

DISSERTATION

submitted to the

**Combined Faculty for the Natural Sciences
and Mathematics**

of the

University of Heidelberg, Germany

for the degree of
Doctor of Natural Sciences

put forward by
Dipl. Stat. Eva-Maria Didden
born in Memmingen

Date of oral examination:

Modeling of Locally Scaled Spatial Point Processes, and Applications in Image Analysis

Advisors: Dr. Thordis Linda Thorarinsdottir
Prof. Dr. Christoph Schnörr

Zusammenfassung

Räumliche Punktprozesse sind ein bewährtes statistisches Analysewerkzeug, welches in zahlreichen Wissenschaftsbereichen – wie beispielsweise der Ökologie, der Epidemiologie oder der Werkstoffkunde – zum Einsatz kommt. Intensive Forschung wird vor allem dahingehend betrieben, mit Hilfe von Punktprozessen systematische Veränderungen in räumlichen Strukturen zu beschreiben sowie innovative Methoden der Parameterschätzung zu entwickeln. Das Anwendungsspektrum räumlicher Punktprozessmodelle ist insbesondere in der Bildanalyse groß. Typische Beispiele sind Aufnahmen von Baumbeständen in der Fernerkundung, von Zellen in der Biologie oder von Verbundkonstruktionen in den Materialwissenschaften. Auf Grund ihrer Praxisrelevanz und vielfältigen Einsetzbarkeit erscheint vor allem die Modellklasse der lokal-skalierten Punktprozesse geeignet zur Beschreibung räumlicher Objektanordnungen. Eine unbekannte Normalisierungskonstante in der Likelihood erschwert allerdings die statistische Inferenz und verlangt nach ausgeklügelten Simulations- und Schätzstrategien. In dieser Arbeit wird ein bayesianischer Ansatz zur Modellierung lokal-skalierten Punktprozesse eingeführt und anschließend unter anderem dazu verwendet, Mais-Genotypen anhand der Gefäßstrukturen in den Halmen zu klassifizieren. Es liegen hierfür Querschnittsaufnahmen der Halme vor. Weitere räumliche Punktprozessmodelle sind speziell für die bild-basierte Beschreibung der Ausrichtung texturierter Oberflächen im dreidimensionalen Raum vorgesehen, welche allgemein unter den Begriff *Shape-from-Texture-Analyse* fällt. In den Bereichen Mustererkennung und Bildverarbeitung sind das Verstehen von Textur sowie die Erkennung und Quantifizierung von Geometrien zwei wichtige Problemstellungen. Häufig geht es darum, die Fotografie einer bestimmten Szene zunächst auf die lokalen geometrischen Bildstrukturen hin zu untersuchen und darauf basierend die Szenenausrichtung im Dreidimensionalen zu erschließen. Somit lassen sich u.a. Kameraeinstellungen rekonstruieren. Die in dieser Arbeit vorgestellten statistischen Methoden zur Shape-from-Texture-Analyse umfassen lokal-skalierte Punktprozessverfahren sowie den Entwurf eines bayesianischen markierten Punktprozessmodells.

Abstract

Spatial point processes provide a statistical framework for modeling random arrangements of objects, which is of relevance in a variety of scientific disciplines, including ecology, spatial epidemiology and material science. Describing systematic spatial variations within this framework and developing methods for estimating parameters from empirical data constitute an active area of research. Image analysis, in particular, provides a range of scenarios to which point process models are applicable. Typical examples are images of trees in remote sensing, cells in biology, or composite structures in material science. Due to its real-world orientation and versatility, the class of the recently developed locally scaled point processes appears particularly suitable for the modeling of spatial object patterns. An unknown normalizing constant in the likelihood, however, makes inference complicated and requires elaborate techniques. This work presents an efficient Bayesian inference concept for locally scaled point processes. The suggested optimization procedure is applied to images of cross-sections through the stems of maize plants, where the goal is to accurately describe and classify different genotypes based on the spatial arrangement of their vascular bundles. A further spatial point process framework is specifically provided for the estimation of shape from texture. Texture learning and the estimation of surface orientation are two important tasks in pattern analysis and computer vision. Given the image of a scene in three-dimensional space, a frequent goal is to derive global geometrical knowledge, e.g. information on camera positioning and angle, from the local textural characteristics in the image. The statistical framework proposed comprises locally scaled point process strategies as well as the draft of a Bayesian marked point process model for inferring shape from texture.

Acknowledgements

I gratefully acknowledge the financial support by the German Research Foundation (DFG) within the program “Spatio-/Temporal Graphical Models and Applications in Image Analysis”, and the scientific support by the Combined Faculty for the Natural Sciences and Mathematics of the University of Heidelberg.

Particularly, I would like to express my gratitude to my advisors Thordis Thorarinsdottir and Christoph Schnörr who have helped me during my PhD studies with their expertise, scientific know-how, patience, and warm and friendly nature. My deep appreciation also goes to Alex Lenkoski, Donald Richards, Tilmann Gneiting and David Legland for their valuable feedback and our fruitful discussions.

I owe an important debt to my friends & (ex-)colleagues in Heidelberg. Thanks for the great time I have/had with you, Andreas, Andreea, Annette, Arati, Barbara, Bernhard, Bogdan, Boris, Christoph, Dominic, Ecatarina, Elena, Evelyn, Fabian, Fabian (I will miss our library sessions), Florian, Frank, Gabriel, Jasmin, Johannes, Jörg, Karsten, Katharina, Kira, Markus, Martin, Mercedes, Michael, Mattia, Niko, Paul, Robert, Stefania, Tabea, Tobias, Vera,...!

Despite the distance, I have always benefited from the lively scientific and non-scientific conversations with Elisabeth, Birgit, Mirjam and the rest of my former “gang” from Munich, and I am very proud of having such faithful friends.

Last but not least, I want to thank MaPa, Chrissi, Roni and Jakob for their never-ending mental support and encouragement.

Contents

List of Algorithms	iv
List of Figures	v
List of Tables	viii
1 Introduction	1
1.1 Overview	1
1.2 Related Work	4
1.3 Contributions and Organization	10
2 Preliminaries on Spatial Point Processes	13
2.1 Point Process Theory	14
2.2 Poisson Processes	19
2.3 Markov/Gibbs Processes	22
2.4 Sampling Techniques	26
2.5 Parameter Estimation	35
3 Preliminaries on Image Analysis	45
3.1 Camera Projection	45
3.2 Detection of Local Symmetry Elements	54
4 Locally Scaled Spatial Point Processes	61
4.1 Locally Scaled Point Process Models	62
4.2 Proper Scaling Functions	64
4.3 Bayesian Inference Framework	67
4.4 Simulation Study	73
5 Analysis of Biological Image Data	77
5.1 Data Preparation	78
5.2 Step-wise Scaling Constraints	79
5.3 Two-Stage Inference Procedure	83
5.4 Conclusion and Outlook	90
6 Shape from Texture using Locally Scaled Point Processes	95
6.1 Motivation	95
6.2 Intuitive Preliminary Framework	100
6.3 Estimation of the Latent Point Process	105
6.4 Model Construction and Inference	110
6.5 Case Studies	113
6.6 Outlook	119

7	Analysis of Textured 3D Scenes using Marked Point Processes	121
7.1	Motivation	121
7.2	Geometric Feature Extraction	122
7.3	Shape from Texture via Marked Point Processes	128
7.4	Case Studies	136
7.5	Outlook	141
8	Discussion	143
	Bibliography	149
	Subject Index	157
	Notation Index	163

List of Algorithms

2.1	Gibbs sampler for a fixed number of variables	29
2.2	Metropolis-Hastings sampler	30
2.3	Sampling from a homogeneous Poisson process.	31
2.4	Sampling from an inhomogeneous Poisson process.	32
2.5	Sampling from a fix-n Gibbs process	33
2.6	Birth-death-move-M-H algorithm	43
2.7	Exchange algorithm.	44
4.1	Sketch: Dominated coupling from the past	72
6.1	Inferring a point process realization and its parameters from a probability map	103
6.2	Neighborhood-based point process estimation	108
6.3	Point process estimation and Voronoi tessellation	109
7.1	Estimating the posterior distribution of a marked point process .	135

List of Figures

1.1	Cellular structures in maize stems of a certain genotype	2
1.2	Examples of textured 3D scenes	3
2.1	Simulation of a forest area	18
2.2	Samples from point processes with different structural properties	19
2.3	Simulations from different Strauss type models	25
3.1	Process of image formation	46
3.2	Perspective projections of a 3D scene point	51
3.3	Original image of a brick wall	57
3.4	Brick wall image after median filtering	58
3.5	Brick wall image after median and DoG filtering	58
3.6	Distance transformation of a brick wall image after median and DoG filtering	59
4.1	Simulations from locally scaled Strauss models	66
4.2	Exponential scaling effects of varying strength	67
4.3	Simulations from an exponentially scaled Strauss model	74
4.4	Full conditionals of the model parameters	75
4.5	Optimization of the pseudo log-likelihood subject to the nuisance parameter R and the previously estimated scaling effects	75
5.1	Cellular structures in cross-sections through maize stems	77
5.2	Cross-sections through representative stems before and after a circular normalization	79
5.3	Step scaling effects in a point pattern in $W = D_1(0)$	81
5.4	Determination of the inner stem radius (genotype I)	82
5.5	Determination of the inner stem radius (genotype II)	83
5.6	Full conditional posteriors with gray bars marking the priors . .	86
5.7	Results from PIT calibration diagnostics	87
5.8	Full conditionals describing $p(\boldsymbol{\theta} \mathbf{x})$	89
5.9	Full conditionals describing $p(\boldsymbol{\theta} \mathbf{x}^{(1)})$ and $p(\boldsymbol{\theta} \mathbf{x}^{(1)}, \mathbf{x}^{(2)})$	92
6.1	Sequence of operations proposed for inferring shape from texture	97
6.2	3D camera model	98
6.3	Mappings of regular homogeneous point patterns in 3D space onto a 2D plane	99
6.4	Exponentially scaled distances from the point $(0, 0)^\top$	102
6.5	Learning a point process realization from a brick image	104

6.6	Model validation I	105
6.7	Model validation II and estimation of δ	106
6.8	Voronoi tessellation framework	110
6.9	Examples of point distances under perspective scaling	112
6.10	Simulated Poisson point patterns with 3D shape determined by the outer normal δ	114
6.11	Point process estimation in scenes with different amounts of non- convex texture elements	115
6.12	Estimation of shape from texture.	117
6.13	Estimation of the unit normal δ with respect k_1	118
6.14	Log-likelihood ratios with respect to k_1	118
7.1	Finite library of geometric marks	123
7.2	Images of interest	129
7.3	Euclidean similarity transformations of a dictionary atom	130
7.4	Euclidean similarity transformations of a square	131
7.5	Modeling of a latent marked point process	136
7.6	Estimates of the full conditionals of the geometric parameters	137
7.7	Pixel-wise MAP estimates of the geometric parameters	138
7.8	Convergence of the algorithm and overall model performance	139
7.9	Estimation of a marked point process realization	140
7.10	Empirical distributions of the geometric posterior estimates	141

List of Tables

5.1	Results from the first stage of the analysis	91
5.2	Results from the second stage of the analysis	91
6.1	True angles and composite likelihood estimates	114
6.2	Composite likelihood estimates of the surface normals	116
6.3	Estimates of the perspective scaling parameters	119

1 Introduction

God used beautiful mathematics in creating the world.

– Paul A. M. Dirac (August 8th, 1902 – October 20th, 1984)

This dissertation is interdisciplinarily oriented in that it applies mathematical principles and spatial statistics to demanding real-world problems related to biology and computer science. To be precise, it introduces newly developed inference strategies for spatial point processes on the one hand, and couples spatial statistics with the geometric analysis of image data on the other hand. Therefore, a detailed treatment of both statistical theory and the relevant principles on imaging geometry will be provided. Our work is intended to be accessible by a readership particularly interested in spatial statistics and/or computer vision. In the following section, we briefly summarize our research objectives and thus give an overview of the essential items of this thesis.

1.1 Overview

Many classes of spatial point processes have been developed in the past decades, among which the class of locally scaled point processes introduced by Hahn et al. (2003) appears particularly appealing in terms of further advancements and practical applications. This new type of inhomogeneous point process stems from a local adjustment of a homogeneous template process, such that regions of dissimilar packing density differ only by a scale factor. As a consequence, local geometries are kept constant. Locally scaled point processes prove to be promising for various fields of study, including biology, demography and astronomy. Due to their complexity and numerical intractability, however, the development of efficient inference strategies requires sophistication, which motivates our first research question [Didden et al. (2015)].

How to Implement a Bayesian Hierarchical Inference Framework for Locally Scaled Point Processes?

Although work on simulation-based inference techniques and, in particular, on the development of flexible Markov chain Monte Carlo (MCMC) methods is one of the major research topics in the spatial statistics community, modeling concepts for locally scaled point processes are rare. In Chapter 4, we therefore introduce a Bayesian alternative to the frequentist composite likelihood approach presented by Prokešová et al. (2006). Owing to the coupling of inhomogeneity and interaction constraints, the likelihood of a locally scaled point

process is usually not available in complete form, which makes the posterior distribution of the model parameters doubly intractable. This problem cannot be tackled via common MCMC algorithms, such as standard Gibbs sampling or Metropolis-Hastings routines. We propose an appropriate implementation of the exchange algorithm by Murray et al. (2012), which includes perfect sampling from the unnormalized likelihood performed via dominated coupling from the past [Berthelsen and Møller (2003)].

A properly scaled point process model relies on a scaling function which is identifiable and allows for the exact computation of locally scaled distances. Depending on the point pattern, the definition of a proper scaling function poses minor or major difficulties. From our collaborators David Legland¹ and Marie-Françoise Devaux², we are given images of cross-sections through maize stems (see Fig. 1.1), where the spatial arrangement of the vascular bundles can be seen as a realization of an inhomogeneous point process. Our second research question therefore concerns the modeling of the bundle distribution under suitably specified scaling assumptions [Didden et al. (2015)].

How to Classify Genotypes of Maize Plants through a Locally Scaled Point Process Approach?




Figure 1.1: Cellular structures in maize stems of a certain genotype.

Two series of maize data from two different genotypes are at our disposal. With the objective of detecting differences and similarities in the spatial distribution of the vascular bundles, we first project the bundle coordinates to circular discs in order to transform the original stem contours into simpler geometric shapes that can be handled easier. Since it appears that the bundle intensity is denser in the outer than in the inner parts of the stems, we develop a proper and intuitively comprehensible step scaling function. Via a two-stage procedure, we first model the homogeneous bundle distribution in the inner parts of the stems in a Bayesian manner and then enter the posterior information obtained on the unknown parameters as prior knowledge into the second modeling stage.

¹INRA & AgroParisTech, *UMR 782 Food Process Engineering and Microbiology*, Thiverval-Grignon; INRA & AgroParisTech, *UMR 1318 Institut Jean-Pierre Bourgin*, Versailles

²INRA, UR 1268 *Biopolymers, Interactions and Assemblies*, Nantes

The second stage comprises the analysis of the full datasets, i.e. the locations of the vascular bundles in the inner and outer stem sections. It is performed via our newly developed Bayesian inference technique (see Chapter 4), under inhomogeneity constraints induced through the proposed step scaling function. The whole biological project is presented in Chapter 5 of this thesis.

An entirely different type of image data forms the basis for our third research project dealing with the question of how to estimate geometric attributes, such as camera positioning and angle, from the two-dimensional (2D) image of a textured three-dimensional (3D) scene [Didden et al. (2013)].

How to Use Locally Scaled Point Processes for the Estimation of Shape from Texture?

We work with images of textured scenes, where the texture elements are near-regular in shape (see Fig. 1.2). Chapter 6 and Chapter 7 are concerned with




Figure 1.2: Images of textured 3D scenes.

the introduction of two suitable concepts for inferring 3D geometries from such images. The objective of the first concept is to learn 3D shape from the inhomogeneous spatial distribution of the texture elements in the 2D image plane (see Chapter 6). For this purpose, texture is regarded as a realization of a locally scaled point process where each point is associated with exactly one texture element. We propose two alternative approaches, a Gibbs model with hardcore interaction constraints and a Strauss model. In both cases, appropriately specified scaling assumptions based on spherical coordinates allow us to reconstruct the orientation of the camera towards the 3D space. The respective optimization algorithms are of simple structure. However, they require some image preprocessing facilitating the learning of a point process realization from the given image. We apply smoothing and distance transformation mechanisms to translate the image into a probability map. From the probability map, we can read how likely each single pixel corresponds to the symmetry center of a texture element and thus to a point of the latent point process.

Our second modeling concept does not require any image preprocessing. It is based on the idea to infer 3D shape from the local 2D deformations of the texture elements in the image plane (see Chapter 7). We propose a Bayesian marked

point process framework, where the marks correspond to geometric transformations of distinctive patches of texture learned from a training image without 3D effect. Assuming that the spatial pixel value distribution follows a mixture of Gaussians, we define appropriate priors for the latent point process, the mark attributes and the remaining unknown parameters contained in the Gaussian likelihood, and implement a reversible jump birth-death-move Metropolis-Hastings algorithm. Although our hierarchical inference framework performs promising with regard to the identification and estimation of local 2D geometries, we have not yet established a connection between this location-dependent knowledge and a global 3D geometric description for the entire scene.

Summing up, this dissertation comprises three principal research topics linked through their methodological similarities, while somewhat differing in their objective targets. Our first research question is of a theoretical nature, whereas the two other problems have an application-oriented and interdisciplinary character. The proposed modeling concepts are therefore inspired by related literature from different scientific disciplines, particularly from spatial statistics and image analysis. Section 1.2 presents a selection of research articles and books, while further helpful references will be provided throughout this thesis.

1.2 Related Work

First focusing on the statistical aspects of our work, we present a selection of related research publications in which the relevant statistical methodology is discussed, statistical optimization techniques are provided, and approaches to data problems similar to our case studies are suggested. We then move on to literature on image analysis and discuss projects dealing with the question of how to learn 3D shape from a 2D image with depth effect. Finally, an overview is given on marked point process approaches to the extraction of geometrical features from images.

Spatial Point Processes

Profound mathematical basic knowledge on point processes in metric spaces of arbitrary dimension can be gathered from the books by Daley and Vere-Jones (2003, 2007). For a less formal and more compact introduction to spatial point processes, we recommend the contributions available in Baddeley et al. (2006) as well as in Illian et al. (2008). The collected edition published by Gelfand et al. (2010) is concerned with various fields of spatial statistics. A broad overview of spatial statistics methodology can also be gathered from Baddeley and Turner (2005) who have collected and implemented several relevant algorithms and made these routines available in the R-package ‘`spatstat`’. For a detailed

in-depth introduction to simulation and inference techniques for spatial point processes, we refer to Møller and Waagepetersen (2003).

The majority of the statistical approaches introduced or discussed in our work are based on the locally scaled point processes developed by Hahn et al. (2003). Technical explanations concerning the modeling of this new class of spatial point processes are provided by Prokešová et al. (2006). A locally scaled point process corresponds to an inhomogeneous Markov process [Kendall (1990)] where the variation in the point packing density is driven by a location-dependent scale factor. Other previously established approaches to the inclusion of heterogeneity into a Markov model are, for instance, presented by Ogata and Tanemura (1986) and Baddeley et al. (2000). A modeling class which is similar to the locally scaled point processes by Hahn et al. is the class of the transformation inhomogeneous point processes introduced by Jensen and Nielsen (2001) [see also Nielsen and Jensen (2004)]. In contrast to Jensen and Nielsen, Hahn et al. consider isotropic neighborhoods for evaluating interaction. The location-dependent scaling due to Hahn et al. implicates that local geometry is not affected by inhomogeneity, which makes densely packed areas look like scaled versions of regions covered with fewer points. This specific property is of practical relevance in view of various real-world data problems. Shimatani and Kubota (2004) and Eckel et al. (2009), for example, discuss the necessity to model tree populations or root networks under the consideration of location-dependent variations in the soil conditions. Despite their real-world orientation, previous work on locally scaled point processes still leaves plenty of room for methodological and technical development. We use this rather uninvestigated point process class as a primary statistical tool and ingredient for our interdisciplinary research study.

