one of the two matrices $\widehat{\Gamma}_1(\boldsymbol{\xi})$ or $\widehat{\Gamma}_2(\boldsymbol{\xi})$ is of rank at most one (which is the case when we compute the distance to a model $\text{ADSN}(\mathbf{u})$), then $\widehat{\Gamma}_1(\boldsymbol{\xi})^{1/2} \widehat{\Gamma}_2(\boldsymbol{\xi}) \widehat{\Gamma}_1(\boldsymbol{\xi})^{1/2}$ is also of rank at most one and therefore we can use the fact that if A is a positive symmetric matrix of rank at most one, then $\text{tr}(A^{1/2}) = \text{tr}(A)^{1/2}$. Moreover to have more stable numerical results, we also use the identity: $\text{tr}(A_1) + \text{tr}(A_2) - 2\text{tr}(A_1 A_2)^{1/2} = \text{tr}(A_2 - A_1) + 2\text{tr}(A_1(A_1 - A_2)) / (\text{tr}(A_1) + \text{tr}(A_1 A_2)^{1/2})$.

To conclude this section on the matricial texton we point out that, unlike the α -color texton, the matricial texton is defined for any color Gaussian texture. For real textures, it is practically equivalent (in terms of concentration property) to the $\tilde{\alpha}$ -color texton as shown on the following figures.

2.5 Appendix

2.5.1 Proof of Theorem 2.1

Proof. Since Fourier transform is an invertible linear operator, $(U(\mathbf{x}))_{\mathbf{x} \in \Omega}$ is a centered Gaussian vector if and only if $\{\widehat{U}(\mathbf{0}), \text{Re} \widehat{U}(\boldsymbol{\xi}), \text{Im} \widehat{U}(\boldsymbol{\xi}); \boldsymbol{\xi} \in \Omega_+\}$ is a centered Gaussian vector. We

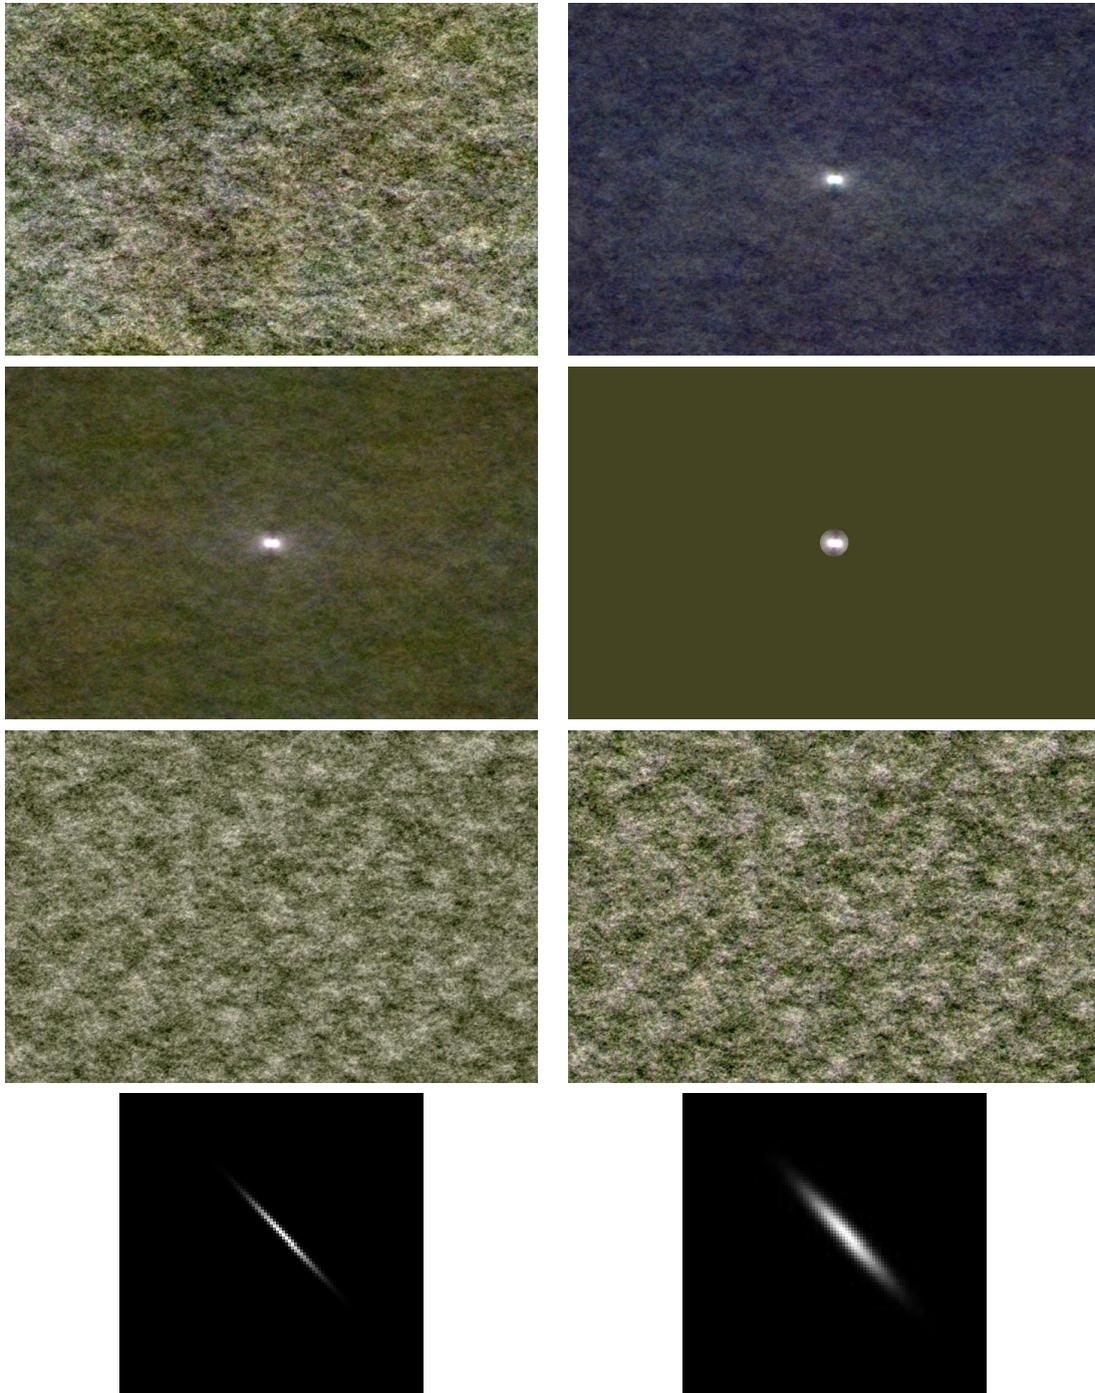


Figure 2.6: First line: the original texture “wall” (size 512×512) and the off-diagonal matricial texton. Second line: the diagonal of the matricial texton and diagonal of the matricial texton after a crop with a disk of radius 20 pixels. Third line: texture synthesized with the cropped matricial texton (notice how the blue-grey shades of the original texture are almost lost) and same texture after covariance equalization. Fourth line: joint distribution of the values of the red and blue channels of the synthesized texture before covariance equalization (on the left) and after (on the right). This second joint distribution is exactly equal to the one of the original texture.

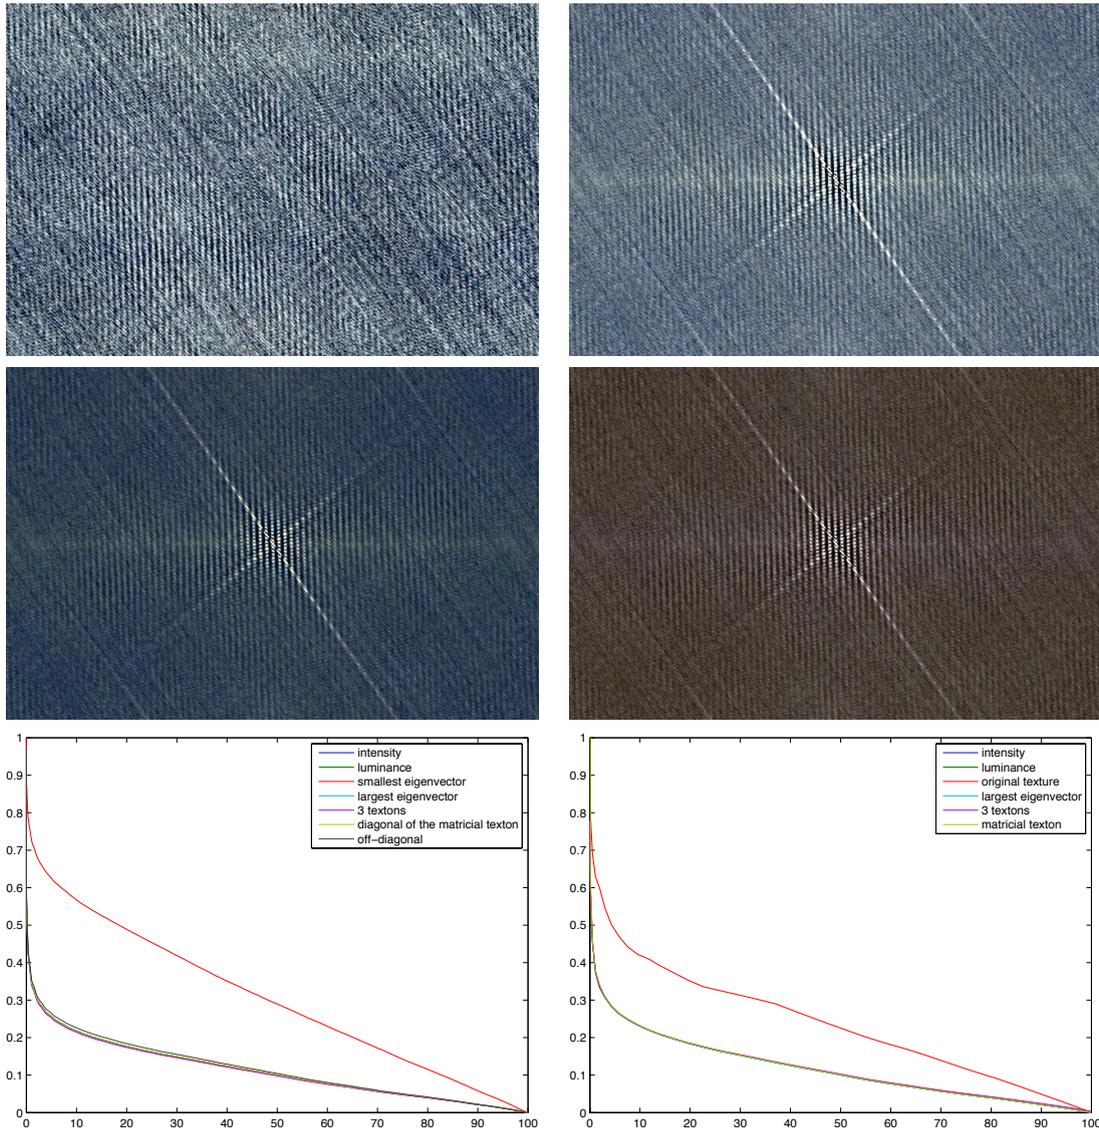


Figure 2.7: First line: left, the original texture image; right, the $\tilde{\alpha}$ -color texton. Second line: the matricial texton visualized as two color images: the diagonal elements on the left, and the off-diagonal elements on the right. Third line: left, energy of the different textons as a function of the percentage of pixels in a disk centered at $\mathbf{0}$; right, square OT distance of the different cropped textons after covariance equalization.

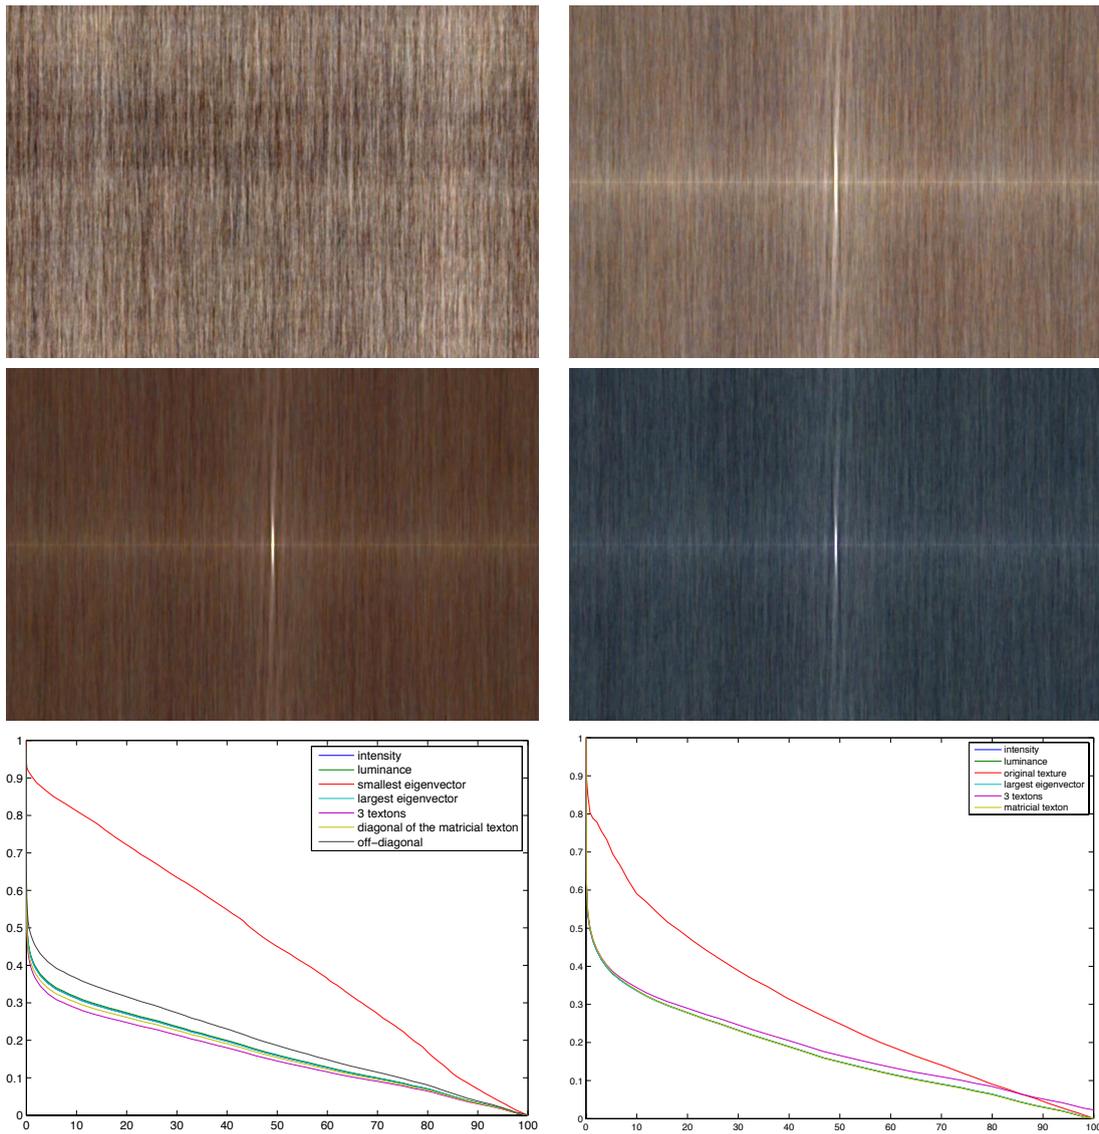


Figure 2.8: First line: left, the original texture image; right, the $\tilde{\alpha}$ -color texton. Second line: the matricial texton visualized as two color images: the diagonal elements on the left, and the off-diagonal elements on the right. Third line: left, energy of the different textons as a function of the percentage of pixels in a disk centered at $\mathbf{0}$; right, square OT distance of the different cropped textons after covariance equalization.

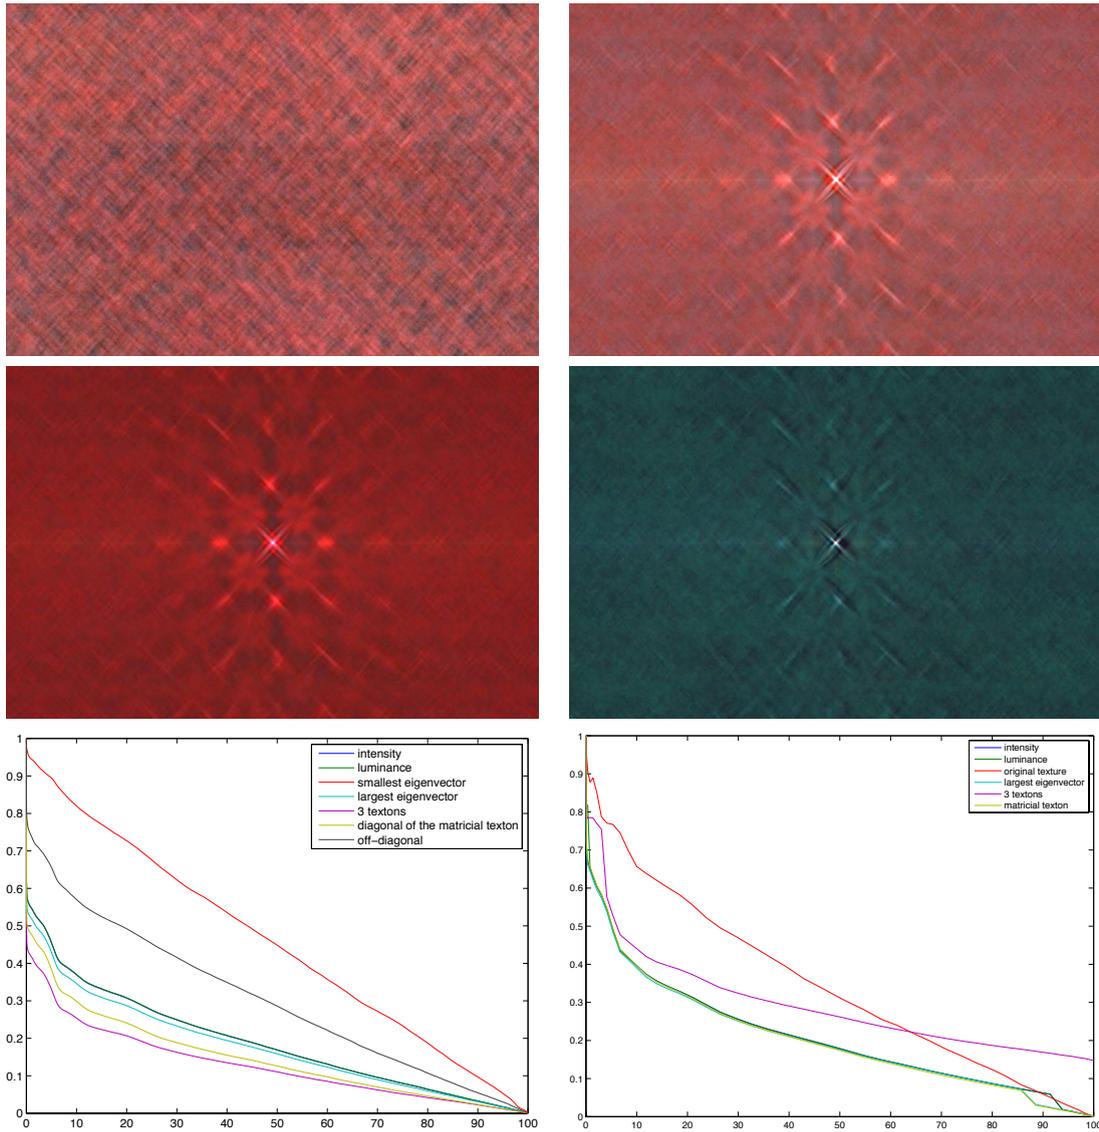


Figure 2.9: First line: left, the original texture image; right, the $\tilde{\alpha}$ -color texton. Second line: the matricial texton visualized as two color images: the diagonal elements on the left, and the off-diagonal elements on the right. Third line: left, energy of the different textons as a function of the percentage of pixels in a disk centered at $\mathbf{0}$; right, square OT distance of the different cropped textons after covariance equalization.

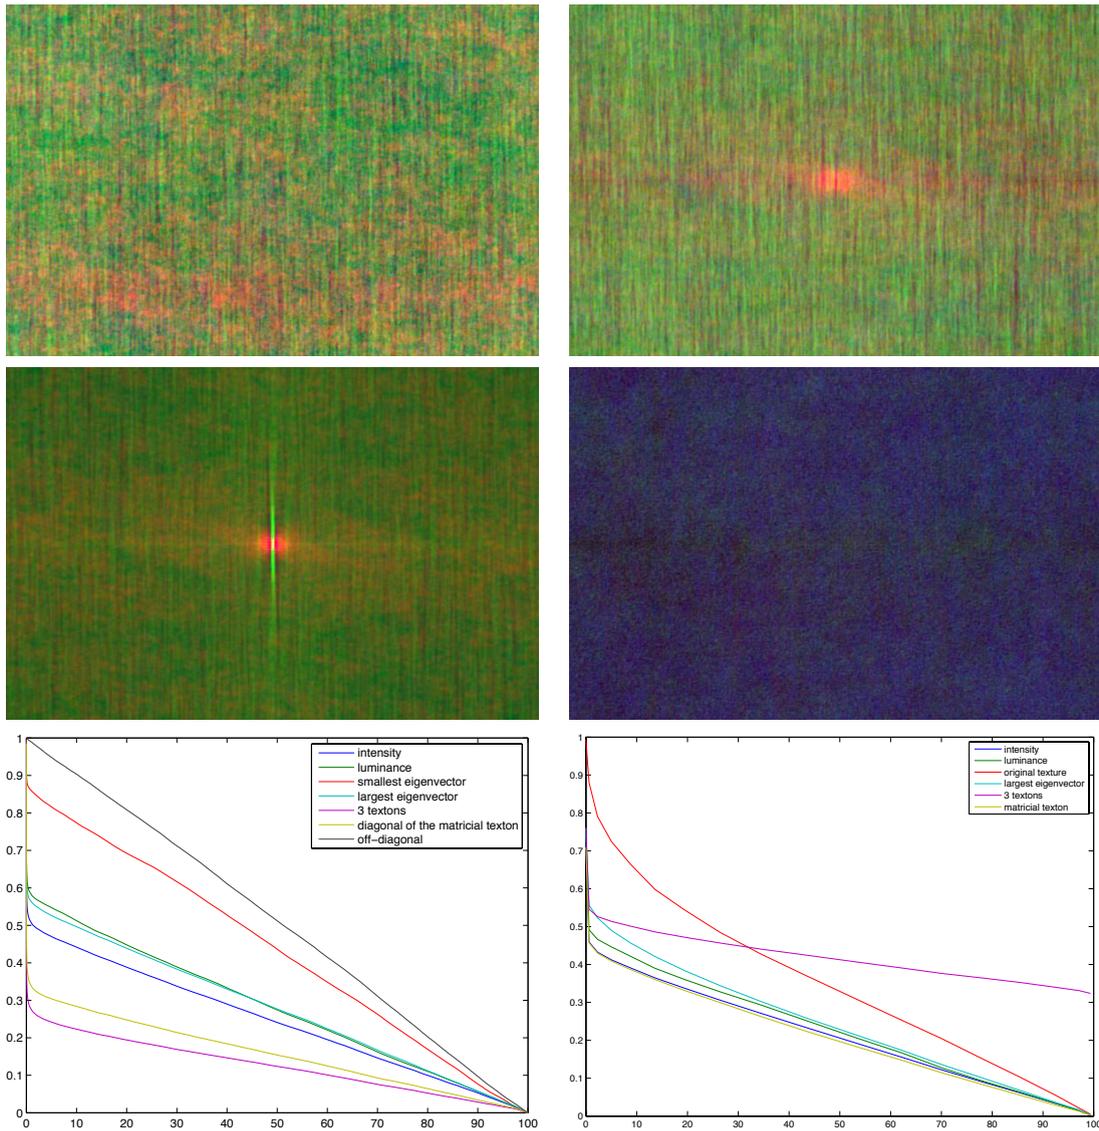


Figure 2.10: First line: left, the original texture image; right, the $\tilde{\alpha}$ -color texton. Second line: the matricial texton visualized as two color images: the diagonal elements on the left, and the off-diagonal elements on the right. Third line: left, energy of the different textons as a function of the percentage of pixels in a disk centered at $\mathbf{0}$; right, square OT distance of the different cropped textons **without** covariance equalization. Here, the diagonal matrices in (2.56) are not numerically invertible.

need now to compute how their covariances are related. Now, assume that $(U(\mathbf{x}))_{\mathbf{x} \in \Omega}$ is a centered Gaussian periodic stationary random field, and we denote $\Gamma(\mathbf{x}) = \text{Cov}(U(\mathbf{x}), U(\mathbf{0}))$ for all $\mathbf{x} \in \Omega$. Because U is periodic stationary we have that $\Gamma(\mathbf{x}) = \text{Cov}(U(\mathbf{x} + \mathbf{y}), U(\mathbf{y}))$ for all $\mathbf{x}, \mathbf{y} \in \Omega$. Moreover $\Gamma(\mathbf{x}) = \Gamma(-\mathbf{x})$ and as a consequence, $\widehat{\Gamma}$ is real. Let $\boldsymbol{\xi}$ and $\boldsymbol{\xi}'$ in Ω_+ , we start for instance by computing $\text{Cov}(\text{Re } \widehat{U}(\boldsymbol{\xi}), \text{Re } \widehat{U}(\boldsymbol{\xi}'))$. We get

$$\begin{aligned} \text{Cov}(\text{Re } \widehat{U}(\boldsymbol{\xi}), \text{Re } \widehat{U}(\boldsymbol{\xi}')) &= \mathbb{E} \left(\left(\sum_{\mathbf{x} \in \Omega} U(\mathbf{x}) \cos(2\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle) \right) \left(\sum_{\mathbf{x} \in \Omega} U(\mathbf{x}) \cos(2\pi \langle \mathbf{x}, \boldsymbol{\xi}' \rangle) \right) \right) \\ &= \sum_{\mathbf{x} \in \Omega} \sum_{\mathbf{y} \in \Omega} \mathbb{E}(U(\mathbf{x})U(\mathbf{x} + \mathbf{y})) \cos(2\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle) \cos(2\pi \langle \mathbf{x} + \mathbf{y}, \boldsymbol{\xi}' \rangle) \\ &= \sum_{\mathbf{y} \in \Omega} \Gamma(\mathbf{y}) \sum_{\mathbf{x} \in \Omega} \cos(2\pi \langle \mathbf{x}, \boldsymbol{\xi} \rangle) \cos(2\pi \langle \mathbf{x} + \mathbf{y}, \boldsymbol{\xi}' \rangle) \\ &= \sum_{\mathbf{y} \in \Omega} \Gamma(\mathbf{y}) \frac{MN}{2} \cos(2\pi \langle \mathbf{y}, \boldsymbol{\xi} \rangle) \delta_{\boldsymbol{\xi} = \boldsymbol{\xi}'} = \frac{MN}{2} \widehat{\Gamma}(\boldsymbol{\xi}) \delta_{\boldsymbol{\xi} = \boldsymbol{\xi}'} \end{aligned}$$

Similar computations can be made for the imaginary parts of $\widehat{U}(\boldsymbol{\xi})$ and $\widehat{U}(\boldsymbol{\xi}')$, and also for the covariance between the real and the imaginary parts.

The converse part of the theorem works also in a very similar way. □

2.5.2 Proof of Proposition 2.10

Proof. Let $\|S\|_\infty$ denote the upper bound of S over $[-\pi, \pi]^2$. Thanks to periodicity, this is also the upper bound of S over \mathbb{R}^2 . Let $\mathbf{x}_1, \dots, \mathbf{x}_l$ denote l distinct points in \mathbb{Z}^2 . Take (M, N) in \mathbb{N}^2 and consider $(\widetilde{U}_{M,N}(\mathbf{x}_1), \dots, \widetilde{U}_{M,N}(\mathbf{x}_l))$ and the joint characteristic function:

$$\Phi_{M,N}(t_1, \dots, t_l) = \mathbb{E}[\exp(it_1 \widetilde{U}_{M,N}(\mathbf{x}_1) + \dots + it_l \widetilde{U}_{M,N}(\mathbf{x}_l))]$$

defined for any (t_1, \dots, t_l) in \mathbb{R}^l . For any $\mathbf{x} \in \mathbb{Z}^2$, we have, using the hypothesis on S and the symmetry condition on ϕ , that

$$\begin{aligned} \widetilde{U}_{M,N}(\mathbf{x}) &= \frac{1}{\sqrt{MN}} \sum_{\boldsymbol{\xi} \in \Omega_{MN}} S(\boldsymbol{\xi}) e^{i\phi(\boldsymbol{\xi})} e^{2i\pi \mathbf{x} \cdot \boldsymbol{\xi}} \\ &= \frac{2}{\sqrt{MN}} \sum_{\boldsymbol{\xi} \in \Omega_{MN+}} S(\boldsymbol{\xi}) \cos(\phi(\boldsymbol{\xi}) + 2\pi \mathbf{x} \cdot \boldsymbol{\xi}). \end{aligned}$$

Then, using the independance of the $\phi(\boldsymbol{\xi})$ for $\boldsymbol{\xi}$ in Ω_{MN+} , we get

$$\log \Phi_{M,N}(t_1, \dots, t_l) = \sum_{\boldsymbol{\xi} \in \Omega_{MN+}} \log \mathbb{E} \left(e^{i \frac{2}{\sqrt{MN}} \sum_{j=1}^l t_j S(\boldsymbol{\xi}) \cos(\phi(\boldsymbol{\xi}) + 2\pi \mathbf{x}_j \cdot \boldsymbol{\xi})} \right).$$

Now, for $\boldsymbol{\xi} \in \Omega_{MN+}$ let us denote

$$X_{\boldsymbol{\xi}} = \frac{2}{\sqrt{MN}} \sum_{j=1}^l t_j S(\boldsymbol{\xi}) \cos(\phi(\boldsymbol{\xi}) + 2\pi \mathbf{x}_j \cdot \boldsymbol{\xi}).$$

It is straightforward to see that $X_{\boldsymbol{\xi}}$ is bounded independently of $\boldsymbol{\xi}$ by

$$|X_{\boldsymbol{\xi}}| \leq \frac{2}{\sqrt{MN}}(|t_1| + \dots + |t_l|)\|S\|_{\infty}.$$

We can also compute the first two moments of $X_{\boldsymbol{\xi}}$ and we obtain, using the property that for any uniform random variable ϕ over $\mathbb{R}/2\pi\mathbb{Z}$ then $\mathbb{E}(\cos(\phi + u)) = 0$ and $\mathbb{E}(\cos(\phi + u)\cos(\phi + v)) = \frac{1}{2}\cos(u - v)$,

$$\mathbb{E}(X_{\boldsymbol{\xi}}) = 0 \quad \text{and} \quad \mathbb{E}(X_{\boldsymbol{\xi}}^2) = \frac{2}{MN} \sum_{j,k=1}^l t_j t_k S(\boldsymbol{\xi})^2 \cos(2\pi(\mathbf{x}_j - \mathbf{x}_k) \cdot \boldsymbol{\xi}).$$

Now, for any ε small enough in \mathbb{R} , we have by Taylor formula that

$$|e^{i\varepsilon} - 1 - i\varepsilon + \frac{1}{2}\varepsilon^2| \leq \varepsilon^3 \quad \text{and} \quad |\log(1 + \varepsilon) - \varepsilon| \leq 2\varepsilon^2.$$

Therefore, there exists a constant C independant of M, N and $\boldsymbol{\xi}$ such that

$$\forall \boldsymbol{\xi} \in \Omega_{MN+}, \quad |\log \mathbb{E}(e^{iX_{\boldsymbol{\xi}}}) + \frac{1}{MN} \sum_{j,k=1}^l t_j t_k S(\boldsymbol{\xi})^2 \cos(2\pi(\mathbf{x}_j - \mathbf{x}_k) \cdot \boldsymbol{\xi})| \leq \frac{C}{(MN)^{3/2}}.$$

Then

$$|\log \Phi_{M,N}(t_1, \dots, t_l) + \frac{1}{MN} \sum_{\boldsymbol{\xi} \in \Omega_{MN+}} \sum_{j,k=1}^l t_j t_k S(\boldsymbol{\xi})^2 \cos(2\pi(\mathbf{x}_j - \mathbf{x}_k) \cdot \boldsymbol{\xi})| \leq \frac{C}{\sqrt{MN}}.$$

Thus, as (M, N) grows to infinity, since S is assumed to be piecewise continuous, the Riemann sum converges and we finally obtain

$$\log \Phi_{M,N}(t_1, \dots, t_l) \xrightarrow{M,N \rightarrow \infty} -\frac{1}{4\pi^2} \sum_{j,k=1}^l t_j t_k \int_{[-\pi,\pi]^2} S(\boldsymbol{\xi})^2 \cos((\mathbf{x}_j - \mathbf{x}_k) \cdot \boldsymbol{\xi}) d\boldsymbol{\xi}.$$

That concludes the proof. □

Chapter 3

Large Error Approximations of Images

We revisit the importance of phase of the Fourier transform of signals through a paradigm of “large error approximations” (LEA) of signals. We investigate variants of this approximation that either challenge or comfort the commonly accepted interpretation that the shapes of an image are coded within the phase. We also discuss the importance of the discrete Fourier transform itself, as a choice among other unitary operators.

3.1 Introduction

Many tasks in signal processing can be described as finding a signal, within a given set, close to an initial signal in either a quantifiable or a perceptual way. For instance, lossy signal compression problems can be stated in the following way:

- the initial signal is a raw collection of sampled measurements
- the approximation set is characterised by the sparsity of its elements in some representation (a given ratio of non-zero entries with respect to some representation, see *e.g.* [62] for the DCT and [91] for orthogonal wavelets)
- the error can be measured by the distance induced by the $\|\cdot\|_2$ norm, but is often assessed by tests with respect to human perception.

Such examples share the characteristic that the approximations they render are of somewhat “small error”, typically with a relative distance less than 10%.