In general, point patterns with a constant intensity and without repulsion or clustering effects are mathematically easier to handle than heterogeneous patterns with interacting points. Based on the test principles presented by Dwass (1957), a set of Monte Carlo approaches has been made available for determining whether a homogeneous or inhomogeneous point process contains repulsive, attractive, or independent points [see e.g. Ripley (1977), Besag and Diggle (1977), Besag and Clifford (1989)]. The most common Monte Carlo tests rely either on second-order summary statistics, including the widely used K-, L-, and g-functions [see e.g. Ripley (1976), Baddeley et al. (2000)] and the directional K-functions [see Stoyan and Stoyan (2000), Brix and Møller (2001)], or on distance-based summary statistics, including the F-, G-, and J-functions introduced by Van Lieshout and Baddeley (1996) and discussed by Bedford and Van den Berg (1997). To assess whether a point process model captures spatial variations and dependency structures to a satisfactory extent, residuals may be calculated and examined [Baddeley et al. (2005)]. Thorarinsdóttir (2013) recommends calibration diagnostics based on the probability integral

transform (PIT) [Dawid (1984)] for model validation and selection purposes [see also Diebold et al. (1998), Gneiting et al. (2007)].

Due to a numerically inaccessible normalizing constant, the likelihood of interaction point processes is usually not available in complete form. Optimization procedures therefore require special diligence. The most common frequentist inference approaches are based on approximations of the likelihood. The composite likelihood [Lindsay (1988)], for instance, is derivable from the first-order intensity function of the point process, whereas the pseudo likelihood [Besag (1977)] also accounts for second-order structures. Several follow-up methods are based on the pseudo likelihood principle [see e.g. Huang and Ogata (1999, 2002)], while effort has also been made on the development of alternative modeling strategies, such as approximate maximum likelihood techniques [see e.g. Ogata and Tanemura (1981, 1984)], or Monte Carlo likelihood simulation [Geyer (1999)]. The latter belongs to the class of the Markov chain Monte Carlo (MCMC) concepts [Meyn and Tweedie (2009)], among which the Gibbs and the Metropolis-Hasting (M-H) samplers [Casella and George (1992), Metropolis et al. (1953), Hastings (1970)] are the most common tools. If the set of the unknown parameters is not of fixed dimension, enhanced M-H algorithms accounting for reversible jumps between different parameter spaces need to be considered [see e.g. Green (1995)]. This is often the case when marked point processes are being modeled. Another MCMC routine is provided in the work by Murray et al. (2008) who introduce a Gaussian process density sampler as a tool for non-parametric Bayesian inference [see also Adams et al. (2009)], whereas Walker (2011) suggests a Bayesian procedure relying on a latent-variable-based approximation of the integral that formally determines the normalizing constant.

There is some gap in the statistical literature in view of inference strategies for locally scaled interaction point processes. Rajala and Penttinen (2012) propose a Markovian modeling framework for hardcore Gibbs processes. A pseudo likelihood approach to the analysis of exponentially scaled point patterns of Strauss type [Strauss (1975)] is discussed in the work by Prokešová et al. (2006), and Bognar (2005) presents an M-H framework where the intractable acceptance ratio is estimated via importance sampling. This latter method, however, suffers from computational inefficiency. We present an alternative Bayesian inference method to estimate the posterior distributions of the model parameters, despite the unknown normalizing constant in the likelihood and the resulting double-intractability of the posterior model.

To deal with double-intractable distributions, Møller et al. (2006) and Murray et al. (2012) have developed M-H algorithms with an additional auxiliary variable scheme that makes the unknown normalizing constant in the likelihood cancel out. We decide for the exchange principle by Murray et al., since it is slightly more straightforward than the single-auxiliary variable method proposed

by Møller et al.. To collect the auxiliary samples, we make use of the perfect simulation framework for spatial point processes introduced by Berthelsen and Møller (2003) [see also Berthelsen and Møller (2002)] which proves to be applicable under the consideration of a location-dependent scaling. Further perfect simulation concepts for interaction point processes are explained by Kendall (1998), Kendall and Møller (2000), and Fernández et al. (2002). We prefer exact sampling procedures to approximate samplers [Liang et al. (2007), Jin and Liang (2012)], such as the double-M-H sampler by Liang (2010), even though these MCMC techniques are normally favorable in terms of computer runtime.

As already stated in Section 1.1, we apply our newly established Bayesian inference framework to images of cross-sections through the stems of maize plants, in order to analyze genotype-specific characteristics in the packing density of the vascular bundles. The data have been preprocessed and made available by David Legland³ and Marie-Françoise Devaux⁴. Our collaborators have previously been working on the maize images and successfully developed a normalization mechanism as well as a bundle intensity estimator allowing for direct visual comparisons between different stems [Legland et al. (2014)]. This framework, however, analyzes the packing density of the vascular bundles in a very general manner and neither accounts for repulsions between the bundles nor provides a concrete statistical model describing the inhomogeneity in their spatial arrangement.

Shape-from-Texture

Besides the biological image data, images of textured 3D scenes attract our research attention. To be precise, the learning of 3D shape from the 2D textural information contained in the images is a key topic of our interdisciplinary research work. Therefore, mathematical basic knowledge on camera projection is indispensable, including backgrounds on projective geometry as well as on camera modeling and calibration. For an overview on computer graphics and classical geometries, we refer to the standard works by Hughes et al. (2014) and Ramírez Galarza and Seade (2007). Hartley and Zisserman (2000) and Faugeras and Luong (2001) lay particular stress on multiple view set-ups and interrelations between camera projections. For a brief and concise mathematical introduction to optical flow, we recommend the article contributed by Becker et al. (2014). Being familiar with digital image processing [see e.g. Jähne (1989), Mather and Koch (2010)] and computational symmetry [Liu et al. (2009)] is a further advantage with regard to the image preparation and texture identification methods used in this work.

Tuceryan and Jain (1998) review and discuss texture analysis in several respects, particularly focusing on mechanisms to detect textural features in real-

³INRA & AgroParisTech, *UMR 782 Food Process Engineering and Microbiology*, Thiverval-Grignon; INRA & AgroParisTech, *UMR 1318 Institut Jean-Pierre Bourgin*, Versailles

⁴INRA, UR 1268 *Biopolymers, Interactions and Assemblies*, Nantes

world scenes projected onto image planes. Stevens (1980) explains how surface orientation affects the appearance of texture in images basically in that the texture elements feature foreshortening effects, geometric deformations, and variations in their packing density. The author uses the slant [see also Gibson (1950)] and the tilt to measure surface orientation, where the slant is defined as the angle between a normal to the surface and a normal to the image plane, and the tilt corresponds to the angle between the projection of the surface normal onto the image plane and a predetermined coordinate axis in the image. Gibson (1950) discusses, from a psychological as well as from a mathematical standpoint, to which extent natural 2D images provide immediate perceptual knowledge on the 3D shapes of objects and surfaces.

Several statistical strategies have been established for inferring the orientation of the camera towards a textured 3D scene from one single image. Most of these modeling techniques, however, require strict symmetry assumptions or a very regular arrangement of the texture elements. Blostein and Ahuja (1989), for instance, only deal with uniformly distributed texture elements of circular shape. Under the assumption of an almost undistorted camera projection, Witkin (1981) proposes to learn surface orientation from the edge locations and directions observed in the image. Aloimonos (1988), and Warren and Mamasian (2010) additionally consider perspective projection. The inference concept by Blake and Marinos (1990) is applicable to images of 3D planes containing isotropically oriented line elements. On the backgrounds of Witkin’s statistical approach and Kanatani’s texture moments [Kanatani (1989)], the authors estimate surface orientation along with a spatial error distribution which is important for integrating shape information, and furthermore present a tool for testing hypotheses about intrinsic texture attributes.

Focusing on curved surfaces and assuming strict homogeneity, Malik and Rosenholtz (1997) take affine transformation of adjacent image patches as a basis to locally describe and quantify texture distortion. Building on the work by Malik and Rosenholtz, Clerc and Mallat (2002) propose a framework for modeling 3D orientation under less restrictive homogeneity constraints that, however, are hard to verify in practice. A critical overview on the strong and practically unsuitable assumptions prior work has been based on is provided in the research article by Loh and Hartley (2005). The authors introduce a methodological framework for dealing with inhomogeneous, non-stationary, anisotropic and perspective texture. This framework is closely related to the inference procedure suggested by Forsyth (2006). Both methods rely on estimating deformations of individual texture elements, while strongly relaxing restrictions on the global textural appearance.

In contrast to the above-referenced research on the learning of shape from texture, we regard the image of a textured plane in 3D space as a realization of

a latent point process where each point represent one texture element. This allows us to quantify perspective foreshortening and inhomogeneity of the texture through the location-dependent scaling properties of the point process, provided that the texture elements are near-regular and close to convex in shape.

Image Analysis via Marked Point Processes

Another way to look at texture by means of point processes is to regard the texture elements as latent marked points. In the literature, several such marked point process approaches are discussed, including their applications to real-world data samples. The modeling framework by Descombes and Zerubia (2008), for example, allows for the identification of rectangular buildings in remotely sensed images of densely populated areas, whereas Tournaire et al. (1965) propose an inference technique specifically designed to model the dashed lines of road markings. While both approaches have a very limited scope of application, Lacoste et al. (2005) introduce a marked point process concept for analyzing more general types of line networks on aerial photographs, such as hydrographic networks. Sun et al. (2007) provide a helpful two-stage inference framework to search for vascular trees on angiograms. An ecological data example of particular importance for animal scientists is discussed in the research work by Descamps et al. (2008). To count individual flamingos in remotely sensed images of their colonies, the authors propose a suitable Gibbs model and a simulated annealing procedure [Brooks and Morgan (1995), Van Laarhoven and Aarts (1987), Salmon et al. (2002)] with fast birth-death dynamics. Another ecological problem concerns forestry and has statistically been approached by Perrin et al. (2005) who introduce a simulated annealing strategy similar to the optimization algorithm by Descamps et al.. This strategy is specifically applicable to the estimation of tree populations from aerial photographs of the tree crowns.

The research works on marked point processes referenced so far are focused on one specific type of object and therefore not very flexibly employable. A more general type of model has been developed by Ge and Collins (2010). Arguing that different human shapes can be seen as geometric transformations of each other, the authors introduce a Bayesian marked point process tool for identifying and counting individuals in noisy scenes. The inference procedure discussed by Lafarge et al. (2010) is even more universal. Lafarge et al. use a sophisticated Gibbs model to match objects from a library of elementary geometric shapes with the given image, and implement a jump diffusion process for optimization. This approach is particularly appealing in that it is capable of detecting structures of interest in the image, delivering a representation of texture by simpler geometric shapes, or estimating crowd densities. Due to these advantages, Section 7.2 is devoted to a more detailed description of the inference framework by Lafarge et al.. We work on a technically similar marked

point approach which is intended to allow for the learning of 3D shape from the local 2D deformations of texture patches on the image plane.

Summing up, we first conclude that, even though locally scaled point processes appear to be highly suitable for modeling real-world phenomena, the range of the existing inference techniques is limited. Second, we realize that most of the available methods for estimating shape from texture are based on strict assumptions on the geometrical form or the arrangement of the texture elements and thus only applicable to a very specific type of texture. Our third and last remark concerns the geometric description of images of textured scenes via marked point processes. Although a considerable amount of marked point processes carefully designed for specific data problems has been made available in the past decades, there is, to our knowledge, no general modeling framework to locally describe geometric texture deformations in a 2D image with depth effect and to simultaneously infer knowledge on geometries in the original 3D scene. These gaps in the literature have motivated our research questions. In the next section, we present the organization of our work, with particular emphasis on our own contributions.

1.3 Contributions and Organization

As stated above, this thesis establishes a connection between spatial point process theory and selected problems related to image analysis. Consequently, methodological and technical tools from both scientific disciplines need to be understood, appropriately implemented, and coupled with each other. To clearly distinguish between the presentation of the relevant existing concepts and the introduction of our newly developed modeling strategies and approaches, we briefly summarize the contents of each chapter in what follows.

The next two chapters impart very basic knowledge on spatial point processes and statistical inference techniques (see Chapter 2), as well as on camera projection and image analysis (see Chapter 3). Chapter 2 starts with a theoretical introduction of *spatial point processes* in Section 2.1, where several fundamental definitions are given and the relevant notational conventions are explained. The most common class of point processes, the class of the *Poisson processes*, is discussed in Section 2.2. Provided that the spatial *Markov property* is fulfilled, patterns of interacting points are usually termed *Markov* or *Gibbs processes* (see Section 2.3). The last sections of Chapter 2 are concerned with the presentation of well-established inference concepts for spatial point processes. Section 2.4 points out the principle of *Markov chain Monte Carlo* (MCMC) simulation before it describes some concrete MCMC methods to sample from spatial point process distributions. Mainly referring to the models introduced in Section 2.2 – 2.3, Section 2.5 gives an overview of widely used frequentist and Bayesian

approaches to parameter estimation. Throughout Chapter 2, we distinguish between *homogeneous* and *inhomogeneous* point processes, involving *marked* point processes. Extensive and detailed in-depth information is provided to ensure the traceability of the strategies and algorithms presented later in Chapter 4 – 7.

Having discussed all statistical basics that are relevant for a better understanding of our work and contributions, Chapter 3 introduces some fundamental concepts related to *camera projections* and the detection of symmetric structures and *near-regular* objects in images. These concepts form an important methodological basis for Chapter 6 – 7. Section 3.1 presents a commonly used camera model based on which the meaning of the terms *assignment field* and *two-view geometry* is explained and graphically demonstrated. Thereby, the interrelation between a *homogeneous* and an *inhomogeneous coordinate representation* is pointed out. Detached from Section 3.1, Section 3.2 deals with the question of how to identify symmetry elements in the image of a textured scene. Based on the necessary mathematical background, it describes a typical sequence of smoothing and distance transformation steps which turn the given image into a *probability map* facilitating the detection of local symmetry centers.

Chapter 4 can be seen as the key chapter of this thesis. In Section 4.1, the concept of inhomogeneous spatial point processes by *location-dependent scaling* is explained, and subsequently, in Section 4.2, the meaning and construction of a *proper scaling function* is discussed. In Section 4.3, we introduce our newly established Bayesian inference framework for the analysis of *locally scaled Strauss processes* [see also Didden et al. (2015)]. The performance of the proposed optimization algorithm is evaluated and discussed in Section 4.4 where we simulate and analyze point patterns under *exponential scaling* assumptions. We compare our results to the results generated by an alternative frequentist optimization algorithm.

Afterwards, we apply our Bayesian inference method to the classification of maize plants by genotype [Didden et al. (2015)], which we perform based on the locations of the vascular bundles in cross-sectional images of the stems (see Chapter 5). Besides the preprocessing carried out by our collaborators, David Legland⁵ and Marie-Françoise Devaux⁶, we first project the point data, i.e. the bundle coordinates, to circular discs of fixed radii (see Section 5.1). Based on the resulting coherent and simplified representation of the spatial bundle distribution, we have developed a new proper scaling function, the *step scaling function* discussed in Section 5.2. The particular properties of this function suggest the implementation of two-stage inference framework as proposed in Section 5.3. Our final results are discussed in Section 5.4 where similarities and differences between the genotypes are worked out, a general evaluation of our

⁵INRA & AgroParisTech, UMR 782 Food Process Engineering and Microbiology, Thiverval-Grignon; INRA & AgroParisTech, UMR 1318 Institut Jean-Pierre Bourgin, Versailles

⁶INRA, UR 1268 Biopolymers, Interactions and Assemblies, Nantes

two-stage procedure is given, and possible further advancements are mentioned.

We provide another newly and specifically designed scaling function in Chapter 6 that deals with the estimation of *shape from texture*. At this point, spatial statistics becomes connected to image analysis in a narrow sense, which we discuss in detail in Section 6.1. Section 6.2 presents a first straightforward and rather intuitive *Gibbs modeling approach* to the estimation of shape from texture under exponential scaling constraints, linked to *spherical coordinates*. To simplify this preliminary concept by immediately employing spherical coordinates as perspective scaling parameters, we have developed an alternative inference procedure [Didden et al. (2013)], where the symmetry centers in the image are estimated and replaced by points in a first step (see Section 6.3), and the resulting pattern is modeled as a realization of a locally scaled Strauss process afterwards (see Section 6.4). In Section 6.5, the suggested inference framework is critically examined and evaluated by means of several case studies. We conclude Chapter 6 with an outlook on potential for future advancement (see Section 6.6).

Chapter 7 is also concerned with the development of a statistical shape-from-texture approach. In contrast to before, however, we here focus on marked point processes. Referring to related work, Section 7.1 surveys how marked point processes have emerged as a versatile tool for the geometric analysis of textured scenes. Since, from a technical and methodological point of view, our research is strongly related to the modeling framework by Lafarge et al. (2010), we shortly summarize the authors' work in Section 7.2. In Section 7.3, we present our own inference strategy which is still in progress, but already comprises a marked point process tool for estimating local 2D geometric deformations in the image of a textured 3D scene. Having analyzed both a manually constructed and a real-world image of a brick wall, the outcomes delivered by this newly proposed tool are summarized and graphically illustrated in Section 7.4. We discuss possible further steps in Section 7.5, particularly focusing on the learning of global 3D geometries from the local 2D information contained in the marked points and their spatial distribution in the image.

Chapter 8 concludes this work with a summary of the proposed methods, our research results, advantages and open problems, and an outlook on possible advancements.

2 Preliminaries on Spatial Point Processes

Spatial point processes provide a statistical framework for modeling random arrangements of objects, which is of relevance in a variety of scientific disciplines including ecology, spatial epidemiology and material science. Describing systematic spatial variations within this framework and developing methods for estimating parameters from empirical data constitute an active area of research. Typically, the locations or centers of objects in a two- or three-dimensional observation space are being modeled, and additional object attributes are induced through marks, if available. Since point process theory and methodology is a very lively field of research with a broad range of possible applications, newly emerging questions become more and more demanding. Depending on the observation space as well as on the types and interrelations of the objects of interest, different inference strategies prove useful. The more indications for spatial inhomogeneity or object interactions a point pattern shows, the more complex suitable point process models are and the more sophistication inference requires. We concentrate on *simple* spatial point processes in \mathbb{R}^2 , where each point is assumed to represent the location of exactly one object. That is, no two points have identical coordinates.

The books by Daley and Vere-Jones (2003, 2007) as well as Part IV of the comprehensive handbook of spatial statistics by Gelfand et al. (2010) provide a broad range of theoretical as well as technical details on point processes. A more data-oriented introduction to the modeling of spatial point patterns can be gathered from Illian et al. (2008). Over the past years, great importance has been attached to the development of efficient optimization techniques constituting a comprehensive inference framework. A compact overview of different classes of spatial point processes and appropriate simulation routines is given in the standard work by Møller and Waagepetersen (2003).

In this chapter, an introduction to spatial point process theory is presented, along with the modeling and inference schemes that are relevant for our research studies. We deal with fundamental definitions related to spatial point processes and their mathematical properties in Section 2.1. Section 2.2 is devoted to the class of Poisson point processes, addressing both modeling approaches to point patterns showing a homogeneous packing density and patterns containing heterogeneously distributed points. Since Poisson models do not capture potential point interactions, Section 2.3 goes over to Markovian point processes also referred to as Gibbs processes. We particularly focus on pairwise interaction models and, amongst these, on Strauss models as well as on Gibbs models with

a fixed number of points. These approaches account for interaction in the form of pairwise point repulsions. The more complex a point process model is the more sophisticated sampling routines are required and the more elaborate tools for parameter estimation have to be considered. In the concluding Sections 2.4 – 2.5, we therefore introduce a selection of appropriate inference schemes.

2.1 Point Process Theory

As mentioned above, spatial point processes are of relevance in many scientific disciplines. In ecology, for instance, researchers might be interested in the spatial distribution of soil bacteria populations [Nunan et al. (2003)], whereas in geographical sciences, the occurrences of sand dunes in a desert or earthquakes in a certain danger area might be among the relevant topics [Eberhard et al. (2012), Møller and Sørensen (1994)]. A medical question addresses the modeling of patterns of nerve fibers in the human skin tissue [Waller (2005)]. In zoology, the predominant locations of the territories of rare animal species might attract the researchers' attention [Edelman (2012), Klaver et al. (2012)], and the spatial modeling of road accidents is a crucial task in traffic engineering [Yamada and Thill (2004)].