In different contexts however, the need for such a small distance from the original image is irrelevant to the processing task. Histogram equalization, as highlighted by the morphological model [57], may yield results that are, by design, far from the initial image – the error can often be higher than 50%. However, the resulting image remains “faithful” to the original, in the sense that the former can easily be recognized in the latter by human vision. We call such processing tasks “large error approximations” (LEA). Examples of LEA are countless and comprise approximations of images with constant patches (down-sampling or

pixelation), dominos or ASCII characters (see Knowlton and Harmon in [74] and [58]), photographic patches from a library (photomosaics as pioneered by Silvers in [123]), patches from a texturing image (texture transfer, see Efros and Freeman [41]), stipples (see *e.g.* Deussen in [37], and Secord [121]), dots of different sizes (halftones, see *e.g.* Ostromoukhov [107]) and even “Traveling Salesman Problem” paths (see Bosch and Kaplan [71]). We refer to Figure 3.1 for a few graphic examples.



Figure 3.1: “Lena” and three large error approximations. Top left: pixellized 64×64 sub-sampling. Top right: photomosaics as pioneered in [123] (credit: [1]). Bottom left: stippling (6000 stipples). Bottom right: TSP art (credit Seb Perez-Duarte via Flickr).

The fact that the initial signal can be recognized in large error approximations can be of great importance for vision psychologists and physiologist, neuroscientists, brain modelers etc. For instance, the recognition of photomosaics or downsamplings of an image (see Figure 3.1), along with the experiments of Oliva and Torralba (see *e.g.* [102]) are quite well explained, *e.g.* thanks to the findings of Campbell and Robson [20]: visual neurons respond to specific spatial

frequencies. Navon [101] offered a concurrent explanation: in certain conditions (especially scene vision), perception follows a hierarchy from global to local.

Many LEA can be obtained as the solution of a projection problem. More precisely, let D denote sampling set. Let \mathcal{S} be a subset of \mathbb{R}^D (resp. \mathbb{C}^D) – the “constraint set” – *e.g.* the set of images with some given histogram in the task of histogram equalization. Let u – the “target image” – be an image, modeled by an element in \mathbb{R}^D . We consider images s^* in the “best approximation set” defined by

$$\arg \min_{s \in \mathcal{S}} \|s - u\|_2. \quad (3.1)$$

We shall refer to LEA obtained through projections as “large error projections” (LEP). For instance, the operator that maps images onto the set of images with a prescribed histogram is a projection and finds the closest image, within the (non-convex) set of images with a prescribed histogram, from the initial image.

This chapter deals with a particular LEP that was discussed by Oppenheim and Lim in a celebrated article [106]. This classical phenomenon in signal processing, namely “the importance of phase in signals”, is still quite puzzling for the scientific understanding of human vision. Morrone and Burr have suggested that the phase of signals could be somehow “hard-wired” in human perception system [100], but this hypothesis has turned out to be hard to test.

The purpose of this chapter is to revisit the findings described by Oppenheim and Lim in [106], with the LEA/LEP paradigm in mind. We argue that the importance of phase is deeply connected to the fact that the phase characterizes the solutions of a LEP under constraints on the mean and the (periodical) auto-correlation, which is detailed in Section 2. In Section 3, we support this argument by providing experiments of LEA/LEP where the constraints on the modulus are replaced by constraints on the phase itself, which surprisingly lead to the recognition of the target image. In Section 4, we provide further experiments in order to investigate the role of the Fourier transform itself and discuss an argument as for its specificity among other unitary transforms. In Section 5, we conclude by discussing the connection between LEA and LEP.

3.2 The importance of phase in signals as a LEA problem

In this section, we detail a LEP problem that is deeply connected to Oppenheim and Lim observations in [106] where they pointed out the importance of phase of the discrete Fourier transform of signals. One common interpretation is that some critical information of a signal (*e.g.* the edges and contours of an image) is roughly coded within its phase. The rationale for this interpretation is that the phase is essential to form edges from the summation of planar sine waves. Under some over-sampling hypotheses, the information in the phase alone is accurate enough to exactly retrieve the signal (see Section 5 in [106]).



Figure 3.2: Top left: “Lena”. Top right: “clouds”. Bottom left: the Oppenheim-Lim image w as defined by (3.2), with “Lena” as u (phase) and “clouds” as v (modulus). Bottom right: the Oppenheim-Lim image w as defined by (3.2) with “clouds” as u (phase) and “Lena” as v (modulus).

3.2.1 The LEP problem

Let u and v be two grey-scale images defined over the domain D (*i.e.* two elements in \mathbb{R}^D), and let \hat{u} and \hat{v} denote their discrete Fourier transforms. In order to point the importance of the phase of the Fourier transform, Oppenheim and Lim consider the signal w in \mathbb{R}^D defined by its Fourier transform as follows: for all ξ in D

$$\hat{w}(\xi) = \mathbf{1}_{\{\hat{u}(\xi) \neq 0\}} \frac{\hat{u}(\xi)}{|\hat{u}(\xi)|} |\hat{v}(\xi)|. \quad (3.2)$$

In other words, w is the image that has the phase of u , $\phi_u \left(\frac{\hat{w}(\xi)}{|\hat{w}(\xi)|} \right) = \frac{\hat{u}(\xi)}{|\hat{u}(\xi)|} = e^{i\phi_u(\xi)}$ – where $\phi_u(\xi)$ is in $\mathbb{R}/2\pi\mathbb{Z}$ – for all ξ such that $\hat{u}(\xi) \neq 0$ and the modulus of v ($|\hat{w}(\xi)| = |\hat{v}(\xi)|$) for

all ξ in D), which has been documented as being “similar” to the image u . Let us show that w solves a very particular LEA problem of the form defined by (3.1).

Proposition 3.1. *Define*

$$\mathcal{M}_v := \{s \in \mathbb{C}^D; |\widehat{s}(\xi)| = |\widehat{v}(\xi)| \ (\forall \xi \in D)\}. \quad (3.3)$$

Then

$$w \in \arg \min_{s \in \mathcal{M}_v} \|s - u\|_2. \quad (3.4)$$

Moreover, w is the unique point in $\arg \min_{s \in \mathcal{M}_v} \|s - u\|_2$ if and only if

$$\forall \xi \in D \ (\widehat{u}(\xi) = 0) \Rightarrow (\widehat{v}(\xi) = 0). \quad (3.5)$$

Proof. For simplicity we adopt the unitary definition of the discrete Fourier transform. Thanks to Parseval identity, $\|s - u\|_2 = \|\widehat{s} - \widehat{u}\|_2$, and since minimizing $\|s - u\|_2$ is equivalent to minimizing $\|s - u\|_2^2$, we have the following identity

$$\arg \min_{s \in \mathcal{M}_v} \|s - u\|_2 = \arg \min_{s \in \mathcal{M}_v} \sum_{\xi \in D} |\widehat{s}(\xi) - \widehat{u}(\xi)|^2. \quad (3.6)$$

Now notice that for one fixed non-zero complex number a , and a variable complex number z with a fixed modulus $|z| = |b| > 0$, the quantity $|z - a|$ is minimal if and only if $z = |b| \frac{a}{|a|}$, i.e. when a and z are on the same half-line starting at 0 on the complex plane (see Figure 3.3).

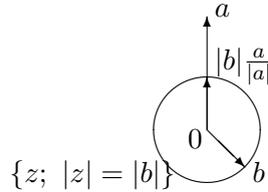


Figure 3.3: Orthogonal projection onto a circle.

Hence, for each ξ in D , for $s \in \mathcal{M}_v$, $|\widehat{s}(\xi) - \widehat{u}(\xi)|$ is minimal when $\widehat{s}(\xi) = |\widehat{v}(\xi)| e^{i\phi_u(\xi)}$. Thus w , as defined by (3.2), is in $\arg \min_{s \in \mathcal{M}_v} \sum_{\xi \in D} |\widehat{s}(\xi) - \widehat{u}(\xi)|^2$. Notice that w is indeed a real image since u is itself a real image and thus for all ξ in D , $e^{i\phi_u(-\xi)} = e^{-i\phi_u(\xi)}$ which implies in turn that $\widehat{w}(-\xi) = \overline{\widehat{w}(\xi)}$.

The uniqueness condition is straightforward. \square

Remark 3.1. *The error of the LEA problem (3.4), defined by $\min_{s \in \mathcal{M}_v} \|s - u\|_2$, is exactly $\| |\widehat{v}| - |\widehat{u}| \|_2$.*

Remark 3.2. *The choice of the $\|\cdot\|_2$ distance in (3.4) shows that w defined by (3.2) is an “orthogonal projection” of u onto the set \mathcal{M}_v . We write*

$$w = P_{\mathcal{M}_v}(u). \quad (3.7)$$

Remark 3.3. The space \mathcal{M}_v is generally non-convex and $\arg \min_{s \in \mathcal{M}_v} \|s - u\|_2$ is generally not reduced to a singleton. More precisely if there exists ξ_0 in $D \setminus \{0\}$ such that $\widehat{u}(\xi_0) = 0$ and $\widehat{v}(\xi_0) \neq 0$, replacing the value of $P_{\mathcal{M}_v}(u)(\xi_0)$ (resp. $P_{\mathcal{M}_v}(u)(-\xi_0)$) by $\widehat{v}(\xi_0)e^{i\phi}$ (resp. $\widehat{v}(\xi_0)e^{-i\phi}$) for any phase ϕ in $\mathbb{R}/2\pi\mathbb{Z}$ provides another point in $\arg \min_{s \in \mathcal{M}_v} \|s - u\|_2$. However, if u is assumed to be such that for all ξ in D , $\widehat{u}(\xi) \neq 0$, then for all v , $\arg \min_{s \in \mathcal{M}_v} \|s - u\|_2$ is a singleton.

Remark 3.4. Notice that $P_{\mathcal{M}_v}(\lambda u) = P_{\mathcal{M}_v}(u)$ for all $\lambda > 0$. However, $P_{\mathcal{M}_v}(\lambda u) = -P_{\mathcal{M}_v}(u)$ for all $\lambda < 0$.

Now, for all s in \mathbb{C}^D , let us write

$$m_s = \frac{1}{D} \sum_{x \in D} s(x) \quad (3.8)$$

the means of s and, for all x in D

$$\Gamma_s(x) = \frac{1}{|D|} \sum_{y \in D} (s(y) - m_s) \overline{(s(y+x) - m_s)} \quad (3.9)$$

the periodical autocorrelation of s ($y+x$ is defined in the finite group D). Recall that

$$\widehat{\Gamma}_s(\xi) = |\widehat{s}(\xi)|^2 \quad (3.10)$$

for all ξ in $D \setminus \{0\}$ and thus for all s in \mathbb{C}^D we have the following

$$\{s \in \mathbb{R}^D; |m_s| = |m_v| \text{ and } \Gamma_s(x) = \Gamma_v(x) (\forall x \in D)\} = \mathcal{M}_v. \quad (3.11)$$

3.2.2 The texton as an approximation of the Dirac mass

In Chapter 2, we defined a representation of Gaussian and Random Phase Noise textures u by defining its *texton* by

$$\widehat{T}(u) = |\widehat{u}|. \quad (3.12)$$

Recall that δ_0 denotes the Dirac at 0 is defined by $\delta_0(x) = 1$ if $x = \mathbf{0}$ and $\delta_0(x) = 0$ otherwise. Proposition 3.1 implies the following result.

Corollary 3.1. For all u

$$T(u) = P_{\mathcal{M}_u}(\delta_0). \quad (3.13)$$

Moreover, $T(u)$ is the only image in \mathbb{R}^D to satisfy

$$T(u) \in \arg \min_{s \in \mathcal{M}_u} \|s - \delta_0\|_2. \quad (3.14)$$

Proof. Simply notice that $\widehat{\delta}_0(\xi) = \frac{1}{\sqrt{D}} \neq 0$ for all ξ in D , so that

$$\frac{\widehat{\delta}_0(\xi)}{|\widehat{\delta}_0(\xi)|} = 1 \quad (3.15)$$

for all ξ in D .

For the uniqueness property, notice that $\widehat{\delta}_0(\xi) \neq 0$ for all ξ in D , so as mentioned in Remark 3.3, condition (3.5) is always satisfied, regardless of u . \square

This result can be linked to the fact that textons are “concentrated” around 0. Indeed, δ_0 is the most concentrated signal around 0 and $T(u)$ is the closest signal to δ_0 according to the $l^2(D)$ distance, with the constraint $\Gamma_{T(u)} = \Gamma_u$ and $|m_u| = |m_{T(u)}|$ (recall identity (3.11)). Notice that Remark 3.4 yields $T(u) = P_{\mathcal{M}_u}(\lambda\delta_0)$ for all $\lambda > 0$.

3.3 Four large error projections with constraints on the phase

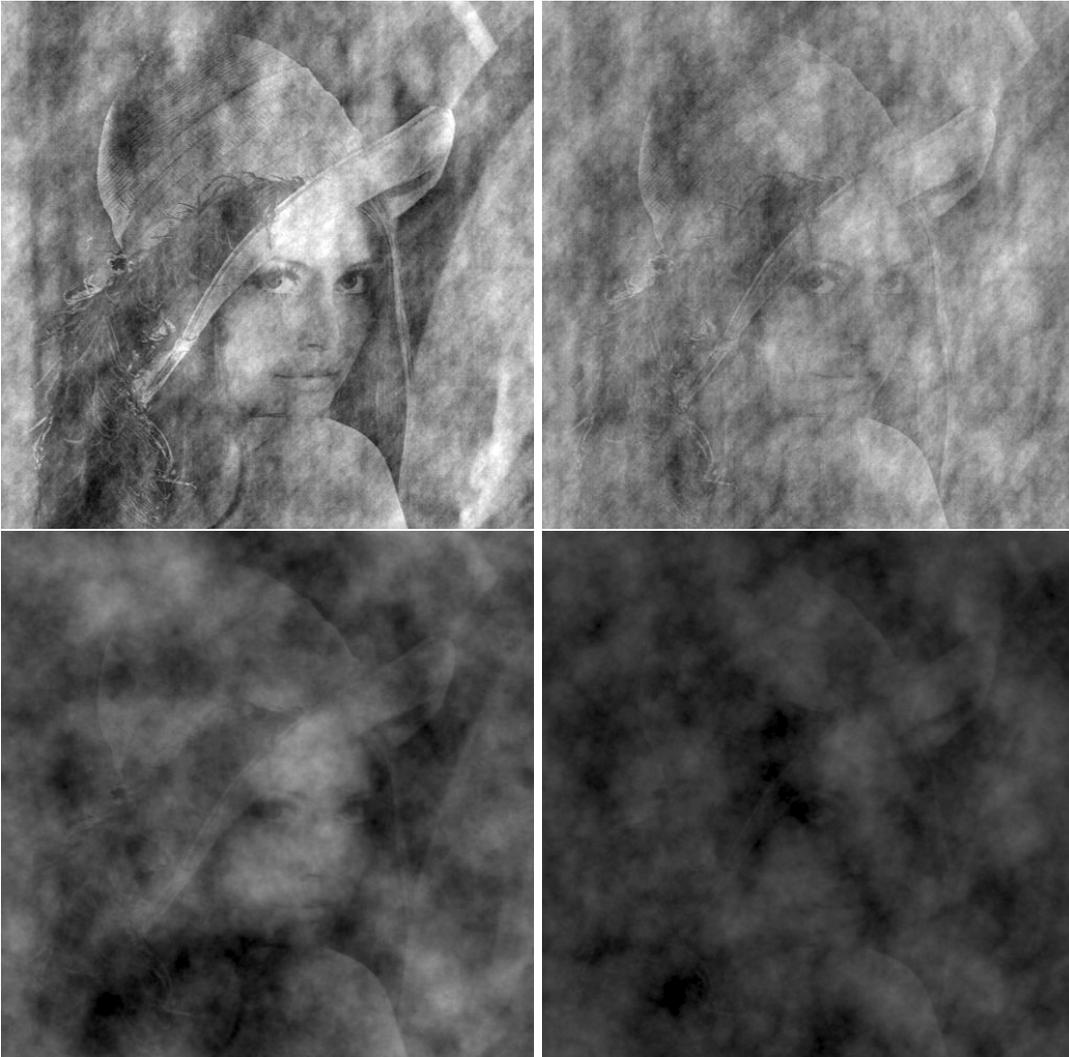


Figure 3.4: Images of the projections of “Lena” (as the target image u) onto different constraint sets defined by “clouds” (as v). In particular, all the images shown here have the phase of “clouds” (modulo π and 2π), but “Lena” is somewhat recognizable in each projection. Top left: $P_{\mathcal{D}_v^{(\pi)}}(u)$ (phase of “clouds” modulo π). Top right: $P_{\mathcal{D}_v^{(2\pi)}}(u)$ (phase of “clouds” modulo 2π). Bottom left: $P_{\mathcal{F}_v^{(\pi)}}(u)$ (phase of “clouds” modulo π , modulus of “clouds”). Bottom right: $P_{\mathcal{F}_v^{(2\pi)}}(u)$ (phase of “clouds” modulo π , modulus of “clouds” or zero).

In this section, we are interested in how well a given image can be approximated, while their phase is strongly constrained, typically by imposing them from a different image. The purpose is to reconsider and nuance the proverbial importance of phase in signals. Thus, we focus on LEP onto sets that are defined by the phase of its elements. We investigate the possibility of preserving the geometry of the target image despite such constraints. We shall first consider the projections onto the set of images which phase modulo π are given (*e.g.* taken from a given different image), and restrict this to the phase modulo 2π . As we shall illustrate in Figure 3.4 and Figure 3.6, these projections allow a fair recognition of the target image. Interestingly, under somewhat weak hypotheses, a small number of these projections allow for exact reconstruction of the entire target image.

Given an image v , we also investigate projections onto the set of images where the whole Fourier coefficients can be either kept identical or gotten rid of, and projections onto the set of images where the Fourier coefficients can only be kept identical or multiplied by -1 .

3.3.1 Projections with constraints on the phase

Let us consider the two following constraint sets

$$\mathcal{D}_v^{(\pi)} = \{s \in \mathbb{R}^D; \widehat{s}(\xi) \in \mathbb{R} \cdot \widehat{v}(\xi) (\forall \xi \in D)\} \quad (3.16)$$

$$\mathcal{D}_v^{(2\pi)} = \{s \in \mathbb{R}^D; \widehat{s}(\xi) \in \mathbb{R}_+ \cdot \widehat{v}(\xi) (\forall \xi \in D)\}. \quad (3.17)$$

Notice that $\mathcal{D}_v^{(2\pi)} \subset \mathcal{D}_v^{(\pi)}$. We are considering the LEA problems based on the target image u and the constraint sets $\mathcal{D}_v^{(\pi)}$ and $\mathcal{D}_v^{(2\pi)}$. These constraint sets are designed to preserve the phase of the image v . Notice that in $\mathcal{D}_v^{(\pi)}$ the constraint holds on the phase modulo π , whereas $\mathcal{D}_v^{(2\pi)}$ the constraint holds on the phase modulo 2π . Thus, in the corresponding projection $P_{\mathcal{D}_v^{(\pi)}}$ (*resp.* $P_{\mathcal{D}_v^{(2\pi)}}$) defined *infra*, the optimization is performed on the signed Fourier modulus of the optimized variable (*resp.* its unsigned modulus). We now derive the solutions of the LEP.

Proposition 3.2. *The set $\mathcal{D}_v^{(\pi)}$ is a linear subspace of \mathbb{R}^D and $\mathcal{D}_v^{(2\pi)}$ is a convex cone. Moreover*

$$\arg \min_{s \in \mathcal{D}_v^{(\pi)}} \|s - u\|_2 = \{P_{\mathcal{D}_v^{(\pi)}}(u)\} \quad (3.18)$$

and

$$\arg \min_{s \in \mathcal{D}_v^{(2\pi)}} \|s - u\|_2 = \{P_{\mathcal{D}_v^{(2\pi)}}(u)\} \quad (3.19)$$

where $P_{\mathcal{D}_v^{(\pi)}}(u)$ (*resp.* $P_{\mathcal{D}_v^{(2\pi)}}(u)$) is the orthogonal projections of u onto $\mathcal{D}_v^{(\pi)}$ (*resp.* $\mathcal{D}_v^{(2\pi)}$). They can be defined through their discrete Fourier transform by

$$\widehat{P_{\mathcal{D}_v^{(\pi)}}(u)}(\xi) = \mathbb{1}_{\{\widehat{v}(\xi) \neq 0\}} \operatorname{Re} \left(\widehat{u}(\xi) \overline{\widehat{v}(\xi)} \right) \frac{\widehat{v}(\xi)}{|\widehat{v}(\xi)|^2} \quad (3.20)$$

and

$$\widehat{P_{\mathcal{D}_v^{(2\pi)}}(u)}(\xi) = \mathbb{1}_{\{\operatorname{Re}(\widehat{u}(\xi) \overline{\widehat{v}(\xi)}) > 0\}} \operatorname{Re} \left(\widehat{u}(\xi) \overline{\widehat{v}(\xi)} \right) \frac{\widehat{v}(\xi)}{|\widehat{v}(\xi)|^2}. \quad (3.21)$$

Proof. Clearly $\mathcal{D}_v^{(\pi)}$ is a \mathbb{R} -linear subspace of \mathbb{C}^D as the image of a \mathbb{R} -linear subspace of \mathbb{C}^D through the inverse Fourier transform, which is \mathbb{R} -linear. The fact that $\mathcal{D}_v^{(\pi)}$ is a subset of \mathbb{R}^D is straightforward thanks to the symmetry condition. To prove that $\mathcal{D}_v^{(\pi)}$ is a convex cone, notice that it is an intersection of convex cones, namely \mathbb{R}^D and $\{s \in \mathbb{C}^D; \widehat{s}(\xi) \in \mathbb{R}_+ \cdot \widehat{v}(\xi) (\forall \xi \in D)\}$. The formulas for the Fourier transforms of the projections of u onto $\mathcal{D}_v^{(\pi)}$ (resp. $\mathcal{D}_v^{(2\pi)}$) $P_{\mathcal{D}_v^{(\pi)}}(u)$ (resp. $P_{\mathcal{D}_v^{(2\pi)}}(u)$) are simple consequence of Parseval identity, since

$$\arg \min_{s \in \mathcal{S}} \|s - u\|_2 = \arg \min_{s \in \mathcal{S}} \|\widehat{s} - \widehat{u}\|_2^2. \quad (3.22)$$

for any subset \mathcal{S} of \mathbb{C}^D .

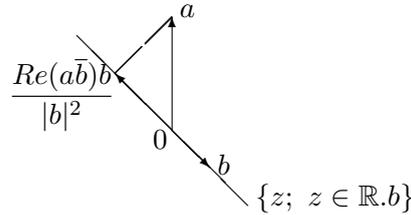


Figure 3.5: Orthogonal projection onto a line.

Now, as in the proof of Theorem 3.1, simply notice that for one fixed non-zero complex number a , and a complex non-zero number z in $\mathbb{R}.b$ (with $b \neq 0$), the quantity $|z - a|$ is minimal if and only if $z = \frac{Re(a\bar{b})b}{|b|^2}$, regardless of the sign of $Re(a\bar{b})$. Moreover, notice that for one fixed non-zero complex number a , and a complex non-zero number z in $\mathbb{R}_+.b$ (with $b \neq 0$), the quantity $|z - a|$ is minimal if and only if $z = \frac{Re(a\bar{b})b}{|b|^2}$ whenever $Re(a\bar{b}) > 0$ and $z = 0$ otherwise. See Figure 3.5 for an illustration. \square

Remark 3.5. *The phase of $P_{\mathcal{D}_v^{(2\pi)}}(u)$ is identical to the phase of v (modulo 2π), except on the subset of the Fourier domain $\{\xi \in D; Re(\widehat{u}(\xi)\overline{\widehat{v}(\xi)}) \leq 0\}$ where the phase is not properly defined. Notice that this set collects the points ξ in the Fourier domain where the complex numbers $\widehat{u}(\xi)$ and $\widehat{v}(\xi)$ form an obtuse angle in the complex plane (see Figure 3.5).*

3.3.2 Reconstruction formulas

We now state a straightforward result of exact reconstruction from the projections. The interest of such a result lies mostly in the comparison with the exact reconstruction algorithm that Oppenheim and Lim used to illustrate the importance of phase. For two images v_1 and v_2 , let us state the following hypothesis

$$(\mathcal{H}): \widehat{v}_1(\xi)\overline{\widehat{v}_2(\xi)} \notin \mathbb{R} \text{ for all } \xi \text{ in } D.$$

In other words, Hypothesis (\mathcal{H}) guarantees that for each spatial frequency ξ , the complex numbers $\widehat{v}_1(\xi)$ and $\widehat{v}_2(\xi)$ are \mathbb{R} -independent. Moreover, Hypothesis (\mathcal{H}) is quite easily satisfied, as illustrates the following result.

Proposition 3.3. *Let v_1 and v_2 be two images such that $\widehat{v}_1(\xi) \neq 0$ (resp. $\widehat{v}_2(\xi) \neq 0$) for all ξ . Define $RPN(v_2)$ the random phase noise image associated with v_2 as defined in [51]. Then, almost surely v_1 and $RPN(v_2)$ satisfy Hypothesis (\mathcal{H}) .*

Proof. For each ξ in D , recall that $\widehat{RPN(v_2)}(\xi) = \widehat{v_2}(\xi)e^{i\Phi(\xi)}$ where $\Phi(\xi)$ is a random variable uniformly distributed on the circle. Thus, for each ξ , $\widehat{RPN(v_2)}(\xi)$ and $\widehat{v}_1(\xi)$ are \mathbb{R} -independent with probability 1. Since there is a finite number of such ξ , the conclusion follows. \square

We now state a result of exact reconstruction based on the projections onto sets of the form $\mathcal{D}_v^{(\pi)}$.

Proposition 3.4. *Let v_1 and v_2 be such that they satisfy Hypothesis (\mathcal{H}) . Then, every image u can be exactly reconstructed from v_1 , v_2 , $P_{\mathcal{D}_{v_1}^{(\pi)}}(u)$ and $P_{\mathcal{D}_{v_2}^{(\pi)}}(u)$ by*

$$\widehat{u}(\xi) = \alpha_1(\xi)\widehat{v}_1(\xi) + \alpha_2(\xi)\widehat{v}_2(\xi) \quad (3.23)$$

where $\alpha_1(\xi)$ and $\alpha_2(\xi)$ are defined by

$$\begin{cases} \alpha_1(\xi) = \frac{|\widehat{v}_2(\xi)|^2 P_{\mathcal{D}_{v_1}^{(\pi)}}(u)(\xi)\widehat{v}_1(\xi) - \operatorname{Re}(\widehat{v}_1(\xi)\widehat{v}_2(\xi))P_{\mathcal{D}_{v_2}^{(\pi)}}(u)(\xi)\widehat{v}_2(\xi)}{|\widehat{v}_1(\xi)|^2|\widehat{v}_2(\xi)|^2 - \operatorname{Re}(\widehat{v}_1(\xi)\widehat{v}_2(\xi))^2} \\ \alpha_2(\xi) = \frac{|\widehat{v}_1(\xi)|^2 P_{\mathcal{D}_{v_2}^{(\pi)}}(u)(\xi)\widehat{v}_2(\xi) - \operatorname{Re}(\widehat{v}_2(\xi)\widehat{v}_1(\xi))P_{\mathcal{D}_{v_1}^{(\pi)}}(u)(\xi)\widehat{v}_1(\xi)}{|\widehat{v}_1(\xi)|^2|\widehat{v}_2(\xi)|^2 - \operatorname{Re}(\widehat{v}_1(\xi)\widehat{v}_2(\xi))^2}. \end{cases}$$

Proof. Hypothesis (\mathcal{H}) ensures that for each ξ , $\widehat{v}_1(\xi)$ and $\widehat{v}_2(\xi)$ are \mathbb{R} -independent vectors of \mathbb{R}^2 . Thus, for each ξ there are a unique couple $(\alpha_1(\xi), \alpha_2(\xi))$ in \mathbb{R}^2 such that $\widehat{u}(\xi) = \alpha_1(\xi)\widehat{v}_1(\xi) + \alpha_2(\xi)\widehat{v}_2(\xi)$. Let us show that $\alpha_1(\xi)$ and $\alpha_2(\xi)$ solve (uniquely) the \mathbb{R} -linear system

$$\begin{cases} \alpha_1(\xi)|\widehat{v}_1(\xi)|^2 + \alpha_2(\xi)\operatorname{Re}(\widehat{v}_1(\xi)\widehat{v}_2(\xi)) = P_{\mathcal{D}_{v_1}^{(\pi)}}(u)(\xi)\widehat{v}_1(\xi) \\ \alpha_1(\xi)\operatorname{Re}(\widehat{v}_1(\xi)\widehat{v}_2(\xi)) + \alpha_2(\xi)|\widehat{v}_2(\xi)|^2 = P_{\mathcal{D}_{v_2}^{(\pi)}}(u)(\xi)\widehat{v}_2(\xi). \end{cases} \quad (3.24)$$

Indeed, recall that by definition $P_{\mathcal{D}_{v_1}^{(\pi)}}(u)(\xi) = \frac{\operatorname{Re}(\widehat{u}(\xi)\widehat{v}_1(\xi))}{|\widehat{v}_1(\xi)|^2}\widehat{v}_1(\xi)$. Hence, on the one hand

$$\operatorname{Re}(\widehat{u}(\xi)\widehat{v}_1(\xi)) = P_{\mathcal{D}_{v_1}^{(\pi)}}(u)(\xi)\widehat{v}_1(\xi)$$

and on the other hand

$$\operatorname{Re}(\widehat{u}(\xi)\widehat{v}_1(\xi)) = \alpha_1(\xi)|\widehat{v}_1(\xi)|^2 + \alpha_2(\xi)\operatorname{Re}(\widehat{v}_1(\xi)\widehat{v}_2(\xi))$$

which yields the first equation of the system. The second equation follows by symmetry. Notice that the system is always determinate since $|\widehat{v}_1|^2|\widehat{v}_2|^2 - \operatorname{Re}(\widehat{v}_1\widehat{v}_2) > 0$ thanks to Hypothesis (\mathcal{H}) . The formulas for $\alpha_1(\xi)$ and $\alpha_2(\xi)$ follow by inverting the 2×2 matrix of the system. \square

We can now prove a similar result of exact reconstruction based on the projections onto sets of the form $\mathcal{D}_v^{(2\pi)}$.

Proposition 3.5. *Let v_1 and v_2 be such that they satisfy Hypothesis (\mathcal{H}). Then, every image u can be exactly reconstructed from $v_1, v_2, w_1^+ = P_{\mathcal{D}_{v_1}^{(2\pi)}}(u), w_1^- = P_{\mathcal{D}_{-v_1}^{(2\pi)}}(u), w_2^+ = P_{\mathcal{D}_{v_2}^{(2\pi)}}(u)$ and $w_2^- = P_{\mathcal{D}_{-v_2}^{(2\pi)}}(u)$ by*

$$\widehat{u}(\xi) = \beta_1(\xi)\varepsilon_1(\xi)\widehat{v}_1(\xi) + \beta_2(\xi)\varepsilon_2(\xi)\widehat{v}_2(\xi) \quad (3.25)$$

where $\varepsilon_1(\xi) = +$ (resp. $\varepsilon_2(\xi) = +$) if $w_1^+ \neq 0$ (resp. if $w_2^+ \neq 0$) and $\varepsilon_1(\xi) = -$ (resp. $\varepsilon_2(\xi) = -$) otherwise, $\beta_1(\xi)$ and $\beta_2(\xi)$ solve (uniquely) the \mathbb{R} -linear system

$$\begin{cases} \beta_1(\xi) = \frac{|\widehat{v}_2(\xi)|^2 \widehat{w_1^{\varepsilon_1(\xi)}}(\xi) \overline{\widehat{v}_1(\xi)} - \varepsilon_1(\xi)\varepsilon_2(\xi) \operatorname{Re}(\widehat{v}_1(\xi)\overline{\widehat{v}_2(\xi)}) \widehat{w_2^{\varepsilon_2(\xi)}}(\xi) \overline{\widehat{v}_2(\xi)}}{|\widehat{v}_1(\xi)|^2 |\widehat{v}_2(\xi)|^2 - \operatorname{Re}(\widehat{v}_1(\xi)\overline{\widehat{v}_2(\xi)})^2} \\ \beta_2(\xi) = \frac{|\widehat{v}_1(\xi)|^2 \widehat{w_2^{\varepsilon_2(\xi)}}(\xi) \overline{\widehat{v}_2(\xi)} - \varepsilon_1(\xi)\varepsilon_2(\xi) \operatorname{Re}(\widehat{v}_2(\xi)\overline{\widehat{v}_1(\xi)}) \widehat{w_1^{\varepsilon_1(\xi)}}(\xi) \overline{\widehat{v}_1(\xi)}}{|\widehat{v}_1(\xi)|^2 |\widehat{v}_2(\xi)|^2 - \operatorname{Re}(\widehat{v}_1(\xi)\overline{\widehat{v}_2(\xi)})^2}. \end{cases}$$

Proof. Let us prove that for each ξ , $\beta_1(\xi)$ and $\beta_2(\xi)$ solve (uniquely) the \mathbb{R} -linear system

$$\begin{cases} \beta_1(\xi)|\widehat{v}_1(\xi)|^2 + \beta_2(\xi)\varepsilon_1(\xi)\varepsilon_2(\xi) \operatorname{Re}(\widehat{v}_1(\xi)\overline{\widehat{v}_2(\xi)}) = \varepsilon_1(\xi) \widehat{w_1^{\varepsilon_1(\xi)}}(\xi) \overline{\widehat{v}_1(\xi)} \\ \beta_1(\xi)\varepsilon_1(\xi)\varepsilon_2(\xi) \operatorname{Re}(\widehat{v}_1(\xi)\overline{\widehat{v}_2(\xi)}) + \beta_2(\xi)|\widehat{v}_2(\xi)|^2 = \varepsilon_2(\xi) \widehat{w_2^{\varepsilon_2(\xi)}}(\xi) \overline{\widehat{v}_2(\xi)}. \end{cases}$$

First, notice that $\widehat{w_j^+}(\xi) + \widehat{w_j^-}(\xi) = \varepsilon_j(\xi)\widehat{w_j}(\xi)$ and that $\widehat{w_j^+}(\xi)$ and $\widehat{w_j^-}(\xi)$ cannot be non-zero at the same time. The remaining of the proof is identical to the proof of Proposition 3.4. \square

Let us compare these results to the result of reconstruction obtained by Oppenheim and Lim [106]. The authors of the latter article describe an algorithm that allows exact reconstruction based on the information of the phase only, under the hypothesis that a large number of coefficients of the image (actually 75%) are zeroes. In other words, a phase-only image, or an image with inaccurate modulus, allows for perfect reconstruction of roughly 25% of the image. Comparatively, according to Proposition 3.4, projections onto phase cones of the form $\mathcal{D}_v^{(\pi)}$ allow for reconstruction of 25% of the signal, since four images of the dimension of u are needed. According to Proposition 3.5, projections onto phase cones of the form $\mathcal{D}_v^{(2\pi)}$ allow for reconstruction of one sixth of the signal. Actually these estimations turn out to be quite conservative. Indeed, under the hypothesis that for each ξ in $D \setminus \{0\}$ both $\operatorname{Re}(\widehat{v}_1(\xi)\widehat{u}(\xi)) \neq 0$ and $\operatorname{Re}(\widehat{v}_2(\xi)\widehat{u}(\xi)) \neq 0$, the information of $\mathcal{D}_{v_1}^{(\pi)}$ and $\mathcal{D}_{v_2}^{(\pi)}$ (resp. $\mathcal{D}_{v_1}^{(2\pi)}, \mathcal{D}_{-v_1}^{(2\pi)}, \mathcal{D}_{v_2}^{(2\pi)}$ and $\mathcal{D}_{-v_2}^{(2\pi)}$) are sufficient for exact reconstruction. These hypothesis are quite easily satisfied as they are clearly satisfied by two independent RPN textures V_1 and V_2 . Under these more favorable hypotheses, projections onto phase cones of the form $\mathcal{D}_v^{(\pi)}$ allow for reconstruction of **one half** of the signal and projections onto phase cones of the form $\mathcal{D}_v^{(2\pi)}$ allow for reconstruction of **one quarter** of the signal, which compares favorably with the Oppenheim-Lim reconstruction framework.