In this thesis, we consider two types of datasets: A framework for modeling the spatial distribution of the vascular bundles in maize plants is discussed in Chapter 4. The database consists of several cross-sections through maize stems and, in particular, through the enclosed vascular bundles that we replace by points. Our second database contains photographs of textured three-dimensional scenes (see Chapters 6 – 7). The key idea here is to assign a point to each texture element, and to draw conclusions on camera positioning and angle from the spatial distribution of the resulting point process realization.

The definitions in the following refer to spatial point processes in \mathbb{R}^2 . However, they can easily be adapted to spaces \mathbb{R}^d of arbitrary dimension $d \in \mathbb{N} = \{1, 2, \dots\}$. Following Møller and Waagepetersen (2003), we characterize a spatial point process \mathbf{X} as a random countable subset of a two-dimensional observation window $W \subseteq \mathbb{R}^2$. For the sake of convenience and practical relevance, most of our definitions refer to bounded windows $B \subseteq W$. Moreover, we assume that each realization $\mathbf{x} \subseteq W$ of \mathbf{X} is simple and corresponds to a locally finite point configuration. This means that no two points have exactly the same coordinates, as previously indicated, and that the number of points $N(\mathbf{x}_B)$ with realizations $n(\mathbf{x}_B)$ is finite on each bounded subset $B \subseteq W$. \mathbf{X} thus takes values in the space

$$O_{\mathbf{X}} = \{\mathbf{x} \subseteq W : n(\mathbf{x}_B) < \infty, \forall B \subseteq W\} . \quad (2.1)$$

The subsequent definitions of a point process are aligned to the definitions in Gelfand et al. (2010, ch. 16), Daley and Vere-Jones (2003, 2007, ch. 3, ch. 15).

Definition 2.1 A *point process* on a complete, separable metric space W with Borel σ -algebra \mathcal{B} is a projection \mathbf{X} from a probability space $(\Omega, \mathcal{A}, \mathcal{P})$ to $O_{\mathbf{X}}$. Moreover, the number of points $N(\mathbf{X}_B)$ in any bounded subset $B \subseteq W$ is a finite random variable.

In other words, \mathbf{X} is a random variable taking values denoted by \mathbf{x} on the measurable space $(O_{\mathbf{X}}, \sigma_O)$, where σ_O is the smallest σ -algebra providing for the measurability of the mapping $\mathbf{x} \mapsto n(\mathbf{x}_B)$, for all $B \subseteq W$ and $\mathbf{x}_B \subseteq B$. The induced probability measure \mathcal{P} in Def. 2.1 relates to the distribution of the point process and hence constitutes its statistical properties.

Definition 2.2 A point process \mathbf{X} due to Def. 2.1 is called *strongly stationary* if the underlying probability measure \mathcal{P} is translation invariant, which means that \mathcal{P} is robust to the shifting of \mathbf{X} by any finite vector h . If, in addition, any unitary rotational transition of the point process \mathbf{X} keeps the probability measure \mathcal{P} unaffected, \mathbf{X} is referred to as *strongly isotropic*.

Assuming that we are given a point process with a proper density $f(\cdot)$ subject to the Lebesgue integral measure [Tao (2011, ch. 1)], we can interpret Def. 2.2 following the lecture notes by Schmid and Feilke (2012). The first part of Def. 2.2 implies that, for all $x = (x_1, x_2)^\top \in \mathbf{x}$ and any $h = (h_1, h_2)^\top \in \mathbb{R}^2$,

$$x \mapsto z = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} \implies f(\mathbf{x}) = f(\mathbf{z})$$

holds for any realization \mathbf{x} of a strongly stationary point process \mathbf{X} . Due the second statement of Def. 2.2, \mathbf{X} is strongly isotropic if

$$x \mapsto z = \begin{pmatrix} \cos \rho & -\sin \rho \\ \sin \rho & \cos \rho \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} h_1 \\ h_2 \end{pmatrix} \implies f(\mathbf{x}) = f(\mathbf{z}),$$

for any angle of rotation $\rho \in [0, 2\pi)$ [see also Hughes et al. (2014, ch. 5)]. Though less compactly presentable for $d > 2$, the interpretations of Def. 2.2 similarly apply to any d -dimensional space \mathbb{R}^d , $d \in \mathbb{N}$.

Instead of perfectly determining a point process by means of its distributional properties, it is possible to incompletely characterize it through its intensity. Therefore, we take the bounded subsets $B_u \subseteq W$ and $B_v \subseteq W$ as given. We assume that u and v are locations in B_u and B_v , surrounded by circular discs du and dv with two dimensional (Lebesgue) volume measures $\nu^2(du)$ and $\nu^2(dv)$. Based on this notation, the so-called (*first-order*) *intensity* function evaluated at u equates to

$$\beta(u) = \lim_{\nu^2(du) \rightarrow 0} \frac{E(N(du))}{\nu^2(du)}, \quad (2.2)$$

whereas its *second-order* equivalent applied to the pair $\{u, v\}$ is of the form

$$\beta_2(\{u, v\}) = \lim_{\substack{\nu^2(du) \rightarrow 0 \\ \nu^2(dv) \rightarrow 0}} \frac{E(N(du) N(dv))}{\nu^2(du) \nu^2(dv)}. \quad (2.3)$$

It is easy to show that $\beta(\cdot)$ and $\beta_2(\cdot)$ fully determine the first two moments, i.e. the expectation $E(\cdot)$ and the covariance $C(\cdot, \cdot)$, of the distribution of the random variable $N(\mathbf{X})$:

$$\begin{aligned} E(N(\mathbf{X})) &= \mu(W) = \int_W \beta(u) du \\ E(N(\mathbf{X}_{B_u}) N(\mathbf{X}_{B_v})) &= \int_{B_u} \int_{B_v} \beta_2(\{u, v\}) du dv - \mu(\{B_u \cap B_v\}) \\ C(N(\mathbf{X}_{B_u}), N(\mathbf{X}_{B_v})) &= E(N(\mathbf{X}_{B_u}) N(\mathbf{X}_{B_v})) - \mu(B_u) \mu(B_v) \end{aligned}$$

The probability of observing one point in an infinitesimally small region du surrounding any location $u \in B_u$ can be approximated through

$$\beta(u) \nu^2(du) \approx E(N(du)) \approx P(\text{"1 point in } du\text{"}) .$$

It is obvious that, for receiving a statistically valid probability $P(\dots) \not\approx 1$, this approximation must be based on $du \rightarrow 0$ to ensure that $E(N(du)) \leq 1$. Both intensity functions (cf. Eq. (2.2) and Eq. (2.3)) form the basis of a weakened version of Def. 2.2.

Definition 2.3 A point process \mathbf{X} due to Def. 2.1 is called (*weakly*) *stationary* if its intensity functions are translation invariant. If the intensity functions are furthermore robust to rotations, \mathbf{X} is called (*weakly*) *isotropic*.

The interpretation of Def. 2.3 is similar to that of Def 2.2, as weak stationarity implies that

$$\beta(u) = \beta(u + h) = \beta = \text{const} \quad \text{and} \quad \beta_2(\{u, v\}) = \beta_2(\{u + h, v + h\}) ,$$

for any finite displacement vector h . Moreover, the second-order intensity of a weakly isotropic point process only depends on the Euclidean distances between pairs of points, and not on the exact point coordinates. That is,

$$\beta_2(\{u, v\}) = \beta_2(\|u - v\|) ,$$

where $\|u - v\| = \sqrt{\sum_{i=1}^d (u_i - v_i)^2}$ denotes the Euclidean distance between $u \in \mathbb{R}^d$ and $v \in \mathbb{R}^d$, $d \in \mathbb{N}$. As we concentrate on two-dimensional spaces $W \subseteq \mathbb{R}^2$, $\|u - v\| = \sqrt{(u_1 - v_1)^2 + (u_2 - v_2)^2}$ goes into the second-order intensity term under the assumption of weak isotropy. Strong stationarity and strong isotropy always imply weak stationarity and weak isotropy, but not vice versa. Cressie (1993) surveys several types of point processes clearly distinguishing between their isotropic and their non-isotropic versions.

In this paragraph, we give a short introduction to the theory of *marked point processes* as surveyed in Daley and Vere-Jones (2003, ch. 16). A point process

is termed marked if each point $X \in \mathbf{X}$ comes along with a well-defined random characteristic $m_X \in M$. Formally, we write

$$\mathbf{Y} = \{(X, m_X) : X \in \mathbf{X}, m_X \in M\},$$

where \mathbf{Y} stands for the marked point process on the domain

$$O_{\mathbf{Y}} = \{\mathbf{y} = \{(x, m_x) : x \in \mathbf{x}, m_x \in M\} \subseteq (W \times M) : n(\mathbf{x}_B) < \infty, \forall B \subseteq W\}. \quad (2.4)$$

Apart from the additional consideration of an individual attribute m_x accompanying each $x \in \mathbf{x}$, the domain of a marked point process is of the same structure as the domain of a pure point process (cf. Eq. (2.1)). It basically conforms to the set of all possible realizations of a finite point process \mathbf{X} , though, supplementally to Eq. (2.1), it also accounts for all possible mark attributions. Stating that the point process has to be finite in $W \subseteq \mathbb{R}^2$ is equivalent to postulating that, with probability 1, it contains a finite number of points [Daley and Vere-Jones (2003, ch. 5)].

In general, the statistical properties of a pure point process \mathbf{X} do not change if it is turned into a marked point process \mathbf{Y} . Particularly, stationarity and isotropy characteristics due to Def. 2.2 and Def. 2.3 remain unaffected. Additionally to describing a marked point process \mathbf{Y} in terms of the underlying ground process \mathbf{X} , important attributes of the marking process should not be left unnoted. Following Daley and Vere-Jones (2003, ch. 16), we outline two statements concerning dependence-structures between marks and (marked) points.

Definition 2.4 A marked point process $\mathbf{Y} = \{(X, m_X) : X \in \mathbf{X}, m_X \in M\}$ in the product space $A := (W \times M)$, $W \subseteq \mathbb{R}^2$, is said to have *independent marks* if, conditional on the basic process X , the marks are mutually independent random variables. That is, for any realization $\mathbf{x} = \{x_1, \dots, x_n\}$ with mark assignments $\{m_{x_1}, \dots, m_{x_n}\}$, each mark m_{x_i} only depends on x_i , for all $i \in \{1, \dots, n\}$.

Definition 2.5 A marked point process $\mathbf{Y} = \{(X, m_X) : X \in \mathbf{X}, m_X \in M\}$ in the product space $A := (\mathbb{R}^2 \times M)$ is said to have *unpredictable marks* if the mark distribution at any point contained in \mathbf{X} does not depend on the other point locations and their markings. That is, for every realization \mathbf{y} of \mathbf{Y} , the mark distribution in any component x_i of the corresponding ground process \mathbf{x} does not depend on the remaining marked points $\mathbf{y}_{-\{i\}} = \{\mathbf{y} \setminus \{y_i\}\}$.

The simplest model for marks is the *independent marks model* where the marks are *independent and identically distributed* (i.i.d.) random variables not depending on the locations of the points. If the marks do not depend on the point coordinates, but correlate with each other, a *random field model* is typically considered. Random field models are also known as *geostatistical marking*.

More general models take correlations between the spatial point distribution and the marking into account. In this work, however, we only consider cases where the marks are assumed to be independent and unpredictable.

It is obvious that marked point processes are of high practical relevance. Assuming, for instance, that \mathbf{X} describes the locations of the trees in a forest, \mathbf{Y} might contain information on the stem diameters. If \mathbf{Y} comprises further details, such as the type, age or height of the trees, we are concerned with a *multi-marked point process*. Fig. 2.1 illustrates a simple example.




Figure 2.1: Simulation of a forest area. (a) corresponds to a realization of a non-marked point process showing the locations of the trees, and (b) additionally depicts the relative stem diameters. (c) provides information on the locations, stem diameters, and types of the trees.

Summing up in terms of notation, \mathbf{X} stands for a non-marked point process and \mathbf{x} for one of its realizations, and X denotes a single point and x one of its realizations. In the marked case, \mathbf{Y} symbolizes the random process, \mathbf{y} the realization, and Y as well as y a single point (realization). $N(\cdot)$ stands for the number of points contained in \mathbf{X} or \mathbf{Y} , whereas $n(\cdot)$ counts the actual number of points in \mathbf{x} or \mathbf{y} . While $W \subseteq \mathbb{R}^2$ denotes an arbitrary observation window, B is always bounded.

We differentiate between three types of point patterns with regard to their structural appearances: In an entirely *random* pattern, points can be located arbitrarily close to each other. The coordinates of one point are thus independent of the coordinates of all the other points. A point configuration is *regular* if the average distance from a point to its nearest neighbor is higher than under the assumption of total randomness. Otherwise, if this distance is shorter than for independently located points, the observed pattern is termed *aggregated*. Fig. 2.2 shows a random (a), a regular (b) and an aggregated (c) point arrangement.

It is possible to check a homogeneous or inhomogeneous point process for clustering or repulsion effects by means of suitably implemented test statistics [see e.g. Ripley (1977), Besag and Diggle (1977), Besag and Clifford (1989)].




Figure 2.2: Samples from point processes featuring different structural properties. In each window, the expected number of points is equal to 50.

Common test statistics are either based on inter-point distances, such as the F-, G- or J-functions discussed by Van Lieshout and Baddeley (1996) and Bedford and Van den Berg (1997), or on second-order measures, such as the class of the K-, L-, and g-functions [see e.g. Ripley (1976), Baddeley et al. (2000)]. Furthermore, directional K-functions have been developed to investigate possible anisotropy [Stoyan and Stoyan (2000), Brix and Møller (2001)].

In the following sections, we discuss some common ways to determine and characterize the distribution of a point process. These distributional principles are needed for clarifying the technical details of our research work in the later chapters. While our main focus will be on Markovian point processes (see Section 2.3), we start with a general description of their superordinate class, the class of Poisson point processes.

2.2 Poisson Processes

A spatial Poisson point process \mathbf{X} on a subset $W \subset \mathbb{R}^2$ can be described as a random arrangement of points, in which the location of each single point does not depend on the coordinates of the other points (cf. Fig. 2.2 (a)). The definition of a Poisson point process and its characteristics can thus be derived from the properties of a binomial point process [see also Gelfand et al. (2010, ch. 17), Illian et al. (2008, ch. 2), and Møller and Waagepetersen (2003, ch. 3)].

Definition 2.6 A *binomial point process* $\mathbf{X} \sim \text{binomial}(W, n, f)$ in a set $W \subset \mathbb{R}^2$ is a random arrangement of $n \in \mathbb{N} = \{1, 2, \dots\}$ i.i.d. points, where the point intensity depends on the underlying density function f on W .

Definition 2.7 A binomial point process $\mathbf{X} \sim \text{binomial}(W, n = 1, f)$ is termed *Bernoulli process* in W subject to f .

Definition 2.8 A *Poisson point process* $\mathbf{X} \sim \text{Poisson}(W, \beta)$ in a set $W \subset \mathbb{R}^2$ with underlying intensity β has the following properties:

- (a) For any $B \subseteq W$, the number of points in B follows a Poisson distribution with mean $\mu(B)$, i.e. $N(\mathbf{X}_B) \sim \text{Pois}(\mu(B))$, under the constraint that $\mu(B) < \infty$. Assuming that B is bounded, $\mu(B) = E(N(\mathbf{X}_B))$.
- (b) For any $n \in \mathbb{N}$ and $B \subseteq W$, given that $N(\mathbf{X}_B) = n$ and $0 < \mu(B) < \infty$, $\mathbf{X}_B \sim \text{binomial}(B, n, f)$ with density $f(x) = \frac{\beta(x)}{\mu(B)}$, $\forall x \in \mathbf{X}_B$.
- (c) For any I disjoint subsets $B_i \subset W$, $N(\mathbf{X}_{B_1}), \dots, N(\mathbf{X}_{B_I})$ are independent random variables.

Definition 2.9 The Poisson process with constant intensity $\beta = 1$, is called the *standard* or *unit rate Poisson process*.

We first assume that the points are uniformly distributed in W . The respective point process is then called a *homogeneous binomial* and a *homogeneous Poisson point process*, respectively. It is stationary and isotropic due to Def. 2.2. Under homogeneity conditions, Def. 2.6 simplifies in the sense that $f(x) = \frac{1}{|B|}$, $x \in \mathbf{x}$, is the uniform density in $B \subseteq W$. Moreover, β becomes a constant factor and the relation $\mu(B) = \beta |B|$ goes into Def. 2.8. It follows from Def. 2.8 (a) that the number of points, $N(\mathbf{X}_B)$, is Poisson distributed,

$$P(N(\mathbf{X}_B) = n) = \frac{(\beta |B|)^n}{n!} \exp\{-\beta |B|\},$$

with $P(N(\mathbf{X}_B) = 0) = \exp\{-\beta |B|\}$. The term *void probability* is widely-used for $P(\cdot = 0)$ in the statistical literature.

Under homogeneity assumptions, the relation between a binomial and a Poisson process, as described in Def. 2.8 (a), results in $f(x) = \frac{\beta}{\mu(B)} = \frac{1}{|B|}$, $x \in \mathbf{x}$. Similarly, the intensity parameter of the Poisson process can be deduced from the expected number of points of the binomial process by computing $\beta = \frac{E(N(\mathbf{X}_B))}{|B|}$, provided that B is bounded implying that $E(N(\mathbf{X}_B)) = \mu(B)$.

If the packing density of the point pattern is not constant, but driven by an intensity function $\beta : W \rightarrow [0, \infty)$, the process corresponds to an *inhomogeneous Poisson process*. For modeling purposes, $\beta(\cdot)$ needs to be locally integrable, meaning that $\int_B \beta(x) dx < \infty$ for all bounded subsets $B \subseteq W$. The equation $\mu(B) = \int_B \beta(x) dx$ has to be solved to quantify the intensity measure $\mu(B)$ in Def. 2.8. Since $\mu(\{x\}) = 0$ for all $x \in \mathbf{X}_B$, we call $\mu(\cdot)$ diffuse. Following Def. 2.8 (a), the number of points in an inhomogeneous pattern is Poisson distributed with density

$$P(N(\mathbf{X}_B) = n) = \frac{(\int_B \beta(x) dx)^n}{n!} \exp\left\{-\int_B \beta(x) dx\right\}.$$

This thesis is focused on parametric intensity functions, meaning that $\beta(\cdot)$ is determined by a set of parameters $\boldsymbol{\theta}$. We thus henceforth use the notation $\beta_{\boldsymbol{\theta}}(\cdot)$. To assess $\boldsymbol{\theta}$ with respect to a given dataset $\boldsymbol{x} = \{x_1, \dots, x_n\}$ on $W \subset \mathbb{R}^2$, we apply the *likelihood*

$$f(\boldsymbol{x}|\boldsymbol{\theta}) = \exp \left\{ - \int_W \beta_{\boldsymbol{\theta}}(x) dx \right\} \prod_{x \in \boldsymbol{x}} \beta_{\boldsymbol{\theta}}(x) . \quad (2.5)$$

Likelihood-based parameter estimation is discussed in more detail in Section 2.5. Formally, the likelihood is equivalent to the discrete or continuous density of the data. By using the term likelihood, however, the fact is stressed that the exact values of $\boldsymbol{\theta}$ are not known and modeled subject to the observed data \boldsymbol{x} , although $f(\boldsymbol{x}|\boldsymbol{\theta})$ itself suggests conditionality on $\boldsymbol{\theta}$. In Daley and Vere-Jones (2003, ch. 7), the principle and motivation of the likelihood representation of a point process with proper density $f(\cdot)$ is step-wisely developed and explained based on all relevant theoretical fundamentals.

In case of analyzing a homogeneous point pattern, Eq. (2.5) reduces to

$$f(\boldsymbol{x}|\boldsymbol{\theta}) = \exp \{ -\beta |W| \} \beta^n , \quad (2.6)$$

which facilitates inference considerably.