3.3.3 Constraints on both the phase and the modulus

The relative quality of the results of the projections onto $\mathcal{D}_v^{(\pi)}$ and $\mathcal{D}_v^{(2\pi)}$ (see discussion in the next subsection along with Figure 3.4 and Figure 3.6) led us to investigate the conservation

of shapes of the target image u through the projections onto smaller sets. In particular we wonder whether one can constrain both the phase and the modulus through a projection and still recognize the target image u . Hence, we now consider the following constraint sets

$$\mathcal{F}_v^{(\pi)} := \{s \in \mathbb{R}^D; \widehat{s}(\xi) \in \{\widehat{v}(\xi), -\widehat{v}(\xi)\} (\forall \xi \in D)\} \quad (3.26)$$

$$\mathcal{F}_v^{(2\pi)} := \{s \in \mathbb{R}^D; \widehat{s}(\xi) \in \{\widehat{v}(\xi), 0\} (\forall \xi \in D)\}. \quad (3.27)$$

Notice that $\mathcal{F}_v^{(\pi)} \subset \mathcal{D}_v^{(\pi)}$ and $\mathcal{F}_v^{(2\pi)} \subset \mathcal{D}_v^{(2\pi)}$.

Notice furthermore that both sets $\mathcal{F}_v^{(\pi)}$ and $\mathcal{F}_v^{(2\pi)}$ are finite sets. It is straightforward to compute the cardinalities of $\mathcal{F}_v^{(\pi)}$ and $\mathcal{F}_v^{(2\pi)}$ along with the topological dimensions of $\mathcal{D}_v^{(\pi)}$ and $\mathcal{D}_v^{(2\pi)}$, which satisfy

$$|\mathcal{F}_v^{(\pi)}| = |\mathcal{F}_v^{(2\pi)}| = 2^{\dim(\mathcal{D}_v^{(\pi)})} = 2^{\dim(\mathcal{D}_v^{(2\pi)})} = 2^{|\{\xi \in D \setminus \{0\}; \widehat{v}(\xi) \neq 0\}|/2}. \quad (3.28)$$

Notice that $|\mathcal{F}_v^{(\pi)}| = |\mathcal{F}_v^{(2\pi)}| < 2^D$, which is the size of the set of binary images over D . Thus, projecting onto $\mathcal{F}_v^{(\pi)}$ and $\mathcal{F}_v^{(2\pi)}$ allows for even fewer bits than the binary projection. We now turn to the solution of the optimization problem for $\mathcal{F}_v^{(\pi)}$ and $\mathcal{F}_v^{(2\pi)}$.

Proposition 3.6.

$$\arg \min_{s \in \mathcal{F}_v^{(\pi)}} \|s - u\|_2 = \{w_{\mathcal{F}^{(\pi)}}\} \quad (3.29)$$

and

$$\arg \min_{s \in \mathcal{F}_v^{(2\pi)}} \|s - u\|_2 = \{w_{\mathcal{F}^{(2\pi)}}\}. \quad (3.30)$$

where $w_{\mathcal{F}^{(\pi)}}$ and $w_{\mathcal{F}^{(2\pi)}}$ are defined through their discrete Fourier transform by

$$\widehat{w_{\mathcal{F}^{(\pi)}}}(\xi) = \operatorname{sgn}(\operatorname{Re}(\widehat{u}(\xi)\overline{\widehat{v}(\xi)}))\widehat{v}(\xi) \quad (3.31)$$

and

$$\widehat{w_{\mathcal{F}^{(2\pi)}}}(\xi) = \mathbf{1}_{\{2\operatorname{Re}(\widehat{u}(\xi)\overline{\widehat{v}(\xi)}) > |\widehat{v}(\xi)|^2\}}\widehat{v}(\xi). \quad (3.32)$$

Proof. The proof is almost identical to the proof of Proposition 3.2, with different set of constraints on the frequency domain. Simply notice that $|a - b| \leq |b|$ if and only if $|a|^2 \leq 2\operatorname{Re}(a\bar{b})$ and $|a - b| \leq |-a - b|$ if and only if $\operatorname{Re}(a\bar{b}) \geq 0$. \square

3.3.4 Experiments

In Figure 3.4, we can observe that all of the projections $P_{\mathcal{D}_v^{(\pi)}}(u)$, $P_{\mathcal{D}_v^{(2\pi)}}(u)$, $P_{\mathcal{F}_v^{(\pi)}}(u)$ and $P_{\mathcal{F}_v^{(2\pi)}}(u)$ exhibit shapes from the target image u . This is surprising since these images are designed to have the phase of v , an image with radically different shapes than u . More precisely, in this figure the phase of v is chosen randomly like in a Gaussian or Random Phase Noise texture – see [51]. Arguably, the degree of recognition grows with the size (or the inclusion relationships) of the sets onto which “Lena” is projected: $P_{\mathcal{D}_v^{(\pi)}}(u)$ is the best approximation of u , followed by $P_{\mathcal{D}_v^{(2\pi)}}(u)$, $P_{\mathcal{F}_v^{(\pi)}}(u)$ and $P_{\mathcal{F}_v^{(2\pi)}}(u)$.

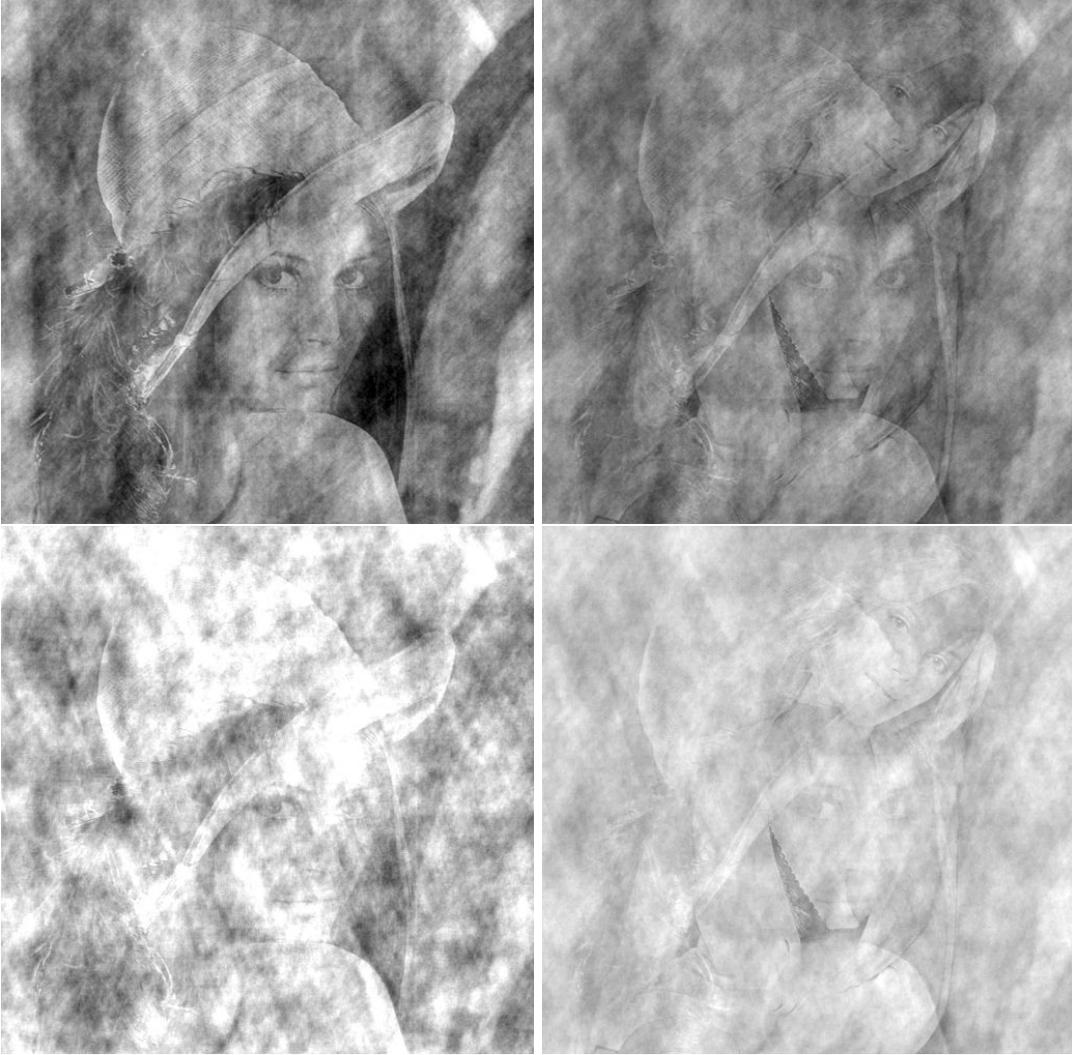


Figure 3.6: Images of the projections of “Lena” (as the target image u) onto different constraint sets defined by “Scarlett Johansson” (as v). In particular, all the images shown here have the phase of “Scarlett Johansson” (modulo 2π or π), but “Lena” is always somewhat recognizable. Top left: $P_{\mathcal{D}_v^{(\pi)}}(u)$ (phase of “Scarlett Johansson” modulo π , projected modulus of “Lena”). Top right: $P_{\mathcal{D}_v^{(2\pi)}}(u)$ (phase of “Scarlett Johansson” modulo 2π , projected modulus of “Lena”). Bottom left: $P_{\mathcal{F}_v^{(\pi)}}(u)$ (phase of “Scarlett Johansson” modulo π , modulus of “Lena”). Bottom right: $P_{\mathcal{F}_v^{(2\pi)}}(u)$ (phase of “Scarlett Johansson” modulo 2π , modulus of “Lena” or zero).

In Figure 3.6, we can also observe that $P_{\mathcal{D}_v^{(\pi)}}(u)$, $P_{\mathcal{D}_v^{(2\pi)}}(u)$, $P_{\mathcal{F}_v^{(\pi)}}(u)$ and $P_{\mathcal{F}_v^{(2\pi)}}(u)$ exhibit shapes from the target image u . However, both $P_{\mathcal{D}_v^{(2\pi)}}(u)$ and $P_{\mathcal{F}_v^{(2\pi)}}(u)$ concurrently exhibit shapes from the “Scarlett Johansson” image v that defines the constraint sets $\mathcal{D}_v^{(2\pi)}$ and $\mathcal{F}_v^{(2\pi)}$. Hence, when the phase of v is chosen to be kept modulo 2π , its shapes still appear in the projections onto $\mathcal{D}_v^{(2\pi)}$ (resp. $\mathcal{F}_v^{(2\pi)}$). Interestingly, the imprints of original images u and v in $P_{\mathcal{D}_v^{(2\pi)}}(u)$ and $P_{\mathcal{F}_v^{(2\pi)}}(u)$ seem to be comparable on the perceptual level. Let us also notice

that these experiments highlight that the phase taken modulo π alone is quite irrelevant for image approximations *à la* Oppenheim and Lim.

3.4 The importance of the Fourier transform

In this section, we are interested in the role of the Fourier transform in the LEA/LEP experiments discussed in previous sections. We first discuss an experiment where the Fourier transform is replaced by a unitary transform chosen at random. We then address the specificity of the Fourier transform by highlighting the relationship of the Fourier transform with respect to derivation operators.

3.4.1 Replacing the Fourier transform: experiments

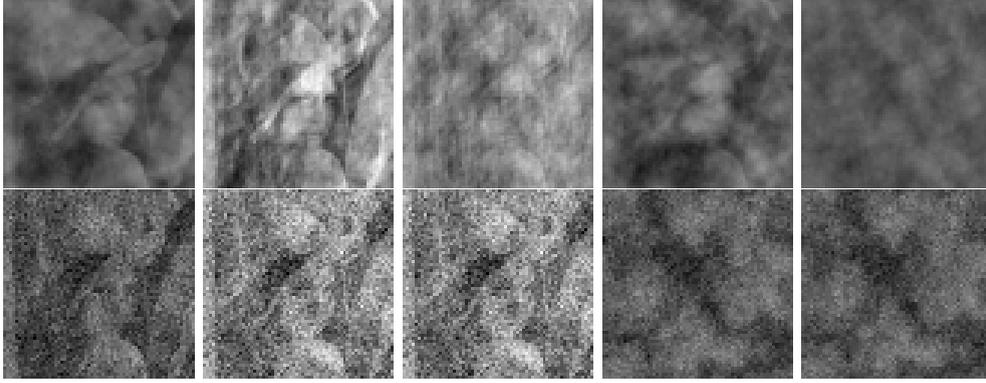


Figure 3.7: Images of the projections of “Lena” (as u) onto sets defined by “clouds” (as v). Top row: projections onto sets that are defined by the Fourier transform – from left to right \mathcal{M}_v , $\mathcal{D}_v^{(\pi)}$, $\mathcal{D}_v^{(2\pi)}$, $\mathcal{F}_v^{(\pi)}$ and $\mathcal{F}_v^{(2\pi)}$. Bottom row: the real part of their equivalent where the Fourier transform is replaced by a random unitary transform U .

Let u and v be two images (in \mathbb{R}^D). For any unitary transform U over D ($U : \mathbb{C}^D \rightarrow \mathbb{C}^D$ such that $U^*U = I_D$), the Fourier transform can be replaced by U in the sets \mathcal{M}_v , $\mathcal{D}_v^{(\pi)}$, $\mathcal{D}_v^{(2\pi)}$, $\mathcal{F}_v^{(\pi)}$ and $\mathcal{D}_v^{(2\pi)}$. More precisely, we define the sets

$$\mathcal{M}_{U,v} = \{s \in \mathbb{C}^D; |Us(\xi)| = |Uv(\xi)| \text{ for all } \xi \text{ in } D\} \quad (3.33)$$

$$\mathcal{D}_{U,v}^{(\pi)} = \{s \in \mathbb{C}^D; Us(\xi) \in Uv(\xi) \cdot \mathbb{R} \text{ for all } \xi \text{ in } D\} \quad (3.34)$$

$$\mathcal{D}_{U,v}^{(2\pi)} = \{s \in \mathbb{C}^D; Us(\xi) \in Uv(\xi) \cdot \mathbb{R}_+ \text{ for all } \xi \text{ in } D\} \quad (3.35)$$

$$\mathcal{F}_{U,v}^{(\pi)} = \{s \in \mathbb{C}^D; Us(\xi) \in \{Uv(\xi), -Uv(\xi)\} \text{ for all } \xi \text{ in } D\} \quad (3.36)$$

and

$$\mathcal{F}_{U,v}^{(2\pi)} = \{s \in \mathbb{C}^D; Us(\xi) \in \{Uv(\xi), 0\} \text{ for all } \xi \text{ in } D\} \quad (3.37)$$

and consider the optimization problem

$$\arg \min_{s \in \mathcal{S}} \|s - u\|_2$$

for each of these sets. Solutions are easily computed based on Propositions 3.1, 3.2 and 3.6. Let us remark that without further assumption on U , images in the sets have no reason to belong to \mathbb{R}^D . However, for the sake of the experiment, we compute the real parts of these images in \mathbb{C}^D . Let us notice that the imaginary parts seem to convey no particular visual information distinguishable from noise.

We choose U randomly with the Haar probability measure over the compact group $\mathcal{U}(|D|)$ of unitary matrices of size $|D|$ (recall that $|D|$ denotes the cardinality of D). Let us recall that $\mathcal{U}(|D|)$ is a compact topological group and thus admits a single probability law, the Haar probability law P_{Haar} , such that for each measurable subset \mathcal{S} of $\mathcal{U}(|D|)$, $P_{\text{Haar}}(U\mathcal{S}) = P_{\text{Haar}}(\mathcal{S})$ for all U in $\mathcal{U}(|D|)$. A simple and well known way to get unitary matrices randomly according to the Haar measure is described in [40] and is based on the Gram-Schmidt algorithm. More precisely, it can be proven that the matrix obtained by

1. drawing a square matrix N of size D with *i.i.d.* centered complex Gaussian entries with a non-zero variance $\sigma > 0$
2. performing the Gram-Schmidt Hermitian orthonormalization algorithm upon N

follows the Haar probability over $\mathcal{U}(|D|)$. Unfortunately, this algorithm does not scale very well for a large number of entries, so limited ourselves to 64×64 images – which already requires computing the Hermitian orthonormalization of a 4096×4096 complex matrix.

Let us comment the results of the experiment as shown in Figure 3.7. The projection onto the random unitary set $\mathcal{M}_{U,v}$ appears to be more noisy than the projection onto \mathcal{M}_v . However, the extent to which “Lena” can be recognized in the (real part of the) projection is somehow comparable in both cases. The same observations hold for the set $\mathcal{D}_v^{(\pi)}$ and its random unitary counterpart $\mathcal{D}_{U,v}^{(\pi)}$. However, in the case of the random unitary equivalents of $\mathcal{F}_v^{(\pi)}$ and $\mathcal{F}_v^{(2\pi)}$, which are finite sets, “Lena” cannot be distinguished in the results of the projections.

3.4.2 Fourier transformation and the gradient

The gradient of an image plays a fundamental role in computer vision, *e.g.* for edge and contour detection (see Canny [23] and Kass [72]). A point with a high gradient norm is likely to be sampled from an image that exhibits an edge in the direction orthogonal to the direction of the gradient.

We consider the periodical partial derivatives of an image $\partial_j s$ (for $j = 1$ or 2) defined by $\partial_1 s(x_1, x_2) = s(x_1 + 1, x_2) - s(x_1, x_2)$ (resp. $\partial_2 s(x_1, x_2) = s(x_1, x_2 + 1) - s(x_1, x_2)$) for all (x_1, x_2) in the domain D (recall that we consider D as the group $\mathbb{Z}/N_1\mathbb{Z} \times \mathbb{Z}/N_2\mathbb{Z}$). For $j = 1$ or 2 , let us consider the following LEP

$$\arg \min_{s \in \mathcal{S}} \|\partial_j s - \partial_j u\|_2 \quad (3.38)$$

for each of the Fourier-based constraint sets \mathcal{S} that we have defined *supra*. We have the following result.

Proposition 3.7. For all $j = 1$ or 2 , and for any constraint set \mathcal{S} in $\{\mathcal{M}_v, \mathcal{D}_v^{(\pi)}, \mathcal{D}_v^{(2\pi)}, \mathcal{F}_v^{(\pi)}, \mathcal{F}_v^{(2\pi)}\}$ the following holds

$$P_{\mathcal{S}}(u) \in \arg \min_{s \in \mathcal{S}} \|\partial_j s - \partial_j u\|_2. \quad (3.39)$$

Proof. First, recall that for each j in $\{1, 2\}$,

$$\widehat{\partial_j s}(\xi) = 2i \sin(\pi \xi_j / N_j) \widehat{s}(\xi) \quad (3.40)$$

for all ξ in D . From here, the argument is very similar to the proof of Propositions 3.1, 3.2 and 3.6. Parseval identity yields that $\|\partial_j s - \partial_j u\|_2 = \|\widehat{\partial_j s} - \widehat{\partial_j u}\|_2$ and thus

$$\begin{aligned} \arg \min_{s \in \mathcal{S}} \|\partial_j s - \partial_j u\|_2 &= \arg \min_{s \in \mathcal{S}} \sum_{\xi \in D} |\widehat{\partial_j s}(\xi) - \widehat{\partial_j u}(\xi)|^2 = \arg \min_{s \in \mathcal{S}} \sum_{\xi \in D} |2i \sin(\pi \xi_j / N_j) \widehat{s}(\xi) - 2i \sin(\pi \xi_j / N_j) \widehat{u}(\xi)|^2 \\ &= \arg \min_{s \in \mathcal{S}} \sum_{\xi \in D} |2 \sin(\pi \xi_j / N_j)|^2 |\widehat{s}(\xi) - \widehat{u}(\xi)|^2 \end{aligned} \quad (3.41)$$

for any constraint set \mathcal{S} in $\{\mathcal{M}_v, \mathcal{D}_v^{(\pi)}, \mathcal{D}_v^{(2\pi)}, \mathcal{F}_v^{(\pi)}, \mathcal{F}_v^{(2\pi)}\}$.

The arguments that prove that $P_{\mathcal{S}}(u) \in \arg \min_{s \in \mathcal{S}} \|\widehat{s} - \widehat{u}\|_2$ in Propositions 3.1, 3.2 and 3.6 also prove that $P_{\mathcal{S}}(u) \in \arg \min_{s \in \mathcal{S}} \|\partial_j s - \partial_j u\|_2$. \square

Notice that, as opposed to Propositions 3.1, 3.2 and 3.6, $\arg \min_{s \in \mathcal{S}} \|\partial_j s - \partial_j u\|_2$ is never a singleton. Indeed, for any constant signal s_0 , $\partial_j s_0(x) = 0$ for all x in D , so one can add a non-zero constant image s_0 to s while the distance $\|\partial_j(\widehat{s + s_0}) - \widehat{\partial_j u}\|_2$ remains unchanged.

Let us define the gradient operator ∇ classically by $\nabla s = (\partial_1 s, \partial_2 s) : D \rightarrow \mathbb{R}^2$.

Corollary 3.2. For any constraint set \mathcal{S} in $\{\mathcal{M}_v, \mathcal{D}_v^{(\pi)}, \mathcal{D}_v^{(2\pi)}, \mathcal{F}_v^{(\pi)}, \mathcal{F}_v^{(2\pi)}\}$ the following holds

$$P_{\mathcal{S}}(u) \in \arg \min_{s \in \mathcal{S}} \|\nabla s - \nabla u\|_2. \quad (3.42)$$

Proof. First, notice that

$$\arg \min_{s \in \mathcal{S}} \|\nabla s - \nabla u\|_2 = \arg \min_{s \in \mathcal{S}} \sum_{\xi \in D} \sum_{j=1}^2 |\partial_j s(\xi) - \partial_j u(\xi)|^2. \quad (3.43)$$

To conclude, recall that

$$P_{\mathcal{S}}(u) \in \arg \min_{s \in \mathcal{S}} \sum_{\xi \in D} |\partial_j s(\xi) - \partial_j u(\xi)|^2 \quad (3.44)$$

for each j in $\{1, 2\}$ and thus

$$\begin{aligned} P_{\mathcal{S}}(u) &\in \arg \min_{s \in \mathcal{S}} \sum_{j=1}^2 \sum_{\xi \in D} |\partial_j s(\xi) - \partial_j u(\xi)|^2 = \arg \min_{s \in \mathcal{S}} \sum_{\xi \in D} \sum_{j=1}^2 |\partial_j s(\xi) - \partial_j u(\xi)|^2 \\ &= \arg \min_{s \in \mathcal{S}} \|\nabla s - \nabla u\|_2^2 = \arg \min_{s \in \mathcal{S}} \|\nabla s - \nabla u\|_2. \end{aligned} \quad (3.45)$$

\square

The validity of the inclusion

$$\arg \min_{s \in \mathcal{S}} \|s - u\| \subset \arg \min_{s \in \mathcal{S}} \|\nabla s - \nabla u\| \quad (3.46)$$

is clearly not true for all sets \mathcal{S} . In particular, the fact that the Fourier operator diagonalizes the gradient operator is crucial in the proof of Corollary 3.2. Notice that the argument in the proof of Proposition 3.7 still holds if one replaces the Fourier transform by any unitary transform that diagonalizes the translations, such as *e.g.* F_{k_1, k_2} defined for any k_1 in $\{1, \dots, N_1\}$ and k_2 in $\{1, \dots, N_2\}$ by

$$(F_{k_1, k_2} u)(\xi_1, \xi_2) = \frac{1}{\sqrt{N_1 N_2}} \sum_{0 \leq x_1 < N_1} \sum_{0 \leq x_2 < N_2} u(x_1, x_2) e^{i(\frac{k_1 x_1}{N_1} + \frac{k_2 x_2}{N_2})2\pi} \quad (3.47)$$

for all u in \mathbb{C}^D . It is also worth noting that the gradient operator in Corollary 3.2 can be replaced by other edge detector such as a smoothed gradient, since the smooth operator is a convolution and can thus also be expressed diagonally in the Fourier domain.

3.5 Concluding remarks

3.5.1 The l^2 norm

Let us briefly discuss the role of the $l^2(D)$ norm in LEP problems (3.1). The point of this subsection is to show that the $l^2(D)$ norm that has been used so far in our discussion on LEP can be replaced by other norms while conserving visually compelling results. First, notice that many other norms than the canonical $l^2(D)$ norm allow to formulate problems that are equivalent to (3.1). Indeed, for all real $p \geq 1$, changing the $l^2(D)$ norm with the norm $\nu_p : u \mapsto \|\hat{u}\|_p$ shall not change the solutions of the LEP discussed in Sections 2 and 3.

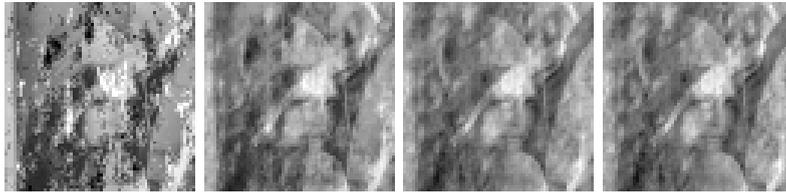


Figure 3.8: Images of the solutions of Problem (3.48) with various choices of parameter p , “Lena” (as u) and “clouds” (as v). From left to right, respectively l^1 , $l^{3/2}$, $l^{5/2}$ and l^3 minimisations.

We now briefly discuss variants of the orthogonal projection onto $\mathcal{D}_v^{(\pi)}$. More precisely, for any choice of $p \geq 1$, we consider the set

$$\arg \min_{s \in \mathcal{D}_v^{(\pi)}} \|s - u\|_p. \quad (3.48)$$

Unless $p = 2$, solutions to Problem (3.48) are not *a priori* easily solved by a formula. However, since the constraint set is a \mathbb{R} -linear sub-space of the \mathbb{R} -linear space \mathbb{C}^D , Problem (3.48) is

a convex optimization problem where algorithmic solutions are well known. Note that a non-convex projection set such as \mathcal{M}_v would yield a much harder problem.

Figure 3.8 shows the results of Problem (3.48), with “Lena” as target (u) and a constant set \mathcal{D}_v defined by v chosen to be the “clouds” image. Interestingly, the choice of p appears to be connected to the nature of the noise of the projected image. For instance, the l^1 optimization that has been documented to lead to results close to l^0 (support) optimization, produces a noise that leaves whole areas identical from the projection to the target image and other areas at extreme values.

3.5.2 LEA and LEP

In this chapter, we have mostly been focusing on LEP solutions for LEA problems. However, LEP might not be appropriate to find visually compelling approximations of images. For instance, let u be some image (*e.g.* “Lena”) with real values, m_u be the empirical mean of u , σ_u the distance to the constant image m_u , $\sigma_u = \|u - m_u\|_2$. Consider the set $\mathcal{S} = \mathbb{R} \cdot \mathbf{1}_D \cup \{s \in \mathbb{R}^D; \|s - u\|_2 > \sigma_u\}$. Clearly, \mathcal{S} is a connected set and $\arg \min_{s \in \mathcal{S}_\varepsilon} \|s - u\|_2 = \{m_u\}$, so the solution to the LEP defined by \mathcal{S} is a constant image. However, for any $\varepsilon > 0$, $u + \sigma_u + \varepsilon$ also belongs to \mathcal{S} and shall always reproduce the shapes from u better than a constant image. Thus, the LEA problem of finding an image in \mathcal{S} that is “faithful” to u is not solved by solving the associated LEP.

This example raises broader questions about the properties of sets that contain faithful approximations of a wide range target images. Indeed, checking that such approximations exist within a given set becomes a much harder task if it cannot be accomplished by simply checking that the closest approximation of a target image (the projection) is faithful. Thus we are left with two puzzling questions. First, when can the existence check of a LEA in a set \mathcal{S} be done through the LEP onto \mathcal{S} ? Second, how can such an approximation be found if the projection is excluded?

Chapter 4

The Billard Theorem for Multiple Random Fourier Series

We propose a generalization of a classical result on random Fourier series, namely the Billard Theorem, for random Fourier series over the d -dimensional torus. We provide an investigation of the independence with respect to a choice of a sequence of partial sums (or *method of summation*). We also study some probabilistic properties of the resulting sum field such as stationarity and characteristics of the marginal distribution.

4.1 Introduction

In this Chapter and the following, we study some local asymptotic properties of microtextures as defined in the Introduction and Chapter 2. We are particularly interested in the definition of RPN models with infinite spatial frequencies: in a sense, our analysis consists in taking into account signals with an arbitrary (maybe not compact or band-limited) Fourier spectrum support. This allows the study of the regularity of our models, an important feature for some classes of stationary signals. As our proofs seldom depend on the dimensionality of the definition domain of the functions, we consider signals defined over the d -dimensional torus for an arbitrary integer d . Our purpose is thus to define and study the limit of the finite random sums of the RPN model when we consider infinitely many terms, *i.e.* random multiple Fourier series.

Random Fourier series have a long and rich history. First introduced by Paley and Zygmund in a series of papers [108] [109] and [110] in the 1930's, the subject has been drawing attention ever since. The most prominent work on the matter, along with many applications to harmonic analysis, has been synthesized by Kahane in [69] and Marcus and Pisier in [96], and many problems are still open as of today.

The purpose of this chapter is to prove the equivalence between different important properties for multiple random Fourier series. In dimension 1, the celebrated Billard Theorem (as stated in Kahane's famous book [69], Theorem 3 p. 58 — the original article by Billard [15] attempts to prove a slightly weaker result) claims a chain of equivalences between almost sure continuity, uniform convergence, uniform boundedness, and pointwise convergence of random

Fourier series. This chain is very surprising since it connects properties that are obviously non-equivalent under general non-probabilistic hypotheses. One interesting point of such a result is that it allows one to define 2π -periodical processes through the law of their Fourier coefficients. Moreover, continuity is proven to hold over the pointwise limit of the series, and not over some modified processes as it is often the case in probability theory.

Interestingly, the Billard Theorem has been partially extended in different directions. The equivalence between almost sure uniform convergence and almost sure uniform boundedness for Gaussian random Fourier series is well known (see *e.g.* [80] Theorem 13.4). Most notably, the same equivalence for Fourier series on any compact group has been proven in [47], without assuming the Fourier coefficients to be Gaussian. However, a proof of an extension of the whole chain of equivalences to the case of Fourier series on the d -dimensional torus (for $d > 1$) was missing. This chapter proposes to extend the techniques introduced by Kahane in order to provide such a proof.

In order to state such an extension, we shall write our hypotheses in Section 2. In Section 3, we introduce a notion of convergence for non absolutely summable sums taken over \mathbb{Z}^d and claim a result of independence with respect to the variations of this notion. This independence is largely based on the Itô-Nisio Theorem [60]. In Section 4, we state and prove an extension of Billard Theorem to the d -dimensional torus. Moreover, we discuss direct generalizations which include the Gaussian case. In Section 5, we study the law of the resulting process under the hypothesis of uniform convergence of the partial sums of the random Fourier series.

4.2 Notations and Hypotheses

4.2.1 Notations

Throughout this chapter and the following, we consider $(\Omega, \mathcal{F}, \mathbb{P})$ a complete probability space, and $\mathbb{T}^d := \mathbb{R}^d / 2\pi\mathbb{Z}^d$ the d -dimensional torus over which we consider the usual Lebesgue measure. We are interested in real stationary centered second-order processes defined on \mathbb{T}^d . Our purpose in this chapter is to define such processes through the law of their random Fourier representation. For any function f in $L^2(\mathbb{T}^d)$ with real or complex values, let us write the Fourier coefficients

$$\hat{f}_n := \langle f, e_n \rangle = \int_{\mathbb{T}^d} f(t) e^{-in \cdot t} dt \quad (4.1)$$

where $e_n : t \mapsto e^{in \cdot t}$ for all n in \mathbb{Z}^d ($a \cdot b$ denotes the canonical inner product in \mathbb{R}^d and $\langle g, h \rangle$ the canonical Hermitian inner product in $L^2(\mathbb{T}^d)$).

Let $X : \Omega \times \mathbb{T}^d \rightarrow \mathbb{R}$ denote a second-order process, *i.e.* such that $X(\cdot, t)$ (often written $X(t)$ in the following) is in $L^2(\Omega)$ for all t in \mathbb{T}^d . Moreover, X is assumed to be centered, *i.e.* such that $\mathbb{E}[X(t)] = 0$ for all t , and weakly stationary, *i.e.* such that $\text{Cov}(X(s), X(t))$ only depends on $t - s$. In particular, $\mathbb{E}[X(t)^2] < \infty$ and this quantity does not depend on t . Thus, thanks to Fubini-Tonelli Theorem, it follows that the sample paths of X belong almost surely to the space $L^2(\mathbb{T}^d)$ of square integrable functions. Hence, for some real non-negative random variables $(A_n)_{n \in \mathbb{Z}^d}$ that are almost surely in $l^2(\mathbb{Z}^d)$ and $(\Phi_n)_{n \in \mathbb{Z}^d}$ random variables

in $\mathbb{R}/2\pi\mathbb{Z}$, one can write

$$X(\omega, \cdot) \stackrel{L^2(\mathbb{T}^d)}{=} \sum_{n \in \mathbb{Z}^d} A_n(\omega) e^{i\Phi_n(\omega)} e_n \quad (4.2)$$

almost surely (recall that $(e_n)_{n \in \mathbb{Z}^d}$ is a Hilbert basis of $L^2(\mathbb{T}^d)$). In other words $\widehat{X}_n(\omega) = A_n(\omega) e^{i\Phi_n(\omega)}$.