As stated in the previous section, a point dataset does not necessarily contain spatial point coordinates only. Frequently, additional parameters in the form of marks provide more detailed information on the objects the points are representing. We define the marked version $\boldsymbol{Y} = \{(X, m_X) : X \in \boldsymbol{X}, m_X \in M\}$ of a Poisson process on the state space $(W \times M)$ according to Møller and Waagepetersen (2003, ch. 3).

Definition 2.10 Let $\boldsymbol{X} \sim \text{Poisson}(W, \beta)$ be a Poisson process with locally integrable intensity function β , and let M be a given mark space. If, conditional on \boldsymbol{X} , the marks $\{m_X : X \in \boldsymbol{X}\}$ are mutually independent (cf. Def. 2.4), \boldsymbol{Y} is called a *marked Poisson process*. Furthermore, in the case that all marks have the same distributional properties, their distribution is called *mark distribution*. Assuming that a proper density p_M is derivable from the distribution of the marks, i.e. $\sum_{m \in M} p_M(m) = 1$ or $\int_M p_M(m) dm = 1$, $p_M(\cdot)$ is termed the discrete or continuous *mark density* associated with the Poisson process \boldsymbol{X} .

Following Daley and Vere-Jones (2003, ch. 7), an independent marks model (cf. Def. 2.4 and discussion) of Poisson type is defined as

$$f(\boldsymbol{y}|\boldsymbol{\theta}) = \exp \left\{ - \int_W \beta_{\boldsymbol{\theta}}(x) dx \right\} \prod_{(x, m_x) \in \boldsymbol{y}} \left[\beta_{\boldsymbol{\theta}}(x) p_M(m_x) \right], \quad (2.7)$$

provided that $p_M(\cdot)$ is proper according to Def. 2.10.

Poisson models per se do not account for potential point dependencies, which means that the conditional distribution of one point $X_i \in \mathbf{X}$ given the remaining pattern $\mathbf{X}_{-\{i\}}$ does not depend on the point locations in the remaining pattern. For a Poisson model with a parametric density following Eq. (2.5), it therefore holds that

$$f(X_i | \mathbf{X}_{-\{i\}}, \boldsymbol{\theta}) = f(X_i | \boldsymbol{\theta}) .$$

In many practical applications, however, the assumption of independent point coordinates proves too simplistic and unrealistic. Nevertheless, by using Poisson models as templates, while imposing additional properly formulated interaction conditions, we derive models that allow to deal with point interactions. The following section gives an overview of the well-established Markovian interaction point processes.

2.3 Markov/Gibbs Processes

The most common and statistically relevant overall class of interaction point processes is the class of *Markov point processes* also known as *Gibbs processes* [Møller and Waagepetersen (2003, ch. 6)]. We call a point process a Markov process if its distribution is affected by the *Markov property*. This means that, considering a point process in $W \subseteq \mathbb{R}^2$, the conditional distribution of any single point given all the other point locations only depends on the locations in the immediate surrounding of the point. In other words, the probability of assigning the position x_i to X_i , while knowing the remaining coordinates $\mathbf{x}_{-\{i\}}$ of $\mathbf{X}_{-\{i\}}$, depends on the point positions in the close proximity of x_i only. The circumference of the “close proximity” needs to be ascertained based on a thoroughly chosen symmetric and reflexive neighborhood relation that we symbolize by \sim_i . Then, the spatial Markov property formally implies that

$$P(X_i = x_i | X_1 = x_1, \dots, X_{i-1} = x_{i-1}, X_{i+1} = x_{i+1}, \dots, X_n = x_n) = P(X_i = x_i | \mathbf{X}_{\sim_i} = \mathbf{x}_{\sim_i}). \quad (2.8)$$

If the point process considered has a proper density f with respect to the Lebesgue integral measure [Tao (2011, ch. 1)], it furthermore holds that

$$f(X_i | \mathbf{X}_{-\{i\}}) = f(X_i | \mathbf{X}_{\sim_i}) .$$

The Markov property can similarly be interpreted with regard to point processes in any d -dimensional subspace of \mathbb{R}^d , where $d \in \mathbb{N}$. Temporal point processes in \mathbb{R} , for instance, are Markovian if the conditional distribution of the occurrence time of the next event only depends on the occurrence time(s) of the immediately preceding event(s) (see also Section 2.4). More precise and technically sophisticated explanations of the meaning of the Markov property

in the context of point process modeling are provided by Kendall (1990) and Daley and Vere-Jones (2007, ch. 10).

Coming back to the modeling of spatial patterns in \mathbb{R}^2 , we limit our discussion to the most typical class of Markov point process models composed of a point process density with respect to the unit rate Poisson process and additional constraints involving interaction. We confine ourselves to *isotropic pairwise interaction point processes* the densities of which depend on Euclidean point distances, but not on exact point coordinates (see also Def. 2.3 and its interpretation). Among that subclass, in turn, our main focus is on *Strauss processes* as well as on Gibbs processes with a known number n of points. For the sake of simplicity, we call the latter *fix- n Gibbs processes*. Both Strauss and fix- n Gibbs models capture interactions in the form of pairwise point repulsions. A comprehensive methodological background as well as supplementary technical details can be gathered from Møller and Waagepetersen (2003, ch. 6).

In general terms, the likelihood of a pairwise interaction point process in \mathbb{R}^2 is specified by its intensity and interaction functions, $\phi_1(\cdot)$ and $\phi_2(\cdot)$, and takes the form

$$f(\mathbf{x}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta})^{-1} \prod_{u \in \mathbf{x}} \phi_1(u) \prod_{\substack{\neq \\ \{u,v\} \subseteq \mathbf{x}}} \phi_2(\{u,v\}) , \quad (2.9)$$

where $\boldsymbol{\theta}$ denotes the set of unknown parameters and $\{u,v\}$ any pair of dissimilar points [Møller and Waagepetersen (2003)]. To ensure that the likelihood corresponds to a proper density, i.e. $\int_{\mathbb{R}^2} f(\mathbf{x}|\boldsymbol{\theta}) d\mathbf{x} = 1$, the normalizing constant $\mathcal{Z}(\cdot)^{-1}$ needs to be defined as

$$\mathcal{Z}(\boldsymbol{\theta}) = \int_{\mathbb{R}^2} f^{(*)}(\mathbf{x}|\boldsymbol{\theta}) d\mathbf{x} , \quad \text{with} \quad f^{(*)}(\mathbf{x}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta}) f(\mathbf{x}|\boldsymbol{\theta}) . \quad (2.10)$$

Since this integral is often not solvable, many point processes suffer from a likelihood not available in complete form.

A proper interaction function is nonnegative and furthermore guarantees that $f(\cdot|\boldsymbol{\theta})$ is integrable with respect to the unit rate Poisson process (cf. Def. 2.9). If and only if $\phi_2(\cdot) \leq 1$ holds for all pairs of points, a repulsive point configuration is modeled (cf. Fig. 2.2 (b)). The point process is furthermore locally stable if $\int_W \phi_1(u) du < \infty$. A Poisson process can be regarded as the limiting case of a repulsive point process, where $\phi_2(\{u,v\}) = 1$ for all pairs of points, and where the Likelihood is available in complete form (cf. Eq. (2.5)). An incomplete likelihood makes inference difficult and requires the implementation of approximate methods. This is why we devote Section 2.4 and Section 2.5 to the introduction of well-established approaches to simulation and parameter estimation in numerically challenging point process frameworks.

Assuming that $\phi_2(\{u,v\}) > 1$, we are concerned with an attractive density measure modeling aggregated patterns (cf. Fig. 2.2 (c)). Most attractive models

are poorly defined and suffer from a lack of accessibility [Kelly and Ripley (1976)]. Therefore as well as due to the fact that our own research deals with repulsive patterns only, we will not go into any more detail at this point.

Given a homogeneous point pattern with pairwise repulsion effects, Eq. (2.9) simplifies to

$$f(\mathbf{x}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta})^{-1} \beta^{n(\mathbf{x})} \prod_{\{u,v\} \subseteq \mathbf{x}} \phi_2(\{u,v\}) ,$$

that is, $\phi_1(u)$ turns into a constant factor $\phi_1 = \beta > 0$, for all $u \in \mathbf{x} = \{x_1, \dots, x_n\}$. Based on the assumption of homogeneity, we characterize the well-known class of *homogeneous Strauss processes* first introduced by Strauss (1975). In a Strauss process model, interactions between pairs of points are included through their Euclidean distances, which means that $\phi_2(\{u,v\}) = \phi_2(\|u - v\|)$. More concretely, the term of interaction takes the form

$$\phi_2(r) = \gamma^{\mathbb{1}[r \leq R]} , \quad (2.11)$$

where $0 \leq \gamma \leq 1$ and $R \geq 0$. We call γ the weight of interaction and R the interaction radius or range. Furthermore, we assume that $0^0 = 1$. In terms of inference, Strauss models come along with a likelihood according to

$$f(\mathbf{x}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta})^{-1} \beta^{n(\mathbf{x})} \gamma^{s_R(\mathbf{x})} , \quad (2.12)$$

where $s_R(\mathbf{x}) = \sum_{\substack{\neq \\ \{u,v\} \subseteq \mathbf{x}}} \mathbb{1}[\|u - v\| \leq R] ,$

and $\mathcal{Z}(\cdot)^{-1}$ is usually intractable. In this formal representation, $\boldsymbol{\theta} = \{\beta, \gamma, R\}$. Turning our attention to the borders of the domain of γ , both limiting cases have a noteworthy meaning: $\gamma = 1$ corresponds to a homogeneous Poisson model (cf. Eq. (2.2)), whereas a Strauss model with $\gamma = 0$ is called a *hardcore model*. The designation “hardcore” is self-explanatory, since the expression on the right-hand side of Eq. (2.12) is equal to 0 if $\gamma = 0$ and at least one pairwise point distance is shorter than R .

We see that γ and R describe the regularity of a point pattern. The question comes up why we abstract away from higher interaction weights, i.e. $\gamma > 1$, although they are expected to capture potential clustering effects. The problem consists, in line with our discussion above, in the non-integrability of the right-hand side of Eq. (2.12). We therefore make reference to Kelly and Ripley (1976) once again.

We move on to Gibbs processes containing a known number of points. As clarified previously, we call these processes fix- n Gibbs processes for convenience. Since n is deterministic, the first product on the right-hand side of the generic likelihood in Eq. (2.9) results in a constant factor commonly equated with an




Figure 2.3: Simulations in $W = [0, 1] \times [0, 1]$ from different Strauss type models with intensity $\beta = 100$ and interaction range $R = 0.05$. The point process realizations in (b) – (c) have been generated via perfect simulation, through progressive thinning of the template Poisson process in (a) (cf. Section 2.4).

expression of the form $\exp\{-\alpha\}$. The interaction condition is set to

$$\prod_{\{u,v\} \in \mathbf{x}}^{\neq} \phi_2(\{u, v\}) = \prod_{\{u,v\} \in \mathbf{x}}^{\neq} \exp\{-\varphi(\|u - v\|)\} ,$$

where $\varphi(\cdot)$ is a suitably designed function evaluating interaction. On the whole, the likelihood of a fix-n Gibbs process is presentable as

$$f(\mathbf{x}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta})^{-1} \exp\left\{-\alpha - \sum_{\{u,v\} \in \mathbf{x}}^{\neq} \varphi(\|u - v\|)\right\} = \mathcal{Z}(\boldsymbol{\theta})^{-1} \exp\{-U(\mathbf{x}|\boldsymbol{\theta})\} . \quad (2.13)$$

$U(\cdot)$ is termed the *total Gibbs energy*, and $\mathcal{Z}(\cdot)^{-1}$ follows Eq. (2.10). If, for all pairs $\{u, v\}$, $\varphi(\|u - v\|) = 0$, the fix-n Gibbs process reduces to a simple binomial process (cf. Def. 2.6), whereas $\varphi(\|u - v\|) = \gamma^{1[\|u - v\| > R]}$ complies with a Strauss process density due to Eq. (2.12).

In conformity with the previous paragraphs, we supplement our description of interaction point processes by addressing their marked versions. The likelihood of a *marked pairwise interaction point process* $\mathbf{Y} = \{(X, m_X) : X \in \mathbf{X}, m_X \in M\}$ is given by

$$f(\mathbf{y}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta})^{-1} \prod_{(u, m_u) \in \mathbf{y}} \phi_1((u, m_u)) \prod_{\{(u, m_u), (v, m_v)\} \subseteq \mathbf{y}}^{\neq} \phi_2(\{(u, m_u), (v, m_v)\}) , \quad (2.14)$$

for any realization \mathbf{y} of \mathbf{Y} . Again, $\phi_1(\cdot)$ and $\phi_2(\cdot)$ have to be positive and integrable. All unknown parameters describing the point locations and the marks go into $\boldsymbol{\theta}$ in this formulation. A common type of marked interaction point process is the so-called *Strauss type disc process*, in which the marks

correspond to disc radii. These disc radii might, for example, quantify the sizes of the root networks in a forest, whereas the points themselves indicate the actual tree locations (cf. Fig. 2.1 (a) – (c)). The trees cannot stand arbitrarily close to each other. In a Strauss type disc process,

$$\phi_1((u, m_u)) = \beta, \quad \text{and} \quad \phi_2(\{(u, m_u), (v, m_v)\}) = \gamma^{1[\|u-v\| \leq m_u + m_v]}.$$

As for the non-marked case, we require that $\beta > 0$, $0 \leq \gamma \leq 1$, and $0^0 = 1$.

So far, we have only discussed homogeneous Markov processes. It is self-evident, however, that there are also inhomogeneous versions of general pairwise interaction point processes. Especially, the consideration of local scaling effects has intensely been studied over the past years. While object configurations showing spatially varying packing densities and interaction behaviors appear in most fields of statistical application, analyzing them is quite challenging. Apart from the fact that the normalizing constant of a Gibbs model is usually intractable rendering exact inference impossible (cf. Eq. (2.10) and discussion), additional scaling constraints complicate matters, mainly for identifiability reasons. We therefore devote the separate Chapter 4 to the introduction of locally scaled spatial point processes and, in particular, to the proposition of a suitable and efficient Bayesian inference framework.

As just stated, most Gibbs models suffer from a lack of completeness in the sense that their normalizing constants are not available in closed form. For the same reason, it is often not easy to model inhomogeneous Poisson processes. Only homogeneous binomial and homogeneous Poisson point process distributions ensure the applicability of exact inference procedures without fail. In Section 2.4, we give an introduction to the Markov chain Monte Carlo simulation framework which has become a crucial foundation for solving complicated modeling problems. Based on this background, we discuss convenient sampling principles related to the point process models described so far. Section 2.5 finally surveys how to perform parameter estimation under these modeling conditions, addressing both frequentist and Bayesian concepts of inference.

2.4 Sampling Techniques

Our discussion on suitable sampling techniques for different types of spatial point processes is mainly based on the books by Gelfand et al. (2010, ch. 12 & 19) and Møller and Waagepetersen (2003, ch. 7 & 11). In the following, we begin with a general description of what is meant by *Markov chain Monte Carlo* simulation.

Markov Chain Monte Carlo

Markov chain Monte Carlo (MCMC) algorithms are iterative stochastic simulation techniques that yield an approximate solution when exact or straightforward solvers are not available. Building on the early seminal works by Metropolis et al. (1953) and Hastings (1970), a huge range of MCMC(-related) methods has been developed during the past decades. Particularly in computer-based sciences as well as in Bayesian statistics, MCMC approaches are among the gold standards.

One fundamental MCMC principle traces back to the substitutability of unsolvable integrals by empirical means, that is, by *Monte Carlo integration* [see e.g. Gelfand et al. (2010, ch. 12)]. We assume that we are given a random variable \mathbf{X} with a density $f(\cdot)$ from which we can simulate realizations denoted by \mathbf{x} . For any suitable function $g(\cdot)$ with unknown mean $E(\cdot)$, it holds that

$$E(g(\mathbf{x})) = \int_W g(\mathbf{x}) f(\mathbf{x}) d\mathbf{x} \approx \frac{1}{T} \sum_{t=1}^T g(\mathbf{x}^{(t)}) ,$$

provided that enough drawings $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(T)}\}$ from $f(\cdot)$ are available and W is the known domain of \mathbf{X} . In case of independent sampling, the central limit theorem [Cam (1986)] yielding

$$\frac{1}{\sqrt{T}} \sum_{t=1}^T g(\mathbf{x}^{(t)}) \stackrel{T \rightarrow \infty}{\rightsquigarrow} N\left(\sqrt{T} E(g(\mathbf{x})), V(g(\mathbf{x}))\right) ,$$

where $V(\cdot)$ can be replaced by the empirical variance, furthermore holds under quite general assumptions.

The second motivation for the development of MCMC methodology results from the problem that exact sampling from an incomplete density $f(\cdot)$ is usually not possible. *Monte Carlo sampling*, however, allows to implicitly draw realizations from $f(\cdot)$. As discussed in Gelfand et al. (2010, ch. 12), Monte Carlo sampling builds on the key idea to develop a *Markov chain*. This chain is constructed such that its stationary distribution accords to the target model, and all its elements take values in the target domain. It is important that the Markov chain is sufficiently long. As soon as it is in its equilibrium regime, all previous stages, i.e. the entire *burn-in phase*, become eliminated, and samples from the remaining chain are regarded as drawings from the intractable target model $f(\cdot)$. In order to avoid *autocorrelation*, not all states of the remaining chain should be taken into account. Autocorrelation effects are dependence structures between successively generated samples [see also Gelfand et al. (2010, ch. 1)].

The chain of states obtained from an MCMC simulation conforms to a succession of outcomes with the Markov property as its characteristic feature. In contrast to the spatial Markov property discussed in Section 2.3 and formally

defined by means of Eq. (2.8), we are now concerned with its discrete temporal version. This means that element $(t + 1)$ of a Markovian sequence so far consisting of the states $\mathbf{X}^{(0)} = \mathbf{x}^{(0)} \rightarrow \mathbf{X}^{(1)} = \mathbf{x}^{(1)} \rightarrow \dots \rightarrow \mathbf{X}^{(t)} = \mathbf{x}^{(t)}$, where $\mathbf{X}^{(0)} = \mathbf{x}^{(0)}$ is a suitably chosen baseline, only depends on the last realization $\mathbf{X}^{(t)} = \mathbf{x}^{(t)}$. Technically,

$$\begin{aligned} P(\mathbf{X}^{(t+1)} = \mathbf{x}^{(t+1)} | \mathbf{X}^{(t)} = \mathbf{x}^{(t)}, \dots, \mathbf{X}^{(0)} = \mathbf{x}^{(0)}) &= P(\mathbf{X}^{(t+1)} = \mathbf{x}^{(t+1)} | \mathbf{X}^{(t)} = \mathbf{x}^{(t)}) \\ &\text{or} \\ f(\mathbf{x}^{(t+1)} | \mathbf{x}^{(t)}, \dots, \mathbf{x}^{(0)}) &= f(\mathbf{x}^{(t+1)} | \mathbf{x}^{(t)}) . \end{aligned} \quad (2.15)$$

The whole Markov process is therefore uniquely determined by its transition distributions.

We let $f(\cdot | \cdot)$ denote a discrete or continuous density which models transitions between any two consecutive states \mathbf{x} and \mathbf{x}' . Under weak regularity assumptions, the whole simulation process converges towards its stationary and unique target distribution $f(\cdot)$. Therefore, it is important and essential that the *detailed balance condition*

$$f(\mathbf{x}) f(\mathbf{x}' | \mathbf{x}) = f(\mathbf{x}') f(\mathbf{x} | \mathbf{x}') \quad (2.16)$$

is fulfilled [see also Møller and Waagepetersen (2003, ch. 7 & 11)]. For more information, we recommend the standard work by Meyn and Tweedie (2009) which provides a broad and sound theoretical knowledge on the properties and the construction of Markovian chains.