However, (5.1) does not hold *a priori* in the sense “almost surely for all t in \mathbb{T}^d ”. Thus, defining a second-order process over \mathbb{T}^d through the law of its random Fourier coefficients is generally not straightforward. Indeed, two second-order processes Y and Z that have the same random Fourier representations ($\widehat{Y}_n = \widehat{Z}_n$ for all n in \mathbb{Z}^d almost surely) do not necessarily satisfy finite-dimensional distribution equality (e.g., one could have $Y(0) = 0$ a.s. and $Z(0) = 1$ a.s.). Moreover, for any set of null Lebesgue-measure $\mathcal{N} \subset \mathbb{T}$, there exists a function f in $\mathcal{C}^0(\mathbb{T}^d)$ such that the Dirichlet sums $\sum_{n=-N}^N \widehat{f}_n e_n$ diverge, as proven by Kahane and Katznelson [70]. Thus, defining a process unambiguously only through its Fourier coefficients can turn out to be difficult.

4.2.2 Two Strategies to Define a Process through its Fourier Coefficients

There are several ways to overcome these difficulties – e.g. random general functions as defined in [53]. Let us propose two radically different elementary strategies, that will turn out to be equivalent in our context. The first strategy consists in restricting our study to processes with continuous sample paths, since continuous functions with identical Fourier coefficients (hence $L^2(\mathbb{T}^d)$ equivalent) are equal everywhere. Since the inclusion of $\mathcal{C}^0(\mathbb{T}^d)$ into $L^2(\mathbb{T}^d)$ is strict, we shall seek conditions for a random family of Fourier coefficients to represent a continuous function almost surely. Another advantage of this strategy is that processes with sample paths that are almost surely in $\mathcal{C}^0(\mathbb{T}^d)$ are Radon random variables, and thus there is equivalence between *equality in law as random variables in the Banach space $\mathcal{C}^0(\mathbb{T}^d)$* and *equality in finite-dimensional law* (see Ledoux and Talagrand [80] p. 46).

Remark 4.1. *In the following, when considering the law of a random function that is almost surely in $\mathcal{C}^0(\mathbb{T}^d)$, we shall consider its finite dimensional law or the law of the entire process indifferently.*

The second strategy consist in considering the pointwise convergence of partial sums and focus on random Fourier coefficients that yield convergence everywhere almost surely. For that matter, a sequence of partial sums or “method of summation” needs to be specified. Indeed, it can be the case that for the same Fourier coefficients $(a_n)_{n \in \mathbb{Z}^d}$, a sequence of partial sums $(\sum_{n \in \mathcal{A}_k} a_n e_n)_{k \in \mathbb{N}}$ is convergent almost everywhere and another sequence $(\sum_{n \in \mathcal{B}_k} a_n e_n)_{k \in \mathbb{N}}$ is divergent on a set of positive measure. This has been pointed out by Fefferman in [44] and [43] for the case $d = 2$.

4.2.3 The Billard Theorem in Dimension 1

In this chapter, we focus on random Fourier coefficients that have the following properties:

- \mathcal{H}_1 : “ $(A_n)_{n \in \mathbb{Z}^d} = (a_n)_{n \in \mathbb{Z}^d}$ is a deterministic, non-negative, even ($a_{-n} = a_n$ for all n in \mathbb{Z}^d) square summable family with $a_0 = 0$.”
- \mathcal{H}_2 : “ $(\Phi_n)_{n \in \mathbb{Z}^d}$ is a pure phase noise field, that is for all n in \mathbb{Z}^d , $\Phi_{-n} = -\Phi_n$ (modulo 2π) almost surely, Φ_n has uniform distribution over $\mathbb{R}/2\pi\mathbb{Z}$ and $(\Phi_n)_{n \in \mathcal{A}}$ are independent for all $\mathcal{A} \subset \mathbb{Z}^d$ such that \mathcal{A} and $-\mathcal{A}$ do not intersect.”

Notice that \mathcal{H}_1 is equivalent to restrict ourselves to real functions. Interestingly, under the hypotheses \mathcal{H}_1 and \mathcal{H}_2 , the two strategies turn out to be equivalent, as we shall see in the Section 4. This generalizes the Billard Theorem (see Billard [15] and Kahane [69] p. 58), obtained in the case where $d = 1$, stating that under hypotheses \mathcal{H}_1 and \mathcal{H}_2 , the conditions

- (i) ω -a.s. convergence everywhere of the Dirichlet sums $(\sum_{n=-k}^k A_n(\omega) e^{i(n \cdot t + \Phi_n(\omega))})_{k \in \mathbb{N}}$
 - (ii) ω -a.s. uniform convergence of the Dirichlet sums
 - (iii) ω -a.s. boundedness of the Dirichlet sums
 - (iv) ω -a.s. existence of a continuous function X with Fourier coefficients $(A_n(\omega) e^{i(n \cdot t + \Phi_n(\omega))})_{n \in \mathbb{Z}}$
- are equivalent.

4.2.4 Discussion on the hypotheses

The evenness hypothesis in \mathcal{H}_1 is equivalent to considering only real-valued processes. As we shall see in Section 5, assuming that $(A_n)_{n \in \mathbb{Z}^d}$ are deterministic is equivalent to considering only second-order processes with a deterministic autocorrelation function. We shall also consider relaxations of this hypothesis in Section 3.

Furthermore, the set of hypotheses \mathcal{H}_1 and \mathcal{H}_2 can also be interpreted as an asymptotical extension of an image processing model ($d = 2$) for texture synthesis. Indeed the random phase noise model (first introduced by van Wijk [127]) has recently drawn a lot of attention, see *e.g.* Galerne *et al.* [51]. In a nutshell, this model of texture synthesis defines a random field over the discrete 2-dimensional torus $T = (\mathbb{Z}/M_1\mathbb{Z}) \times (\mathbb{Z}/M_2\mathbb{Z})$ by taking the inverse discrete Fourier transform of $(a_n e^{i\Phi_n})_{n \in T}$ where the following hypotheses are satisfied:

- \mathcal{H}_1^* : “ $(a_n)_{n \in T}$ is non-negative and even ($a_{-n} = a_n$ for all n in T)”;
- \mathcal{H}_2^* : “ $(\Phi_n)_{n \in T}$ is a finite pure phase noise field: for all n in T , $\Phi_{-n} = -\Phi_n$ (modulo 2π) almost surely, Φ_n has uniform distribution over $\mathbb{R}/2\pi\mathbb{Z}$ and $(\Phi_n)_{n \in \mathcal{A}}$ are independent if \mathcal{A} and $-\mathcal{A}$ do not intersect.”

4.3 Methods of summation in \mathbb{Z}^d

As we shall see in the next section, one of the difficulty in extending Billard Theorem to the case where $d \geq 2$ is that there is no straightforward equivalent of the canonical Dirichlet sums. In other words, if $d \geq 2$, there is no increasing sequence of subsets of \mathbb{Z}^d , say $(\mathcal{A}_k)_{k \in \mathbb{N}}$, such that any other increasing sequence of subsets of \mathbb{Z}^d , say $(\mathcal{B}_k)_{k \in \mathbb{N}}$, is also a subsequence

of $(\mathcal{A}_k)_{k \in \mathbb{N}}$. This has been a major difficulty for generalizing Carleson's theorem in all finite dimensions, as discussed by Ash and Welland in [4], Fefferman in [44] and [43] (see also [126]), and more broadly for the study of Fourier series in multiple dimensions. In the following, we shall focus on increasing sequences of *finite and symmetrical* subsets of \mathbb{Z}^d . This assumption combined with hypotheses \mathcal{H}_1 and \mathcal{H}_2 allows us to focus on real functions.

Definition 4.1. $(\mathcal{A}_k)_{k \in \mathbb{N}}$ is said to be a (symmetrical) method of summation over \mathbb{Z}^d if

1. for each k , \mathcal{A}_k is a finite subset of \mathbb{Z}^d (such that $-\mathcal{A}_k = \mathcal{A}_k$);
2. for each k , $\mathcal{A}_k \subset \mathcal{A}_{k+1}$;
3. $\bigcup_{k \in \mathbb{N}} \mathcal{A}_k = \mathbb{Z}^d$.

Given a Banach space B , a family $(x_n)_{n \in \mathbb{Z}^d}$ of elements in B is said to be summable according to $(\mathcal{A}_k)_{k \in \mathbb{N}}$ if $\sum_{n \in \mathcal{A}_k} x_n$ converges in B as $k \rightarrow \infty$.

Remark 4.2. A method of summation can be seen as a subsequence of an ordering sequence over \mathbb{Z}^d .

Let us consider $(\mathcal{C}^0(\mathbb{T}^d), \|\cdot\|_\infty)$ the Banach space of all continuous functions over \mathbb{T}^d endowed with the uniform convergence topology.

Remark 4.3. Under Hypothesis \mathcal{H}_2 , for each $n \neq 0$ in \mathbb{Z}^d , $t \mapsto \cos(n \cdot t + \Phi_n)$ is a symmetrically distributed random variable in $\mathcal{C}^0(\mathbb{T}^d)$. So Hypothesis \mathcal{H}_2 implies that for any symmetrical method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, the incremental partial sums $t \mapsto \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} a_n e^{i(n \cdot t + \Phi_n)}$ are independent and symmetrically distributed.

The following result builds upon this remark and allows us to overcome the difficulties that arise with sums over \mathbb{Z}^d .

Proposition 4.1. Let $(\mathcal{A}_k)_{k \in \mathbb{N}}$ be a symmetrical method of summation in \mathbb{Z}^d . Assume that, almost surely, the sequence of functions

$$S_{\mathcal{A}_k} : t \mapsto \sum_{n \in \mathcal{A}_k} a_n e^{i(n \cdot t + \Phi_n)} \quad (4.3)$$

converges uniformly (resp. is uniformly bounded) on \mathbb{T}^d as $k \rightarrow \infty$ and call $S_{\mathcal{A}}$ its limit. Then, under the hypotheses \mathcal{H}_1 and \mathcal{H}_2 , for any other method of summation $(\mathcal{B}_k)_{k \in \mathbb{N}}$, the sequence of functions

$$S_{\mathcal{B}_k} : t \mapsto \sum_{n \in \mathcal{B}_k} a_n e^{i(n \cdot t + \Phi_n)} \quad (4.4)$$

converges uniformly to $S_{\mathcal{A}}$ (resp. is uniformly bounded) on \mathbb{T}^d as $k \rightarrow \infty$.

Proof. We first prove the claim for uniform convergence. Notice that each sum over a symmetrical subset $\mathcal{E} \subset \mathbb{Z}^d$ such that $0 \notin \mathcal{E}$ satisfies

$$\sum_{n \in \mathcal{E}} a_n e^{i(n \cdot t + \Phi_n)} = \sum_{n \in \mathcal{E}} a_n \cos(n \cdot t + \Phi_n) \quad (4.5)$$

for every t in \mathbb{T}^d , thanks to \mathcal{H}_1 and \mathcal{H}_2 . Moreover, thanks to \mathcal{H}_2 , $t \mapsto \cos(n \cdot t + \Phi_n)$ are symmetrical random variables and thus so are $t \mapsto \sum_{n \in \mathcal{E}} a_n e^{i(n \cdot t + \Phi_n)}$. Hence, for each $k \geq 2$, $S_{\mathcal{A}_k} = S_{\mathcal{A}_0} + \sum_{p=0}^{k-1} S_{\mathcal{A}_{p+1} \setminus \mathcal{A}_p}$ is a sum of independent symmetrical random variables in the Banach space $\mathcal{C}^0(\mathbb{T}^d)$.

Proposition 4.1 can be deduced as from a well known consequence of the Lévy-Itô-Nisio Theorem (see [60] or [80]) that we recall here. If $(Y_k)_{k \in \mathbb{N}}$ is a sequence of independent symmetrical random variables in some Banach space $(B, \|\cdot\|)$, and if S_k denotes $\sum_{l=1}^k Y_l$, then (see *e.g.* [80] p. 48 and Theorem 1 in [69] p. 13), the following conditions are equivalent:

- (i) $(S_k)_{k \in \mathbb{N}}$ converges almost surely
- (ii) $(S_k)_{k \in \mathbb{N}}$ converges in probability
- (iii) there exists some subsequence $(S_{k_p})_{p \in \mathbb{N}}$ that converges almost surely.

We apply this result to the Banach space of continuous function $(\mathcal{C}^0(\mathbb{T}^d), \|\cdot\|_\infty)$.

Notice that, since $\mathcal{A}_k \subset \mathcal{A}_{k+1}$ and $\bigcup_{k \in \mathbb{N}} \mathcal{A}_k = \mathbb{Z}^d$, for any finite subset $\mathcal{E} \subset \mathbb{Z}^d$ one has $\mathcal{E} \subsetneq \mathcal{A}_k$ for k large enough. Let us define a new method of summation $(\mathcal{AB}_k)_{k \in \mathbb{N}}$ by induction. $\mathcal{AB}_0 = \mathcal{A}_0$, $\mathcal{AB}_1 = \bigcap_{l, \mathcal{AB}_0 \subsetneq \mathcal{B}_l} \mathcal{B}_l$, $\mathcal{AB}_2 = \bigcap_{l, \mathcal{AB}_1 \subsetneq \mathcal{A}_l} \mathcal{A}_l$, and by induction $\mathcal{AB}_{2k} = \bigcap_{l, \mathcal{AB}_{2k-1} \subsetneq \mathcal{A}_l} \mathcal{A}_l$ (resp. $\mathcal{AB}_{2k+1} = \bigcap_{l, \mathcal{AB}_{2k} \subsetneq \mathcal{B}_l} \mathcal{B}_l$) for all k in \mathbb{N} . Notice that this reasoning provides us $(p_k)_{k \in \mathbb{N}}$ and $(q_k)_{k \in \mathbb{N}}$, two strictly increasing sequences of integers such that $\mathcal{AB}_{2k} = \mathcal{A}_{p_k}$ and $\mathcal{AB}_{2k+1} = \mathcal{B}_{q_k}$. Moreover, $(\mathcal{AB}_k)_{k \in \mathbb{N}}$ is clearly a symmetrical method of summation.

Since $(S_{\mathcal{AB}_{2k}})_{k \in \mathbb{N}}$ is a subsequence of $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$, it converges almost surely in $\mathcal{C}^0(\mathbb{T}^d)$ to $S_{\mathcal{A}}$. Hence, thanks to the consequence of Lévy-Itô-Nisio Theorem mentioned earlier, $(S_{\mathcal{AB}_k})_{k \in \mathbb{N}}$ converges also almost surely in $\mathcal{C}^0(\mathbb{T}^d)$ to $S_{\mathcal{A}}$ thanks to the uniqueness of limits. It follows that $(S_{\mathcal{AB}_{2k+1}})_{k \in \mathbb{N}}$ converges also almost surely in $\mathcal{C}^0(\mathbb{T}^d)$ to $S_{\mathcal{A}}$, as a subsequence of $(S_{\mathcal{AB}_k})_{k \in \mathbb{N}}$. Thus, since $(S_{\mathcal{AB}_{2k+1}})_{k \in \mathbb{N}}$ is a subsequence of $(S_{\mathcal{B}_k})_{k \in \mathbb{N}}$, the latter converges also almost surely to $S_{\mathcal{A}}$ in $\mathcal{C}^0(\mathbb{T}^d)$, thanks to the same consequence of Lévy-Itô-Nisio Theorem. Thus $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$, $(S_{\mathcal{B}_k})_{k \in \mathbb{N}}$ and $(S_{\mathcal{AB}_k})_{k \in \mathbb{N}}$ converge simultaneously to the same limit almost surely.

The proof for boundedness uses a slightly different consequence of the Lévy-Itô-Nisio Theorem. Namely with the same hypotheses and notations, the following propositions are equivalent (see *e.g.* Theorem 1 in [69] p. 13):

- (vi) $(S_k)_{k \in \mathbb{N}}$ is bounded almost surely
- (v) there exists some subsequence $(S_{k_p})_{p \in \mathbb{N}}$ that is bounded almost surely.

Assume that $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is almost surely bounded in $\mathcal{C}^0(\mathbb{T}^d)$. Then, so is $(S_{\mathcal{AB}_{2k}})_{k \in \mathbb{N}}$ as a subsequence of $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$. Hence, thanks to the consequence of Lévy-Itô-Nisio mentioned earlier, $(S_{\mathcal{AB}_k})_{k \in \mathbb{N}}$ is also bounded in $\mathcal{C}^0(\mathbb{T}^d)$ almost surely. Hence, $(S_{\mathcal{AB}_{2k+1}})_{k \in \mathbb{N}}$ is also bounded in $\mathcal{C}^0(\mathbb{T}^d)$ almost surely, as a subsequence of $(S_{\mathcal{AB}_k})_{k \in \mathbb{N}}$. Thus, since $(S_{\mathcal{AB}_{2k+1}})_{k \in \mathbb{N}}$ is a subsequence of $(S_{\mathcal{B}_k})_{k \in \mathbb{N}}$, the latter is also bounded in $\mathcal{C}^0(\mathbb{T}^d)$ almost surely, thanks to the same consequence of Lévy-Itô-Nisio Theorem. \square

An important consequence of Proposition 4.1 is that the choice of a method of summation does not matter for the uniform convergence or for the uniform boundedness. As long as uniform convergence (resp. uniform boundedness) happens almost surely for some method of convergence, it also happens almost surely for any other method of convergence and the limit is the same.

4.4 Billard's theorem in arbitrary finite dimension

We can now turn to an extension of Billard's theorem to the case where $d \geq 2$ (recall that the sequence $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ has been defined by equation (4.33)).

Theorem 4.1. *Under the hypotheses \mathcal{H}_1 and \mathcal{H}_2 , the following conditions are equivalent:*

- (i) *almost surely, there exists a continuous function X , such that $(a_n e^{i\Phi_n})_{n \in \mathbb{Z}^d}$ are the Fourier coefficients of X ;*
- (ii) *there exists a method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$ such that, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ converges uniformly;*
- (iii) *for all methods of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ converges uniformly;*
- (iv) *there exists a method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$ such that, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is bounded;*
- (v) *for all methods of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is bounded;*
- (vi) *for all methods of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, almost surely, for all t in \mathbb{T}^d , $(S_{\mathcal{A}_k}(t))_{k \in \mathbb{N}}$ converges.*

The fact that (ii) \Leftrightarrow (iii) (resp. (iv) \Leftrightarrow (v)) follows from Proposition 4.1. Moreover, (iii) implies clearly all the other statements.

Remark 4.4. *A somewhat weaker equivalence between boundedness and continuity, which depends on a method of summation, was proven with much more generality for any compact group instead of \mathbb{T}^d by Figa-Talamanca in [47].*

Definition 4.2. *Under any of the equivalent conditions of Theorem 4.1, the limit X in $\mathcal{C}^0(\mathbb{T}^d)$ is called a random phase noise (RPN) process.*

The remaining of this section is dedicated to the proof of Theorem 4.1.

4.4.1 Proof of (v) \Rightarrow (iii)

Proposition 4.2. *Let $(Y_k)_{k \in \mathbb{N}}$ be a sequence of independent random variables with value in $\mathcal{C}^0(\mathbb{T}^d)$. Assume that*

1. *for each $k \in \mathbb{N}$, Y_k is symmetrically distributed i.e. $-Y_k$ and Y_k have the same law ;*

2. for each $k \in \mathbb{N}$, Y_k is stationary i.e. $Y_k(\cdot - \tau)$ and Y_k have the same law for any τ in \mathbb{T}^d ;
3. the sequence $(\sum_{k \leq l} Y_k)_{l \in \mathbb{N}}$ is almost surely bounded in $\mathcal{C}^0(\mathbb{T}^d)$, according to the $\|\cdot\|_\infty$ norm.

Then, almost surely, $(\sum_{k \leq l} Y_k)_{l \in \mathbb{N}}$ converges in $\mathcal{C}^0(\mathbb{T}^d)$.

Proposition 4.2 is a straightforward extension of Proposition 13 p. 55 in [69] and the proof is postponed in Appendix. We can now prove that (v) implies (iii). Assume (v), and recall that under Hypothesis \mathcal{H}_2 , Remark 4.3 ensures that the incremental partial sums $Y_k := t \mapsto \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} a_n e^{i(n \cdot t + \Phi_n)}$ satisfy the three hypotheses of Proposition 4.2. Thus, for any symmetrical method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ converges almost surely in $\mathcal{C}^0(\mathbb{T}^d)$ and (iii) holds.

The end of this section is largely built upon ideas found in Kahane [69] pp. 48 and 59-60. However, we found the details of our proof to be significantly different from the case $d = 1$, so we provide them in the core of the text.

4.4.2 Proof of (vi) \Rightarrow (iv)

To prove that (vi) implies (iv) we need to prove more intermediate results. The first one deals with trigonometric polynomials. For a trigonometric polynomial P defined on \mathbb{T}^d by

$$P(t) = \sum_{n \in \mathcal{E}} b_n e^{in \cdot t} \quad (4.6)$$

where $b_n \neq 0$ is in \mathbb{C} for each n in the finite set $\mathcal{E} \subset \mathbb{Z}^d$, we define the degree of P as

$$d(P) := \max_{n \in \mathcal{E}} |n|_\infty \quad (4.7)$$

where $|x|_\infty := \max_i |x_i|$ denotes the max norm for x in \mathbb{R}^d . In the following we denote $\mathbb{B}_\infty(t, r) = \{s \in \mathbb{T}^d; |t - s|_\infty < r\}$ the projection onto \mathbb{T}^d of the \mathbb{R}^d open ball of radius r and center t with respect to $|\cdot|_\infty$ onto \mathbb{T}^d .

Proposition 4.3. *Let \mathcal{E} be a finite subset of \mathbb{Z}^d and $P(t) = \sum_{n \in \mathcal{E}} b_n e^{i(n \cdot t + \phi_n)}$ be a complex trigonometric polynomial defined on \mathbb{T}^d . Assume that there exists $q \geq 1$ in \mathbb{N} and l in \mathbb{Z}^d such that $\mathcal{E} \subset l + q\mathbb{Z}^d$, so $t \mapsto |P(t)|$ is $\frac{2\pi}{q}$ -periodic in every direction. Assume moreover that the degree of P is less than K where $K \geq \frac{q}{2\pi}$.*

Then for every radius $\varepsilon \geq 2\pi/q$, and center t in \mathbb{T}^d , there exists t' in $\mathbb{B}_\infty(t, \varepsilon)$ such that

$$\mathbb{B}_\infty(t', \varepsilon') \subset \mathbb{B}_\infty(t, \varepsilon) \quad (4.8)$$

with $\varepsilon' \geq (2K)^{-1}$ and

$$|P(s)| \geq 1/2 \|P\|_\infty \quad (4.9)$$

for all s in $\mathbb{B}_\infty(t', \varepsilon')$.

The proof is postponed in Appendix. We now state a result of symmetrization, useful for the remainder of the proof of Theorem 4.1.

Lemma 4.1. *Let $(a_n)_{n \in \mathbb{Z}^d}$ and $(\Phi_n)_{n \in \mathbb{Z}^d}$ satisfy to Hypotheses \mathcal{H}_1 and \mathcal{H}_2 . Let $(\mathcal{A}_k)_{k \in \mathbb{N}}$ be any method of summation and \mathcal{B} be a subset of \mathbb{Z}^d . Assume that there exists a random variable T such that with non-zero probability (resp. almost surely) the complex-valued sequence*

$$\left(\sum_{n \in \mathcal{A}_k \cap \mathcal{B}} a_n e^{i(n \cdot T + \Phi_n)} \right)_{k \in \mathbb{N}} \quad (4.10)$$

diverges. Then there exists \mathcal{B}^ a symmetrical subset of \mathbb{Z}^d such that with non-zero probability (resp. almost surely) the real-valued sequence*

$$\left(\sum_{n \in \mathcal{A}_k \cap \mathcal{B}^*} a_n e^{i(n \cdot T + \Phi_n)} \right)_{k \in \mathbb{N}} \quad (4.11)$$

diverges.

We can now prove the implication $(vi) \Rightarrow (iv)$. Let us assume that (vi) holds and that (iv) does not, and let us aim at a contradiction. Let $(\mathcal{A}_k)_{k \in \mathbb{N}}$ be any method of summation. The sequence of partial sums $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is not almost surely bounded in $\mathcal{C}^0(\mathbb{T}^d)$. Hence, the event

$$E = \{\omega \in \Omega; (S_{\mathcal{A}_k}(\omega))_{k \in \mathbb{N}} \text{ is bounded in } \mathcal{C}^0(\mathbb{T}^d)\} \quad (4.12)$$

has probability less than 1. For each k , define the σ -algebra \mathcal{F}_k generated by $\{e^{i\Phi_n}\}_{n \in \mathcal{A}_k}$ and notice that $\mathcal{F}_k \subset \mathcal{F}_{k+1}$. The event E belongs to the asymptotic σ -algebra of $(\mathcal{F}_k)_{k \in \mathbb{N}}$, since E is independent of any finite subset of the random variables $(e^{i\Phi_n})_{n \in \mathbb{Z}^d}$. Thanks to the independance hypothesis in \mathcal{H}_2 , the zero-one law applies and $\mathbb{P}(E) = 0$, which in turns implies that, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is unbounded in $\mathcal{C}^0(\mathbb{T}^d)$.

Symmetrization In order to obtain a contradiction, we shall construct \mathcal{B} a (non-random) subset of \mathbb{Z}^d , a method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$ and a random variable T such that, with non-zero probability,

$$\text{the sequence } \left(\sum_{n \in \mathcal{B} \cap \mathcal{A}_k} a_n e^{i(n \cdot T + \Phi_n)} \right)_{k \in \mathbb{N}} \text{ does not converge as } k \rightarrow \infty. \quad (4.13)$$

Thanks to Lemma 4.1, there shall exist \mathcal{B}^* a (non-random) symmetrical subset of \mathbb{Z}^d such that, with non-zero probability,

$$\text{the sequence } \left(\sum_{n \in \mathcal{B}^* \cap \mathcal{A}_k} a_n e^{i(n \cdot T + \Phi_n)} \right)_{k \in \mathbb{N}} \text{ does not converge as } k \rightarrow \infty. \quad (4.14)$$

Let us consider the random Fourier coefficients $(\varepsilon_n a_n e^{i\Phi_n})_{n \in \mathbb{Z}^d}$ where $\varepsilon_n = 1$ whenever $n \in \mathcal{B}^*$ and $\varepsilon_n = -1$ otherwise. Thanks to \mathcal{H}_2 , this family has the same law as $(a_n e^{i\Phi_n})_{n \in \mathbb{Z}^d}$. Hence,

$$S_{\mathcal{A}_k} : t \mapsto \sum_{n \in \mathcal{A}_k} a_n e^{i(n \cdot t + \Phi_n)} \quad (4.15)$$

has the same law in the Banach space $\mathcal{C}^0(\mathbb{T}^d)$ as

$$S'_{\mathcal{A}_k} : t \mapsto \sum_{n \in \mathcal{A}_k} \varepsilon_n a_n e^{i(n \cdot t + \Phi_n)} \quad (4.16)$$

and since $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is assumed to converge everywhere almost surely, $(S'_{\mathcal{A}_k})_{k \in \mathbb{N}}$ shall also converge everywhere almost surely. Hence, the sum

$$S_{\mathcal{A}_k} + S'_{\mathcal{A}_k} : t \mapsto 2 \sum_{n \in \mathcal{B}^* \cap \mathcal{A}_k} a_n e^{i(n \cdot t + \Phi_n)} \quad (4.17)$$

shall in turn converge everywhere almost surely. This is contradictory with (4.14).

Construction Let us now build such a set \mathcal{B} and a method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$. Let $\mathcal{A}_k = \{n \in \mathbb{Z}^d; |n|_\infty \leq k\}$. Let us define the events

$$E_k^{(1)} := \{\omega \in \Omega; \sup_{j \leq k} \|S_{\mathcal{A}_j}(\omega, \cdot)\|_\infty > 2\} \quad (4.18)$$

for each k in \mathbb{N} and notice that $E_k^{(1)} \subset E_{k+1}^{(1)}$ for each k . Since almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is unbounded in $\mathcal{C}^0(\mathbb{T}^d)$, $\mathbb{P}(E_k^{(1)}) \rightarrow 1$ as $k \rightarrow \infty$, so there is an integer k_1 such that the probability of the event $E_{k_1}^{(1)}$ is larger than $1/2$. Furthermore, whenever ω belongs to $E_{k_1}^{(1)}$, thanks to Proposition 4.3 (with $q = 1$, $K = k_1$ and $\varepsilon > \pi$ so $\mathbb{B}_\infty(t, \varepsilon) = \mathbb{T}^d$), there exists a random ball $U_1(\omega) = \mathbb{B}_\infty(T_1(\omega), \varepsilon_1)$ with radius $\varepsilon_1 = (2k_1)^{-1}$ such that

$$\sup_{j \leq k_1} |S_{\mathcal{A}_j}(\omega, t)| > 1 \quad (4.19)$$

for all t in $\mathbb{B}_\infty(T_1(\omega), \varepsilon_1)$. For $\omega \in \Omega \setminus E_{k_1}^{(1)}$ we set $U_1(\omega) = \mathbb{T}^d$. Finally, we define $\mathcal{B}_1 = \mathcal{A}_{k_1}$.

Define $q_1 = \lceil 2\pi/\varepsilon_1 \rceil = \lceil 4\pi k_1 \rceil$. Let us consider the partition of $\mathbb{Z}^d \setminus \mathcal{A}_{k_1}$ into q_1^d subsets

$$\mathcal{C}_{1,l} = (l + q_1 \mathbb{Z}^d) \setminus \mathcal{A}_{k_1} \quad (4.20)$$

for each l in $\{n \in \mathbb{N}^d; |n|_\infty < q_1\}$. For $k \geq k_1$, there are $(2q_1 + 1)^d$ random sequences of functions $(S_k^{(1,l)})_{k \in \mathbb{N}}$ defined by

$$S_k^{(1,l)}(t) := \sum_{n \in \mathcal{A}_k \cap \mathcal{C}_{1,l}} a_n e^{i(n \cdot t + \Phi_n)} \quad (4.21)$$

for each l in $\{n \in \mathbb{N}^d; |n|_\infty < q_1\}$ and they satisfy

$$\sum_{l \in \{n \in \mathbb{N}^d; |n|_\infty < q_1\}} S_k^{(1,l)} = S_{\mathcal{A}_k \setminus \mathcal{A}_{k_1}} = S_{\mathcal{A}_k} - S_{\mathcal{A}_{k_1}} \quad (4.22)$$

for $k > k_1$. Since $(S_{\mathcal{A}_k \setminus \mathcal{B}_1})_{k \in \mathbb{N}}$ is almost surely unbounded in $\mathcal{C}^0(\mathbb{T}^d)$, so must be at least one of the sequences $(S_k^{(1,l)})_{k \in \mathbb{N}}$. Thus for at least one of these sequences, say $(S_k^{(1,l_1)})_{k \in \mathbb{N}}$,

$$\mathbb{P}((S_k^{(1,l_1)})_{k \in \mathbb{N}} \text{ is unbounded}) > 0. \quad (4.23)$$

Thanks to the zero-one law, $(S_k^{(1,l_1)})_{k \in \mathbb{N}}$ must be unbounded with probability 1. For each $k > k_1$, define the event

$$E_k^{(2)} := \{\omega \in \Omega \mid \sup_{k_1 < j \leq k} \|S_j^{(1,l_1)}(\omega, \cdot)\|_\infty > 2\} \quad (4.24)$$

and notice that $E_k^{(2)} \subset E_{k+1}^{(2)}$ for all $k \geq k_1$ in \mathbb{N} . Since $\mathbb{P}(E_k^{(2)}) \rightarrow 1$ as $k \rightarrow \infty$, there exists an integer k_2 (non-random) such that $\mathbb{P}(E_{k_2}^{(2)}) > 1/2$. Thus, whenever ω belongs to $E_{k_2}^{(2)}$, thanks to Proposition 4.3 (invoked with $q = q_1$, $K = k_2$ and $\varepsilon = \varepsilon_1$) we know that $U_1(\omega) = \mathbb{B}_\infty(T_1(\omega), \varepsilon_1)$ contains a random ball $U_2(\omega) = \mathbb{B}_\infty(T_2(\omega), \varepsilon_2)$ with radius $\varepsilon_2 = (2k_2)^{-1}$ such that

$$\sup_{k_1 < j \leq k_2} \left| \sum_{n \in \mathcal{C}_{1,l_1} \cap \mathcal{A}_j} a_n e^{i(n \cdot t + \Phi_n)} \right| > 1 \quad (4.25)$$

for each t in $B(T_2(\omega), \varepsilon_2)$. For ω in $\Omega \setminus E_{k_2}^{(2)}$, we choose $U_2(\omega) = U_1(\omega)$. Finally, we choose $\mathcal{B}_2 := \mathcal{C}_{1,l_1} \cap \mathcal{A}_{k_2}$.

Induction By induction, using the same arguments (Proposition 4.3 invoked with $q = q_p$, $K = k_{p+1}$ and $\varepsilon = \varepsilon_p$), we construct

- two increasing sequences $(k_p)_{p \in \mathbb{N}}$ and $(q_p)_{p \in \mathbb{N}}$ with values in \mathbb{N} and a real sequence $(\varepsilon_p)_{p \in \mathbb{N}}$ such that

$$\forall p, q_p = \lceil 2\pi/\varepsilon_p \rceil = \lceil 4\pi k_p \rceil \quad (4.26)$$

- a sequence $(l_p)_{p \in \mathbb{N}}$ with values in \mathbb{Z}^d and a sequence $(\mathcal{B}_p)_{p \geq 1}$ of finite subsets of \mathbb{Z}^d such that

$$\mathcal{B}_{p+1} \subset (l_p + q_p \mathbb{Z}^d) \cap (\mathcal{A}_{k_{p+1}} \setminus \mathcal{A}_{k_p}) \quad (4.27)$$

- a sequence of events $(E_{k_p}^{(p)})_{p \in \mathbb{N}}$ with probability at least 1/2 such that

$$\forall \omega \in E_{k_p}^{(p)}, \sup_{k_p \leq j \leq k_{p+1}} \sup_{t \in \mathbb{T}^d} \left| \sum_{n \in \mathcal{B}_p \cap \mathcal{A}_j} a_n e^{i(n \cdot t + \Phi_n)} \right| > 2 \quad (4.28)$$

- a sequence $(T_p)_{p \in \mathbb{N}}$ of random variables with values in \mathbb{T}^d
- a sequence of decreasing random open balls $(U_p)_{p \in \mathbb{N}}$ defined by either $U_p(\omega) = \mathbb{B}_\infty(T_p(\omega), \varepsilon_p)$ if $\omega \in E_{k_p}^{(p)}$, or $U_p(\omega) = U_{p-1}(\omega)$ otherwise. such that

$$\forall \omega \in E_{k_p}^{(p)}, \forall t \in U_p(\omega), \sup_{k_p \leq j \leq k_{p+1}} \left| \sum_{n \in \mathcal{B}_p \cap \mathcal{A}_j} a_n e^{i(n \cdot t + \Phi_n)} \right| > 1. \quad (4.29)$$

Let us denote $U^*(\omega) = \bigcap_p U_p(\omega)$. Since the sets $(\mathcal{B}_p)_{p \in \mathbb{N}}$ are disjoint, and thanks to \mathcal{H}_2 , the events $(E_{k_p}^{(p)})_{p \in \mathbb{N}}$ are independent. Moreover since $\mathbb{P}(E_{k_p}^{(p)}) \geq 1/2$, $\sum_p \mathbb{P}(E_{k_p}^{(p)}) = \infty$ and thanks to the Borel-Cantelli Lemma, almost surely, $\mathbb{P}(\liminf_p E_{k_p}^{(p)}) = 1$. Hence, ω -almost surely, there is one and only one (random) point $T^*(\omega)$ in $U^*(\omega)$. Define $\mathcal{B} = \bigcup_p \mathcal{B}_p$. By construction, almost surely, the complex sequence

$$\left(\sum_{n \in \mathcal{B} \cap \mathcal{A}_k} a_n e^{i(n \cdot T^* + \Phi_n)} \right)_{k \in \mathbb{N}} \quad (4.30)$$

is not Cauchy since $T^*(\omega)$ belongs to each $U_p(\omega)$ and (4.29) holds for all p in \mathbb{N} .

We conclude the proof of the implication $(vi) \Rightarrow (iv)$ by noticing that the method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$, the random variable T^* and the subset $\mathcal{B} \subset \mathbb{Z}^d$ satisfy the condition (4.13).

4.4.3 Proof of $(i) \Rightarrow (ii)$

Define $|x|_1 = \sum_{i=1}^d |x_i|$ (x in \mathbb{R}^d) and let us choose $\mathcal{D}_k := \{n \in \mathbb{Z}^d; |n|_1 \leq k\}$ (k in \mathbb{N}) as a method of summation. We consider the $(2d-1)$ -Cesàro means of the sequence of functions $(S_{\mathcal{D}_k})_{k \in \mathbb{N}}$

$$C_{(2d-1,k)}^{\mathcal{D}} := \frac{1}{\binom{k+2d-1}{k}} \sum_{l=0}^k \binom{k-l+2d-2}{k-l} S_{\mathcal{D}_l} \quad (4.31)$$

for k in \mathbb{N} as introduced in [8]. One easily checks that the sums $S_{\mathcal{D}_{k+1} \setminus \mathcal{D}_k}$ are symmetrically distributed and independent. Moreover, notice that the sums $C_{(2d-1,k)}^{\mathcal{D}}$ can be rewritten as

$$C_{(2d-1,k)}^{\mathcal{D}} = \sum_{l=0}^{k-1} b_{k,l} S_{\mathcal{D}_{l+1} \setminus \mathcal{D}_l} \quad (4.32)$$

for each k , with $b_{k,l} := \frac{\binom{k-l+2d-1}{k-l}}{\binom{k+2d-1}{k}}$ for $l \leq k$ (and $b_{k,l} := 0$ otherwise). The coefficients $(b_{k,l})_{k,l \in \mathbb{N}}$ satisfy the properties of a matrix of summation, namely that $b_{k,l} \rightarrow 0$ as $l \rightarrow \infty$ and $b_{k,l} \rightarrow 1$ as $k \rightarrow \infty$ (see Kahane [69] p. 12). Since (i) implies that $C_{(2d-1,k)}^{\mathcal{D}}$ converges uniformly as proven by Berens and Xu in [8]), Theorem 1 p. 13 in [69] yields that $(S_{\mathcal{D}_k})_{k \in \mathbb{N}}$ converges uniformly almost surely, and thus (ii) holds.

This concludes the proof of Theorem 4.1.

4.4.4 Discussion and Extension

Our extension of the Billard Theorem can be generalized to weaker hypotheses. For instance, consider the Hypothesis

\mathcal{H}_1^{**} : " $(A_n)_{n \in \mathbb{Z}^d}$ is such that $(A_n)_{n \in \mathcal{A}}$ are independant whenever \mathcal{A} and $-\mathcal{A}$ do not intersect ; $(A_n)_{n \in \mathbb{Z}^d}$ is independent of Φ ; $\mathbb{E}[\sum_{\mathbb{Z}^d} A_n^2] < \infty$; $A_0 = 0$ almost surely".

Write

$$S_{\mathcal{A}_k}(\omega, t) = \sum_{n \in \mathcal{A}_k} A_n(\omega) e^{i(\Phi_n(\omega) + n \cdot t)} \quad (4.33)$$

for all ω in Ω and t in \mathbb{T}^d . The following result can be easily deduced from Theorem 4.1.

Corollary 4.1. *Under the hypotheses \mathcal{H}_1^{**} and \mathcal{H}_2 , the chain of equivalence of Theorem 1 holds with A_n instead of a_n and $S_{\mathcal{A}_k}$ defined by (4.33).*

Proof. To prove that, notice that $\mathbb{E}[\sum_{\mathbb{Z}^d} A_n^2] < \infty$ implies that $(A_n)_{n \in \mathbb{Z}^d}$ is almost surely square summable, and thus almost surely, Theorem 4.1 can be applied conditionally on $\mathcal{F}(A_n, n \in \mathbb{Z}^d)$, the σ -algebra generated by $(A_n)_{n \in \mathbb{Z}^d}$, since Φ is independent of $\mathcal{F}(A_n, n \in \mathbb{Z}^d)$. \square

Remark 4.5. *This is of particular interest since Gaussian processes satisfy \mathcal{H}_1^{**} and \mathcal{H}_2 .*

Notice however that Hypothesis \mathcal{H}_2 cannot be much relaxed. As argued by Cohen and Cundy in [25], the symmetry assumption on $A_n e^{i\Phi_n}$ for each n cannot be replaced by $\mathbb{E}[A_n e^{i\Phi_n}] = 0$ for each n .

4.5 Properties of Random Phase Noise Processes

Throughout this section, we assume both hypotheses \mathcal{H}_1 and \mathcal{H}_2 to hold. Moreover, we assume the equivalent hypotheses in Theorem 4.1 to hold, and thus the sample paths of the random phase noise field X are almost surely continuous. Explicit conditions (*e.g.* on the coefficients $(a_n)_{n \in \mathbb{Z}^d}$) have been thoroughly studied in the case $d = 1$, *e.g.* in Kahane [69], Chapter 7.

4.5.1 Stationarity

Proposition 4.4. *A random phase noise (RPN) process X is a centered second-order process, with covariance*

$$c_X(t) = \text{Cov}(X(t+s), X(s)) = \sum_{n \in \mathbb{Z}^d} a_n^2 \cos(n \cdot t) \quad (4.34)$$

for all s and t in \mathbb{T}^d (weak stationarity). Moreover, X is strongly stationary in the sense that $(X(t))_{t \in \mathbb{T}^d}$ and $(X(t+\tau))_{t \in \mathbb{T}^d}$ have the same law for any τ in \mathbb{T}^d .

Finally, the autocorrelation of X defined as

$$R_X(\tau) = \frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} X(t) X(t+\tau) dt, \tau \in \mathbb{T}^d, \quad (4.35)$$

is deterministic and a.s. equal to c_X .

Proof. For each t in \mathbb{T}^d , $X(t)$ is the almost sure limit of a centered martingale $(X_k(t))_{k \in \mathbb{N}}$ ($X_k(t) = \sum_{n \in \mathcal{A}_k} a_n e^{in \cdot t + \Phi_n}$ for any method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$), that is bounded by $\sum_{n \in \mathbb{Z}^d} a_n^2$ in the space $L^2(\Omega)$, so it is a centered random variable in $L^2(\Omega)$. It follows that

$$\mathbb{E}[X(s)X(t)] = \mathbb{E}[X(s)\overline{X(t)}] = \lim_{k \rightarrow +\infty} \mathbb{E}[X_k(s)\overline{X_k(t)}] = \sum_{n \in \mathbb{Z}^d} a_n^2 e^{in \cdot (s-t)} = \sum_{n \in \mathbb{Z}^d} a_n^2 \cos(n \cdot (s-t)) \quad (4.36)$$

holds thanks to \mathcal{H}_1 .

Recall that the Fourier coefficients of $t \mapsto X(t - \tau)$ are $(e^{in \cdot \tau} \widehat{X}_n)_{n \in \mathbb{Z}^d}$. By the definition of \widehat{X}_n and \mathcal{H}_2 , $(e^{in \cdot \tau} \widehat{X}_n)_{n \in \mathbb{Z}^d}$ and $(\widehat{X}_n)_{n \in \mathbb{Z}^d}$ have the same finite dimensional law. Thus, X and $X(\cdot - \tau)$ have the same finite dimensional law and the same law thanks to the almost sure continuity (Remark 4.1).

Finally, thanks to Parseval identity,

$$\frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} X(t)X(t + \tau)dt = \frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} X(t)\overline{X(t + \tau)}dt = \sum_{n \in \mathbb{Z}^d} a_n^2 \cos(n \cdot \tau) \quad (4.37)$$

holds for all τ a.s., so we can conclude that $R_X = c_X$ a.s. \square

Hence, a single sample path contains enough information to fully determine the covariance and the law of the entire process, which can have various applications. For instance, one only needs one sample path to get as many independent sample paths with the same law. Interestingly, a second-order process that has a deterministic autocorrelation also has deterministic Fourier modulus.

Proposition 4.5. *Let $Y : \Omega \times \mathbb{T}^d \mapsto \mathbb{R}$ a centered process with sample paths almost surely in $C^0(\mathbb{T}^d)$. Assume that there exists a (deterministic) continuous even function $\rho : \mathbb{T}^d \rightarrow \mathbb{R}$ satisfying*

$$R_Y = \rho \quad (4.38)$$

almost everywhere, almost surely. Then, there exists a unique sequence of non-negative real numbers $(a_n)_{n \in \mathbb{Z}^d}$ and a random phase field Φ such that

$$Y(\omega, t) = \sum_{n \in \mathbb{Z}^d} a_n e^{i(n \cdot t + \Phi_n(\omega))} \quad (4.39)$$

holds in $L^2(\mathbb{T}^d)$ almost surely.

Proof. Almost surely, we can write

$$Y(\omega, \cdot) \stackrel{L^2(\mathbb{T}^d)}{=} \sum_{n \in \mathbb{Z}^d} A_n(\omega) e^{i\Phi_n(\omega)} e_n \quad (4.40)$$

for some random variables $(A_n)_{n \in \mathbb{Z}^d}$ and $(\Phi_n)_{n \in \mathbb{Z}^d}$, with A_n chosen non-negative for all n . Thanks to Parseval identity, we can rewrite

$$R_{Y(\omega)}(\tau) = \frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} Y(\omega, t)Y(\omega, t + \tau)dt = \sum_{n \in \mathbb{Z}^d} A_n(\omega)^2 e^{in \cdot \tau} \quad (4.41)$$

and

$$\rho(\tau) = \sum_{n \in \mathbb{Z}^d} b_n e^{in \cdot \tau} \quad (4.42)$$

for some non-negative Fourier coefficients $(b_n)_{n \in \mathbb{Z}^d}$ thanks to Herglotz Theorem. Take $a_n = \sqrt{b_n}$ for each n , and conclude thanks to the uniqueness of Fourier coefficients. \square

Remark 4.6. *The result also holds under the assumption that the sample paths are almost surely in $L^2(\mathbb{T}^d)$.*

4.5.2 Marginal laws

The law of the marginal, say $X(0) = \sum_{n \in \mathbb{Z}^d} a_n \cos(\Phi_n)$, has already been studied by Blevins in [16] for series with a finite number of terms. We complete this study to fit our more general case of an infinite series that converges in $L^2(\Omega)$. Let us recall that in \mathcal{H}_1 we assume that $a_0 = 0$. Then, one can compute the normalized Kurtosis. Indeed, $\mathbb{E}[|X(0)|^2] = \sum_{n \in \mathbb{Z}^d} a_n^2$ according to Proposition 4.4 and thus $\mathbb{E}[|X(0)|^2]^2 = (\sum_{n \in \mathbb{Z}^d} a_n^2)^2$. Moreover recall that for Φ uniformly distributed in $\mathbb{R}/2\pi\mathbb{Z}$, $\mathbb{E}[\cos^2(\Phi)] = \frac{1}{2}$ and $\mathbb{E}[\cos^4(\Phi)] = \frac{3}{8}$, and thus for $\mathcal{A} \subset \mathbb{Z}^d$ such that $\mathcal{A} \cap -\mathcal{A} = \emptyset$ and $\mathcal{A} \cup -\mathcal{A} = \mathbb{Z}^d \setminus \{0\}$ (so $X(0) = 2 \sum_{n \in \mathcal{A}} a_n \cos(\Phi_n)$) and

$$\begin{aligned} \mathbb{E}\left[\left|\frac{X(0)}{2}\right|^4\right] &= \sum_{n \in \mathcal{A}} a_n^4 \mathbb{E}[\cos^4(\Phi_n)] + 3 \sum_{(l,m) \in \mathcal{A}^2, l \neq m} a_l^2 a_m^2 \mathbb{E}[\cos^2(\Phi_l)] \mathbb{E}[\cos^2(\Phi_m)] \\ &= \frac{3}{8} \sum_{n \in \mathcal{A}} a_n^4 + \frac{3}{4} \sum_{(l,m) \in \mathcal{A}^2, l \neq m} a_l^2 a_m^2 = \frac{3}{4} \left(\sum_{n \in \mathcal{A}} a_n^2 \right)^2 - \frac{3}{8} \sum_{n \in \mathcal{A}} a_n^4 \end{aligned} \quad (4.43)$$

so the kurtosis β_2 of $X(0)$ is given by

$$\beta_2 = \frac{\mathbb{E}[|X(0)|^4]}{\mathbb{E}[|X(0)|^2]^2} = 3 - \frac{3}{2} \frac{\sum_{n \in \mathcal{A}} a_n^4}{(\sum_{n \in \mathcal{A}} a_n^2)^2} = 3 - 3 \frac{\sum_{n \in \mathbb{Z}^d} a_n^4}{(\sum_{n \in \mathbb{Z}^d} a_n^2)^2} < 3, \quad (4.44)$$

which proves that $X(0)$ is not Gaussian.

Remark 4.7. *Actually $X(0)$ is not infinitely-divisible, and thus not Gaussian. Indeed, thanks to the independence hypothesis on Φ in \mathcal{H}_2 , one easily checks that the characteristic function of $X(0)$ is the (maybe infinite) product*

$$\mathbb{E}[e^{i\xi X(0)}] = \prod_{n \in \mathcal{A}} \mathbb{E}[e^{i2a_n \xi \cos(\Phi_n)}] = \prod_{n \in \mathcal{A}} J_0(2a_n \xi) \quad (4.45)$$

with J_0 the Bessel function of the first kind, which admit zeroes on the real line. Hence, the characteristic function of $X(0)$ cannot be the characteristic function of an infinitely divisible random variable (see Theorem 5.3. p. 108 of [88]).

Proposition 4.6. *$X(0)$ is sub-Gaussian. More precisely, for all λ in \mathbb{R} ,*

$$\mathbb{E}[e^{\lambda X(0)}] \leq e^{\lambda^2 \sum_{n \in \mathbb{Z}^d} a_n^2}. \quad (4.46)$$

Proof. First, notice that a centered random variable Y bounded by one is sub-Gaussian. Indeed, let $\lambda \in \mathbb{R}$, then $e^{\lambda Y} \leq \cosh(\lambda) + Y \sinh(\lambda)$ since $|Y| \leq 1$ and $x \mapsto e^{\lambda x}$ is convex. Then, using the fact that Y is centered we get $\mathbb{E}(e^{\lambda Y}) \leq \cosh(\lambda) \leq e^{\lambda^2/2}$. Now let $(\Phi_n)_{n \in \mathbb{Z}^d}$

be a pure phase noise field. For a finite sum $X_k(0) = \sum_{|n|_\infty \leq k} a_n \cos(\Phi_n)$, thanks to the independence hypothesis for a subset $\mathcal{A} \subset \mathbb{Z}^d$ such that $\mathcal{A} \cap -\mathcal{A} = \emptyset$ and $\mathcal{A} \cup -\mathcal{A} = \mathbb{Z}^d \setminus \{0\}$

$$\mathbb{E}[e^{\lambda X_k(0)}] = \prod_{n \in \mathcal{A}, |n|_\infty \leq k} \mathbb{E}[e^{\lambda 2a_n \cos(\Phi_n)}] \leq \prod_{n \in \mathcal{A}, |n|_\infty \leq k} e^{(\lambda 2a_n)^2/2} = e^{\lambda^2 \sum_{|n|_\infty \leq k} a_n^2} \quad (4.47)$$

holds for all $\lambda \in \mathbb{R}$, since $(\cos(\Phi_n))_{n \in \mathcal{A}}$ are independent centered random variables bounded by one.

For a general sum $X(0) = \sum_{n \in \mathbb{Z}^d} a_n \cos(\Phi_n)$ where $(a_n)_{n \in \mathbb{Z}^d}$ is a square summable family, let λ be any real number and notice that $\mathbb{E}[e^{\lambda X_k(0)}] \leq e^{\lambda^2 \sum_{|n|_\infty \leq k} a_n^2} \leq e^{\lambda^2 \sum_{n \in \mathbb{Z}^d} a_n^2}$. Moreover $X_k(0) \rightarrow X(0)$ almost surely, thus $e^{\lambda X_k(0)} \rightarrow e^{\lambda X(0)}$ almost surely and we can apply Fatou's lemma and conclude that (4.46) holds. \square

Proposition 4.7. *Assume that $(a_n)_{n \in \mathbb{Z}^d}$ is a family satisfying \mathcal{H}_1 and \mathcal{H}_2 , such that*

1. *there exists n_1, n_2, n_3 in \mathbb{Z}^d with $a_{n_1} a_{n_2} a_{n_3} \neq 0$;*
2. *$\{n_1, n_2, n_3\} \cap \{-n_1, -n_2, -n_3\} = \emptyset$.*

Then $X(0)$ admits a density function that is uniformly continuous and bounded over \mathbb{R} .

The proof is postponed to Appendix. Interestingly, in the cases where only one or two coefficients are non-zero, the resulting Random Phase Noise process has an unbounded density function.

4.6 Appendix: Proofs

4.6.1 Proof of Proposition 4.2

Proposition 4.2 is based on Proposition 13 pp. 55-56 in [69], we provide a proof for the sake of completeness. Let us first show a lemma, itself based on Proposition 12 p. 55 in Kahane [69].

Lemma 4.2. *Let $(u_k)_{k \in \mathbb{N}}$ be a sequence in $\mathcal{C}^0(\mathbb{T}^d)$ with real or complex values, such that*

$$\limsup_k \|u_k\|_\infty > 0. \quad (4.48)$$

Let $(\Psi_k)_{k \in \mathbb{N}}$ be a sequence of independent random variables uniformly distributed on \mathbb{T}^d . Then, almost surely, there exists T (random) in \mathbb{T}^d such that

$$\limsup_k |u_k(T - \Psi_k)| > 0. \quad (4.49)$$

Proof. Since $\limsup_k \|u_k\|_\infty > 0$ by assumption, there exists both some $\eta > 0$ and a subsequence $(k_p)_{p \in \mathbb{N}}$ such that

$$\|u_{k_p}\|_\infty > \eta \quad (4.50)$$

for all p . Thanks to continuity, $u_{k_p}(t) > \eta$ for t in an open ball $B_\infty(t_p, \varepsilon_p)$. Thus, $|u_{k_p}(t - \Psi_{k_p})| > \eta$ holds for t in a random open ball $U_p := B_\infty(T_p, \varepsilon_p)$ whose center is a random

variable $T_p := t_p + \Psi_{k_p}$ that is equidistributed on \mathbb{T}^d . Moreover $(T_p)_{p \in \mathbb{N}}$ is *i.i.d.* since $(\Psi_k)_{k \in \mathbb{N}}$ is assumed to be *i.i.d.*

Now, let us show that $\limsup_p U_p$ is almost surely non-empty (it can be shown that it is actually almost surely dense). Let t be any point in \mathbb{T}^d , $\varepsilon > 0$ be a positive number and denote $U := B_\infty(t, \varepsilon)$. For each p , $\mathbb{P}(U_p \cap U \neq \emptyset) \geq \mathbb{P}(T_p \in U) = \text{vol}(U)/(2\pi)^d$ since T_p is equidistributed on \mathbb{T}^d . Thus $\sum_p \mathbb{P}(U_p \cap U \neq \emptyset) = \infty$, and since the events $\{\omega | U_p(\omega) \cap U \neq \emptyset\}$ are independent, it follows thanks to Borel-Cantelli Lemma that almost surely $U_p \cap U \neq \emptyset$ happens for infinitely many p . Thus, almost surely, $\limsup_p U_p \neq \emptyset$.

Let us pick some random T in $\limsup_p U_p$ and notice that $\limsup_p |u_{k_p}(T - \Psi_{k_p})| > \eta$ almost surely since T belongs to infinitely many U_p . This concludes the proof. \square

Let us now prove Proposition 4.2.

First, let us recall that since for all k the random variable Y_k (in $\mathcal{C}^0(\mathbb{T}^d)$) is assumed to be symmetric, Itô-Nisio Theorem applies. Hence, the series $\sum_k Y_k$ converges almost surely in $\mathcal{C}^0(\mathbb{T}^d)$ if and only if any subsequence converges in $\mathcal{C}^0(\mathbb{T}^d)$ in probability.

Let us assume that the conclusion does not hold. Then, there must exist some $\eta > 0$ and two sequences of integers $(k_p)_{p \in \mathbb{N}}$ and $(k'_p)_{p \in \mathbb{N}}$ such that $k_p < k'_p < k_{p+1}$ for each p and

$$\mathbb{P} \left(\left\| \sum_{k=k_p+1}^{k_{p+1}} Y_k \right\|_\infty > \eta \right) > \eta \quad (4.51)$$

for all p .

Let $(\Omega', \mathbb{P}_{\Omega'})$ denote the probability space $\Omega \times \Omega_\Psi$ with $\mathbb{P}_{\Omega'} = \mathbb{P}_\Omega \otimes \mathbb{P}_\Psi$, where Ω_Ψ is a probability space in which there is a sequence $(\Psi_n)_{n \in \mathbb{N}}$ of *i.i.d.* random variables equidistributed on \mathbb{T}^d . Let us write $Z_p = \sum_{k=k_p+1}^{k_{p+1}} Y_k$ for all p , and let us consider the series of functions $\sum Z_p(\cdot)$ and $\sum Z_p(\cdot - \Psi_p)$ as random series (in the probability space Ω') of elements in $\mathcal{C}^0(\mathbb{T}^d)$. Since the Y_k (k in \mathbb{N}) are independent and symmetrical (by 1.), so are the Z_p (p in \mathbb{N}). Since for all k , Y_k and its translates have the same law (by 2.), Z_p and $Z_p(\cdot - \Psi_p)$ have the same law for each p . Moreover, since the sequence $(\sum_{k \leq l} Y_k)_{l \in \mathbb{N}}$ is almost surely bounded in $\mathcal{C}^0(\mathbb{T}^d)$ (by 3.), the series of functions $\sum_p Z_p$ is also almost surely bounded.

Moreover $(Z_p)_{p \in \mathbb{N}}$ is a sequence of independent variables and $\mathbb{P}_\Omega(\|Z_p\|_\infty > \eta) > \eta$ for each p , and thus $\sum_p \mathbb{P}_\Omega(\|Z_p\|_\infty > \eta) = \infty$. Hence Borel-Cantelli lemma applies and, almost surely (in Ω), $\limsup_p \|Z_p\|_\infty > \eta$. As a consequence, almost surely (in Ω'), $\limsup_p \|Z_p\|_\infty > \eta$. Lemma 4.2 yields that almost surely in Ω' , $\limsup_p |Z_p(T - \Psi_p)| > 0$ for some (random) T in \mathbb{T}^d .

Let us introduce another probability space $\Omega'' = \Omega' \times \Omega_\varepsilon$ ($\mathbb{P}_{\Omega''} = \mathbb{P}_\Omega \otimes \mathbb{P}_\Psi \otimes \mathbb{P}_\varepsilon$) and a Rademacher sequence $(\varepsilon_p)_{p \in \mathbb{N}}$. We now consider the random series of functions $\sum_p Z_p(t - \Psi_p)$ and $\sum_p \varepsilon_p Z_p(t - \Psi_p)$ on the space Ω'' . Since the random functions Z_p are symmetric, the partial sums have the same law in Ω'' . Moreover, since $\limsup_p |Z_p(T - \Psi_p)| > 0$ almost surely (in Ω''),

$$\sum_{p=1}^{\infty} |Z_p(T - \Psi_p)|^2 = \infty \quad (4.52)$$

holds almost surely (in Ω''). Thus the sequence $(\sum_1^N \varepsilon_p Z_p(t - \Psi_p))_{N \in \mathbb{N}}$ is almost surely (in Ω'') not bounded for some (random) T , thanks to a classic consequence of Paley-Zygmund inequalities (see [69] Theorem 1 p. 54).

To conclude, recall that $(\sum_{p=1}^N Z_p)_{N \in \mathbb{N}}$ is assumed to be almost surely bounded in $\mathcal{C}^0(\mathbb{T}^d)$ (in the probability space Ω and thus also in Ω''). Finally, notice that Z_p and $\varepsilon_p Z_p(\cdot - \Psi_p)$ have the same law in Ω'' and thus $(\varepsilon_p \sum_{p=1}^N Z_p(\cdot - \Psi_p))_{N \in \mathbb{N}}$ must also be almost surely bounded, which is a contradiction.

4.6.2 Proof of Proposition 4.3

The proof is a generalization to $d \geq 2$ of Kahane's [69] Proposition 5 p. 49. We begin with a lemma that gives a Bernstein's inequality for a multivariate trigonometric polynomial. In the following result, $\|\cdot\|$ denotes the norm on linear forms induced by $|\cdot|_\infty$ the maximum norm over \mathbb{R}^d , and $\nabla P(t)$ denotes the gradient of the trigonometric polynomial P at point t .

Lemma 4.3. *Let K be some positive integer and P a trigonometric polynomial on \mathbb{T}^d with degree less than K defined by $P(t) = \sum_{|n|_\infty \leq K} b_n e^{i(n \cdot t + \phi_n)}$ for all t in \mathbb{T}^d . Then*

$$\sup_{t \in \mathbb{T}^d} \|\nabla P(t)\| \leq K \|P\|_\infty. \quad (4.53)$$

Proof. Let us denote $(\theta_k)_{1 \leq k \leq d}$ the canonical basis of \mathbb{R}^d . Let us introduce for $t \in \mathbb{T}^d$, $1 \leq k \leq d$, the real trigonometric polynomial $Q_k(r) = P(t + r\theta_k)$. According to Bernstein's inequality one has

$$\|Q'_k\|_\infty \leq \max_{|n|_\infty \leq K} |n \cdot \theta_k| \|Q_k\|_\infty,$$

which involves that

$$\left| \frac{\partial P}{\partial t_k}(t) \right| = |Q'_k(0)| \leq K \|P\|_\infty,$$

and proves (4.53). \square

We now turn to the proof of Proposition 4.3. Let $\varepsilon \geq 2\pi/q$ and $t \in \mathbb{T}^d$. The function $s \mapsto |P(s)|$ is $2\pi/q$ -periodic on each component. Indeed, write $\mathcal{E} = \{l + qj\}_{j \in \mathcal{E}'}$ ($\mathcal{E}' \subset \mathbb{Z}^d$) and notice that

$$|P(s)| = \left| \sum_{j \in \mathcal{E}'} b_{l+qj} e^{i((l+qj) \cdot s + \phi_n)} \right| = \left| e^{i(l \cdot s)} \sum_{j \in \mathcal{E}'} b_{l+qj} e^{i(qj \cdot s + \phi_n)} \right| = \left| \sum_{j \in \mathcal{E}'} b_{l+qj} e^{i(qj \cdot s + \phi_n)} \right| \quad (4.54)$$

for each s in \mathbb{T}^d . Let t' in \mathbb{T}^d be such that $|P|$ achieves its global maximum $\|P\|_\infty$ at point t' that may be assumed to be in $\mathbb{B}_\infty(t, \varepsilon/2)$, thanks to the $2\pi/q$ -periodicity of $s \mapsto |P(s)|$. For all s in \mathbb{T}^d ,

$$|P(s) - P(t')| = P(t') - P(s) = \|P\|_\infty - P(s) \leq \sup_u \|\nabla P(u)\| |t' - s| \leq K \|P\|_\infty |t' - s| \quad (4.55)$$

thanks to Lemma 4.3, and thus

$$P(s) \geq \|P\|_\infty / 2 \quad (4.56)$$

for all s in $\mathbb{B}_\infty(t', 1/(2K)) = \mathbb{B}_\infty(t', \varepsilon')$. Since $K \geq \frac{q}{2\pi}$, ε' satisfies

$$\varepsilon' \leq \pi/q \leq \varepsilon/2 \quad (4.57)$$

and thus $s \mapsto |P(s)|$ achieves its global maximum $\|P\|_\infty$ on a point t' such that $\mathbb{B}_\infty(t', \varepsilon') \subset \mathbb{B}_\infty(t, \varepsilon)$.

4.6.3 Proof of Lemma 4.4

Define $\mathcal{B}_+ = \mathcal{B} \cap (\mathbb{N} \times \mathbb{Z}^{d-1})$ and $\mathcal{B}_- = \mathcal{B} \cap (-\mathbb{N} \times \mathbb{Z}^{d-1})$. Notice that, with non-zero probability, at least one of the sequences among $(\sum_{n \in \mathcal{B}_+ \cap \mathcal{A}_k} a_n e^{i(n \cdot T + \Phi_n)})_{k \in \mathbb{N}}$ and $(\sum_{n \in \mathcal{B}_- \cap \mathcal{A}_k} a_n e^{i(n \cdot T + \Phi_n)})_{k \in \mathbb{N}}$ diverges. Thus, we can define \mathcal{B}' , a deterministic subset of \mathbb{Z}^d either equal to \mathcal{B}_+ or equal to $-(\mathcal{B}_-)$, such that, with non-zero probability, the sequence $(Z_k)_{k \in \mathbb{N}}$ defined by

$$Z_k = \sum_{n \in \mathcal{A}_k \cap \mathcal{B}'} a_n e^{i(n \cdot T + \Phi_n)} \quad (4.58)$$

diverges. Define $X_k := \operatorname{Re}(Z_k) = \sum_{n \in \mathcal{A}_k \cap \mathcal{B}'} a_n \cos(n \cdot T + \Phi_n)$ and $Y_k := \operatorname{Im}(Z_k) = \sum_{n \in \mathcal{A}_k \cap \mathcal{B}'} a_n \sin(n \cdot T + \Phi_n)$ for all k . With non-zero probability $(X_k)_{k \in \mathbb{N}}$ or $(Y_k)_{k \in \mathbb{N}}$ diverges. Let us define the events

$$E_{\operatorname{div}}^{\cos} := \{\omega \mid \text{the sequence } (X_k)_{k \in \mathbb{N}} \text{ diverges}\} \quad (4.59)$$

and

$$E_{\operatorname{div}}^{\sin} := \{\omega \mid \text{the sequence } (Y_k)_{k \in \mathbb{N}} \text{ diverges}\}. \quad (4.60)$$

The event

$$E_{\operatorname{div}} := E_{\operatorname{div}}^{\cos} \cup E_{\operatorname{div}}^{\sin} \quad (4.61)$$

happens with non-zero probability.

Since $(\Phi_n)_{n \in \mathbb{N} \times \mathbb{Z}^{d-1}}$ and $(\Phi_n + \frac{\pi}{2})_{n \in \mathbb{N} \times \mathbb{Z}^{d-1}}$ have the same law (direct consequence of \mathcal{H}_2), the events $E_{\operatorname{div}}^{\cos}$ and $E_{\operatorname{div}}^{\sin}$ have the same probability. Thus, the probability of the event $E_{\operatorname{div}}^{\cos}$ is non-zero. We conclude by defining $\mathcal{B}^* = \mathcal{B}' \cup (-\mathcal{B}')$ and noticing that $\sum_{n \in \mathcal{A}_k \cap \mathcal{B}^*} a_n \cos(n \cdot T + \Phi_n) = 2X_k$ for each k .

4.6.4 Proof of Proposition 4.7

Since the law $\mathbb{P}_{X(0)}$ of the limit $X(0)$ does not depend on a method of summation (Proposition 4.1), let us pick one ordering in $\mathbb{N} \times \mathbb{Z}^{d-1}$, $(n_k)_{k \in \mathbb{N}}$, and rewrite $b_k := a_{n_k}$ for each $k \in \mathbb{N}$. For simplicity, let us assume that $b_k \neq 0$ for each k . We may write

$$X(0) = \sum_{k=1}^3 b_k \cos(\Phi_{n_k}) + \sum_{k=4}^{+\infty} b_k \cos(\Phi_{n_k}) = Y + Z.$$

By independence one has $\mathbb{P}_{X(0)} = \mathbb{P}_Y * \mathbb{P}_Z$. Since the convolution of a probability measure with an absolutely continuous measure with uniformly continuous bounded density remains an absolutely continuous measure with uniformly continuous bounded density, it is sufficient

to prove that $Y = \sum_{k=1}^3 b_k \cos(\Phi_{n_k})$ admits a uniformly continuous bounded density.