The *Gibbs sampler* is one of the two very elementary MCMC tools. Given a random variable \mathbf{X} consisting of n components, i.e. $\mathbf{X} = \{X_1, \dots, X_n\}$, it iteratively generates samples $\mathbf{x} = \{x_1, \dots, x_n\}$ from the *full conditionals* $f(x_i | \mathbf{x}_{-\{i\}})$, where $\mathbf{x}_{-\{i\}} = \{x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n\}$. Consequently, Gibbs sampling avoids the necessity to immediately assess the incomplete density $f(\mathbf{x})$, provided that all full conditionals are numerically accessible. In Alg. 2.1, the stages of a Gibbs simulation are formally outlined for a set of n unknown variables and a sufficiently large number T of iterations. A more detailed description of the Gibbs sampler can be gathered from Casella and George (1992).

The question comes up how to proceed if it is not possible to immediately draw realizations from the full conditionals. In such situations, algorithms of the type *Metropolis-Hastings* (M-H), tracing back to the works by Metropolis et al. (1953) and Hastings (1970), often prove expedient. M-H sampling is not the most efficient simulation strategy and there is a lot of research going on, trying to develop more innovative methods. However, M-H and Gibbs algorithms are probably the most wide-spread and well-known techniques. To implement a standard M-H routine, appropriate transition density (or probability) distributions $q(\mathbf{x}^{(t+1)} | \mathbf{x}^{(t)})$ need to be specified in advance. It must be possible to draw exact samples on the target domain from these distributions. Although the

Algorithm 2.1: Gibbs sampler for a fixed number of variables.

Data: Number of variables n , burn-in threshold T_0
Result: Samples $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$

Initialize $\mathbf{x}^{(0)} = \{x_1^{(0)}, \dots, x_n^{(0)}\}$;

for $t \in \{0, \dots, T - 1\}$ **do**

 for $i \in \{1, \dots, n\}$ **do**

 draw $x_i^{(t+1)} \sim f(x_i | \mathbf{x}_{-\{i\}}^{(t)})$;

 end

 set $\mathbf{x}^{(t+1)} = \{x_1^{(t+1)}, \dots, x_n^{(t+1)}\}$;

end

Return $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$;

true transition densities $f(\cdot|\cdot)$ appear numerically intractable or not available in complete form, they are related to the user-defined functions $q(\cdot|\cdot)$ through

$$f(\mathbf{x}^{(t+1)}|\mathbf{x}^{(t)}) = q(\mathbf{x}^{(t+1)}|\mathbf{x}^{(t)}) P_{\text{acc}}(\mathbf{x}^{(t+1)}|\mathbf{x}^{(t)}),$$

where $P_{\text{acc}}(\cdot|\cdot)$ stands for the probability to accept the transition proposed based on $q(\cdot|\cdot)$. The detailed balance condition in Eq. (2.16) can thus be restated as follows:

$$\begin{aligned} f(\mathbf{x}) q(\mathbf{x}'|\mathbf{x}) P_{\text{acc}}(\mathbf{x}'|\mathbf{x}) &= f(\mathbf{x}') q(\mathbf{x}|\mathbf{x}') P_{\text{acc}}(\mathbf{x}|\mathbf{x}') \\ \implies \frac{P_{\text{acc}}(\mathbf{x}'|\mathbf{x})}{P_{\text{acc}}(\mathbf{x}|\mathbf{x}')} &= \frac{f(\mathbf{x}')}{f(\mathbf{x})} \frac{q(\mathbf{x}|\mathbf{x}')}{q(\mathbf{x}'|\mathbf{x})} \end{aligned}$$

This representation provides for the conclusion that

$$P_{\text{acc}}(\mathbf{x}'|\mathbf{x}) = \min \left\{ 1, R_{\text{MH}} := \frac{f(\mathbf{x}')}{f(\mathbf{x})} \frac{q(\mathbf{x}|\mathbf{x}')}{q(\mathbf{x}'|\mathbf{x})} \right\}. \quad (2.17)$$

The calculation and evaluation of R_{MH} , known as *Hastings ratio*, is the central part of the M-H algorithm step-wisely sketched in Alg. 2.2. In spite of the intractability of $f(\cdot)$, the calculation of R_{MH} does not pose any problems, as it makes the unknown normalizing constants cancel out. More precisely,

$$R_{\text{MH}} = \frac{f(\mathbf{x}')}{f(\mathbf{x})} \frac{q(\mathbf{x}|\mathbf{x}')}{q(\mathbf{x}'|\mathbf{x})} = \frac{\mathcal{Z}^{-1} f^{(*)}(\mathbf{x}')}{\mathcal{Z}^{-1} f^{(*)}(\mathbf{x})} \frac{q(\mathbf{x}|\mathbf{x}')}{q(\mathbf{x}'|\mathbf{x})} = \frac{f^{(*)}(\mathbf{x}')}{f^{(*)}(\mathbf{x})} \frac{q(\mathbf{x}|\mathbf{x}')}{q(\mathbf{x}'|\mathbf{x})},$$

where $f^{(*)}(\cdot)$ is the known kernel of $f(\cdot)$ and \mathcal{Z}^{-1} the inaccessible normalizer (see also Eq. 2.10). In general, the suggested transition distribution $q(\cdot|\cdot)$ should not be too restrictive, that is, too much input in terms of vague speculations

Algorithm 2.2: Metropolis-Hastings sampler.

Data: Burn-in threshold T_0 , selected method

Result: Samples $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$

Initialize $\mathbf{x}^{(0)} = \{x_1^{(0)}, \dots, x_n^{(0)}\}$;

for $t \in \{0, \dots, T-1\}$ **do**

if $method == \text{"individual"}$ **then**

for $i \in \{1, \dots, n\}$ **do**

 draw $x'_i \sim q_i(x_i | \mathbf{x}^{(t)})$ and set $\mathbf{x}' = \{\mathbf{x}^{(t)}_{-\{i\}} \cup \{x'_i\}\}$;

 calculate $R_{MH} = \frac{f(\mathbf{x}') q_i(x_i^{(t)} | \mathbf{x}')}{f(\mathbf{x}) q_i(x'_i | \mathbf{x}^{(t)})}$;

 draw $R_t \sim U[0, 1]$;

 set $\mathbf{x}^{(t)} = \begin{cases} \mathbf{x}' & , \text{ if } R_{MH} > R_t \\ \mathbf{x}^{(t)} & , \text{ else} \end{cases}$;

end

 set $\mathbf{x}^{(t+1)} = \{x_1^{(t+1)}, \dots, x_n^{(t+1)}\}$;

end

else

 draw $\mathbf{x}' \sim q(\mathbf{x} | \mathbf{x}^{(t)})$;

 calculate $R_{MH} = \frac{f(\mathbf{x}') q(\mathbf{x}^{(t)} | \mathbf{x}')}{f(\mathbf{x}^{(t)}) q(\mathbf{x}' | \mathbf{x}^{(t)})}$;

 draw $R_t \sim U[0, 1]$;

 set $\mathbf{x}^{(t+1)} = \begin{cases} \mathbf{x}' & , \text{ if } R_{MH} > R_t \\ \mathbf{x}^{(t)} & , \text{ else} \end{cases}$;

end

end

Return $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$;

should be avoided. Flat (truncated) Gaussian or uniform distributions often prove suitable and convenient.

Regarding Alg. 2.2, if $q_i(x'_i | \mathbf{x}^{(t)}) = q_i(x'_i | x_i^{(t)}) = q_i(x_i^{(t)} | x'_i)$ for all i , the algorithm is turned into a *Metropolis algorithm* [Metropolis et al. (1953)]. Moreover, it is obvious that $q_i(x'_i | \mathbf{x}^{(t)}) = f(x'_i | \mathbf{x}^{(t)})$ reduces the MH-routine to a Gibbs sampler as illustrated in Alg. 2.1.

The following section is devoted to a detailed description of how MCMC methodology can be applied to simulate from different spatial point process

models. In addition to the technical explanations provided, we give some tips concerning the implementation of the proposed algorithms. Here, the open source software R serves as a basis¹. We introduce relevant commands from the basic R-toolbox² as well as from the ‘spatstat’-package by Baddeley and Turner (2005).

Sampling from a Spatial Point Process Distribution

We revert to the notation introduced in Section 2.1 and let $\mathbf{x} = \{x_1, \dots, x_n\}$ denote the state of a point process \mathbf{X} in a two-dimensional data space $W \subseteq \mathbb{R}^2$. For a start, we consider a homogeneous Poisson process \mathbf{X} in a bounded set $B \subseteq W$ with known intensity parameter β . From \mathbf{X} , we sample a realization \mathbf{x} via a two-stage procedure as described in Illian et al. (2008, ch. 2). We first draw the total number of points, $n(\mathbf{x}) = n$, from a Poisson distribution with intensity parameter $\beta|B|$. Thereafter, we randomly distribute the points over B by superposing n independent samples from a Bernoulli process. Alg. 2.3 presents the suggested sequence of operations in brief. Using R, the functions `rpois()` and `runif()` allow for its step-wise implementation, whereas the `rpoispp()` provided in the R-package ‘spatstat’ [Baddeley and Turner (2005)] immediately returns a point process realization.

Algorithm 2.3: Sampling from a homogeneous Poisson process.

Data: Intensity parameter β , observation window B

Result: Point process realization $\mathbf{x} = \{x_1, \dots, x_n\}$

Draw $n \sim \text{Pois}(\beta |B|)$;

for $i \in \{1, \dots, n\}$ **do**

 draw $x_i \sim \text{binomial}\left(B, n = 1, f = \frac{1}{|B|}\right)$;

end

Under heterogeneity assumptions, sampling from \mathbf{X} becomes more demanding. Location-dependent thinning according to Illian et al. (2008, ch. 3), however, is an intuitive and feasible method to simulate from an inhomogeneous Poisson process with intensity function $\beta(\cdot)$. The key idea is to take a realization from a densely packed homogeneous Poisson process as a basis, and to successively delete each point with a specific probability defined through $\beta(\cdot)$. Normally, the homogeneous template realization $\mathbf{x}^* = \{x_1^*, \dots, x_{n^*}^*\}$ is sampled from a Poisson process with intensity parameter $\beta^{(*)} \geq \sup_B \{\beta(\cdot)\}$, and the probability of deleting a point x_i^* from \mathbf{x}^* in the thinning phase is defined as

¹<http://www.r-project.org/>

²<https://stat.ethz.ch/R-manual/R-devel/library/base/>

$P(x_i^* \notin \mathbf{x}) = d(x_i^*) = 1 - \frac{\beta(x_i^*)}{\beta^{(*)}}$, where \mathbf{x} is a realization from the inhomogeneous target process and $i \in \{1, \dots, n^*\}$. Alg. 2.4 illustrates this sampling procedure in a nutshell. The R-function `rpoispp()` mentioned previously is not only applicable to homogeneous, but also to inhomogeneous point process models. Via thinning, it automatically generates a point process realization with locally varying packing density if the given intensity is not a constant, but a function. Therefor, $\beta^{(*)}$ has to be determined by the user.

Algorithm 2.4: Sampling from an inhomogeneous Poisson process.

Data: Intensity function $\beta(\cdot)$, observation window B

Result: Point process realization $\mathbf{x} = \{x_1, \dots, x_n\}$

Specify $\beta^{(*)} \geq \sup_B \{\beta(\cdot)\}$;

Draw $\mathbf{x}^* = \{x_1^*, \dots, x_{n^*}^*\} \sim \text{Poisson}(B, \beta^{(*)})$ due to Alg. 2.3;

Set $\mathbf{x} = \mathbf{x}^*$;

for $i \in \{1, \dots, n^*\}$ **do**

 with probability $d(x_i^*) = 1 - \frac{\beta(x_i^*)}{\beta^{(*)}}$, set $\mathbf{x} = \{\mathbf{x} \setminus \{x_i^*\}\} = \mathbf{x}_{-\{i\}}$;

end

Not only with regard to the Poisson case, but also in applications to more sophisticated point process models, thinning proves expedient. Section 4.3 dealing with the class of locally scaled spatial point processes, for instance, shows that the concept of thinning-based inference can be adjusted to inhomogeneous patterns with repulsive points.

The skeletal structure of a sampling routine for homogeneous interaction point processes with a density following Eq. (2.9) basically depends on whether the total number of points in B is fixed or not [see also Møller and Waagepetersen (2003, ch. 7)]. Considering Gibbs processes with a known number of points, which we summarize under the term fix-n Gibbs processes (cf. Section 2.3), a starting state $\mathbf{x}^{(0)}$ has to be generated first, e.g. by placing n points on a regular grid. Thereafter, a randomly selected point $x_i^{(0)} \in \mathbf{x}^{(0)}$, $i \in \{1, \dots, n\}$, is replaced by a new one $x_i^{(1)}$ drawn from the conditional point process density $f(\cdot | \mathbf{x}_{-\{i\}}^{(0)})$, and $\mathbf{x}^{(0)}$ is turned into $\mathbf{x}^{(1)} = \{\mathbf{x}_{-\{i\}}^{(0)} \cup \{x_i^{(1)}\}\}$. This procedure is repeated again and again. After sufficiently many iterations, the states of the point process can be regarded as realizations from the target fix-n Gibbs model. Alg. 2.5 summarizes the proposed sampling stages.

Although the algorithm appears simple at first view, its actual implementation requires some more lines of code if the conditional density $f(x | \mathbf{x}_{-\{i\}}^{(t)})$ is not tractable or difficult to access. Unless given a limiting case in the form of a homogeneous Poisson or a hardcore process, only approximate samplers, such

Algorithm 2.5: Sampling from a fix-n Gibbs process, e.g. from a hardcore process with interaction radius R .

Data: Number of variables n , burn-in threshold T_0 , interaction radius R

Result: Samples $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$

// General algorithm:

Initialize $\mathbf{x}^{(0)} = \{x_1^{(0)}, \dots, x_n^{(0)}\}$;

for $t \in \{1, \dots, T-1\}$ do

 draw $i \sim U[1, n]$;

 replace $x_i^{(t)} \in \mathbf{x}^{(t)}$ by $x_i^{(t+1)} \sim f(x|\mathbf{x}_{-\{i\}}^{(t)})$;

 set $\mathbf{x}^{(t+1)} = \{\mathbf{x}_{-\{i\}}^{(t)} \cup \{x_i^{(t+1)}\}\}$;

end

Return $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$;

// Algorithm for hardcore processes:

Initialize $\mathbf{x}^{(0)} = \{x_1^{(0)}, \dots, x_n^{(0)}\}$;

for $t \in \{1, \dots, T-1\}$ do

 draw $i \sim U[1, n]$;

 draw $u \sim \text{binomial}\left(B, n=1, f = \frac{1}{|B|}\right)$;

 set $x_i^{(t+1)} = \begin{cases} u, & \text{if } \min_{v \in \mathbf{x}_{-\{i\}}^{(t)}} \{\|u - v\|\} > R \\ x_i^{(t)}, & \text{else} \end{cases} ;$

end

Return $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$;

as samplers of the type Metropolis Hastings, provide for a reasonable solution. In an M-H framework for point process models, the point replacement stage is normally called a move step and the entire simulation procedure a *move-M-H algorithm* [Møller and Waagepetersen (2003, ch. 7)]. As outlined by means of Alg. 2.2, the proposed substitution is not necessarily accepted.

To explain the sequence of operations of an M-H move procedure in the context of sampling from a spatial point process, we merge Alg. 2.2 with Alg. 2.5. First, we replace $f(\cdot|\mathbf{x}_{-\{i\}}^{(t)})$ in Alg. 2.5 by predetermined transition densities $q_i(\cdot|\mathbf{x}^{(t)})$ from which we can draw proposals $x'_i \sim q_i(x_i|\mathbf{x}^{(t)})$, such that $\mathbf{x}' = \{\mathbf{x}_{-\{i\}}^{(t)} \cup \{x'_i\}\}$ for $i \in \{1, \dots, n\}$ and $t \in \{1, \dots, T\}$. Following Alg. 2.2, we

implement *move ratios* of the form

$$R_m = \frac{f(\mathbf{x}')}{f(\mathbf{x}^{(t)})} \frac{q_i(x_i^{(t)}|\mathbf{x}')}{q_i(x'_i|\mathbf{x}^{(t)})}, \quad (2.18)$$

which we accept or reject with probability $\min\{1, R_m\}$ and $1 - \min\{1, R_m\}$, respectively. Alg. 2.6 concisely presents the embedding of move steps in an M-H framework.

Again referring to Møller and Waagepetersen (2003, ch. 7), we continue with the introduction of a sampling scheme for Gibbs processes with an unknown number of points. As no information on the exact size of the pattern is available, any point configuration – and therefore even an empty point set – may serve as starting pattern $\mathbf{x}^{(0)}$. Instead of solely conducting move steps in the form of iterative point replacements, we furthermore consider “pure” point additions and deletions, called births and deaths. It is possible to completely leave the move steps out, but this usually reduces speed of convergence considerably. In M-H terminology, the described approaches are referred to as *birth-death-move* or *birth-death-M-H algorithms*. Preliminarily to the implementation of a birth-death routine, the probability $p_b(\mathbf{x})$ of proposing to add a new point to the current state \mathbf{x} of the Markovian chain needs to be specified. A point deletion is therefore suggested with probability $1 - p_b(\mathbf{x})$. For transition purposes, we need to take numerically tractable birth and death kernels, $q_b(\cdot|\mathbf{x})$ and $q_d(\cdot|\mathbf{x})$, as a basis.

Having decided for a birth proposal by means of $p_b(\cdot)$, we first draw a new point location $x'_b \sim q_b(\cdot|\mathbf{x})$ and set $\mathbf{x}' = \{\mathbf{x} \cup \{x'_b\}\}$. We evaluate the transition probability $q_b(x'_b|\mathbf{x})$ as well as the reversal probability $q_d(x'_b|\mathbf{x}')$ of deleting x'_b from \mathbf{x}' . Building on these measures, the *birth ratio* results in

$$R_b = \frac{f(\mathbf{x}')}{f(\mathbf{x})} \frac{q_d(x'_b|\mathbf{x}')}{q_b(x'_b|\mathbf{x})} \frac{1 - p_b(\mathbf{x}')}{p_b(\mathbf{x})}. \quad (2.19)$$

A death proposal, in contrast, starts with the random selection of an existing point $x'_d \in \mathbf{x}$ subject to $q_d(x|\mathbf{x})$. The new candidate state of the point process becomes $\mathbf{x}' = \{\mathbf{x} \setminus \{x'_d\}\}$. Just as in the context of a birth proposal, we evaluate all functions that are relevant for the suggested modification as well as for its reversal, and thus end up with a *death ratio* of the form

$$R_d = \frac{f(\mathbf{x}')}{f(\mathbf{x})} \frac{q_b(x'_d|\mathbf{x}')}{q_d(x'_d|\mathbf{x})} \frac{p_b(\mathbf{x}')}{1 - p_b(\mathbf{x})}. \quad (2.20)$$

Concerning the inclusion of move steps in an M-H simulation, we make reference to Eq. (2.18) and the accompanying explanations. We adopt the birth, death, or move transition proposed in the t -th iteration of the algorithm if $R_{\text{MH}} > R_t \sim U[0, 1]$. In this case, the new state of the chain becomes $\mathbf{x}^{(t+1)} = \mathbf{x}'$, whereas otherwise, the old state is retained, i.e. $\mathbf{x}^{(t+1)} = \mathbf{x}^{(t)}$. A step-by-step

illustration of an M-H algorithm with birth, death and move steps is provided in Alg. 2.6.

The R-function `rmh()` contained in the ‘`spatstat`’-package allows for simulations from numerous common point process distributions with densities according to Eq. (2.9) [Baddeley and Turner (2005)]. The command works for both processes with a fixed and processes with a random number of points. The default method is an M-H sampling routine that automatically adapts to the attributes of the model specified by the user.

Apart from the M-H framework, a toolbox of alternative sampling approaches summarized under the term *perfect simulation techniques* has been developed and made available in the past decades [see e.g. Berthelsen and Møller (2002)]. It comprehends, amongst others, the research studies by Kendall (1998), Kendall and Møller (2000), and Fernández et al. (2002). Perfect simulation methodology is applicable under local stability assumptions and proves beneficial particularly in applications to interaction point processes. In Chapter 4 of this thesis, we describe *dominated coupling from the past* (CFTP) in terms of its technical principles (see Section 4.3). Per default, dominated CFTP is carried out when simulations from a Strauss model are initiated via the function call `rStrauss()` from the ‘`spatstat`’-package.