Let us recall that the density function of $b \cos(\Phi)$, where Φ is a uniform random variable over $\mathbb{R}/2\pi\mathbb{Z}$ and b a non-zero real number, is $t \mapsto \frac{1}{b} f\left(\frac{t}{b}\right)$, with $f(t) = \mathbb{1}_{(-1,1)} \frac{1}{\pi\sqrt{1-t^2}}$. Moreover, one easily checks that f is in $L^p(\mathbb{R})$ for every p in $[1, 2)$ and hence so is the density function of the random variable $b \cos(\Phi)$. Let f_1 (resp. f_2, f_3) denote the density function of $b_1 \cos(\Phi_{n_1})$ (resp. $b_2 \cos(\Phi_{n_2}), b_3 \cos(\Phi_{n_3})$) such that the density of Y is given by $f_1 * f_2 * f_3$. Recall that Young's inequalities (see *e.g.* [86] p. 99) state that if p, q, r are in $[1, \infty]$, such that

$$\frac{1}{p} + \frac{1}{q} = 1 + \frac{1}{r} \quad (4.62)$$

and if $f \in L^p(\mathbb{R})$ and $g \in L^q(\mathbb{R})$ one has $f * g \in L^r(\mathbb{R})$ with $\|f * g\|_r \leq \|f\|_p \|g\|_q$. It follows that the convolution $f_1 * f_2$ belongs to $L^p(\mathbb{R})$ for all p in $[1, \infty)$, since f_1 and f_2 are in $L^p(\mathbb{R})$ for every p in $[1, 2)$. In particular it belongs to $L^3(\mathbb{R})$. Moreover, f_3 belongs to $L^{3/2}(\mathbb{R})$ and since 3 and 3/2 are conjugate exponents ($1/3 + 2/3 = 1$), $(f_1 * f_2) * f_3$ is uniformly continuous and bounded (see *e.g.* [86] p. 70).

Chapter 5

On the Regularity of some Multiple Random Fourier Series

We propose a generalization of classical one-dimensional conditions for the convergence of random Fourier series over the d -torus \mathbb{T}^d . We provide an investigation of the regularity of the random infinite sum and discuss sufficient and necessary conditions on the modulus of Fourier coefficients. An emphasis is put onto anisotropic geometries that distinguish the multi-dimensional case from the classical one-dimensional case.

5.1 Introduction

In this chapter, we pursue our discussion on Random Fourier series, with a focus on conditions for regularity. Such conditions, especially for continuity and uniform convergence have been one of the main focus of interest on this topic, with the first results obtained by Paley and Zygmund in their seminal series of papers [108], [109] and [110]. Major contributions were made by Salem and Zygmund in [119], Kahane in [69], Marcus in [93], along with Jain [61] and Pisier [96]. Fernique introduced metric entropy methods to the field in [45] and [46] and found the first necessary and sufficient conditions for Gaussian periodic stationary processes. These results led to generalizations by Marcus and Pisier in [95] and [96] to compact groups (both abelian and non-abelian). Marcus and Pisier used such necessary and sufficient conditions to derive sufficient conditions based on the Fourier coefficients of random Fourier series ([96] chapter 7 and [97] section 4 – see also [26] section 3) ⁽¹⁾. In this chapter, we obtain similar conditions through a direct method inspired by Kahane [69] chapter 7 that doesn't involve metric entropy. Conditions for Hölder regularity have also been investigated in Kahane's book [69], chapter 7. In this chapter, we focus on similar conditions on the Fourier coefficients and extend existing results for usual random Fourier series (defined on the circle $\mathbb{R}/2\pi\mathbb{Z}$) to multiple random Fourier series defined on the d -torus \mathbb{T}^d , using an anisotropic notion of Hölder regularity.

⁽¹⁾We would like to thank C. Cuny for pointing to the references [97] and [26].

In order to articulate such an extension, we write our main hypotheses and recall some basic results in Section 2. In Section 3, we introduce a fundamental lemma and state sufficient conditions for continuity. In Section 4, we focus on the converse problem and state necessary conditions for almost sure unboundedness. In Section 5, we focus on anisotropic Hölder regularity and state sufficient conditions as well as necessary conditions in the isotropic framework. In Section 6 we discuss extensions such as the Gaussian case.

5.2 Notations and Hypotheses

Recall that $(\Omega, \mathcal{F}, \mathbb{P})$ denotes a complete probability space, $\mathbb{T}^d := \mathbb{R}^d/2\pi\mathbb{Z}^d$ the d -dimensional torus and $(e_n)_{n \in \mathbb{Z}^d}$ the trigonometric Hilbert basis of $L^2(\mathbb{T}^d)$. In Chapter 4 we considered random continuous functions X defined by their random Fourier coefficients by

$$X(\omega, \cdot) \stackrel{L^2(\mathbb{T}^d)}{=} \sum_{n \in \mathbb{Z}^d} A_n(\omega) e^{i\Phi_n(\omega)} e_n(\cdot) \quad (5.1)$$

where the sum in $L^2(\mathbb{T}^d)$ can also be taken in the sense of uniform convergence.

Let us recall the following hypotheses that we made on the random Fourier coefficients:

- \mathcal{H}_1 : “ $(A_n)_{n \in \mathbb{Z}^d} = (a_n)_{n \in \mathbb{Z}^d}$ is a deterministic, non-negative, even ($a_{-n} = a_n$ for all n in \mathbb{Z}^d) square summable family with $a_0 = 0$.”
- \mathcal{H}_2 : “ $(\Phi_n)_{n \in \mathbb{Z}^d}$ is a pure phase noise field, that is for all n in \mathbb{Z}^d , $\Phi_{-n} = -\Phi_n$ (modulo 2π) almost surely, Φ_n has uniform distribution over $\mathbb{R}/2\pi\mathbb{Z}$ and $(\Phi_n)_{n \in \mathcal{A}}$ are independent for all $\mathcal{A} \subset \mathbb{Z}^d$ such that \mathcal{A} and $-\mathcal{A}$ do not intersect.”

Recall also that $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ denotes the sequence of random trigonometric functions defined for t in \mathbb{T}^d and k in \mathbb{N} by

$$S_{\mathcal{A}_k}(t) = \sum_{n \in \mathcal{A}_k} A_n e^{in \cdot t + \Phi_n}. \quad (5.2)$$

For any trigonometric polynomial

$$p : t \mapsto \sum_{n \in \mathcal{A}} a_n e^{i(n \cdot t + \phi_n)} \quad (5.3)$$

we consider its degree defined as

$$d_\infty(p) := \max\{|n|_\infty; n \in \mathcal{A}, a_n \neq 0\}, \quad (5.4)$$

with $|n|_\infty = \max_{1 \leq i \leq d} |n_i|$.

In the remaining of this chapter, we shall seek explicit conditions on $(a_n)_{n \in \mathbb{Z}^d}$ for the set equivalent conditions in Theorem 4.1 to be satisfied. We name this set of equivalent conditions (\star) . We shall also seek conditions for stronger notions of regularity, such as Hölder regularity. Recall that for α in $(0, 1)$, a continuous function $f : \mathbb{T}^d \mapsto \mathbb{R}$ is said to be α -Hölder over \mathbb{T}^d if

$$\sup_{t \neq s \in \mathbb{T}^d} \frac{|f(t) - f(s)|}{|t - s|^\alpha} < \infty, \quad (5.5)$$

where $|\cdot|$ is the Euclidean norm on \mathbb{R}^d . For $p \in [1, +\infty]$, we may consider the p -norm defined by $|x|_p = (\sum_{1 \leq i \leq d} |x_i|^p)^{1/p}$ for $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ and their corresponding closed balls $B_p(r) = \{x \in \mathbb{R}^d; |x|_p \leq r\}$, $r > 0$. Choosing $\mathcal{A}_k(p) = B_p(r_k)$ for $k \in \mathbb{N}$, with an increasing positive sequence $(r_k)_k$, we obtain different conditions, relying on sums of the a_n^2 on subsets of \mathbb{Z}^d delimited by these sets, implying all Hölder regularity, in view of the equivalence of norms on \mathbb{R}^d . Following [11] (see also *e.g.* [7]), we also consider an anisotropic generalization of the Hölder regularity property using quasi-norms instead of norms. More precisely, if $E = \text{diag}(\mu_1, \dots, \mu_d)$ is a diagonal matrix with positive eigenvalues $\mu_1, \dots, \mu_d \in (0, +\infty)$, we consider $\tau_E : \mathbb{R}^d \rightarrow \mathbb{R}^+$ a continuous even function such that

- i) for all $x \neq 0$, $\tau_E(x) > 0$;
- ii) for all $r > 0$ and all $x \in \mathbb{R}^d$, $\tau_E(r^E x) = r \tau_E(x)$ with $r^E = \exp((\log r)E) = \text{diag}(r^{\mu_1}, \dots, r^{\mu_d})$.

Note that τ_E remains faithful by i), the homogeneity property of norms is replaced by the E -homogeneity property ii) and τ_E satisfies a quasi-triangular inequality (by Lemma 2.2 of [13]): there exists $\kappa_E \geq 1$ such that

$$\forall x, y \in \mathbb{R}^d, \tau_E(x + y) \leq \kappa_E (\tau_E(x) + \tau_E(y)). \quad (5.6)$$

Then, for $\alpha > 0$, a continuous function $f : \mathbb{T}^d \mapsto \mathbb{R}$ is said to be (α, τ_E) -Hölder over \mathbb{T}^d if

$$\sup_{t \neq s \in \mathbb{T}^d} \frac{|f(t) - f(s)|}{(\tau_E(t - s))^\alpha} < \infty. \quad (5.7)$$

Hence we may consider closed balls $B_{\tau_E}(r) = \{x \in \mathbb{R}^d; \tau_E(x) \leq r\}$, $r > 0$ and a corresponding method of summation $\mathcal{A}_k(\tau_E) = B_{\tau_E}(r_k)$, for $k \in \mathbb{N}$, with an increasing positive sequence $(r_k)_{k \in \mathbb{N}}$. Let us quote that this framework generalizes the previous one since any norm is a quasi-norm for $E = I_d$ the identity matrix. A simple example of a quasi-norm is given by $\tau_E(x) = \sum_{1 \leq i \leq d} |x_i|^{H_i}$ with $H_i = 1/\mu_i$ for $1 \leq i \leq d$ and $x = (x_1, \dots, x_d) \in \mathbb{R}^d$. Moreover, when $H_i \in (0, 1]$ for all $1 \leq i \leq d$, this quasi-norm satisfies the triangular inequality, meaning that $\kappa_E = 1$.

In the following, C_d (resp. C'_d etc.) denotes a universal constant. However, its value may vary across different contexts, but as long as ambiguity is avoided, we shall keep the same notations purposely.

5.3 Conditions for continuity

We are interested in sufficient conditions on the modulus sequence $(a_n)_{n \in \mathbb{Z}^d}$ for a RPN process to be well defined, that is, to satisfy the equivalent conditions of Billard's theorem named (\star) . Sufficient conditions for continuity have been thoroughly studied in the one-dimensional case (see Marcus *et al.* [98], [93], [94], Kahane [69]) and methods introduced by Fernique [45], led to the generalization of necessary and sufficient conditions to any compact group as shown by Marcus and Pisier in [96].

Our problem is to effectively derive whether or not a given random Fourier series satisfying hypotheses \mathcal{H}_1 and \mathcal{H}_2 converges uniformly, and we thus seek conditions directly on the coefficients $(a_n)_{n \in \mathbb{Z}^d}$ as in [96], [97] and [26], with elementary techniques, as found in [69].

5.3.1 Preliminary results

In the following we denote $\mathbb{B}_\infty(t, r) = \{s \in \mathbb{T}^d; |t - s|_\infty < r\}$ the projection onto \mathbb{T}^d of the \mathbb{R}^d open ball of radius r and center t with respect to $|\cdot|_\infty$ onto \mathbb{T}^d . Let us first recall a result inspired by [69], Chapter 5.

Most of the results in this chapter rely on the following proposition, which is inspired by Theorem 2 in [69], Chapter 6.

Lemma 5.1. *Let $K \geq 1$ and $(\varepsilon_k)_{1 \leq k \leq K}$ be a finite family of independent real random variables that are sub-normal, i.e. sub-Gaussian and such that*

$$\mathbb{E}[e^{\lambda \varepsilon_k}] \leq e^{\lambda^2/2} \quad (5.8)$$

for each k . Let $(p_k)_{1 \leq k \leq K}$ be a finite family of real trigonometric polynomials with degree less than N , where $N \geq 1$.

Define $P(\omega, t) := \sum_{1 \leq k \leq K} \varepsilon_k(\omega) p_k(t)$ as a random trigonometric polynomial. Then

$$\mathbb{P} \left(\|P\|_\infty \geq C_d \cdot \left(\log(N) \cdot \sum_{k=1}^K \|p_k\|_\infty^2 \right)^{1/2} \right) \leq \frac{(2\pi)^d}{N} \quad (5.9)$$

holds for some universal constant $C_d > 0$.

Proof. Let us define $r^2 := \sum_{k=1}^K \|p_k\|_\infty^2$. Thanks to the independence of the $(\varepsilon_k)_{1 \leq k \leq K}$, for $t \in \mathbb{T}^d$,

$$\forall \lambda \in \mathbb{R}, \quad \mathbb{E}[e^{\lambda P(t)}] = \mathbb{E}[e^{\lambda \sum_{k=1}^K \varepsilon_k p_k(t)}] = \prod_{k=1}^K \mathbb{E}[e^{\lambda \varepsilon_k p_k(t)}]. \quad (5.10)$$

Due to the subnormality of the $(\varepsilon_k)_{1 \leq k \leq K}$,

$$\mathbb{E}[e^{\lambda P(t)}] = \prod_{k=1}^K \mathbb{E}[e^{\lambda \varepsilon_k p_k(t)}] \leq \prod_{k=1}^K e^{\lambda^2 |p_k(t)|^2/2} \leq e^{\lambda^2 r^2/2}, \quad (5.11)$$

which holds for every t in \mathbb{T}^d .

Thanks to Proposition 4.3 (with $q = 1$), we know that (surely) there exists a random ball $B(T(\omega), \varepsilon)$, with $\varepsilon \geq (2N)^{-1}$, such that $|P(\omega, t)| \geq \|P(\omega, \cdot)\|_\infty/2$ for every t in $B_\infty(T(\omega), \varepsilon)$.

Since $\text{vol}(B_\infty(T(\omega), \varepsilon)) = (2\varepsilon)^d \geq N^{-d}$,

$$\mathbb{E}[e^{\lambda \|P\|_\infty/2}] \leq N^d \mathbb{E} \left[\int_{B_\infty(T, \varepsilon)} (e^{\lambda P(t)} + e^{-\lambda P(t)}) dt \right] \leq N^d \mathbb{E} \left[\int_{\mathbb{T}^d} (e^{\lambda P(t)} + e^{-\lambda P(t)}) dt \right] \quad (5.12)$$

holds for every λ in \mathbb{R} . Thanks to Fubini-Tonelli, (5.11) yields for every $\lambda \in \mathbb{R}$

$$\mathbb{E}[e^{\lambda \|P\|_\infty/2}] \leq 2(2\pi N)^d e^{\lambda^2 r^2/2} \quad (5.13)$$

which can be rewritten for $\lambda > 0$ as

$$\mathbb{E} \left[e^{\lambda/2(\|P\|_\infty - \lambda r^2 - \frac{2}{\lambda} \log(2(2\pi N)^d s))} \right] \leq 1/s \quad (5.14)$$

for every $s > 0$. By Markov inequality, it follows that

$$\mathbb{P} \left(\|P\|_\infty \geq \lambda r^2 + \frac{2}{\lambda} \log(2(2\pi N)^d s) \right) \leq 1/s \quad (5.15)$$

for any $s > 0$ and $\lambda > 0$. Let us fix $s^* = N/(2\pi)^d$ and $\lambda^* = (2 \log(2(2\pi N)^d s^*)/r^2)^{1/2}$. From (5.15), it follows that

$$\mathbb{P} \left(\|P\|_\infty \geq C_d (\log(N)r^2)^{1/2} \right) \leq (2\pi)^d/N \quad (5.16)$$

holds for some universal constant C_d (depending only on the dimension d). \square

Let us now discuss a very particular and useful case of sub-normal random variable. Recall that ε is a Rademacher random variable if

$$\mathbb{P}(\varepsilon = 1) = \mathbb{P}(\varepsilon = -1) = 1/2. \quad (5.17)$$

A Rademacher random variable is subnormal since, thanks to monotone convergence,

$$\mathbb{E}[e^{\lambda\varepsilon}] = \sum_{k \in \mathbb{N}} \mathbb{E} \left[\frac{\lambda^k \varepsilon^k}{k!} \right] = \sum_{k \in \mathbb{N}} \frac{\lambda^{2k}}{(2k)!} \leq \sum_{k \in \mathbb{N}} \frac{\lambda^{2k}}{2^k k!} = e^{\lambda^2/2} \quad (5.18)$$

for all λ in \mathbb{R} . In the following we shall make an intense use of Rademacher random variables that are independent modulo symmetry.

Definition 5.1. $(\varepsilon_n)_{n \in \mathbb{Z}^d}$ is a symmetric Rademacher random field if

- ε_n is a Rademacher random variable for each n
- $\varepsilon_{-n} = \varepsilon_n$ almost surely for each n
- $(\varepsilon_n)_{n \in \mathcal{A}}$ are independent if $\mathcal{A} \subset \mathbb{Z}^d$ is such that $(\mathcal{A} \setminus \{0\}) \cap (-\mathcal{A}) = \emptyset$.

As we stated Proposition 5.1 for a sequence of real trigonometric polynomials, it will be useful to consider sums on halves of \mathbb{Z}^d so we denote $(\mathbb{Z}^d)^+$ the set

$$\{(n_1, \dots, n_d); (n_1 > 0) \text{ or } (n_1 = 0 \text{ and } n_2 > 0) \dots \text{ or } (n_1 = 0, n_2 = 0 \dots, n_{d-1} = 0 \text{ and } n_d > 0)\} \quad (5.19)$$

such that $(\mathbb{Z}^d)^+ \cup -(\mathbb{Z}^d)^+ \cup \{0\} = \mathbb{Z}^d$ and $(\mathbb{Z}^d)^+ \cap -(\mathbb{Z}^d)^+ = \emptyset$, and \mathcal{A}^+ the set $\mathcal{A} \cap (\mathbb{Z}^d)^+$ for any symmetrical set \mathcal{A} .

Proposition 5.1. Assume that hypotheses $\mathcal{H}_1, \mathcal{H}_2$ are satisfied. Assume that $\mathcal{A} \subset \mathbb{Z}^d$ is a symmetrical subset such that $\mathcal{A} \subset B_\infty(N)$ for some $N \geq 1$. Then,

$$\mathbb{P} \left(\left\| \sum_{n \in \mathcal{A}} a_n e^{i(n \cdot t + \Phi_n)} \right\|_\infty \geq \sqrt{2} C_d \left(\log(N) \sum_{n \in \mathcal{A}} a_n^2 \right)^{1/2} \right) \leq (2\pi)^d/N. \quad (5.20)$$

Proof. First, recall that $a_0 = 0$ by Hypothesis \mathcal{H}_1 . Define the random polynomial

$$P^\Phi(t) = \sum_{n \in \mathcal{A}} a_n e^{i(n \cdot t + \Phi_n)}. \quad (5.21)$$

Let $(\varepsilon_n)_{n \in \mathbb{Z}^d}$ be a symmetric Rademacher random field, such that $(\varepsilon_n)_{n \in \mathbb{Z}^d}$ is independent of $(\Phi_n)_{n \in \mathbb{Z}^d}$ and define the random trigonometric polynomial

$$P^\varepsilon(t) = \sum_{n \in \mathcal{A}} \varepsilon_n a_n e^{i(n \cdot t + \Phi_n)}. \quad (5.22)$$

Note that since \mathcal{A} is finite, thanks to \mathcal{H}_1 , \mathcal{H}_2 , and the hypothesis that $(\varepsilon_n)_{n \in \mathbb{Z}^d}$ is independent of $(\Phi_n)_{n \in \mathbb{Z}^d}$, P^ε and P^Φ have the same law. Moreover,

$$P^\varepsilon(t) = \sum_{n \in \mathcal{A}^+} \varepsilon_n P_n(t) \quad (5.23)$$

where P_n is the real random trigonometrical polynomials defined by

$$P_n(t) = a_n e^{i(n \cdot t + \Phi_n)} + a_{-n} e^{i(-n \cdot t + \Phi_{-n})} = 2a_n \cos(n \cdot t + \Phi_n) \quad (5.24)$$

for n in \mathcal{A}^+ . Notice that for each n , $\|P_n\|_\infty = 2|a_n|$. Since $(\varepsilon_n)_{n \in \mathcal{A}^+}$ are independent and subnormal, $(\Phi_n)_{n \in \mathcal{A}^+}$ is independent of $(\varepsilon_n)_{n \in \mathcal{A}^+}$, and $N \geq 1$, Lemma 5.1 yields

$$\mathbb{P} \left(\|P^\varepsilon\|_\infty \geq \sqrt{2} C_d \left(\log(N) \sum_{n \in \mathcal{A}} a_n^2 \right)^{1/2} \mid (\Phi_n)_{n \in \mathcal{A}^+} \right) \leq (2\pi)^d / N. \quad (5.25)$$

Notice that neither the upper bound $(2\pi)^d / N$ nor the lower bound $\sqrt{2} C_d (\log(N) \sum_{n \in \mathcal{A}} a_n^2)^{1/2}$ depend on $(\Phi_n)_{n \in \mathcal{A}^+}$. The conditionality can thus be removed and (5.25) yields

$$\mathbb{P} \left(\|P^\varepsilon\|_\infty \geq \sqrt{2} C_d \left(\log(N) \sum_{n \in \mathcal{A}} a_n^2 \right)^{1/2} \right) \leq (2\pi)^d / N \quad (5.26)$$

which in turn yields the desired inequality

$$\mathbb{P} \left(\|P^\Phi\|_\infty \geq \sqrt{2} C_d \left(\log(N) \sum_{n \in \mathcal{A}} a_n^2 \right)^{1/2} \right) \leq (2\pi)^d / N, \quad (5.27)$$

since P^Φ and P^ε have the same law. \square

5.3.2 Sufficient conditions for almost sure continuity

Assume that Hypothesis \mathcal{H}_1 is satisfied and assume that $(N_k)_{k \in \mathbb{N}}$ is an increasing sequence of integers and that $(\mathcal{A}_k)_{k \in \mathbb{N}}$ is a method of summation such that, for each k , $\mathcal{A}_k \subset B_\infty(N_k)$. Let us state the two hypothesis

- \mathcal{H}_3 : $\sum_{k \in \mathbb{N}} \left(\log(N_{k+1}) \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} a_n^2 \right)^{1/2} < \infty$.

- \mathcal{H}_4 : $\sum_{k \in \mathbb{N}} \frac{1}{N_k} < \infty$.

Theorem 5.1. *Assume that hypotheses \mathcal{H}_1 , \mathcal{H}_2 , \mathcal{H}_3 and \mathcal{H}_4 are satisfied. Then (\star) holds.*

Proof. For each k , define the random polynomial

$$P_k^\Phi(t) = \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} a_n e^{i(n-t+\Phi_n)}. \quad (5.28)$$

Proposition 5.1 yields

$$\mathbb{P} \left(\|P_k^\Phi\|_\infty \geq \sqrt{2}C_d \left(\log(N_{k+1}) \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} a_n^2 \right)^{1/2} \right) \leq (2\pi)^d / N_{k+1}. \quad (5.29)$$

Since $\sum_{k \in \mathbb{N}} 1/N_k < \infty$ by Hypothesis \mathcal{H}_4 , Borel-Cantelli Lemma applies, and thus almost surely

$$\|P_k^\Phi\|_\infty \leq \sqrt{2}C_d \left(\log(N_{k+1}) \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} a_n^2 \right)^{1/2} \quad (5.30)$$

holds for all k in \mathbb{N} except maybe for a finite random number of them. Thanks to the Hypothesis \mathcal{H}_3 , it follows that the series of functions $\left(\sum_{l \leq k} P_l^\Phi \right)_{k \in \mathbb{N}}$ converges normally, almost surely. Hence, almost surely, $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ converges uniformly, which in turn implies condition (ii) in Theorem 4.1 for the method of convergence $(\mathcal{A}_k)_{k \in \mathbb{N}}$. Hence, condition (\star) holds. \square

Remark 5.1. *One possible choice of N_k (k in \mathbb{N}) is to take $\text{diam}_\infty(\mathcal{A}_k)$, the diameter of \mathcal{A}_k with respect to the $|\cdot|_\infty$ norm, as defined by*

$$\text{diam}_\infty(\mathcal{A}_k) := \min\{N \in \mathbb{N}; \mathcal{A}_k \subset B_\infty(N)\}. \quad (5.31)$$

Furthermore, notice that, thanks to equivalence of norms in \mathbb{R}^d , the diameter $\text{diam}_\infty(\cdot)$ can actually be replaced with the diameter with respect to any norm.

5.3.3 Examples

Conditions with norms and quasi-norms

Theorem 5.1 yields many practical and somewhat tractable conditions. We state conditions that deal with norms and quasi-norms.

Let $E = \text{diag}(\mu_1, \dots, \mu_d)$ be a diagonal matrix with positive eigenvalues $\mu_1, \dots, \mu_d \in (0, +\infty)$ and τ_E be a quasi-norm associated with E . Note that for $E = I$ we may choose τ_E as a p -norm on \mathbb{R}^d for any $p \in [1, +\infty]$. For $k \in \mathbb{N}$ we consider the symmetrical subset

$$\mathcal{A}_k(\tau_E) = B_{\tau_E}(2^k) = \{x \in \mathbb{R}^d; \tau_E(x) \leq 2^k\}. \quad (5.32)$$

Let $\underline{\mu} = \min_{1 \leq i \leq d} \mu_i$ and $\bar{\mu} = \max_{1 \leq i \leq d} \mu_i$. We also note $\underline{H} = \bar{\mu}^{-1} = \min_{1 \leq i \leq d} \mu_i^{-1}$ and $\bar{H} = \underline{\mu}^{-1} = \max_{1 \leq i \leq d} \mu_i^{-1}$. By Proposition 3.5 of [12] we may find $c_1, c_2 > 0$ such that for all $x \in \mathbb{R}^d$

$$c_1 \min(|x|_{\infty}^{\underline{H}}, |x|_{\infty}^{\bar{H}}) \leq \tau_E(x) \leq c_2 \max(|x|_{\infty}^{\underline{H}}, |x|_{\infty}^{\bar{H}}). \quad (5.33)$$

It follows that

$$B_{\infty} \left(c_2^{-\underline{\mu}} 2^{k\underline{\mu}} \right) \subset \mathcal{A}_k(\tau_E) \subset B_{\infty} \left(c_1^{-\bar{\mu}} 2^{k\bar{\mu}} \right). \quad (5.34)$$

Assume that Hypothesis \mathcal{H}_1 is satisfied and, for all k in \mathbb{N} , define $s_{\tau_E, k} \geq 0$ by

$$s_{\tau_E, k}^2 := 2^k \sum_{2^{2k} < \tau_E(n) \leq 2^{2k+1}} a_n^2 = 2^k \sum_{n \in \mathcal{A}_{2^{k+1}}(\tau_E) \setminus \mathcal{A}_{2^k}(\tau_E)} a_n^2. \quad (5.35)$$

Corollary 5.1. *Assume that hypotheses $\mathcal{H}_1, \mathcal{H}_2$ and*

$$\mathcal{H}_5(\tau_E) : \sum_{k \in \mathbb{N}} s_{\tau_E, k} < \infty$$

are satisfied, with $s_{\tau_E, k}^2$ defined by (5.35). Then condition (\star) holds.

Proof. Notice that by (5.34) we get $\mathcal{A}_{2^k}(\tau_E) \subset B_{\infty}(N_k)$ with

$$c_2^{-\bar{\mu}} 2^{2k\bar{\mu}} \leq N_k \leq c_1^{-\underline{\mu}} 2^{2k\underline{\mu}}.$$

It follows that \mathcal{H}_4 is satisfied and $2^{-k} \log(N_{k+1})$ remains bounded so that $\mathcal{H}_5(\tau_E)$ implies \mathcal{H}_3 . Then, the conclusion follows from Theorem 5.1. \square

Remark 5.2. *Kahane (in chapter 7 of [69]) states similar conditions for the case $d = 1$, namely*

$$\sum_{k \in \mathbb{N}} s_k < \infty \quad (5.36)$$

with $s_k^2 = 2^k \sum_{2^{2k} < n \leq 2^{2k+1}} a_n^2$ for k in \mathbb{N} . In the case where $d > 1$, the topology of \mathbb{R}^d allows for a greater variety of conditions, not only defined by norms.

Conditions with integrals

We assume in this section that the quasi-norm τ_E satisfies the triangular inequality, meaning $\kappa_E = 1$ in (5.6). In particular, this is the case when $E = I$ with τ_E a norm on \mathbb{R}^d . When E is a diagonal matrix with eigenvalues $\mu_1, \dots, \mu_d \geq 1$ we may choose τ_E defined by $\tau_E(x) = \sum_{i=1}^d |x_i|^{1/\mu_i}$. In this case $q_E := \text{tr}(E) = \sum_{i=1}^d \mu_i \geq d \geq 1$.

Proposition 5.2. *Let E and τ_E such that $\kappa_E = 1$ in (5.6) and $q_E = \text{tr}(E) \geq 1$. Let $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ with $r \mapsto r^{q_E-1} f^2(r)$ non-increasing. Let us consider \mathcal{H}_1 with $a_n = f(\tau_E(n))$. Assume that \mathcal{H}_2 and*

$$\mathcal{H}_6(\tau_E) : \sum_{k \in \mathbb{N}} \left(2^k \int_{2^{2k}}^{2^{2k+1}} r^{q_E-1} f^2(r) dr \right)^{1/2} < \infty$$

are satisfied. Then condition (\star) holds.

Proof. We first need the following sum-integral comparison:

Lemma 5.2. *Let E and τ_E such that $\kappa_E = 1$ and $q_E \geq 1$. Let $g : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be a function such that $r \mapsto r^{q_E-1}g(r)$ is non-increasing on \mathbb{R}_+ . Then there exists a constant $C > 0$ such that, for any two integers $1 \leq N_1 < N_2$,*

$$\sum_{N_1 < \tau_E(n) \leq N_2} g(\tau_E(n)) \leq C \int_{N_1-1}^{N_2-1} r^{q_E-1} g(r) dr. \quad (5.37)$$

Proof. Let $N \geq 1$. We want to give an upper bound for

$$\gamma_N := \sum_{N < \tau_E(n) \leq N+1} g(\tau_E(n)). \quad (5.38)$$

Notice that since $r \mapsto r^{q_E-1}g(r)$ is non-increasing, g is non-negative and $q_E \geq 1$, g must be non-increasing. Hence,

$$\gamma_N \leq \#\{n \in \mathbb{Z}^d; N < \tau_E(n) \leq N+1\}g(N). \quad (5.39)$$

Let us estimate $\#\{n \in \mathbb{Z}^d; N < \tau_E(n) \leq N+1\}$. For each n in \mathbb{Z}^d , we consider the subset of \mathbb{R}^d defined by $\mathcal{C}_n = n + (-\frac{1}{2}, \frac{1}{2}]^d$. Notice that for n_1 and n_2 in \mathbb{Z}^d , $\mathcal{C}_{n_1} \cap \mathcal{C}_{n_2} \neq \emptyset$ implies that $n_1 = n_2$. Hence, for \mathcal{E} subset of \mathbb{Z}^d , the volume of $\bigcup_{n \in \mathcal{E}} \mathcal{C}_n$ is also the cardinality of \mathcal{E} . It follows that,

$$\#\{n \in \mathbb{Z}^d; N < \tau_E(n) \leq N+1\} = \text{vol} \left(\left\{ x \in \mathbb{R}^d; x \in \bigcup_{n; N < \tau_E(n) \leq N+1} \mathcal{C}_n \right\} \right).$$

Using the triangular inequality and (5.33), for all $N > c_2$,

$$\bigcup_{n; N < \tau_E(n) \leq N+1} \mathcal{C}_n \subset B_{\tau_E}(N+1+c_2) \setminus B_{\tau_E}(N-c_2).$$

Now, for $r > 0$, by E -homogeneity, we get

$$\text{vol}(B_{\tau_E}(r)) = \int_{\mathbb{R}^d} \mathbf{1}_{\tau_E(x) \leq r} dx = \int_{\mathbb{R}^d} \mathbf{1}_{\tau_E(r^{-E}x) \leq 1} dx = v_{\tau_E} r^{q_E},$$

with $v_{\tau_E} = \text{vol}(B_{\tau_E}(1)) > 0$. It follows that

$$\#\{n \in \mathbb{Z}^d; N < \tau_E(n) \leq N+1\} \leq v_{\tau_E} ((N+1+c_2)^{q_E} - (N-c_2)^{q_E}).$$

Therefore we may find a constant C such that for all $N \geq 1$

$$\#\{n \in \mathbb{Z}^d; N < \tau_E(n) \leq N+1\} \leq CN^{q_E-1}.$$

Then,

$$\gamma_N \leq CN^{q_E-1}g(N) \leq C \int_{N-1}^N g(r)r^{q_E-1} dr,$$

using the fact that $r \mapsto r^{q_E-1}g(r)$ is non-increasing. Finally, summing up for N in $\{N_1, \dots, N_2-1\}$ yields the announced result. \square

Now let us prove Proposition 5.2. For any $k \in \mathbb{N}$, we apply Lemma 5.2 with $g = f^2$, $N_1 = 2^{2^k}$ and $N_2 = 2^{2^{k+1}}$, so that (5.37) becomes

$$2^{-k} s_{\tau_E, k}^2 \leq C \int_{2^{2^k-1}}^{2^{2^{k+1}}-1} r^{q_E-1} f^2(r) dr,$$

and consequently

$$\sum_{k \geq 1} s_{\tau_E, k} \leq \sum_{k \geq 1} \left(2^k C \int_{2^{2^k-1}}^{2^{2^{k+1}}-1} r^{q_E-1} f^2(r) dr \right)^{1/2} \leq C^{1/2} \sum_{k \geq 1} (I_k + 2I_{k-1})^{1/2}$$

$$\text{with } I_k = 2^k \int_{2^{2^k}}^{2^{2^{k+1}}} r^{q_E-1} f^2(r) dr.$$

Using the inequality $\sqrt{I_k + 2I_{k-1}} \leq \sqrt{I_k} + \sqrt{2I_{k-1}}$, we obtain

$$\sum_{k \geq 1} s_{\tau_E, k} \leq C^{1/2} (1 + \sqrt{2}) \sum_{k \in \mathbb{N}} (I_k)^{1/2}$$

which is finite thanks to $\mathcal{H}_6(\tau_E)$. Hence, $\mathcal{H}_5(\tau_E)$ is satisfied and Corollary 5.1 applies. \square

One important class of examples is obtained by considering the function $f(r) = r^{-\alpha}$ for $\alpha > 0$. Proposition 5.2 implies that the condition (\star) is satisfied for a RPN process with modulus $(a_n)_{n \in \mathbb{Z}^d}$ where $a_0 = 0$ and $a_n = \tau_E(n)^{-\alpha}$ for $n \neq 0$, with $\alpha > q_E/2$. Indeed, one derives easily that for each k in \mathbb{N}

$$\int_{2^{2^k}}^{2^{2^{k+1}}} r^{q_E-1} r^{-2\alpha} dr = \frac{1}{2\alpha - q_E} \left((2^{2^k})^{q_E-2\alpha} - (2^{2^{k+1}})^{q_E-2\alpha} \right) \leq \frac{1}{2\alpha - q_E} 2^{(q_E-2\alpha)2^k}$$

and thus

$$2^k \left(\int_{2^{2^k}}^{2^{2^{k+1}}} r^{q_E-1} r^{-2\alpha} dr \right)^{1/2} \leq \frac{1}{(2\alpha - q_E)^{1/2}} 2^{k - (2\alpha - q_E)2^{k-1}},$$

that is summable since $2\alpha - q_E > 0$.