2.5 Parameter Estimation

We continue with an outline of the most common approaches to parameter estimation in a spatial point process framework. First, we devote our attention to frequentist inference principles, that is, to parameter estimation based on the maximization of the log-likelihood or a feasible alternative construct with similar properties [see also Møller and Waagepetersen (2003, ch. 7), Gelfand et al. (2010, ch. 19)]. Afterwards, we survey the respective Bayesian inference schemes [see also Møller and Waagepetersen (2003, ch. 9), Gelfand et al. (2010, ch. 19)]. Parameter estimation based on the Bayes theorem requires information on the joint posterior which, as discussed in detail below, is defined as the conditional distribution of the parameters given the data.

Frequentist Inference Framework

This section covers frequentist inference techniques for spatial point processes and is mainly based on the definitions provided by Møller and Waagepetersen (2003, ch. 7), and Gelfand et al. (2010, ch. 19). The contents are structured according to Section 2.2 – 2.3, starting with the treatment of homogeneous and inhomogeneous Poisson processes and passing on to pairwise interaction point processes. In conformance with above, we consider a bounded observation

window $B \subseteq W$, where $W \subseteq \mathbb{R}^2$. We let \mathbf{X} denote a point process in B and $n(\mathbf{x}) = n$ the number of points contained in the realization \mathbf{x} .

As the likelihood of a homogeneous Poisson process is presentable in exact form, and as it only depends on the unknown intensity parameter β , maximum likelihood (ML) estimation is possible without any difficulty. Complying with the well-known standard ML process sequence [Aldrich (1997)], we differentiate the log-likelihood with respect to β , equate the resulting score function with zero and solve the system for β . This yields the ML estimator

$$\hat{\beta}_{\text{ML}} = \frac{n}{|B|},$$

which is equivalent to the expected number of points per unit of the observation window B (see also Def. 2.8 and its interpretation).

Concerning Poisson processes with an inhomogeneous intensity (cf. Eq. (2.5)), there is no unique way to present the ML estimator. The variety of potential intensity functions with different parameterizations, mathematically expressed by $\beta_{\boldsymbol{\theta}}(\cdot)$, requires individual treatments of the integral $\int_B \beta_{\boldsymbol{\theta}}(u) du$ as well as individually adapted optimization techniques.

An explicit ML formulation is also not available if the given point process shows pairwise point repulsions. In such case, direct ML inference fails because of the inaccessible normalizing constant in the data model. A well-proven approach to this problem has been introduced by Besag (1977). Bringing the conditional intensity of the point process into focus, Besag proposes a *pseudo likelihood* framework. The conditional intensity $\beta_{\boldsymbol{\theta}}(x_{n+1}|\{x_1, \dots, x_n\})$, also referred to as *Papangelou (conditional) intensity* [Papangelou (1974)], can informally be perceived as the probability that a point pattern with known point coordinates $\{x_1, \dots, x_n\}$ also contains a point located in x_{n+1} . It holds that

$$\begin{aligned} \beta_{\boldsymbol{\theta}}(x_{n+1}|\{x_1, \dots, x_n\}) &= \frac{f(\{x_1, \dots, x_{n+1}\}|\boldsymbol{\theta})}{f(\{x_1, \dots, x_n\}|\boldsymbol{\theta})} \\ &= \frac{\mathcal{Z}(\boldsymbol{\theta})^{-1} f^{(*)}(\{x_1, \dots, x_{n+1}\}|\boldsymbol{\theta})}{\mathcal{Z}(\boldsymbol{\theta})^{-1} f^{(*)}(\{x_1, \dots, x_n\}|\boldsymbol{\theta})}, \quad x_{n+1} \notin \{x_1, \dots, x_n\}. \end{aligned} \tag{2.21}$$

As the the normalizing constant $\mathcal{Z}(\cdot)^{-1}$ cancels out and only the tractable kernels $f^{(*)}(\cdot)$ remain, the conditional intensity is calculable in a straightforward and exact manner.

Briefly returning to models not accounting for point interactions, it is evident that $\beta_{\boldsymbol{\theta}}(x_{n+1}|\{x_1, \dots, x_n\}) = \beta_{\boldsymbol{\theta}}(x_{n+1})$. Under independence assumptions, the Papangelou conditional intensity thus equates to the first-order intensity introduced by means of Eq. (2.2).

For pairwise interaction models following Eq. (2.9), the conditional intensity results in

$$\beta(x_{n+1}|\{x_1, \dots, x_n\}) = \phi_1(x_{n+1}) \prod_{i=1}^n \phi_2(\{x_i, x_{n+1}\}) .$$

We recall that $\phi_1(\cdot)$ is a function measuring intensity, whereas interaction effects are imposed through $\phi_2(\cdot)$.

From the introduction of the Papangelou conditional intensity, we revert to its application in pseudo likelihood inference and assume a point process realization $\mathbf{x} = \{x_1, \dots, x_n\}$ on $B \subseteq W$. The conditional intensity imposed by the parametric model considered is denoted by $\beta_{\boldsymbol{\theta}}(x_i|\mathbf{x}_{-\{i\}})$, for $i \in \{1, \dots, n\}$. Based on this notation, the pseudo likelihood due to Besag (1977) takes the form

$$L_P(\boldsymbol{\theta}) = \exp \left\{ - \int_B \beta_{\boldsymbol{\theta}}(u|\mathbf{x}) \, du \right\} \prod_{i=1}^n \beta_{\boldsymbol{\theta}}(x_i|\mathbf{x}_{-\{i\}}) . \quad (2.22)$$

The exponential term corresponds to the normalizing constant. It ensures that the pseudo likelihood is presentable in closed form and hence identifiable. We obtain a function comprehending the same structural components as a Poisson model according to Eq. (2.5). If the integral in the exponent is not solvable in an exact manner, we replace it by a feasible approximate operator, e.g. a Riemann sum [Daley and Vere-Jones (2003)].

Gibbs modeling approaches, as recalled above, with intensity and interaction conditions factorized in line with Eq. (2.9) have a pseudo likelihood conforming to

$$L_P(\boldsymbol{\theta}) = \exp \left\{ - \int_B \phi_1(u|\boldsymbol{\theta}) \prod_{i=1}^n \phi_2(\{u, x_i\}|\boldsymbol{\theta}) \, du \right\} \prod_{i=1}^n \phi_1(x_i|\boldsymbol{\theta}) \prod_{i \neq j} \phi_2(\{x_i, x_j\}|\boldsymbol{\theta}) .$$

For a Strauss model (cf. Eq. (2.11)), in particular, we obtain

$$L_P(\boldsymbol{\theta}) = \exp \left\{ - \beta \int_B \gamma^{\sum_{i=1}^n \mathbb{1}[\|x_i - u\| \leq R]} \, du \right\} \beta^n \gamma^{\sum_{i \neq j} \mathbb{1}[\|x_i - x_j\| \leq R]} .$$

The build-up of the pseudo likelihood is strongly related to that of the *composite likelihood* by Lindsay (1988). In contrast to the pseudo likelihood, however, the composite likelihood is a first-order-moment-based construct and thus not appropriate for modeling interaction. It is determined through the first-order intensity defined in Section 2.1 (cf. Eq. (2.2)) and can be denoted as

$$L_C(\boldsymbol{\theta}) = \exp \left\{ - \int_B \beta_{\boldsymbol{\theta}}(u) \, du \right\} \prod_{i=1}^n \beta_{\boldsymbol{\theta}}(x_i) . \quad (2.23)$$

We see that the composite likelihood looks identical to the density of an inhomogeneous Poisson process with intensity $\beta_{\boldsymbol{\theta}}(\cdot)$ (cf. Eq. (2.5)). Moreover, it is

obvious that, under the assumption of independent points, the pseudo likelihood takes exactly the form of the composite likelihood. For a more detailed mathematical derivation of the composite likelihood principle, we refer to Lindsay (1988) and Gelfand et al. (2010, ch. 19).

To sum up, the pseudo likelihood is a widely-used tool for modeling Gibbs point processes. Given the intensity of a point process where the points do not interact, a composite likelihood approach is usually taken as a basis. In a Poisson framework, the exact likelihood available in complete form is the foundation for frequentist inference. Based on an appropriate implementation of the likelihood, that is, based on a tractable representation of the data model as a function of the unknown parameters, the ML estimator can be determined as the mode of the likelihood. Bayesian approaches, in contrast, allow to explore the *posterior (distribution)* of the parameters. This means that, instead of immediately assessing the optimal parameter values subject to a closed-form variant of the likelihood, the focus is rather on the distributional properties of the parameters. These properties are deducible from the data model merged with all available prior information.

Bayesian Inference Framework

In this section, we first give a short reminder of the Bayes rule and its application to statistical inference. Our theoretical introduction is mainly based on the book by Bernardo and Smith (2009). Following Waller (2005), we afterwards discuss the implementation of Bayesian inference schemes for homogeneous and inhomogeneous Poisson processes. Finally, we survey how inference relying on an M-H routine with auxiliary variables allows us to analyze point processes under the restrictive assumption of an incomplete likelihood, e.g. owing to the incorporation of important interaction constraints [see also Møller et al. (2006), Murray et al. (2012)].

Referring to Bernardo and Smith (2009, ch. 2), we recall that Bayesian inference is motivated by the Bayes theorem stating that

$$\begin{aligned} p(\boldsymbol{\theta}|\mathbf{x}) &= \frac{f(\mathbf{x}, \boldsymbol{\theta})}{\mathcal{C}(\mathbf{x})} = \frac{f(\mathbf{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta})}{\mathcal{C}(\mathbf{x})} \\ &\propto f(\mathbf{x}|\boldsymbol{\theta}) \times p(\boldsymbol{\theta}) . \end{aligned} \quad (2.24)$$

In words, the posterior of the parameters $\boldsymbol{\theta}$ given the data \mathbf{x} is, up to a normalizing constant $\mathcal{C}(\cdot)^{-1}$, equal to the product of the likelihood $f(\cdot|\boldsymbol{\theta})$ and some distributional information on $\boldsymbol{\theta}$ included in the *prior* $p(\cdot)$. To ensure that $p(\cdot|\mathbf{x})$ is a proper density, $\mathcal{C}(\cdot)$ must conform to

$$\mathcal{C}(\mathbf{x}) = \begin{cases} \sum_{\boldsymbol{\theta} \in \Theta} f(\mathbf{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta}) , & p(\cdot) \text{ discrete} , \\ \int_{\Theta} f(\mathbf{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta} , & p(\cdot) \text{ continuous} , \end{cases} \quad (2.25)$$

which is often not exactly calculable in the continuous case. The choice of the prior consequently influences the form and compactness of the posterior. Previous knowledge of θ may be imposed through a *conjugate prior*, which makes $p(\cdot|\mathbf{x})$ take the same form as $p(\cdot)$. Conjugate priors are mathematically convenient, though not always the best way to convey prior information. A necessary but not sufficient prerequisite for $p(\cdot)$ representing a conjugate prior is that both $p(\cdot)$ and the likelihood $f(\cdot|\theta)$ conform to *exponential families*, whereas it is essential that $p(\cdot)$ and the posterior $p(\cdot|\mathbf{x})$ belong to the same exponential family [see also Bernardo and Smith (2009, ch. 4 – 5)]. In other words, the key characteristic of a conjugate prior is that its multiplication with the likelihood results in a posterior kernel of identical structure as the prior kernel, and that the domain of θ is the same subject to both the given prior and the estimated posterior model. Considering a homogeneous Poisson process density according to Eq. (2.2), for instance, the choice of a Gamma prior for $\theta = \beta$ brings about a Gamma posterior for $\beta|\mathbf{x}$.

To obviate biases resulting from an over-specification, the prior should be kept sufficiently flat. Often, *non-informative priors* prove most suitable [Bernardo and Smith (2009, ch. 5)]. A non-informative prior, such as e.g. a uniform density, yields a posterior which is proportional to the data distribution and hence conforms to $p(\theta|\mathbf{x}) = f(\mathbf{x}|\theta) \times \text{const.}$ When deciding for a non-informative prior, it is important that the posterior is still well-defined. That is, $p(\cdot)$ must be chosen such that the integral or sum determining $\mathcal{C}(\cdot)$ in Eq. (2.24) converges. This ensures that the discrete or continuous posterior density $p(\cdot|\mathbf{x})$ is proper in that it sums or integrates to 1. Moreover, the posterior is only precisely definable if the data model is available in complete form.

Problems come up if, irrespective of the choice of the prior, the posterior is not accessible in a straightforward manner. A common reason is that the data distributional assumptions are not in line with any exponential family. If the structure of the posterior is not directly derivable from the product of the prior and the data distribution, MCMC simulation techniques may be used. This means that, depending on whether a Gibbs or an M-H framework proves appropriate, parameter samples are indirectly drawn from $p(\cdot|\mathbf{x})$ based on either the full conditionals (cf. Alg. 2.1) or feasible transition conditions (cf. Alg. 2.2). A chain of states $\theta^{(0)} \rightarrow \theta^{(1)} \rightarrow \dots \rightarrow \theta^{(T)}$ is thus computed. As soon as it is in its equilibrium regime, the statistical properties of $p(\cdot|\mathbf{x})$ can be characterized through the parameter samples.

MCMC-based parameter estimation requires further explanation if the data model has an incomplete density, such as e.g. a common Gibbs process density (cf. Eq. (2.9)). This confronts us with the problem of *double intractability*, since

$$p(\theta|\mathbf{x}) = \frac{f(\mathbf{x}|\theta) p(\theta)}{\mathcal{C}(\mathbf{x})} = \mathcal{C}(\mathbf{x})^{-1} \times \mathcal{Z}(\theta)^{-1} f^{(*)}(\mathbf{x}|\theta) \times p(\theta) \quad (2.26)$$

contains two unknown normalizing terms, $\mathcal{C}(\cdot)^{-1}$ and $\mathcal{Z}(\cdot)^{-1}$. We denote the tractable kernel of the data distribution by $f^{(*)}(\cdot|\boldsymbol{\theta})$ (cf. Eq. (2.10)). In a standard M-H framework according to Alg. 2.2, the M-H ratio yields the elimination of $\mathcal{C}(\cdot)^{-1}$. To get rid of $\mathcal{Z}(\cdot)^{-1}$, too, advanced M-H methodology is needed. Aiming at maintaining the structure of the M-H algorithm, Møller et al. (2006) propose the addition of a cleverly devised auxiliary variable scheme. Murray et al. (2012) introduce the so-called *exchange algorithm* which is strongly related to the approach by Møller et al. in terms of its motivation and conception.

The key idea behind the exchange procedure developed by Murray et al. is to extend the M-H ratio by a statistically valid quotient that makes $\mathcal{Z}(\cdot)^{-1}$ cancel out. For a more detailed description, we let $\boldsymbol{\theta} = \{\theta_1, \dots, \theta_J\}$ denote the set of model parameters and $p(\cdot)$ their joint prior. We assume a point process realization \mathbf{x} with density $f(\cdot|\boldsymbol{\theta})$ on $W \subseteq \mathbb{R}^2$ and predefine transition kernels $q(\cdot|\boldsymbol{\theta})$ and $q_j(\cdot|\theta_j)$, for $j \in \{1, \dots, J\}$. From these kernels, proposals $\boldsymbol{\theta}'$ and, respectively, θ'_j must be drawable. Unlike in Alg. 2.2, the exchange principle requires the simulation of an additional set of data, $\mathbf{w} \sim f(\mathbf{w}|\boldsymbol{\theta}')$, within each inner loop of the algorithm. Defining $f(\mathbf{w}|\boldsymbol{\theta}')$ such that its normalizing constant is equal to the inaccessible normalizing constant in $f(\mathbf{x}|\boldsymbol{\theta}')$, the common M-H ratio in Alg. 2.2 can be extended to

$$\begin{aligned} R_{\text{eMH}} &= \frac{\mathcal{Z}(\boldsymbol{\theta}')^{-1} f^{(*)}(\mathbf{x}|\boldsymbol{\theta}') p(\boldsymbol{\theta}')}{\mathcal{Z}(\boldsymbol{\theta})^{-1} f^{(*)}(\mathbf{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta})} \frac{q_j(\theta_j|\boldsymbol{\theta}', \mathbf{x})}{q_j(\theta'_j|\boldsymbol{\theta}, \mathbf{x})} \frac{\mathcal{Z}(\boldsymbol{\theta})^{-1} f^{(*)}(\mathbf{w}|\boldsymbol{\theta})}{\mathcal{Z}(\boldsymbol{\theta}')^{-1} f^{(*)}(\mathbf{w}|\boldsymbol{\theta}')} \\ &= \frac{f^{(*)}(\mathbf{x}|\boldsymbol{\theta}') p(\boldsymbol{\theta}')}{f^{(*)}(\mathbf{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta})} \frac{q_j(\theta_j|\boldsymbol{\theta}', \mathbf{x})}{q_j(\theta'_j|\boldsymbol{\theta}, \mathbf{x})} \frac{f^{(*)}(\mathbf{w}|\boldsymbol{\theta})}{f^{(*)}(\mathbf{w}|\boldsymbol{\theta}')} , \end{aligned} \quad (2.27)$$

where, again, $f^{(*)}(\cdot|\cdot)$ symbolizes tractable density kernels. A rearrangement of the quotient allows for plausible explanatory statements referring to the auxiliary variable scheme. $\frac{f(\mathbf{w}|\boldsymbol{\theta})}{f(\mathbf{x}|\boldsymbol{\theta})}$ indicates, whether and how strongly $\boldsymbol{\theta}$ prefers \mathbf{w} to \mathbf{x} under the data model considered, whereas $\frac{f(\mathbf{x}|\boldsymbol{\theta}')}{f(\mathbf{w}|\boldsymbol{\theta}')}$ similarly compares the strength of the statistical link between \mathbf{x} and $\boldsymbol{\theta}'$ to that between \mathbf{w} and $\boldsymbol{\theta}'$.

Alg. 2.7 shows the step sequence of the exchange algorithm by Murray et al. (2012). If not realizable in a direct manner, we suggest to conduct the auxiliary sampling step via one of the simulation techniques explained in Section 4.3 [see also Berthelsen and Møller (2003), Liang (2010)]. The exchange principle provides a crucial basis for our research. In Chapter 4, we show that a sophisticated implementation of the exchange algorithm enables Bayesian inference in a Strauss process framework with location-dependent scaling attributes (see Section 4.1 – 4.3).

In general, diverse posteriori estimators are derivable from a posterior distribution. Apart from the classical measures – including posterior mean, variance, modus and median –, posterior moments of higher order, posterior quantiles and many further distributional characteristics of interest may be determined.

Remark: Including Marks

To conclude this section, we briefly discuss how to estimate the parameters of a marked point process. Referring to the book by Møller and Waagepetersen (2003), we sketch the necessary adjustments of both the frequentist and the Bayesian algorithms explained above. We use the notation introduced in Section 2.1, where we define a marked point process as $\mathbf{Y} = \{(X, m_X) : X \in \mathbf{X}, m_X \in M\}$ and denote a realization of \mathbf{Y} by \mathbf{y} . The domain of a marked point process thus conforms to the domain of a pure point process (cf. Eq. (2.1)) under additional consideration of all possible mark assignments subject to the given mark space M (cf. Eq. (2.4)). Without loss of generality, the following explanations apply to any type of point process model presentable by a tractable or incomplete density $f(\mathbf{y}|\boldsymbol{\theta})$ with respect to the unit rate Poisson process. The marked point process is assumed to act on the product space $A = (W \times M)$. Furthermore, an appropriately specified discrete or continuous density $p_M(\cdot)$ describes the marginal mark distribution on M .