Remark 5.3. Notice that without the assumption concerning the quasi-triangular inequality (5.6), one could not compare the sum with the integral.

5.4 Almost sure unboundedness

We may obtain a partial converse in the isotropic case for $E = I$, considering $\tau_E = |\cdot|_\infty$. Actually, in this setting, we can extend the original Paley-Zygmund approach in the case where $d = 1$ to the general case where d is any positive integer.

For all k in \mathbb{N} , let us define $\sigma_k \geq 0$ by

$$\sigma_k^2 := \sum_{2^k < |n|_\infty \leq 2^{k+1}} a_n^2. \quad (5.40)$$

Theorem 5.2. *Assume hypotheses \mathcal{H}_1 , \mathcal{H}_2 and*

$$\mathcal{H}_7 : \sum_{k \in \mathbb{N}} \sigma_k = \infty$$

are satisfied, with σ_k^2 defined by (5.40). Then (\star) does not hold, and $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is almost surely unbounded for every method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$.

This theorem is stated and proven for $d = 1$ in [69], chapter 8. Our proof is very similar to the one-dimensional case and extends the argument to any dimension.

Proof. First note that (\star) does not hold implies that $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is unbounded for every method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$ almost surely, thanks to the zero-one law.

We prove the result by contradiction. Let us assume that (\star) holds. Thus, almost surely, uniform convergence occurs for any method of summation and the limit X is bounded. We shall prove that this implies that $\sum_k \sigma_k < \infty$.

Let us define $\chi(x) = \max(1 - |x|, 0)$ and $\theta(x) = \prod_{k=0}^{\infty} \chi\left(\frac{x}{2^k}\right)$ for all x in \mathbb{R} as in Lemma 3 p.105 of [69]. By Lemma 1 p.105 of [69], it follows that the trigonometric polynomial

$$\varphi(x) = \sum_{n \in \mathbb{Z}} \psi(n) e^{inx},$$

is a kernel function on \mathbb{T} meaning that $\varphi \geq 0$ and $\frac{1}{2\pi} \int_{\mathbb{T}} \varphi(x) dx = 1$ with both $\psi(\cdot) = \chi\left(\frac{\cdot}{2^k}\right)$ and $\psi(\cdot) = \theta\left(\frac{\cdot}{2^k}\right)$ for any $k \in \mathbb{N}$. For such a function ψ , we define a function on \mathbb{R}^d by

$$\psi_d(t) = \prod_{j=1}^d \psi(t_j) \text{ for } t = (t_1, \dots, t_d) \in \mathbb{R}^d,$$

so that the trigonometric polynomial

$$\varphi_d(t) = \sum_{n \in \mathbb{Z}^d} \psi_d(n) e^{in \cdot t} = \prod_{j=1}^d \psi(n_j) e^{in_j t_j} = \prod_{j=1}^d \varphi(t_j), \quad (5.41)$$

is a kernel function on \mathbb{T}^d meaning that $\varphi_d \geq 0$ and $\frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} \varphi_d(t) dt = 1$. Recall that for all continuous function $f : \mathbb{T}^d \rightarrow \mathbb{R}$ the convolution product $\varphi_d * f$ satisfies for all $t \in \mathbb{T}^d$

$$\varphi_d * f(t) := \frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} \varphi_d(u) f(t - u) du \leq \frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} \varphi_d(u) (\sup f) du \leq \sup f. \quad (5.42)$$

Now, let $k \in \mathbb{N}$ and define the following random trigonometric polynomials

$$P_k(t) = \sum_{n \in \mathbb{Z}^d} a_n \theta_d\left(\frac{n}{2^k}\right) e^{i(n \cdot t + \Phi_n)} = \sum_{|n|_{\infty} \leq 2^k} a_n \theta_d\left(\frac{n}{2^k}\right) e^{i(n \cdot t + \Phi_n)} \quad (5.43)$$

$$Q_k(t) = \sum_{|n|_{\infty} \leq 2^k} a_n \theta_d\left(\frac{n}{2^{k+1}}\right) e^{i(n \cdot t + \Phi_n)} \quad (5.44)$$

and

$$R_k(t) = \sum_{2^k < |n|_\infty \leq 2^{k+1}} a_n \theta_d \left(\frac{n}{2^{k+1}} \right) e^{i(n \cdot t + \Phi_n)} \quad (5.45)$$

so

$$R_k + Q_k = P_{k+1}. \quad (5.46)$$

Let us remark that $\theta_d(\frac{n}{2^k}) = \theta_d(\frac{n}{2^{k+1}})\chi_d(\frac{n}{2^k})$ so that $P_k = Q_k * \varphi_d^k$ and $P_k = P_{k+1} * \varphi_d^k$ with φ_d^k defined using $\psi_d(\cdot) = \chi_d(\cdot/2^k)$ in (5.41). Thanks to (5.42), we get $P_k(t) \leq \sup Q_k$ for all t , so that $M_k \leq \sup Q_k$, where we define the random variable $M_k := \sup P_k$. Then, we can define a random variable T_k , independent of R_k by \mathcal{H}_2 , such that

$$Q_k(T_k) \geq M_k. \quad (5.47)$$

Notice that

$$M_{k+1} - M_k \geq P_{k+1}(T_k) - Q_k(T_k) = R_k(T_k). \quad (5.48)$$

Moreover, thanks to (5.42) we also get $P_k(t) \leq \sup P_{k+1}$ so

$$M_{k+1} \geq M_k \quad (5.49)$$

and we conclude that

$$M_{k+1} - M_k \geq \sup(R_k(T_k), 0). \quad (5.50)$$

Now recall the Paley-Zygmund inequality for Y , a positive random variable satisfying $\mathbb{E}(Y^2) < +\infty$: for all $0 < \lambda < 1$

$$\mathbb{P}(Y > \lambda^2 \mathbb{E}(Y)) \geq (1 - \lambda^2)^2 \frac{\mathbb{E}(Y)^2}{\mathbb{E}(Y^2)}.$$

We consider this inequality for the random variable $R_k(t)^2$. By (4.44),

$$\frac{\mathbb{E}(R_k(t)^2)^2}{\mathbb{E}(R_k(t)^4)} \geq \frac{1}{3},$$

with

$$v_k^2 := \mathbb{E}(R_k(t)^2) = \sum_{n \in B_\infty(2^{k+1}) \setminus B_\infty(2^k)} a_n^2 \theta_d \left(\frac{n}{2^{k+1}} \right)^2. \quad (5.51)$$

It follows that

$$\mathbb{P}(|R_k(t)|^2 > \lambda^2 v_k^2) \geq \frac{1}{3}(1 - \lambda^2)^2. \quad (5.52)$$

Thanks to the fact that $R_k(t)$ is symmetrically distributed, this yields

$$\mathbb{P}(\sup(R_k(t), 0) > \lambda v_k) \geq \frac{1}{6}(1 - \lambda^2)^2. \quad (5.53)$$

As a consequence, for all t in \mathbb{T}^d ,

$$\mathbb{E}[\sup(R_k(t), 0)] \geq \frac{\lambda}{6}(1 - \lambda^2)^2 v_k. \quad (5.54)$$

Since T_k is independent of R_k , we can write that

$$\mathbb{E}[\sup(R_k(T_k), 0)t] = \mathbb{E}[\mathbb{E}[\sup(R_k(T_k), 0)|T_k]] \geq \frac{\lambda}{6}(1 - \lambda^2)^2 v_k. \quad (5.55)$$

Moreover, thanks to the symmetry of R_k and the independence with T_k

$$\mathbb{E}[\sup(R_k(T_k), 0)^2] = \frac{1}{2}\mathbb{E}[R_k(T_k)^2] = \frac{1}{2}\mathbb{E}[\mathbb{E}[R_k(T_k)^2|T_k]] = \frac{1}{2}v_k^2. \quad (5.56)$$

Thus, renaming $Y_k = \sup(R_k(T_k), 0)$, both

$$\mathbb{E}\left[\sum_{j=1}^k Y_j\right] \geq \frac{\lambda}{6}(1 - \lambda^2)^2 \sum_{j=1}^k v_j \quad (5.57)$$

and

$$\mathbb{E}\left[\left(\sum_{j=1}^k Y_j\right)^2\right]^{1/2} \leq \sum_{j=1}^k \mathbb{E}[Y_j^2]^{1/2} \leq \frac{1}{\sqrt{2}} \sum_{j=1}^k v_j \quad (5.58)$$

hold. Thus, thanks to Paley-Zygmund inequality, for $0 < \alpha < 1$,

$$\mathbb{P}\left(\sum_{j=1}^k Y_j > \alpha \frac{\lambda}{6}(1 - \lambda^2)^2 \sum_{j=1}^k v_j\right) \geq \mathbb{P}\left(\sum_{j=1}^k Y_j > \alpha \mathbb{E}\left(\sum_{j=1}^k Y_j\right)\right) \geq (1 - \alpha)^2 \frac{\lambda^2(1 - \lambda^2)^4}{18}. \quad (5.59)$$

Now, notice that adding (5.50) from 1 to k yields

$$\sum_{j=0}^k Y_j \leq M_k. \quad (5.60)$$

Moreover, notice that $P_k = X * \varphi_d^k$ (recall that φ_d^k is defined using $\psi_d(\cdot) = \theta_d(\cdot/2^k)$ in (5.41)) so that by (5.42) $M_k \leq \sup X$. It follows that

$$\mathbb{P}\left(\|X\|_\infty > \alpha \frac{\lambda}{6}(1 - \lambda^2)^2 \sum_{j=1}^k v_j\right) > (1 - \alpha)^2 \frac{\lambda^2(1 - \lambda^2)^4}{18} > 0. \quad (5.61)$$

Thus, since X is assumed to be almost surely bounded, the sequence $(\sum_{j=1}^k v_j)_{k \in \mathbb{N}}$ is bounded and thus converges. Now, notice that

$$\begin{aligned} v_k^2 &= \sum_{2^k < |n|_\infty \leq 2^{k+1}} \theta_d\left(\frac{n}{2^{k+1}}\right)^2 a_n^2 \geq \sum_{2^k < |n|_\infty \leq 3 \cdot 2^{k-1}} \theta_d\left(\frac{n}{2^{k+1}}\right)^2 a_n^2 \\ &\geq \inf_{\frac{1}{2} \leq |t|_\infty \leq \frac{3}{4}} \theta_d(t) \sum_{2^k < |n|_\infty \leq 3 \cdot 2^{k-1}} a_n^2. \end{aligned}$$

Notice that $\inf_{\frac{1}{2} \leq |t|_\infty \leq \frac{3}{4}} \theta_d(t) \geq \inf_{|t|_\infty \leq \frac{3}{4}} \theta_d(t) = \theta\left(\frac{3}{4}\right)^d$. Moreover, $\theta\left(\frac{3}{4}\right) > 0$ since clearly

$\sum_{k=0}^{\infty} \log\left(1 - \frac{3}{4} \frac{1}{2^k}\right) > -\infty$. Thus,

$$\sum_k \left(\sum_{2^k < |n|_\infty \leq 3 \cdot 2^{k-1}} a_n^2\right)^{1/2} < \infty. \quad (5.62)$$

Moreover, replacing X by $t \mapsto X(3t)$ in (5.62) yields

$$\sum_k \left(\sum_{2^k < |3n|_\infty \leq 3 \cdot 2^{k-1}} a_n^2 \right)^{1/2} < \infty. \quad (5.63)$$

Since

$$\sigma_k^2 \leq \sum_{2^k < |n|_\infty \leq 3 \cdot 2^{k-1}} a_n^2 + \sum_{2^{k+2} < |3n|_\infty \leq 3 \cdot 2^{k+1}} a_n^2, \quad (5.64)$$

(5.62) and (5.63) imply that

$$\sum_k \sigma_k < \infty \quad (5.65)$$

which concludes the proof. \square

Let us illustrate this result in the isotropic power spectrum case ($a_n = |n|_2^{-\alpha}$ for $n \neq 0$). Thanks to the equivalence of norms in finite-dimensional normed spaces, there exists a constant C_d such that

$$a_n \geq C_d |n|_\infty^{-\alpha} \quad (5.66)$$

for all $n \neq 0$ in \mathbb{Z}^d . Thus

$$\begin{aligned} \sum_{2^k < |n|_\infty \leq 2^{k+1}} a_n^2 &\geq C_d^2 \sum_{2^k < |n|_\infty \leq 2^{k+1}} |n|_\infty^{-2\alpha} = C_d^2 \sum_{p=2^k+1}^{2^{k+1}} \sum_{|n|_\infty=p} p^{-2\alpha} \\ &\geq C_d^2 \sum_{p=2^k+1}^{2^{k+1}} p^{d-1-2\alpha} \geq C'_d 2^{k(d-2\alpha)}, \end{aligned}$$

for some constant $C'_d > 0$. Hence,

$$\sigma_k = \left(\sum_{2^k < |n|_\infty \leq 2^{k+1}} a_n^2 \right)^{1/2} \geq (C'_d)^{1/2} 2^{k(d-2\alpha)/2}. \quad (5.67)$$

and $\sum_k \sigma_k = \infty$ as soon as $\alpha \leq d/2$.

Remark 5.4. Notice that thanks to the equivalence of norms in \mathbb{R}^d , the same reasoning applies to $a_n = \nu(n)^{-\alpha}$ for any norm ν on \mathbb{R}^d , and thus (\star) does not hold as soon as $\alpha \leq d/2$.

5.5 Hölder regularity

In this section, we focus on stronger assumptions on coefficients $(a_n)_{n \in \mathbb{Z}^d}$ implying that the sample paths of a RPN process are not only continuous but uniformly (α, τ_E) -Hölder on \mathbb{T}^d for some $\alpha > 0$ and τ_E a quasi-norm associated with a diagonal matrix $E = \text{diag}(\mu_1, \dots, \mu_d)$ with positive eigenvalues or simply a norm when $E = I$. Classical α -Hölder regularity can provide an interesting measure of smoothness (resp. roughness) of a texture. In particular, α -Hölder regularity is deeply connected to the fractal dimension of the sample paths graph. However,

when considering stationary processes this property does not allow to reveal anisotropy (see Theorem 1 of [33]). We may get around this drawback using a quasi-norm τ_E associated with the diagonal matrix $E = \text{diag}(\mu_1, \dots, \mu_d)$. Actually, using equivalence between quasi-norm given in Proposition 3.3 of [24] we may find $c_3 > 0$ such that for all $x \in \mathbb{R}^d$,

$$c_3^{-1} \tau_E(x) \leq \sum_{j=1}^d |x_j|^{1/\mu_j} \leq c_3 \tau_E(x). \quad (5.68)$$

It follows that when a continuous function f is (α, τ_E) -Hölder on \mathbb{T}^d , one can find constants $C, C' > 0$ such that

$$\forall t, s \in \mathbb{T}^d, |f(t) - f(s)| \leq C \tau_E(t - s)^\alpha \leq C' \left(\sum_{j=1}^d |t_j - s_j|^{1/\mu_j} \right)^\alpha.$$

This means in particular that the function f is α/μ_j -Hölder along all the straight lines directed by e_j with $(e_j)_{1 \leq j \leq d}$ the canonical basis of \mathbb{R}^d . Note also that this requires that $0 < \alpha \leq \mu_j$ for all $1 \leq j \leq d$ for non constant functions.

5.5.1 Sufficient conditions

Let $E = \text{diag}(\mu_1, \dots, \mu_d)$ be a diagonal matrix with positive eigenvalues $\mu_1, \dots, \mu_d \in (0, +\infty)$ and τ_E be a quasi-norm associated with E . Let us recall our notation $\underline{\mu} = \min_{1 \leq j \leq d} \mu_j$. Assume that Hypothesis \mathcal{H}_1 is satisfied and, for all k in \mathbb{N} , define $\sigma_{\tau_E, k} \geq 0$ by

$$\sigma_{\tau_E, k}^2 := \sum_{2^k < \tau_E(n) \leq 2^{k+1}} a_n^2 = \sum_{n \in \mathcal{A}_{k+1}(\tau_E) \setminus \mathcal{A}_k(\tau_E)} a_n^2. \quad (5.69)$$

Theorem 5.3. *Assume that hypotheses \mathcal{H}_1 , \mathcal{H}_2 and*

$$\mathcal{H}_8(\tau_E) : \exists \beta \in (0, \underline{\mu}) \text{ and } C > 0; \forall k \in \mathbb{N}, \sigma_{\tau_E, k} \leq C 2^{-\beta k}$$

are satisfied, with $\sigma_{\tau_E, k}^2$ defined by (5.69). Then, almost surely, (\star) holds and, for any method of summation, the limit function X satisfies that almost surely, there exists a constant $C > 0$ such that

$$\forall t, s \in \mathbb{T}^d, |X(t) - X(s)| \leq C \tau_E(t - s)^\beta \log(1 + \tau_E(t - s)^{-1})^{1/2}.$$

In particular, almost surely, X is uniformly (α, τ_E) -Hölder on \mathbb{T}^d for all $\alpha < \beta$

Proof. First of all, notice that Hypothesis $\mathcal{H}_8(\tau_E)$ implies that

$$s_{\tau_E, k}^2 = 2^k \sum_{j=2^k}^{2^{k+1}} \sigma_{\tau_E, j}^2 \leq 2^k \sum_{j=2^k}^{2^{k+1}} C^2 2^{-2\beta j} \leq C^2 \frac{2^{-\beta 2^{k+1}}}{1 - 2^{-2\beta}} \quad (5.70)$$

for all k and thus $\sum_k s_{\tau_E, k} < \infty$, so $\mathcal{H}_5(\tau_E)$ is satisfied. Then Corollary 5.1 applies and (\star) holds. We write X the associated RPN process.

Let $h_q = 2^{-q}$, q being some positive integer (a parameter). We consider the random polynomial

$$P_{0,q}^\Phi(t) = \sum_{\tau_E(n) \leq 2^q} a_n e^{i(n \cdot t + \Phi_n)} = \sum_{n \in \mathcal{A}_q(\tau_E)} a_n e^{i(n \cdot t + \Phi_n)},$$

where $\mathcal{A}_l(\tau_E)$ is the symmetrical subset defined by (5.32) for $l \in \mathbb{N}$. The gradient of $P_{0,q}^\Phi$ is given by $\nabla P_{0,q}^\Phi = (\partial_1 P_{0,q}^\Phi, \dots, \partial_d P_{0,q}^\Phi)$ with, for $1 \leq j \leq d$,

$$\partial_j P_{0,q}^\Phi(t) = \sum_{\tau_E(n) \leq 2^q} i n_j a_n e^{i(n \cdot t + \Phi_n)} = \sum_{n \in \mathcal{A}_q(\tau_E)} i n_j a_n e^{i(n \cdot t + \Phi_n)}. \quad (5.71)$$

For $k \geq 1$ let us denote $l_{k,q} = q2^{k-1}$ and $N_{k,q} = 2^{l_{k,q}}$. We introduce the random polynomials

$$P_{k,q}^\Phi(t) = \sum_{N_{k,q} < \tau_E(n) \leq N_{k+1,q}} a_n e^{i(n \cdot t + \Phi_n)} = \sum_{n \in \mathcal{A}_{l_{k+1,q}}(\tau_E) \setminus \mathcal{A}_{l_{k,q}}(\tau_E)} a_n e^{i(n \cdot t + \Phi_n)}. \quad (5.72)$$

Note that by (5.34), for all $k \geq 1$, $\mathcal{A}_{l_{k,q}}(\tau_E) \subset B_\infty(c_1^{-\bar{\mu}} N_{k,q}^{\bar{\mu}})$. Then, choosing $q \geq q_d$ large enough such that $c_1^{-\bar{\mu}} N_{1,q}^{\bar{\mu}} \geq 1$, Proposition 5.1 yields

$$\forall 1 \leq j \leq d, \mathbb{P} \left(\left\| \partial_j P_{0,q}^\Phi \right\|_\infty \geq \sqrt{2} C_d \left(\log \left(c_1^{-\bar{\mu}} N_{1,q}^{\bar{\mu}} \right) \sum_{\tau_E(n) \leq N_{1,q}} n_j^2 a_n^2 \right)^{1/2} \right) \leq c_1^{\bar{\mu}} (2\pi)^d / N_{1,q}^{\bar{\mu}}, \quad (5.73)$$

and

$$\mathbb{P} \left(\left\| P_{k,q}^\Phi \right\|_\infty \geq \sqrt{2} C_d \left(\log \left(c_1^{-\bar{\mu}} N_{k+1,q}^{\bar{\mu}} \right) \sum_{N_{k,q} < \tau_E(n) \leq N_{k+1,q}} a_n^2 \right)^{1/2} \right) \leq c_1^{\bar{\mu}} (2\pi)^d / N_{k+1,q}^{\bar{\mu}}. \quad (5.74)$$

Let us define for $q \geq q_d$ the event

$$E_q := \left\{ \omega \in \Omega; \left\| \partial_j P_{0,q}^\Phi(\omega) \right\|_\infty < \sqrt{2} C_d \left(\log \left(c_1^{-\bar{\mu}} N_{1,q}^{\bar{\mu}} \right) \sum_{\tau_E(n) \leq N_{1,q}} n_j^2 a_n^2 \right)^{1/2} \text{ for all } 1 \leq j \leq d \right. \\ \left. \text{and } \left\| P_{k,q}^\Phi(\omega) \right\|_\infty < \sqrt{2} C_d \left(\log \left(c_1^{-\bar{\mu}} N_{k+1,q}^{\bar{\mu}} \right) \sum_{N_{k,q} < \tau_E(n) \leq N_{k+1,q}} a_n^2 \right)^{1/2} \right\}. \quad (5.75)$$

Thanks to (5.73) and (5.74), $\mathbb{P}(E_q) \geq 1 - c_1^{\bar{\mu}} (2\pi)^d \left(1/N_{1,q}^{\bar{\mu}} + \sum_{k \geq 1} 1/N_{k+1,q}^{\bar{\mu}} \right)$. On the one hand, for all $k \in \mathbb{N}$, by definition of $N_{k,q}$ we have

$$\sum_{N_{k,q} < \tau_E(n) \leq N_{k+1,q}} a_n^2 = \sum_{2^{l_{k,q}} < \tau_E(n) \leq 2^{l_{k+1,q}}} a_n^2 = \sum_{m=l_{k,q}}^{l_{k+1,q}} \sigma_{\tau_E, m}^2,$$

in view of (5.69). Under $\mathcal{H}_8(\tau_E)$, it follows that

$$\sum_{N_{k,q} < \tau_E(n) \leq N_{k+1,q}} a_n^2 \leq C^2 \sum_{m=l_{k,q}}^{l_{k+1,q}} 2^{-2\beta m} \leq \frac{C^2}{1 - 2^{-2\beta}} 2^{-2\beta l_{k,q}}.$$

On the other hand, using equivalence between quasi-norm (5.68), we get $|x_j| \leq c_3^{\mu_j} \tau_E(x)^{\mu_j}$, for all $1 \leq j \leq d$. Therefore, for all $1 \leq j \leq d$,

$$\sum_{\tau_E(n) \leq N_{1,q}} n_j^2 a_n^2 = \sum_{\tau_E(n) \leq 2^{l_{1,q}}} n_j^2 a_n^2 \leq (2c_3^2)^{\mu_j} \sum_{m=0}^{l_{1,q}} 2^{2m\mu_j} \sigma_{\tau_E, m}^2 + c_1^{-2} \sum_{|n|_\infty \leq c_1^{-1}} a_n^2,$$

using (5.33) for the last sum corresponding to $\sum_{\tau_E(n) \leq 1} n_j^2 a_n^2$. Under $\mathcal{H}_8(\tau_E)$ and since $\beta < \underline{\mu}$, it follows that

$$\sum_{m=0}^{l_{1,q}} 2^{m\mu_j} \sigma_{\tau_E, m}^2 \leq \frac{C^2}{1 - 2^{-2(\mu_j - \beta)}} 2^{2l_{1,q}(\mu_j - \beta)}.$$

Hence we may find a constant $C > 0$ such that for all $q \geq q_d$ and $\omega \in E_q$, for all $k \in \mathbb{N}$ and $1 \leq j \leq d$,

$$\|P_{k,q}(\omega)\|_\infty \leq C l_{k,q}^{1/2} 2^{-\beta l_{k,q}} \quad \text{and} \quad \|\partial_j P_{0,q}(\omega)\|_\infty \leq C l_{1,q}^{1/2} 2^{l_{1,q}(\mu_j - \beta)}.$$

Thus, on E_q , using the mean value Theorem, for all $t, \theta \in \mathbb{T}^d$

$$\begin{aligned} |X(t + h_q^E \theta) - X(t)| &\leq \sum_{k=0}^{\infty} |P_{k,q}(t + h_q^E \theta) - P_{k,q}(t)| \\ &\leq |\theta|_\infty \sum_{j=1}^d \|\partial_j P_{0,q}\|_\infty h_q^{\mu_j} + 2 \sum_{k=1}^{\infty} \|P_{k,q}(\omega)\|_\infty \\ &\leq 2\pi C \left(l_{1,q}^{1/2} 2^{-l_{1,q}\beta} + 2 \sum_{k=1}^{\infty} l_{k,q}^{1/2} 2^{-\beta l_{k,q}} \right), \end{aligned}$$

since $h_q = 2^{l_{1,q}}$. Then, since $l_{k,q} = 2^{k-1} l_{1,q}$ and $l_{1,q}^{1/2} 2^{-l_{1,q}\beta} = \log(h_q^{-1})^{1/2} h_q^\beta$, it follows that we may find a constant $C > 0$ such that on E_q

$$\sup_{t, \theta \in \mathbb{T}^d} |X(t + h_q^E \theta) - X(t)| \leq C \log(h_q^{-1})^{1/2} h_q^\beta.$$

Moreover,

$$\begin{aligned} \sum_{q \geq q_d} \mathbb{P}(\Omega \setminus E_q) &\leq c_1^{\bar{\mu}} (2\pi)^d \left(\sum_{q \geq 1} 1/N_{1,q}^{\bar{\mu}} + (2\pi)^d \sum_{q \geq 1} \sum_{k \geq 2} 1/N_{k,q}^{\bar{\mu}} \right) \\ &\leq c_1^{\bar{\mu}} (2\pi)^d \left(\sum_{q=1}^{\infty} 2^{-q\bar{\mu}} + (2\pi)^d \sum_{k \geq 2} \sum_{q \geq 1} 2^{-q\bar{\mu} 2^{k-1}} \right) \end{aligned}$$

$\leq c_1^{\bar{\mu}}(2\pi)^d \left(\frac{1}{2^{\bar{\mu}-1}} + \sum_{k \geq 2} \frac{1}{2^{\bar{\mu}2^{k-1}} - 1} \right) < \infty$ thanks to Fubini and the definition of $N_{k,q}$.

Borel-Cantelli Lemma applies and thus we may find a random constant C' such that almost surely for all $h > 0$ small enough

$$\sup_{t, \theta \in \mathbb{T}^d} |X(t + h^E \theta) - X(t)| \leq C' h^\beta \log(1 + h^{-1})^{1/2}.$$

Now, for $s, t \in \mathbb{T}^d$ with $s \neq t$, let us remark that $s = t + \tau_E(t-s)^E \theta$ with $\theta = \tau_E(s-t)^{-1}(s-t)$ in the compact set $\tau_E^{-1}(\{1\})$. Since X is continuous a.s, we may therefore find a random constant C'' such that almost surely

$$\sup_{t, s \in \mathbb{T}^d} |X(s) - X(t)| \leq C'' \tau_E(t-s)^\beta \log(1 + \tau_E(t-s)^{-1})^{1/2}.$$

It follows that almost surely X is (α, τ_E) -Hölder on \mathbb{T}^d for all $\alpha \in (0, \beta)$. \square

5.5.2 A necessary condition

We may obtain a partial converse in the isotropic case for $E = I$, considering $\tau_E = |\cdot|_\infty$. Actually, in this setting, we extend the classical condition for the case $d = 1$ (see [69], Chapter 7, Theorem 3).

Proposition 5.3. *Let $f : \mathbb{T}^d \mapsto \mathbb{R}$ be a α -Hölder function. Then for all k in \mathbb{N}*

$$\sum_{2^k < |n|_\infty \leq 2^{k+1}} |\hat{f}_n|^2 \leq C_f 2^{-2\alpha k} \quad (5.76)$$

holds for some constant C_f , where \hat{f}_n denotes the n -th Fourier coefficient given by

$$\hat{f}_n = \frac{1}{(2\pi)^d} \int_{\mathbb{T}^d} f(t) e^{-in \cdot t} dt.$$

Proof. For any q , define $h_q = \frac{\pi}{3} 2^{-q}$. Assume f to be α -Hölder, and denote e_j the j -th vector of the canonical basis. There is some constant C such that

$$\sup_{t \in \mathbb{T}^d} |f(t + h_q e_j) - f(t)|^2 \leq C h_q^{2\alpha} \quad (5.77)$$

holds for all $1 \leq j \leq d$. Integrating the last inequality over \mathbb{T}^d and applying Parseval's identity yields

$$\sum_{n \in \mathbb{Z}^d} |\hat{f}_n|^2 \sin^2(n \cdot h_q e_j / 2) = \int_{\mathbb{T}^d} |f(t + h_q e_j) - f(t)|^2 dt \leq (2\pi)^d C h_q^{2\alpha} \quad (5.78)$$

for all j and all q . Denote

$$B^{(q,j)} = B_\infty(2^{q+1}) \setminus B_\infty(2^q) \cap \{n \in \mathbb{Z}^d; 2^q < |n_j| \leq 2^{q+1}\}. \quad (5.79)$$

One derives

$$\sum_{n \in \mathbb{Z}^d} |\hat{f}_n|^2 \sin^2(n \cdot h_q e_j / 2) \geq \sum_{n \in B^{(q,j)}} |\hat{f}_n|^2 \sin^2(h_q n_j / 2) \quad (5.80)$$

and since $\sin^2(h_q n_j/2) \geq \frac{3}{4}$ for all $n \in B^{(q,j)}$,

$$\frac{3}{4} \sum_{n \in B^{(q,j)}} |\widehat{f}_n|^2 \leq (2\pi)^d C h_q^{2\alpha} \quad (5.81)$$

holds for each j in $\{1, \dots, d\}$. To conclude, notice that $B_\infty(2^{q+1}) \setminus B_\infty(2^q) = \bigcup_{i=1}^d B^{(q,i)}$ so that

$$\sum_{n \in B_\infty(2^{q+1}) \setminus B_\infty(2^q)} |\widehat{f}_n|^2 \leq \sum_j \sum_{n \in B^{(q,j)}} |\widehat{f}_n|^2 \leq d(2\pi)^d C h_q^{2\alpha} = d(2\pi)^d C 2^{-2\alpha q} \quad (5.82)$$

holds for all q , which concludes the proof. \square

Let us study the example of the isotropic power spectrum, say $a_n = |n|_2^{-\alpha}$ for some α in $(0, 1)$. From Section 3, we know that if $\alpha > d/2$, the condition (\star) holds. Moreover, in this case,

$$C_d 2^{k(d-2\alpha)/2} \leq \sigma_k \leq C'_d 2^{k(d-2\alpha)/2}. \quad (5.83)$$

The upper bound is obtained by Lemma 5.2 with $E = I$ and $q_E = \text{tr}(E) = d$. The lower bound follows from (5.67). Thus, the RPN associated with $a_n = |n|_2^{-\alpha}$ and $\alpha \in (d/2, d/2+1)$, is almost surely

- β -Hölder for all $\beta < \alpha - \frac{d}{2}$ (Theorem 5.3)
- not β -Hölder for $\beta > \alpha - \frac{d}{2}$ (Proposition 5.3).

Hence $\alpha - d/2 \in (0, 1)$ appears as a critical Hölder exponent as defined in [17].

5.6 Extensions

We discuss extensions of the main results of this chapter, which are designed to apply to the Gaussian case. Recall that the Fourier coefficients of a real center stationary Gaussian process on \mathbb{T}^d are complex Gaussian random variables that are independent on every half-space of \mathbb{Z}^d . Let us consider the hypothesis

\mathcal{H}_1^* : “ $(A_n)_{n \in \mathbb{Z}^d}$ is such that $(A_n)_{n \in \mathcal{A}}$ are independent whenever \mathcal{A} and $-\mathcal{A}$ do not intersect ; $(A_n)_{n \in \mathbb{Z}^d}$ is independent of Φ ; $\mathbb{E}[\sum_{\mathbb{Z}^d} A_n^2] < \infty$; $A_0 = 0$ almost surely”.

Let us write

$$S_{\mathcal{A}_k}(\omega, t) = \sum_{n \in \mathcal{A}_k} A_n(\omega) e^{i(\Phi_n(\omega) + n \cdot t)} \quad (5.84)$$

for all ω in Ω and t in \mathbb{T}^d . As discussed in Chapter 4 (Corollary 4.1) the chain of equivalences in Theorem 1 holds with random variables $(A_n)_{n \in \mathbb{Z}^d}$ instead of deterministic sequence $(a_n)_{n \in \mathbb{Z}^d}$ as long as Hypothesis \mathcal{H}_1^* is fulfilled. Let us name $(\star\star)$ this chain of equivalences. The following results can be straightforwardly deduced from Theorem 5.1, with a modification of Hypothesis \mathcal{H}_3 : $(\mathcal{A}_k)_{k \in \mathbb{N}}$ is a method of summation such that, for each k , $\mathcal{A}_k \subset B_\infty(N_k)$ for some N_k in \mathbb{N} . Let us state the two hypotheses

- \mathcal{H}_3^* : $\sum_k \left(\log(N_{k+1}) \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} A_n^2 \right)^{1/2} < \infty$ almost surely
- \mathcal{H}_4 : $\sum_k \frac{1}{N_k} < \infty$

Corollary 5.2. *Assume that hypotheses \mathcal{H}_1^* , \mathcal{H}_2 , \mathcal{H}_3^* and \mathcal{H}_4 are satisfied. Then $(\star\star)$ holds.*

Remark 5.5. \mathcal{H}_3^* is clearly implied by \mathcal{H}_3^{**} : $\sum_k \mathbb{E} \left[\left(\log(N_{k+1}) \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} A_n^2 \right)^{1/2} \right] < \infty$,

which is itself implied by \mathcal{H}_3^{***} : $\sum_k \left(\log(N_{k+1}) \sum_{n \in \mathcal{A}_{k+1} \setminus \mathcal{A}_k} \mathbb{E}[A_n^2] \right)^{1/2} < \infty$ thanks to Jensen inequality.

We now turn to extending Theorem 5.2. For all k in \mathbb{N} , let us define the random variable $\sigma_k^* \geq 0$ by

$$\sigma_k^{*2} := \sum_{n \in B_\infty(2^{k+1}) \setminus B_\infty(2^k)} A_n^2. \quad (5.85)$$

Corollary 5.3. *Assume hypotheses \mathcal{H}_1^* , \mathcal{H}_2 and*

$$\mathcal{H}_7^* : \sum_{k \in \mathbb{N}} \sigma_k^* = \infty \text{ almost surely}$$

are satisfied, with σ_k^ defined by (5.85). Then (\star) does not hold, and $(S_{\mathcal{A}_k})_{k \in \mathbb{N}}$ is almost surely unbounded for every method of summation $(\mathcal{A}_k)_{k \in \mathbb{N}}$.*

Finally, let us extend Theorem 5.3. Let $E = \text{diag}(\mu_1, \dots, \mu_d)$ be a diagonal matrix with positive eigenvalues $\mu_1, \dots, \mu_d \in (0, +\infty)$, $\underline{\mu} = \min_{1 \leq j \leq d} \mu_j$, and τ_E be a quasi-norm associated with E . Assume that Hypothesis \mathcal{H}_1^* is satisfied and, for all k in \mathbb{N} , define the random variable $\sigma_{\tau_E, k}^* \geq 0$ by

$$\sigma_{\tau_E, k}^{*2} := \sum_{2^k < \tau_E(n) \leq 2^{k+1}} A_n^2 = \sum_{n \in \mathcal{A}_{k+1}(\tau_E) \setminus \mathcal{A}_k(\tau_E)} A_n^2. \quad (5.86)$$

Corollary 5.4. *Assume that hypotheses \mathcal{H}_1^* , \mathcal{H}_2 and*

$$\mathcal{H}_8^*(\tau_E) : \exists \beta \in (0, \underline{\mu}) \text{ and } C > 0; \forall k \in \mathbb{N}, \sigma_{\tau_E, k}^* \leq C 2^{-\beta k} \text{ almost surely}$$

*are satisfied, with $\sigma_{\tau_E, k}^{*2}$ defined by (5.86). Then, almost surely, (\star) holds and, for any method of summation, the limit function X satisfies that almost surely, there exists a constant $C > 0$ such that*

$$\forall t, s \in \mathbb{T}^d, |X(t) - X(s)| \leq C \tau_E(t-s)^\beta \log(1 + \tau_E(t-s)^{-1})^{1/2}.$$

Remark 5.6. $\mathcal{H}_8^*(\tau_E)$ is clearly implied by $\mathcal{H}_8^{**}(\tau_E)$: $\sum_{k \in \mathbb{N}} 2^{\beta k} \mathbb{E} \left[\left(\sum_{2^k < \tau_E(n) \leq 2^{k+1}} A_n^2 \right)^{1/2} \right] < \infty$

which is itself implied by \mathcal{H}_8^{***} : $\sum_{k \in \mathbb{N}} 2^{\beta k} \left(\sum_{2^k < \tau_E(n) \leq 2^{k+1}} \mathbb{E}[A_n^2] \right)^{1/2} < \infty$ thanks to Jensen inequality.

Let us conclude by remarking that in the case where X is a real centered stationary Gaussian process almost surely continuous on \mathbb{T}^d , its Fourier coefficients may be written as

$$\hat{X}_n = \frac{1}{2\pi^d} \int_{\mathbb{T}^d} X(t) e^{-in \cdot t} dt = A_n e^{i\Phi_n},$$

with $(A_n)_{n \in \mathbb{Z}^d}$ and $(\Phi_n)_{n \in \mathbb{Z}^d}$ satisfying hypotheses \mathcal{H}_1^* and \mathcal{H}_2 . In order to deal with (α, τ_E) -Hölder regularity the classical assumption is that

$$\forall t, s \in \mathbb{T}^d, \mathbb{E}((X(t) - X(s))^2)^{1/2} \leq C \tau_E (t - s)^\beta \log(1 + \tau_E (t - s)^{-1})^\eta,$$

for some $\beta \in (0, \underline{\mu})$ and $\eta \in \mathbb{R}$. Actually, under these assumptions, by Proposition 5.3 of [10] one may prove that, for all $\varepsilon > 0$, there exists a random constant C such that almost surely,

$$|X(t) - X(s)| \leq C \tau_E (t - s)^\beta \log(1 + \tau_E (t - s)^{-1})^{\eta+1/2+\varepsilon}.$$

The link between this two kinds of assumption is given by the fact that

$$\mathbb{E}((X(t) - X(s))^2) = 4 \sum_{n \in \mathbb{Z}^d} \mathbb{E}(A_n^2) \sin(n \cdot (t - s)/2)^2.$$

We may conjecture that the power 1/2 on the logarithmic term is sharp (for $\eta = 0$) in view of exact modulus of continuity obtained for operator scaling Gaussian random fields in [85].

Chapter 6

Conclusion

Let us conclude by suggesting a few ideas for future work connected to this thesis.

An interesting extension of our work in Chapter 2 would be, for textures actually based on localized sparse spots, to be able to find this original spot based on a single texture sample, or a few samples. A case where this might be conveniently achieved is for symmetrical spots, which already have a discrete Fourier transform that is entirely real. Thus, starting from the texton of this spot, the right choice of signs for each coefficient would yield the original sparse symmetrical spot. This choice of sign could rely on the fact that the Fourier transform of a localized spot is smooth, and thus the connected components of identically signed neighbors have fair chances to be quite large. Determining these connected components could in turn rely on the Dijkstra algorithm or graph cut techniques [56].

The idea of a texton is itself quite general, and has been introduced long before our work [65]. There is however, no reason to believe that it is only relevant to vision. Thus the notion of texton might be extended to stationary noisy signals of any dimensions. For instance, a sparse representation of sound noise might benefit from our approach in images. Another and arguably more challenging extension would be to extract a somehow minimal set of information that would allow for other texture synthesis algorithms than RPN or Gaussian synthesis, to perform somewhat faithful synthesis.

Our discussion on the importance of phase in signals in Chapter 3 seems to be mostly relevant for human vision experts. Although we only show that the importance of the phase can be traced back to a higher level of abstraction – namely the projection of signals onto sets of signals – the conditions for such projections to yield a visually compelling result (at least recognizable) are quite unclear.

Textures don't need to be periodical, and it would thus be interesting to extend our periodical random phase fields to non-periodical fields. A convenient way to do this while keeping the random phase approach would be to relax the constraint imposing the spectrum of the field to be contained within a lattice while keeping a discrete spectrum approach. Continuity results have recently been proven in this direction [27], and our results on regularity might also be extended in this more general framework.

Defining an equivalent of the Random Phase field with a continuous spectrum is a somewhat puzzling task. Asymptotics results proven in Chapter 2 seem to hint that Gaussian fields

might be the only qualifying fields, but this remains to be proven. However, quasi-periodic functions might be a relevant approach to define a non-Gaussian non-periodical random phase field.

Last but not least, our approach to regularity was mainly dominated by global Hölder exponents, somehow backed by the stationarity of the fields introduced. However, a more thorough study of the local Hölder exponents (and their probabilistic distribution) is still lacking and would make an interesting extension of our work.

Bibliography

- [1] A. Aaron, M. Hill, and A. Srivatsa. Mosmat 5000, a Photomosaic Generator. <http://scien.stanford.edu/pages/labsite/2000/ee368/projects2000/project12>, 2000. Accessed: 2013-11-05.
- [2] C. Aguerrebere, Y. Gousseau, and G. Tartavel. Exemplar-Based Texture Synthesis: The Efros-Leung Algorithm. *Image Processing On Line*, 3:223–241, 2013.
- [3] K.-I. Anjyo. A Simple Spectral Approach to Stochastic Modeling for Natural Objects. In *Proceedings of Eurographics*, volume 88, pages 285–296, 1988.
- [4] J. M. Ash and G. V. Welland. Convergence, Uniqueness, and Summability of Multiple Trigonometric Series. *Transactions of the American Mathematical Society*, 163:401–436, 1972.
- [5] F. Attneave. Some Informational Aspects of Visual Perception. *Psychological review*, 61(3):183, 1954.
- [6] H. B. Barlow. Possible Principles Underlying the Transformation of Sensory Messages. *Sensory communication*, pages 217–234, 1961.
- [7] M. Ben Slimane. Multifractal Formalism and Anisotropic Selfsimilar Functions. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 124, pages 329–363. Cambridge Univ Press, 1998.
- [8] H. Berens and Y. Xu. Fejèr Means for Multivariate Fourier Series. *Mathematische Zeitschrift*, 221:449–465, 1996.
- [9] J. R. Bergen and E. H. Adelson. Early Vision and Texture Perception. *Nature*, 333(6171):363–364, 1988.
- [10] H. Biermé and C. Lacaux. Hölder Regularity for Operator Scaling Stable Random Fields. *Stochastic Processes and their Applications*, 119(7):2222–2248, 2009.
- [11] H. Biermé and C. Lacaux. Modulus of Continuity of Conditionally Sub-Gaussian Random Series, Application to Stable Random Fields. Preprint, 2013.
- [12] H. Biermé, C. Lacaux, and H.-P. Scheffler. Multi-Operator Scaling Random Fields. *Stochastic Processes and their Applications*, 121(11):2642–2677, 2011.

-
- [13] H. Biermé, M. M. Meerschaert, and H.-P. Scheffler. Operator Scaling Stable Random Fields. *Stochastic Processes and their Applications*, 117(3):312–332, 2007.
- [14] H. Biermé, L. Moisan, and S. Ronsin. On the Regularity of some Multiple Random Fourier Series. *Preprint*, 2014.
- [15] P. Billard. Séries de Fourier Aléatoirement Bornées, Continues, Uniformément Convergentes. *Studia Mathematica*, 22:310–330, 1963.
- [16] R.D. Blevins. Probability Density of Finite Fourier Series with Random Phases. *Journal of Sound and Vibration*, 208(4):617 – 652, 1997.
- [17] A. Bonami and A. Estrade. Anisotropic Analysis of some Gaussian Models. *Journal of Fourier Analysis and Applications*, 9:215–236, 2003.
- [18] B. Brunet-Imbault, G. Lemineur, C. Chappard, R. Harba, and C.-L. Benhamou. A New Anisotropy Index on Trabecular Bone Radiographic Images Using the Fast Fourier Transform. *BMC Medical Imaging*, 5(1):4, 2005.
- [19] J. W. Byng, M. J. Yaffe, G. A. Lockwood, L. E. Little, D. L. Tritchler, and N. F. Boyd. Automated Analysis of Mammographic Densities and Breast Carcinoma Risk. *Cancer*, 80(1):66–74, 1997.
- [20] F. W. Campbell and J. G. Robson. Application of Fourier Analysis to the Visibility of Gratings. *The Journal of Physiology*, 197(3):551, 1968.
- [21] E. J. Candès, J. Romberg, and T. Tao. Robust Uncertainty Principles: Exact Signal Reconstruction from Highly Incomplete Frequency Information. *IEEE Transactions on Information Theory*, 52(2):489–509, 2006.
- [22] E. J. Candes, J. Romberg, and T. Tao. Stable Signal Recovery from Incomplete and Inaccurate Measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [23] J. Canny. A Computational Approach to Edge Detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, (6):679–698, 1986.
- [24] M. Clausel and B. Vedel. An Optimality Result about Sample Path Properties of Operator Scaling Gaussian Random Fields. *arXiv preprint arXiv:1302.0818*, 2013.
- [25] G. Cohen and C. Cuny. On Billard’s Theorem for Random Fourier Series. *Bulletin of the Polish Academy of Sciences, Mathematics*, 2005.
- [26] G. Cohen and C. Cuny. On Random Almost Periodic Series and Random Ergodic Theory. *Ergodic Theory and Dynamical Systems*, 26(03):683–709, 2006.
- [27] G. Cohen and C. Cuny. On Random Almost Periodic Trigonometric Polynomials and Applications to Ergodic Theory. *The Annals of Probability*, pages 39–79, 2006.

-
- [28] R. Coifman, F. Geshwind, and Y. Meyer. Noiselets. *Applied and Computational Harmonic Analysis*, 10(1):27–44, 2001.
- [29] R. L. Cook and T. DeRose. Wavelet Noise. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 803–811. ACM, 2005.
- [30] J. W. Cooley, P. A. Lewis, and P. D. Welch. Historical Notes on the Fast Fourier Transform. *Proceedings of the IEEE*, 55(10):1675–1677, 1967.
- [31] J. W. Cooley and J. W. Tukey. An Algorithm for the Machine Calculation of Complex Fourier Series. *Mathematics of computation*, 19(90):297–301, 1965.
- [32] J. Cuzick and T. L. Lai. On Random Fourier Series. *Transactions of the American Mathematical Society*, 261(1):53–80, 1980.
- [33] S. Davies and P. Hall. Fractal Analysis of Surface Roughness by Using Spatial Data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(1):3–37, 1999.
- [34] A. Desolneux, L. Moisan, and S. Ronsin. Vers un Texton pour les Micro-Textures. In *Actes du Groupe d’Etudes du Traitement du Signal et des Images (GRETSI)*, 2011.
- [35] A. Desolneux, L. Moisan, and S. Ronsin. A Compact Representation of Random Phase and Gaussian Textures. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 1381–1384, 2012.
- [36] A. Desolneux, L. Moisan, and S. Ronsin. A Texton for Random Phase and Gaussian Textures. *Preprint*, 2014.
- [37] O. Deussen, S. Hiller, C. Van Overveld, and T. Strothotte. Floating Points: A Method for Computing Stipple Drawings. In *Computer Graphics Forum*, volume 19, pages 41–50. Wiley Online Library, 2000.
- [38] D. L. Donoho. Compressed Sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [39] D. C. Dowson and B. V. Landau. The Fréchet Distance between Multivariate Normal Distributions. *Journal of multivariate analysis*, 12(3):450–455, 1982.
- [40] M. L. Eaton. *Multivariate Statistics: a Vector Space Approach*. Wiley New York, 1983.
- [41] A. A. Efros and W. T. Freeman. Image Quilting for Texture Synthesis and Transfer. In *Proceedings of the 28th annual Conference on Computer Graphics and Interactive Techniques*, pages 341–346. ACM, 2001.
- [42] A. A. Efros and T. K. Leung. Texture Synthesis by Non-Parametric Sampling. *Proceedings of the Seventh IEEE International Conference on Computer Vision*, 2:1033–1038, 1999.

- [43] C. Fefferman. On the Convergence of Multiple Fourier Series. *Bulletin of the American Mathematical Society*, 77(5):pp. 744–745, 1971.
- [44] C. Fefferman. On the Divergence of Multiple Fourier Series. *Bulletin of the American Mathematical Society*, 77(2):pp. 191–195, 1971.
- [45] X. Fernique. Continuité des Processus Gaussiens. *Comptes Rendus de l'Académie des Sciences*, 258:pp. 6058–6060, 1964.
- [46] X. Fernique. Régularité de Processus Gaussiens. *Inventiones mathematicae*, 12(4):304–320, 1971.
- [47] A. Figà-Talamanca. Bounded and Continuous Random Fourier Series on Noncommutative Groups. *Proceedings of the American Mathematical Society*, 22:pp. 573–578, 1969.
- [48] I. K. Fodor. A Survey of Dimension Reduction Techniques, 2002.
- [49] M. Fréchet. Sur la Distance de Deux Lois de Probabilité. *Comptes Rendus de l'Académie des Sciences*, 244(6):689–692, 1957.
- [50] B. Galerne, Y. Gousseau, and J.-M. Morel. Micro-Texture Synthesis by Phase Randomization. *Image Processing On Line*, 1, 2011.
- [51] B. Galerne, Y. Gousseau, and J.-M. Morel. Random Phase Textures: Theory and Synthesis. *IEEE Transactions on Image Processing*, 20(1):257–267, 2011.
- [52] B. Galerne, A. Lagae, S. Lefebvre, and G. Drettakis. Gabor Noise by Example. *ACM Transactions on Graphics (TOG)*, 31(4):73, 2012.
- [53] I. M. Gelfand and N. Y. Vilenkin. *Generalized Functions. Vol. 4, Applications of Harmonic Analysis*. Academic Press, 1964.
- [54] G. Gilet, J.-M. Dischler, and L. Soler. Procedural Descriptions of Anisotropic Noisy Textures by Example. In *Eurographics 2010 - Short Papers*, pages 77–80. The Eurographics Association, 2010.
- [55] A. Goldberg, M. Zwicker, and F. Durand. Anisotropic Noise. In *ACM Transactions on Graphics (TOG)*, volume 27, page 54. ACM, 2008.
- [56] DM Greig, BT Porteous, and Allan H Seheult. Exact Maximum a Posteriori Estimation for Binary Images. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 271–279, 1989.
- [57] R. M. Haralick, S. R. Sternberg, and X. Zhuang. Image Analysis Using Mathematical Morphology. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, (4):532–550, 1987.
- [58] L. D. Harmon and K. C. Knowlton. Picture Processing by Computer. *Science*, 164:19–29, 1969.

- [59] D. J. Heeger and J. R. Bergen. Pyramid-Based Texture Analysis/Synthesis. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, pages 229–238. ACM, 1995.
- [60] K. Itó and M. Nisio. On the Convergence of Sums of Independent Banach Space Valued Random Variables. *Osaka Journal of Mathematics*, 5:35–48, 1968.
- [61] N. C. Jain and M. B. Marcus. Sufficient Conditions for the Continuity of Stationary Gaussian Processes and Applications to Random Series of Functions. In *Annales de l'Institut Fourier*, volume 24, pages 117–141. Institut Fourier, 1974.
- [62] N. Jayant, J. Johnston, and R. Safranek. Signal Compression Based on Models of Human Perception. *Proceedings of the IEEE*, 81(10):1385–1422, 1993.
- [63] B. Julesz. Visual Pattern Discrimination. *Information Theory, IRE Transactions on*, 8(2):84–92, 1962.
- [64] B. Julesz. *Foundations of Cyclopean Perception*. MIT Press, 1971.
- [65] B. Julesz. Textons, the Elements of Texture Perception, and their Interactions. *Nature*, 1981.
- [66] B. Julesz. A theory of preattentive texture discrimination based on first-order statistics of textons. *Biological Cybernetics*, 41:131–181, 1981.
- [67] B. Julesz and J. R. Bergen. Textons, the Fundamental Elements in Preattentive Vision and Perception of Textures. *Bell System Technical Journal*, 1983.
- [68] B. Julesz, E. N. Gilbert, L. A. Shepp, and H.L. Frisch. Inability of Humans to Discriminate between Visual Textures that Agree in Second-Order Statistics-Revisited. *Perception*, 2(4):391–405, 1973.
- [69] J.-P. Kahane. *Some Random Series of Functions*. Cambridge studies in advanced mathematics. Cambridge University Press, 1994.
- [70] J.-P. Kahane and Y. Katznelson. Sur les Ensembles de Divergence des Séries Trigonométriques. *Studia Mathematica*, 26:305–306, 1966.
- [71] C. S. Kaplan and R. Bosch. TSP Art. In *Renaissance Banff: Mathematics, Music, Art, Culture*, pages 301–308. Canadian Mathematical Society, 2005.
- [72] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active Contour Models. *International Journal of Computer Vision*, 1(4):321–331, 1988.
- [73] J. J. Knierim and D. C. Van Essen. Neuronal Responses to Static Texture Patterns in Area V1 of the Alert Macaque Monkey. *Journal of Neurophysiology*, 67(4):961–980, 1992.
- [74] K. Knowlton and L. Harmon. Computer-Produced Grey Scales. *Computer Graphics and Image Processing*, 1(1):1–20, 1972.

- [75] V. Kwatra, A. Schödl, I. Essa, G. Turk, and A. Bobick. Graphcut Textures: Image and Video Synthesis Using Graph Cuts. *ACM Transactions on Graphics (TOG)*, 22(3):277–286, 2003.
- [76] A. Lagae, S. Lefebvre, R. Cook, T. DeRose, G. Drettakis, D. S. Ebert, JP. Lewis, K. Perlin, and M. Zwicker. A Survey of Procedural Noise Functions. In *Computer Graphics Forum*, volume 29, pages 2579–2600. Wiley Online Library, 2010.
- [77] A. Lagae, S. Lefebvre, G. Drettakis, and P. Dutré. Procedural Noise Using Sparse Gabor Convolution. In *ACM Transactions on Graphics (TOG)*, 2009.
- [78] M. S. Landy and N. Graham. Visual Perception of Texture. *The Visual Neurosciences*, 2:1106–1118, 2004.
- [79] S. Lazebnik, C. Schmid, and J. Ponce. A Sparse Texture Representation Using Local Affine Regions. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(8):1265–1278, 2005.
- [80] M. Ledoux and M. Talagrand. *Probability in Banach Spaces*. Ergebnisse der Mathematik und ihrer Grenzgebiete. Springer, 1991.
- [81] S. Lefebvre and H. Hoppe. Parallel Controllable Texture Synthesis. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 777–786. ACM, 2005.
- [82] T. Leung and J. Malik. Representing and Recognizing the Visual Appearance of Materials using Three-Dimensional Textons. *International Journal of Computer Vision*, 43(1):29–44, 2001.
- [83] J.-P. Lewis. Texture Synthesis for Digital Painting. In *ACM SIGGRAPH Computer Graphics*, volume 18, pages 245–252. ACM, 1984.
- [84] M. Li and P. M. B. Vitányi. *An Introduction to Kolmogorov Complexity and its Applications*. Springer, 2009.
- [85] Y. Li, W. Wang, and Y. Xiao. Exact Moduli of Continuity for Operator-Scaling Gaussian Random Fields. *Preprint*, 2013.
- [86] E. H. Lieb and M. Loss. *Analysis: Second Edition*. Graduate Studies in Mathematics. AMS, 2001.
- [87] S. Lisberger (director, screenplay, story), B. MacBird (story), and D. Kushner (producer). Tron. *Walt Disney Productions*, 1982.
- [88] E. Lukacs. *Characteristic Functions*. Griffin, 1970.
- [89] J. Malik, S. Belongie, J. Shi, and T. Leung. Textons, contours and regions: Cue integration in image segmentation. In *Proceedings IEEE of the 7th International Conference on Computer Vision, Crfu, Grece - Volume 2*, pages 918–925, 1999.

- [90] J. Malik and P. Perona. Preattentive Texture Discrimination with Early Vision Mechanisms. *Journal of the Optical Society of America A*, 7(5):923–932, 1990.
- [91] S. Mallat and F. Falzon. Analysis of Low Bit Rate Image Transform Coding. *IEEE Transactions on Signal Processing*, 46(4):1027–1042, 1998.
- [92] B. B. Mandelbrot. Intermittent Turbulence in Self-Similar Cascades: Divergence of High Moments and Dimension of the Carrier. *Journal of Fluid Mechanics*, 62(02):331–358, 1974.
- [93] M. B. Marcus. Continuity of Gaussian Processes and Random Fourier Series. *The Annals of Probability*, 1(6):pp. 968–981, 1973.
- [94] M. B. Marcus. Uniform Convergence of Random Fourier Series. *Arkiv för Matematik*, 13:107–122, 1975.
- [95] M. B. Marcus and G. Pisier. *Necessary and Sufficient Conditions for the Uniform Convergence of Random Trigonometric Series: Lectures 1977/78*. Number 50. Aarhus Universitet, Matematisk Institut, 1978.
- [96] M. B. Marcus and G. Pisier. *Random Fourier Series with Applications to Harmonic Analysis*. Annals of Mathematics Studies. Princeton University Press, 1981.
- [97] M. B. Marcus and G. Pisier. Characterizations of Almost Surely Continuous p -Stable Random Fourier Series and Strongly Stationary Processes. *Acta Mathematica*, 152(1):245–301, 1984.
- [98] M. B. Marcus and L. A. Shepp. Continuity of Gaussian Processes. *Transactions of the American Mathematical Society*, 151:377–391, 1970.
- [99] L. Moisan. Periodic Plus Smooth Image Decomposition. *Journal of Mathematical Imaging and Vision*, 39(2):161–179, 2011.
- [100] M. C. Morrone and D. C. Burr. Feature Detection in Human Vision: A Phase-Dependent Energy Model. *Proceedings of the Royal Society of London. Series B, Biological Sciences*, pages 221–245, 1988.
- [101] D. Navon. Forest before Trees: The Precedence of Global Features in Visual Perception. *Cognitive Psychology*, 9(3):353–383, 1977.
- [102] A. Oliva and A. Torralba. Building the Gist of a Scene: The Role of Global Image Features in Recognition. *Progress in Brain Research*, 155:23–36, 2006.
- [103] B. A. Olshausen. Emergence of Simple-Cell Receptive Field Properties by Learning a Sparse Code for Natural Images. *Nature*, 381(6583):607–609, 1996.
- [104] B. A. Olshausen and D. J. Field. Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1? *Vision Research*, 37(23):3311–3325, 1997.

- [105] B. A. Olshausen and D. J. Field. Sparse Coding of Sensory Inputs. *Current Opinion in Neurobiology*, 14(4):481–487, 2004.
- [106] A. V. Oppenheim and J. S. Lim. The Importance of Phase in Signals. *Proceedings of the IEEE*, 69(5):529–541, 1981.
- [107] V. Ostromoukhov and R. D. Hersch. Artistic Screening. In *Proceedings of the 22nd annual Conference on Computer Graphics and Interactive Techniques*, pages 219–228. ACM, 1995.
- [108] R. E. A. C. Paley and A. Zygmund. On Some Series of Functions, part I. *Mathematical Proceedings of the Cambridge Philosophical Society*, 26:337–357, 1930.
- [109] R. E. A. C. Paley and A. Zygmund. On Some Series of Functions, part II. *Mathematical Proceedings of the Cambridge Philosophical Society*, 26:458–474, 1930.
- [110] R. E. A. C. Paley and A. Zygmund. On Some Series of Functions, part III. *Mathematical Proceedings of the Cambridge Philosophical Society*, 28:190–205, 1932.
- [111] R. Penrose. Pentaplexity a Class of Non-Periodic Tilings of the Plane. *The Mathematical Intelligencer*, 2(1):32–37, 1979.
- [112] K. Perlin. An Image Synthesizer. *ACM SIGGRAPH Computer Graphics*, 19(3):287–296, 1985.
- [113] B. Pesquet-Popescu and J. Lévy-Véhel. Stochastic Fractal Models for Image Processing. *Signal Processing Magazine, IEEE*, 19(5):48–62, 2002.
- [114] G. Peyré. Sparse Modeling of Textures. *Journal of Mathematical Imaging and Vision*, 34(1):17–31, 2009.
- [115] J. Portilla and E. P. Simoncelli. A Parametric Texture Model Based on Joint Statistics of Complex Wavelet Coefficients. *International Journal of Computer Vision*, 40(1):49–70, 2000.
- [116] A. R. Rao and G. L. Lohse. Identifying High-Level Features of Texture Perception. In *SPIE/IS&T 1992 Symposium on Electronic Imaging: Science and Technology*, pages 424–435. International Society for Optics and Photonics, 1992.
- [117] S. Ronsin. Synthèse de Textures par Randomisation de Phases. *Master's Thesis, ENS Cachan*, 2010.
- [118] S. Ronsin, H. Biermé, and L. Moisan. The Billard Theorem for Multiple Random Fourier Series. *To appear in the Journal of Fourier Analysis and Applications*, 2014.
- [119] R. Salem and A. Zygmund. Some Properties of Trigonometric Series whose Terms Have Random Signs. *Acta Mathematica*, 91:245–301, 1954.
- [120] D. Saupe. Algorithms for Random Fractals. In *The Science of Fractal Images*, pages 71–136. Springer, 1988.

- [121] A. Secord. Weighted Voronoi Stippling. In *Proceedings of the 2nd International Symposium on Non-Photorealistic Animation and Rendering*, pages 37–43. ACM, 2002.
- [122] C. E. Shannon. A Mathematical Theory of Communication. *The Bell System Technical Journal*, 27(1):379–423 and 623–656, 1948.
- [123] R. Silvers. *Photomosaics: Putting Pictures in their Place*. PhD thesis, Massachusetts Institute of Technology, 1996.
- [124] E. P. Simoncelli and B. A. Olshausen. Natural Image Statistics and Neural Representation. *Annual Review of Neuroscience*, 24(1):1193–1216, 2001.
- [125] G. Svaetichin. Spectral Response Curves from Single Cones. *Acta Physiologica Scandinavica*, 39(134):17, 1956.
- [126] N. Tevzadze. On the Convergence of Double Fourier Series of Quadratic Summable Functions. *Sakharth. SSR Mecn. Akad. Moambe*, 77(5):pp. 277–279, 1970.
- [127] J. van Wijk. Spot Noise Texture Synthesis for Data Visualization. *ACM SIGGRAPH Computer Graphics*, 25(4):309–318, July 1991.
- [128] Richard F Voss. Fractals in Nature: from Characterization to Simulation. In *The Science of Fractal Images*, pages 21–70. Springer, 1988.
- [129] L. Y. Wei, S. Lefebvre, V. Kwatra, and G. Turk. State of the Art in Example-Based Texture Synthesis. *Eurographics 2009, State of the Art Report, EG-STAR*, pages 93–117, 2009.
- [130] L.-Y. Wei and M. Levoy. Fast Texture Synthesis Using Tree-Structured Vector Quantization. In *Proceedings of the 27th annual Conference on Computer Graphics and Interactive Techniques*, pages 479–488. ACM Press/Addison-Wesley Publishing Co., 2000.
- [131] H. Wendt, S. G. Roux, S. Jaffard, and P. Abry. Wavelet Leaders and Bootstrap for Multifractal Analysis of Images. *Signal Processing*, 89(6):1100–1114, 2009.
- [132] G. S. Xia, S. Ferradans, G. Peyré, and J.-F. Aujol. Synthesizing and Mixing Stationary Gaussian Texture Models. *to appear in SIAM Journal on Imaging Sciences*, 2013.
- [133] S. C. Zhu, X. W. Liu, and Y. N. Wu. Exploring Texture Ensembles by Efficient Markov Chain Monte Carlo. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(6):554–569, 2000.
- [134] S. C. Zhu, Y. N. Wu, and D. Mumford. Minimax Entropy Principle and its Application to Texture Modeling. *Neural computation*, 9(8):1627–1660, 1997.
- [135] S. C. Zhu, Y. N. Wu, and D. Mumford. Filters, Random Fields and Maximum Entropy (FRAME): Towards a Unified Theory for Texture Modeling. *International Journal of Computer Vision*, 27(2):107–126, 1998.

- [136] S.C. Zhu, C.E. Guo, Y.Z. Wang, and Z.J. Xu. What are Textons? *International Journal of Computer Vision*, 62(1/2):121–143, 2005.

Résumé

Le travail présenté dans cette thèse porte sur différents aspects de l'analyse et de la synthèse de textures. Plus particulièrement, nous nous intéressons à des modèles de “micro-textures”, c'est-à-dire de textures dépourvues de motifs, dont les phases de la transformée de Fourier sont aléatoires.

La première partie de cette thèse étudie quelques propriétés d'un représentant particulier de chaque classe de micro-texture, que nous appelons son *texton*. Un résultat prouve l'optimalité de la concentration du *texton* autour de l'origine (zéro spatial). Nous tirons parti de ce phénomène de concentration pour proposer des représentations parcimonieuses des micro-textures, approchées et exactes sous certaines hypothèses. Nous discutons différentes généralisations du *texton* au cas des textures en couleurs et nous efforçons d'étendre les approximations parcimonieuses définies dans le cadre d'images à niveaux de gris.

Nous proposons ensuite d'interpréter l'optimalité de la concentration du *texton* comme un résultat de projection. Nous présentons plusieurs simulations de projection sur différents espaces d'images. Ces expériences numériques montrent que l'hypothèse, largement répandue en traitement du signal, selon laquelle “la géométrie des images est codée dans leur phase”, mérite d'être nuancée.

Dans la dernière partie de cette thèse, nous étudions certaines propriétés asymptotiques de micro-textures du modèle de synthèse à phases aléatoires. Après nous être intéressés à la convergence vers un champ Gaussien de ce modèle dans son extension sur le plan discret (non-périodique) dans la première partie de cette thèse, nous nous intéressons à la convergence et aux propriétés locales (continuité et régularité) de sommes de Fourier aléatoires infinies multi-dimensionnelles. Nous étendons au cas de la dimension quelconque un théorème de Billard et Kahane prouvant l'équivalence, pour les sommes aléatoires considérées, entre convergence uniforme p.s., convergence partout p.s. et continuité de la somme p.s. Nous étendons au cadre multi-dimensionnel des conditions suffisantes et des conditions nécessaires pour la continuité et la régularité Hölderienne de ces sommes, dans un cadre d'analyse anisotropique.

Abstract

This dissertation deals with random-phase texture analysis and synthesis – *i.e.* textures without patterns, and with random Fourier phase.

The first part studies properties of a special representant of each class of micro-textures, that we name *texton*. We prove an optimality result with respect to the spatial concentration around the origin. We take advantage of this concentration phenomenon to propose sparse representations of micro-textures, approximate and exact under some hypothesis. We discuss generalizations of the *texton* to color images and extend the sparse approximations developed for gray-scale images.

We interpret the optimality of concentration as a projection result, and discuss several other projection experiments on different image spaces. These numerical experiments show that the hypothesis, widely believed in signal processing, claiming that “the geometry of images is encoded in their phase” deserves further inquiry.

In the last part of this dissertation, we study some asymptotical properties of the random-phase texture model. We proved the convergence to a Gaussian field while extending random-phase textures towards the whole (non-periodic) discrete plane in the first part of the dissertation, and we focus here on convergence and local properties (continuity and regularity) of multi-dimensional infinite random Fourier sums. We extend to the multi-dimensional case a theorem of Billard and Kahane showing the equivalence, for the random sums considered, between a.s. uniform convergence, a.s. pointwise convergence everywhere, and a.s. continuity everywhere. We also extend to the multi-dimensional case, sufficient conditions and necessary conditions for continuity and Hölder regularity of these sums, with an anisotropic framework.