For the purpose of frequentist inference, we implement a pseudo likelihood approach as previously introduced by means of the properties of the Papangelou conditional intensity (cf. Eq. (2.21) – Eq. (2.22)). Accounting for marks does not yield any structural changes of the common form of the pseudo likelihood which can here be written as

$$L_P(\boldsymbol{\theta}) = \prod_{(x, m_x) \in \mathbf{y}} \beta_{\boldsymbol{\theta}}((x, m_x) | \{\mathbf{y} \setminus \{(x, m_x)\}\}) \quad (2.28)$$

$$\times \begin{cases} \exp \left\{ - \int_W \sum_{m_u \in M} \beta_{\boldsymbol{\theta}}((u, m_u) | \mathbf{y}) \mathbb{1}[(u, m_u) \in A] p_M(m_u) du \right\}, & M \text{ discrete,} \\ \exp \left\{ - \int_W \int_M \beta_{\boldsymbol{\theta}}((u, m_u) | \mathbf{y}) \mathbb{1}[(u, m_u) \in A] p_M(m_u) dm_u du \right\}, & \text{else.} \end{cases}$$

More details can be found in Møller and Waagepetersen (2003, ch. 9). In accordance with the unmarked case, conditionality becomes meaningless if interaction is not modeled. The form of the pseudo likelihood is then equivalent to that of the composite likelihood (cf. Eq. (2.23)) which, as discussed above, corresponds to the density of an inhomogeneous Poisson process (cf. Eq. (2.7)).

For estimating the parameters of a marked point process based on the Bayesian inference principle, we again revert to the methods proposed in the context of analyzing non-marked point processes. A discrete or continuous density measure $p_M(\cdot)$ on M needs to be predetermined, which is here regarded as the mark prior. The mark attributes parameterized with \mathbf{m} go into $\boldsymbol{\theta}$ as additional parameters, and accordingly, $p_M(\cdot)$ merges with the prior of the remaining parameters, i.e. $p(\boldsymbol{\theta}) = p_M(\mathbf{m}) \times p(\{\boldsymbol{\theta} \setminus \mathbf{m}\})$. Applying an (exchange) M-H algorithm to the posterior of a marked point process,

$$f(\boldsymbol{\theta}|\mathbf{y}) \propto f(\mathbf{y}|\boldsymbol{\theta}) p(\boldsymbol{\theta}),$$

the M-H ratio takes the form

$$R_{(\text{e})\text{MH}} = \frac{f(\mathbf{y}|\boldsymbol{\theta}')}{f(\mathbf{y}|\boldsymbol{\theta})} \frac{p(\boldsymbol{\theta}')}{p(\boldsymbol{\theta})} \frac{q(\boldsymbol{\theta}|\boldsymbol{\theta}', \mathbf{y})}{q(\boldsymbol{\theta}'|\boldsymbol{\theta}, \mathbf{y})} \left(\times \frac{f(\mathbf{w}|\boldsymbol{\theta})}{f(\mathbf{w}|\boldsymbol{\theta}')} \right), \quad (2.29)$$

provided that $q(\cdot|\cdot)$ is a tractable positive and proper transition density. In an exchange framework, \mathbf{w} denotes a marked auxiliary variable sampled from $f(\mathbf{w}|\boldsymbol{\theta}')$ which contains the same normalizing constant as $f(\mathbf{y}|\boldsymbol{\theta}')$. Just as in the case of non-marked points, it is usually not possible to generate \mathbf{w} in a direct manner. Therefore, we once again make reference to the simulation techniques proposed by Berthelsen and Møller (2003), and Liang (2010), as well as to the related discussions in Section 4.3 of this thesis. For a more detailed description of Eq. (2.29), we refer to Alg. 2.2, Alg. 2.7, and the according explanatory statements in Section 2.4.

We have provided an introductory compendium of definitions and methods related to spatial point processes, which is important for a profound understanding of our research foci. In the later chapters, reference will be made to the tables and equations presented in this section, whenever needed or appropriate. Not all essential theoretical and technical basics have yet been addressed, however. Advanced state-of-the-art methodology of particular relevance for our work will be discussed throughout the respective chapters.

Algorithm 2.6: Birth-death-move-M-H algorithm: Sampling from an interaction point process.

Data: Burn-in threshold T_0

Result: Point process realizations $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$

Initialize $\mathbf{x}^{(0)} = \{x_1^{(0)}, \dots, x_n^{(0)}\}$;

for $t \in \{1, \dots, T - 1\}$ **do**

with probability $p_b(\mathbf{x}^{(t)})$, $p_d(\mathbf{x}^{(t)})$, and $1 - p_b(\mathbf{x}^{(t)}) - p_d(\mathbf{x}^{(t)})$,
propose a birth, a death, or a move;

if $proposal == \text{"birth"}$ **then**

draw $x'_b \sim q_b(x|\mathbf{x}^{(t)})$;

set $\mathbf{x}' = \{\mathbf{x}^{(t)} \cup \{x'_b\}\}$;

set $n' = n^{(t)} + 1$;

calculate $R_{MH} = R_b = \frac{f(\mathbf{x}')}{f(\mathbf{x}^{(t)})} \frac{q_d(x'_b|\mathbf{x}')}{q_b(x'_b|\mathbf{x}^{(t)})} \frac{p_d(\mathbf{x}')}{p_b(\mathbf{x}^{(t)})}$;

end

else

if $proposal == \text{"death"}$ **then**

draw $x'_d \sim q_d(x|\mathbf{x}^{(t)})$;

set $\mathbf{x}' = \{\mathbf{x}^{(t)} \setminus \{x'_d\}\}$;

set $n' = n^{(t)} - 1$;

calculate $R_{MH} = R_d = \frac{f(\mathbf{x}')}{f(\mathbf{x}^{(t)})} \frac{q_b(x'_d|\mathbf{x}')}{q_d(x'_d|\mathbf{x}^{(t)})} \frac{p_b(\mathbf{x}')}{p_d(\mathbf{x}^{(t)})}$;

end

else

draw $i \sim U[1, n^{(t)}]$;

draw $x'_i \sim q_i(x_i|\mathbf{x}^{(t)})$;

set $\mathbf{x}' = \{\mathbf{x}^{(t)}_{-\{i\}} \cup \{x'_i\}\}$;

set $n' = n^{(t)}$;

calculate $R_{MH} = R_m = \frac{f(\mathbf{x}')}{f(\mathbf{x})} \frac{q_i(x'_i|\mathbf{x}')}{q_i(x'_i|\mathbf{x}^{(t)})}$;

end

end

draw $R_t \sim U[0, 1]$;

set $\{\mathbf{x}^{(t+1)}, n^{(t+1)}\} = \begin{cases} \{\mathbf{x}', n'\}, & \text{if } R_{MH} > R_t; \\ \{\mathbf{x}^{(t)}, n^{(t)}\}, & \text{else} \end{cases}$;

end

Return $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$ as drawings from $f(\mathbf{x})$;

Algorithm 2.7: Exchange algorithm.

Data: Burn-in threshold T_0 , method

Result: Samples $\boldsymbol{\theta}^{(T_0)}, \dots, \boldsymbol{\theta}^{(T)}$

Initialize $\boldsymbol{\theta}^{(0)} = \{\theta_1^{(0)}, \dots, \theta_J^{(0)}\}$;

for $t \in \{0, \dots, T - 1\}$ **do**

if $method == \text{"individual"}$ **then**

for $j \in \{1, \dots, J\}$ **do**

 draw $\theta'_j \sim q_j(\theta_j | \boldsymbol{\theta}^{(t)}, \mathbf{x})$;

 set $\boldsymbol{\theta}' = \{\boldsymbol{\theta}^{(t)}_{-\{j\}} \cup \{\theta'_j\}\}$;

 draw $\mathbf{w} \sim f(\mathbf{w} | \boldsymbol{\theta}')$;

 calculate $R_{\text{eMH}} = \frac{f^{(*)}(\mathbf{x} | \boldsymbol{\theta}')}{f^{(*)}(\mathbf{x} | \boldsymbol{\theta})} \frac{p(\boldsymbol{\theta}')}{p(\boldsymbol{\theta})} \frac{q_j(\theta_j | \boldsymbol{\theta}', \mathbf{x})}{q_j(\theta'_j | \boldsymbol{\theta}, \mathbf{x})} \frac{f^{(*)}(\mathbf{w} | \boldsymbol{\theta})}{f^{(*)}(\mathbf{w} | \boldsymbol{\theta}')};$

 draw $R_t \sim U[0, 1]$ and set $\boldsymbol{\theta}^{(t)} = \begin{cases} \boldsymbol{\theta}' & \text{if } R_{\text{eMH}} > R_t \\ \boldsymbol{\theta}^{(t)} & \text{else} \end{cases};$

end

 set $\boldsymbol{\theta}^{(t+1)} = \boldsymbol{\theta}^{(t)}$;

end

else

 draw $\boldsymbol{\theta}' \sim q(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)}, \mathbf{x})$;

 draw $\mathbf{w} \sim f(\mathbf{w} | \boldsymbol{\theta}')$;

 calculate $R_{\text{eMH}} = \frac{f^{(*)}(\mathbf{x} | \boldsymbol{\theta}')}{f^{(*)}(\mathbf{x} | \boldsymbol{\theta})} \frac{p(\boldsymbol{\theta}')}{p(\boldsymbol{\theta})} \frac{q(\boldsymbol{\theta} | \boldsymbol{\theta}', \mathbf{x})}{q(\boldsymbol{\theta}' | \boldsymbol{\theta}, \mathbf{x})} \frac{f^{(*)}(\mathbf{w} | \boldsymbol{\theta})}{f^{(*)}(\mathbf{w} | \boldsymbol{\theta}')};$

 draw $R_t \sim U[0, 1]$;

 set $\boldsymbol{\theta}^{(t+1)} = \begin{cases} \boldsymbol{\theta}' & \text{if } R_{\text{eMH}} > R_t \\ \boldsymbol{\theta}^{(t)} & \text{else} \end{cases};$

end

end

Return $\boldsymbol{\theta}^{(T_0)}, \dots, \boldsymbol{\theta}^{(T)}$ as drawings from $p(\boldsymbol{\theta} | \mathbf{x})$;

3 Preliminaries on Image Analysis

Chapter 6 of this thesis elaborates a point process model that enables the geometric analysis of images of textured scenes with regard to the orientation of the camera. Suitable image preprocessing measures as well as the relations between a three-dimensional (3D) original scene and its two-dimensional (2D) camera projection(s) are the basic prerequisites for the successful implementation and interpretation of our approach. The following sections provide the conceptual and technical preliminaries that are of relevance for developing a method to estimate shape from texture.

Section 3.1 introduces terminological as well as methodological details related to the internal parameters and the Euclidean motion of a perspective camera with respect to a static 3D scene. First, the meaning and the properties of a homogeneous coordinate representation are addressed and the general set-up of a pinhole camera model is illustrated. In what follows, the relations between camera projections from different points of view are discussed. A major task of the image analysis stage concerns the detection of texture elements under weak assumptions. Section 3.2 gives a short account of image processing techniques that are relevant in this connection.

3.1 Camera Projection

Since time immemorial, humans have been engaged with research on geometry and its relations to dimensionality in daily life. The permanently ongoing effort to make the meaning of space mathematically accessible is remarkable. This effort has not only been addressed to describing spaces of a certain dimension, but also to establishing links between geometries in spaces of different dimensions. An obvious example is the projection of objects from a 3D observation space to a 2D image plane using a camera.

In the following, we provide some basic facts that enable the understanding of the interrelations of projections of a fixed scene recorded from different camera positions. We start with some fundamental definitions related to coordinate representations as well as coordinate transformations within and between 2D and 3D spaces. Moreover, we describe the relevant parameters of the standard pinhole camera model. Following the exposition by Becker et al. (2014), we afterwards outline what is meant by *two-view geometry* and introduce the term *assignment field*. For further information, we recommend the standard works by Hartley and Zisserman (2000, Part II), Faugeras and Luong (2001), Ramírez Galarza and Seade (2007), and Hughes et al. (2014).

Basics

Assuming that we are given a 3D scene projected onto a 2D plane $W \subset \mathbb{R}^2$, such that both a 3D scene point $\mathcal{X} \in \mathbb{R}^3$ and the according 2D coordinates $x \in \mathbb{R}^2$ are connected to the camera through a straight line defined through $\lambda\vec{x} = \lambda x$, where $\lambda \in \mathbb{R}$ and \vec{x} symbolizes the vector pointing from the origin towards x (see Fig. 3.1 (a)). To begin with, we assume that the camera is located at the origin of the underlying 3D coordinate system the axes of which we denote by \mathcal{X}_{a1} , \mathcal{X}_{a2} , and \mathcal{X}_{a3} . The camera center is often also called the center of projection or the optical center. In accordance with Fig. 3.1 (a) and the respective illustrations in Hartley and Zisserman (2000, ch. 6), \mathcal{X}_3 is chosen to become the principal axis of the projection meaning that it is directed perpendicular to the image plane W . It meets W in the principal point p_0 .




Figure 3.1: Process of image formation. (a) shows the camera ray $\lambda\vec{x} = \lambda x$ that connects the camera with the 3D scene point \mathcal{X} and maps \mathcal{X} to the 2D coordinates $x \in W$. \mathcal{X}_{a3} serves as principal axis intersecting with W in the principal point p_0 which corresponds to the origin of the 2D image coordinate system chosen here. (b) presents the projection of the image point x subject to a pinhole camera model with focal length f .

Any point on λx , including the image point x and excluding the origin itself, is determined to correspond to the point \tilde{x} of the so-called *projective plane* \mathbb{P}^2 which contains all lines in \mathbb{R}^3 that go through the origin. Consequently, \tilde{x} stands for an equivalence class comprising all 3D points on the camera ray λx , where

$\lambda \in \mathbb{R}_{\setminus\{0\}}$. In technical terminology, $\tilde{x} \in \mathbb{P}^2$ denotes the *homogeneous coordinate* representation of λx . For obvious reasons, homogeneous coordinate systems are not influenced by scale factors, or, more formally,

$$\tilde{x} \simeq \tilde{x}' \Leftrightarrow \tilde{x} = \lambda \tilde{x}', \quad \lambda \neq 0,$$

where “ \simeq ” symbolizes equivalence. We see that the numerically valid affine domain $\{\tilde{x} \in \mathbb{P}^2 : \tilde{x}_3 \neq 0\}$ of the projective plane \mathbb{P}^2 corresponds to the set of all lines in 3D through the origin, as stated above.

Hartley and Zisserman (2000, ch. 6) propose camera models of differing complexity. Nevertheless, we limit our focus to the *basic pinhole camera*. The (positive) distance between the camera center and the image plane W with respect to the principal axis \mathcal{X}_{a3} is termed the *focal length* and symbolized by f . Assuming that the camera is located at the origin of the 3D scene coordinate system, and provided that the principal point p_0 accords to the origin of the 2D image coordinate system, $\mathcal{X} = (\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}_3)^\top$ is mapped to $\tilde{x} = (\tilde{x}_1, \tilde{x}_2, \tilde{x}_3)^\top = (f\mathcal{X}_1/\mathcal{X}_3, f\mathcal{X}_2/\mathcal{X}_3, f)^\top$. Irrespective of the third coordinate, projections from scene to image points are thus of the form

$$(\tilde{x}_1, \tilde{x}_2, \tilde{x}_3)^\top \mapsto \left(f \frac{\mathcal{X}_1}{\mathcal{X}_3}, f \frac{\mathcal{X}_2}{\mathcal{X}_3}\right)^\top. \quad (3.1)$$

If the image plane W is located at $f = \mathcal{X}_3 = 1$, for instance, the link between the homogeneous and the inhomogeneous coordinate representations of the points in W corresponds to

$$x = \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \frac{1}{\tilde{x}_3} \begin{pmatrix} \tilde{x}_1 \\ \tilde{x}_2 \end{pmatrix}. \quad (3.2)$$

A visual interpretation of the pinhole camera model can be gathered from Fig. 3.1 (a), where a camera working in a mathematically ideal and correct way is considered.

In the exact same manner as for linearizing transformations of a subset of \mathbb{R}^2 , it is convenient to define a homogeneous coordinate notation of $\mathcal{X} \in \mathbb{R}^3$, i.e. $\tilde{\mathcal{X}} = (\tilde{\mathcal{X}}_1, \tilde{\mathcal{X}}_2, \tilde{\mathcal{X}}_3, \tilde{\mathcal{X}}_4)^\top \in \mathbb{P}^3$. The equivalent of Eq. (3.2) then corresponds to

$$\mathcal{X} = \begin{pmatrix} \mathcal{X}_1 \\ \mathcal{X}_2 \\ \mathcal{X}_3 \end{pmatrix} = \frac{1}{\tilde{\mathcal{X}}_4} \begin{pmatrix} \tilde{\mathcal{X}}_1 \\ \tilde{\mathcal{X}}_2 \\ \tilde{\mathcal{X}}_3 \end{pmatrix}. \quad (3.3)$$

Based on this notation and the analogous statement in Eq. (3.2), we relate homogeneous coordinates \tilde{x} in the projective plane \mathbb{P}^2 to homogeneous coordinates $\tilde{\mathcal{X}} \in \mathbb{P}^3$ via

$$\tilde{x} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \tilde{\mathcal{X}} = (I_{33}, 0_3) \tilde{\mathcal{X}} = P\tilde{\mathcal{X}}. \quad (3.4)$$

For evident reasons, we call P the *projection matrix* of the mapping.

For taking internal camera-specific parameters into account, a *camera calibration matrix* $K \in \mathbb{R}^{3 \times 3}$ with entries approximating the attributes of the camera configuration has to be embedded in the model. Such a matrix extends $P = (I_{33}, 0_3)$ to $P = K (I_{33}, 0_3)$ yielding

$$\tilde{x} = P\tilde{\mathcal{X}} = K (I_{33}, 0_3) \tilde{\mathcal{X}} . \quad (3.5)$$

Inversely,

$$\tilde{x}^c = K^{-1}\tilde{x} \quad (3.6)$$

is termed the according *calibrated/normalized coordinate* representation, which results in an affine transformation of the image plane. The camera model considered determines the degrees of freedom in K . Given a pinhole camera model following Fig. 3.1 (a), K is of the simple form

$$K = \begin{pmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{pmatrix} ,$$

including the focal length f as its only parameter. If, unlike in Fig. 3.1 (a), the origin of the 2D image coordinate system is shifted by $(s_1, s_2)^\top$ from the principal point p_0 , this translation goes into the projection through

$$K = \begin{pmatrix} f & 0 & s_1 \\ 0 & f & s_2 \\ 0 & 0 & 1 \end{pmatrix} ,$$

yielding

$$\begin{pmatrix} \mathcal{X}_1 \\ \mathcal{X}_2 \\ \mathcal{X}_3 \\ 1 \end{pmatrix} \mapsto K(I_{33}, 0_3) = \begin{pmatrix} f\mathcal{X}_1 + s_1 \\ f\mathcal{X}_2 + s_2 \\ \mathcal{X}_3 \end{pmatrix} . \quad (3.7)$$

More general camera models due to Hartley and Zisserman (2000) are characterized by a more complex internal structure and further degrees of freedom. Such models might contain an additional skewness parameter, for instance.

So far, we have limited our focus to the camera-specific effects on the mapping of a 3D space to a 2D image plane. As the modeling of camera motions through external parameters has not yet been addressed, we devote the subsequent paragraphs to the definition and implementation of so-called *rigid transformations*. We base our explanations on Ramírez Galarza and Seade (2007, ch. 1), and use both inhomogeneous as well as homogeneous coordinate notations of the image and the scene points. Building on that first overview, we later embed rigid transformations in the context of modeling camera projections from different points of view.

Rigid transformations are isometric affine mappings preserving single points, straight lines, and planes. Parallel lines, for instance, remain parallel and distances unchanged. In general, rigid transformations comprehend translations, rotations and reflections, while rigid transformations without reflections are referred to as *proper rigid transformations* or *roto-translations*. Proper rigid transformations ensure that entire objects are kept unmodified in shape and size. For explicit graphical illustrations, we refer to Hughes et al. (2014, ch. 5). Applied to spaces \mathbb{R}^d , $d \in \{2, 3, \dots\}$, rigid transformations are also known as *Euclidean transformations*. The following explanations are related to the set of proper Euclidean rigid transformations, also known as the *special Euclidean group* and abbreviated by $SE(d)$.

For describing a proper rigid transformation, we employ a proper rotation matrix Λ of dimension $(d \times d)$ as well as a finite translation vector h of length d . Λ has to be an element of the *special orthogonal group* of dimension d , $SO(d)$, which corresponds to the set of all proper rotation matrices. A rotation matrix is proper if it is orthogonal, i.e. $\Lambda^\top \Lambda = I_{dd}$, where I_{dd} is the $(d \times d)$ identity/unit matrix with 1's on its diagonal and 0's everywhere else. Another necessary condition for the properness of Λ refers to its determinant, $\det(\Lambda)$, that has to equal +1. If $\det(\Lambda) = -1$, the rotation is improper meaning that it includes a reflection.

A proper rigid transformation $\{h, \Lambda\} \in SE(d)$ of a point $\mathcal{X} \in \mathbb{R}^d$ with homogeneous coordinates $\tilde{\mathcal{X}} \in \mathbb{P}^d$ yields mappings $\mathcal{X} \mapsto \mathcal{X}'$ and $\tilde{\mathcal{X}} \mapsto \tilde{\mathcal{X}}'$, where

$$\mathcal{X}' = \Lambda \mathcal{X} + h \quad \text{and} \quad \tilde{\mathcal{X}}' = \begin{pmatrix} \Lambda & h \\ 0_d^\top & 1 \end{pmatrix} \tilde{\mathcal{X}} \quad (3.8)$$

in homogeneous and inhomogeneous coordinates, respectively. In Eq. (3.8), 0_d is a vector of length d consisting of 0's only. The respective inverse transformation $\mathcal{X}' \mapsto \mathcal{X}$ and $\tilde{\mathcal{X}}' \mapsto \tilde{\mathcal{X}}$ via $\{-\Lambda^\top h, \Lambda^\top\} \in SE(d)$ amounts to

$$\mathcal{X} = \Lambda^\top (\mathcal{X}' - h) \quad \text{and} \quad \tilde{\mathcal{X}} = \begin{pmatrix} \Lambda^\top & -\Lambda^\top h \\ 0_d^\top & 1 \end{pmatrix} \tilde{\mathcal{X}}'. \quad (3.9)$$

Proper rotational matrices in \mathbb{R}^2 have one degree of freedom, which can be parametrized e.g. by

$$\Lambda = \begin{pmatrix} \cos \rho & -\sin \rho \\ \sin \rho & \cos \rho \end{pmatrix}. \quad (3.10)$$

That is, $\Lambda \in \mathbb{R}^{2 \times 2}$ yields a rotation by an angle ρ with respect to the origin [Hughes et al. (2014, ch. 5)]. In three-dimensional spaces, a rotation by $\Lambda \in \mathbb{R}^{3 \times 3}$ has up to three degrees of freedom. Single axis rotations by angles ρ_\bullet require the implementation of rotational matrices Λ_\bullet corresponding to

$$\Lambda_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos \rho_1 & -\sin \rho_1 \\ 0 & \sin \rho_1 & \cos \rho_1 \end{pmatrix}, \quad \Lambda_2 = \begin{pmatrix} \cos \rho_2 & 0 & -\sin \rho_2 \\ 0 & 1 & 0 \\ \sin \rho_2 & 0 & \cos \rho_2 \end{pmatrix},$$

$$\text{and } \Lambda_3 = \begin{pmatrix} \cos \rho_3 & -\sin \rho_3 & 0 \\ \sin \rho_3 & \cos \rho_3 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (3.11)$$

Therein, the subscripts 1, 2, and 3 indicate whether a rotation around the first, second or third axis is carried out. Any combination of the formal definitions in Eq. (3.11) via matrix multiplication, e.g. $\Lambda_{12} = \Lambda_1 \Lambda_2$ or $\Lambda_{123} = \Lambda_1 \Lambda_2 \Lambda_3$, results in a proper rotation in \mathbb{R}^3 [Hughes et al. (2014, ch. 5)].

Given the definition of a pinhole camera with internal calibration matrix K (cf. Eq. (3.7)) and being familiar with proper rigid transformations by external motion parameters $\{h, \Lambda\} \in SE(3)$ (cf. Eq. (3.8) – Eq. (3.9)), we can model the mathematical relationship between camera projections from different points of view. For describing the required theoretical fundamentals, we refer to the definitions in Becker et al. (2014), Faugeras and Luong (2001, ch. 5), and Hartley and Zisserman (2000, ch. 9).

Two-View Geometry

We assume that two cameras are oriented towards a 3D scene according to Fig. 3.2 (a). The positioning of the first camera follows the canonical mathematical form in Eq. (3.4). Focusing on a point x in an image plane W taken by this camera, the respective camera ray intersects with the unknown scene point \mathcal{X} . It projects in the image W' taken by the second camera to the line l' called *epipolar line*. Somewhere on l' , the equivalent x' of x is located. Translating this initial statement into a homogeneous coordinate environment, the epipolar line $\tilde{l}' \in \mathbb{P}^2$ associated with $l' \in \mathbb{R}^2$ can be defined through $\tilde{x}' \in \mathbb{P}^2$ associated with $x' \in \mathbb{R}^2$ by solving the linear constraint

$$\langle \tilde{x}', \tilde{l}' \rangle = \sum_{i=1}^3 \tilde{x}'_i \tilde{l}'_i = 0. \quad (3.12)$$

Eq. (3.12) is not limited to epipolar geometry only. In a given projection space, it is universally applicable to determine a line based on the homogeneous coordinates of one of the points on this line.

Concerning the positions and configurations of the two cameras, we assume an internal calibration matrix K for the first one. The second camera is assumed to be located in a distance h from the first one, rotated by Λ , and provided with the internal calibration matrix K' . Its projection matrix P' thus conforms to

$$P' = K' \Lambda^\top (I_{3 \times 3}, -h).$$

Fixing the image point $x \in W$ taken by the first camera and regarding its homogeneous coordinates $\tilde{x} \in \mathbb{P}^2$, the projection of the original space point \mathcal{X} to the second image point $\tilde{x}' \in W'$ with $\tilde{x}' \in \mathbb{P}^2$ lies on the epipolar line l' . For




Figure 3.2: Perspective projections of a 3D scene point \mathcal{X} to coordinates x and x' in the 2D camera planes W and W' . (a) refers to the positioning of two different cameras with respect to \mathcal{X} . It illustrates the relations between the two mappings by virtue of the epipolar lines l and l' and the epipoles e and e' . (b) refers to the movement of one single camera. It visualizes the so-called 3D epipolar plane \mathcal{E} relating \mathcal{X} (and any other scene point on \mathcal{E}) to the two image planes W and W' recorded before and after the re-positioning of the camera by $\{h, \Lambda\} \in SE(3)$. \mathcal{E} intersects with W and W' in the epipolar lines l and l' containing the point projections x and x' as well as the foci of expansion x_e and x'_e .

the homogeneous form \tilde{l}' of l' , it holds that

$$\tilde{l}' = F \tilde{x}, \quad (3.13)$$

$$\text{where } F = K'^{-\top} \Lambda^\top [h]_\times K^{-1} \quad \text{with} \quad [h]_\times = \begin{pmatrix} 0 & -h_3 & h_2 \\ h_3 & 0 & -h_1 \\ -h_2 & h_1 & 0 \end{pmatrix}$$

is known as the *fundamental matrix* linking both image planes to each other. Inversely,

$$\tilde{l} = F^\top \tilde{x}'. \quad (3.14)$$

Looking at Fig. 3.2 (a) once again, we shortly characterize the projections of camera center, i.e. the mappings of e on l and e' on l' . The respective homogeneous coordinates \tilde{e} and \tilde{e}' can be computed via

$$\tilde{e} = K h \quad \text{and} \quad \tilde{e}' = K' \Lambda^\top h.$$

Therefore, Eq. (3.13) – Eq. (3.14) yield

$$F\tilde{e} = 0 \quad \text{and} \quad F^\top \tilde{e}' = 0 .$$

We denote e and e' as the *epipoles* related to \tilde{x} and \tilde{x}' . As a result of Eq. (3.12) and Eq. (3.13), it follows that

$$\langle \tilde{x}', F\tilde{x} \rangle = 0 . \quad (3.15)$$

Eq. (3.15) is the core equation connecting any two views of an unidentified scene point \mathcal{X} with each other. In order to work with normalized coordinates, this relation has to be modified properly. More precisely, referring to the definition of F in Eq. (3.13), the adjusted scalar product takes the form

$$\langle \tilde{x}'^c, F\tilde{x}^c \rangle = \langle K'^{-1}\tilde{x}', K'^\top FK(K^{-1}\tilde{x}) \rangle = \langle K'^{-1}\tilde{x}', \Lambda^\top[h]_\times(K^{-1}\tilde{x}) \rangle . \quad (3.16)$$

We see that, using calibrated coordinates, the required adjustment of the fundamental matrix F results in a matrix of the compact form $E := K'^\top FK = \Lambda^\top[h]$, called *essential matrix*. E is parameterized with $\{h, \Lambda\} \in SE(3)$, as the calibration parameters contained in K and K' are assumed to be known. Essential matrices hence correspond to smooth manifolds in $\mathbb{R}^{3 \times 3}$.

In the following section, we explain how the motion of a camera can be modeled. That is, instead of considering two different cameras directed towards a 3D scene (i.e. $K' \neq K$), we notionally change the position as well as the angle of one single camera, and model the interrelations between the resulting scene projections. Alternatively, we could think of two cameras with exactly identical internal configuration parameters (i.e. $K' = K$). The structural arrangement of our definitions and explanations complies with the respective section in Becker et al. (2014), and Hartley and Zisserman (2000, ch. 6).

Assignment Fields

We henceforth take the calibration matrix K of the camera used as given and entirely known, which is why we can easily act on the basis of normalized coordinates. To simplify our notation, we leave out the upper index c employed in Eq. (3.6) for symbolizing the calibration. As all camera-specific parameters are fixed and normalized coordinates are hence used, we set $K = I$.

In line with the previous section, we let x denote the mapping of a 3D scene point \mathcal{X} to a 2D image plane W . The camera used is moved from its original position by h and rotated by $\Lambda = (\Lambda^1, \Lambda^2, \Lambda^3)$, where Λ^i corresponds to the i -th column of Λ . This displacement yields the projection of \mathcal{X} to the point x' in the new observation window W' (see Fig. 3.2 (a)).

We describe the linking between x and x' through the linear equation

$$x' = x + a(x) , \quad (3.17)$$

and call $a(\cdot)$ the underlying *assignment vector*. Referring to Eq. (3.2) and Eq. (3.9), it takes the form

$$a(x) = \frac{1}{\langle \Lambda^3, \mathcal{X} - h \rangle} \begin{pmatrix} \langle \Lambda^1, \mathcal{X} - h \rangle \\ \langle \Lambda^2, \mathcal{X} - h \rangle \end{pmatrix} - \frac{1}{\mathcal{X}_3} \begin{pmatrix} \mathcal{X}_1 \\ \mathcal{X}_2 \end{pmatrix}. \quad (3.18)$$

In case of initiating a pure translation (i.e. $\Lambda = I_{33}$), Eq. (3.18) simplifies to

$$\begin{aligned} a(x) &= \frac{1}{\mathcal{X}_3 - h_3} \begin{pmatrix} \mathcal{X}_1 - h_1 \\ \mathcal{X}_2 - h_2 \end{pmatrix} - \frac{1}{\mathcal{X}_3} \begin{pmatrix} \mathcal{X}_1 \\ \mathcal{X}_2 \end{pmatrix} \\ &\stackrel{\bar{h} := \frac{1}{\mathcal{X}_3} h}{=} \frac{1}{\bar{h}_3 - 1} \left[\begin{pmatrix} \bar{h}_1 \\ \bar{h}_2 \end{pmatrix} - \bar{h}_3 \begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \right]. \end{aligned} \quad (3.19)$$

The point x_e rendering $a(x_e) = 0$ is known as the *focus of expansion*. It conforms to

$$x_e = \frac{1}{\bar{h}_3 - 1} \begin{pmatrix} \bar{h}_1 \\ \bar{h}_2 \end{pmatrix}$$

and equals the epipole $x = e$ discussed previously, that yields $F\tilde{e} = 0 \simeq \Lambda^\top[h]_\times h = 0$ for homogeneous coordinates \tilde{e} (cf. Eq. (3.1)).

Having dealt with relating the projections x and x' of a point \mathcal{X} in an arbitrary 3D scene to each other, we go over to focusing on the special case where the recorded scene is a 3D plane. For this purpose, we first give a reminder of the general definition of planes in a three-dimensional environment. Similarly to vector representations in 2D spaces, the dot product $\langle \boldsymbol{\delta}, \mathcal{X} \rangle = \sum_{i=1}^3 \delta_i \mathcal{X}_i$ goes into the equation of the plane. For the purpose of identifiability and by reason of convention, we postulate that $\boldsymbol{\delta}$ is a unit normal, i.e. $\|\boldsymbol{\delta}\| = 1$. We furthermore need to account for the distance d between the plane and the origin. The entire plane equation finally takes the form

$$\langle \boldsymbol{\delta}, \mathcal{X} \rangle - d = 0 \quad \text{or} \quad \langle \tilde{\boldsymbol{\delta}}, \tilde{\mathcal{X}} \rangle = 0, \quad (3.20)$$

where the second version is the homogeneous coordinate equivalent to the first one, implying that $\tilde{\boldsymbol{\delta}} = (\delta_1, \delta_2, \delta_3, -d)^\top$.

As just stated, we notionally let a camera take two images of one and same scene, but from different positions and angles. We associate the projection matrix $P = (I_{33}, 0)$ with the first situation, whereas $P' = \Lambda^\top(I_{33}, -h)$ relates to the second camera setting. For a moment, we only concentrate on the first setting. Given the homogeneous coordinate representation $\tilde{x} \in \mathbb{P}^2$ of the image point $x \in \mathbb{R}^2$ associated with the unknown scene point $\mathcal{X} \in \mathbb{R}^3$ with homogeneous coordinates $\tilde{\mathcal{X}} \in \mathbb{P}^3$, we can localize the scene point by solving the equation

$$\left\langle \tilde{\boldsymbol{\delta}}, \begin{pmatrix} \lambda \tilde{x} \\ 1 \end{pmatrix} \right\rangle = \langle \tilde{\boldsymbol{\delta}}, \tilde{\mathcal{X}}(\lambda) \rangle = \lambda \langle \boldsymbol{\delta}, \tilde{x} \rangle - d = 0. \quad (3.21)$$

In words, knowing that \mathcal{X} lies on the plane being projected (cf. Eq. (3.20)), we can determine its homogeneous coordinates by intersecting the plane with the ray $\tilde{\mathcal{X}}(\lambda)$, where $\lambda \in \mathbb{R}$. We obtain

$$\lambda = \frac{d}{\langle \boldsymbol{\delta}, \tilde{x} \rangle} \quad \text{and} \quad \tilde{\mathcal{X}} = \begin{pmatrix} \frac{d}{\langle \boldsymbol{\delta}, \tilde{x} \rangle} \tilde{x} \\ 1 \end{pmatrix} \simeq \begin{pmatrix} \tilde{x} \\ \frac{\langle \boldsymbol{\delta}, \tilde{x} \rangle}{d} \end{pmatrix}. \quad (3.22)$$

Having determined $\tilde{\mathcal{X}}$, it is straightforward to compute its mapping onto the second image plane which results from shifting and rotating the camera by $\{h, \Lambda\} \in SE(3)$ (cf. Fig. 3.2). The image point \tilde{x}' derived from the camera displacement is accessible via

$$\tilde{x}' = P' \tilde{\mathcal{X}}(\lambda) = \Lambda^\top \left(\tilde{x} - \frac{\langle \boldsymbol{\delta}, \tilde{x} \rangle}{d} h \right) = \Lambda^\top \left(I_{33} - \frac{h}{d} \boldsymbol{\delta}^\top \right) \tilde{x}. \quad (3.23)$$

In the literature, $\Lambda^\top (I_{33} - \frac{h}{d} \boldsymbol{\delta}^\top) := H$ is known as a *homography* of the projective plane \mathbb{P}^2 , induced by a camera moving relative to a 3D plane.

3.2 Detection of Local Symmetry Elements

This section has to be regarded independently of the previous one, since it is concerned with quite a different problem. It deals with the question of how to identify symmetry elements in the image of a partly or fully textured scene. In the following, we introduce a common strategy that allows us to transform images into probability maps, such that each pixel value becomes replaced by its probability of representing a local symmetry center. Beforehand, we provide the necessary mathematical backgrounds as described in the book by Jähne (1989, ch. 5).

Partial Derivatives and Some Related Operators

To turn texture into an arrangement of clearly separated symmetry elements, contrasts between potential symmetry centers and boundary regions need to be detected and accentuated. This is often performed by computing or approximating partial derivatives of the given image function. General theoretical background knowledge on partial derivatives and related mathematical topics is thus necessary for a better understanding of such image (pre-)processing procedures.

We start with the introduction of some notational conventions. First, we explain what is meant by a *monomial*. A monomial $\boldsymbol{x}^{\boldsymbol{v}}$ with respect to the *multi-index*

$$\boldsymbol{v} = (v_1, \dots, v_n) \in \mathbb{N}_0^n$$

is defined as

$$\boldsymbol{x}^{\boldsymbol{v}} = x_1^{v_1} \cdot x_2^{v_2} \cdots x_n^{v_n}, \quad (3.24)$$

where $\mathbf{x} \in \mathbb{R}^n$. It hence conforms to a product of powers of variables, where the exponents are arbitrary non-negative integers. Given $\mathbf{x} = \{x_1, \dots, x_n\}$, we denote the *partial derivative* with respect to the i -th component by

$$\mathcal{D}_i = \frac{\partial}{\partial x_i} . \quad (3.25)$$

Higher order partial derivatives with respect to the entire variable set or vector \mathbf{x} take the form

$$\mathcal{D}^{\mathbf{v}} = \mathcal{D}_1^{v_1} \dots \mathcal{D}_n^{v_n} = \frac{\partial^{|\mathbf{v}|}}{\partial x_1^{v_1} \dots \partial x_n^{v_n}} , \quad \text{with } |\mathbf{v}| = \sum_{i=1}^n v_i , \quad (3.26)$$

where $\frac{\partial^{v_i}}{\partial x_i^{v_i}}$ denotes the partial differentiation of order v_i with respect to x_i .

We now consider a function $f(\cdot)$, $f : \mathbf{x} \mapsto f(\mathbf{x}) \in \mathbb{R}$, which is differentiable with respect to all elements of \mathbf{x} . The vector of its component-wise derivatives,

$$\nabla f(\mathbf{x}) = (\mathcal{D}_1 f(\mathbf{x}), \dots, \mathcal{D}_n f(\mathbf{x}))^\top , \quad (3.27)$$

is termed the *gradient* of $f(\cdot)$. The matrix of the according second order derivatives, the so-called *Hessian matrix*, conforms to

$$\nabla^2 f(\mathbf{x}) = (\mathcal{D}_i \mathcal{D}_j f(\mathbf{x}))_{i,j} , \quad (3.28)$$

with $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, n\}$ labeling the rows and columns of $\nabla^2 f(\cdot)$. From the gradient of $f(\cdot)$, the *Laplacian* is deducible. It is defined as

$$\Delta f(\mathbf{x}) = \text{div}(\nabla f(\mathbf{x})) = \sum_{i=1}^n \mathcal{D}_i^2 f(\mathbf{x}) . \quad (3.29)$$

The operator $\text{div}(\cdot)$ in Eq. (3.29) measures the divergence of a differentiable vector function $g(\cdot)$, $g : \mathbf{x} \mapsto g(\mathbf{x}) \in \mathbb{R}^n$, and takes the form

$$\text{div}(g(\mathbf{x})) = \sum_{i=1}^n \mathcal{D}_i g_i(\mathbf{x}) ,$$

Regarding images as functions that map each pixel coordinate to a pixel value, the definitions introduced above can be used for purposes of image processing, synthesis and analysis. In the following, we explain how partial derivatives of image functions can be approximated with the help of suitable smoothers. For more details, see Mather and Koch (2010, ch. 7).

Estimating Partial Derivatives of Image Functions

To reduce image noise, a *smoothing kernel* $\xi(\cdot)$ is usually taken as a basis. In general, a properly designed smoothing kernel is an infinitely and continuously

differentiable function on \mathbb{R}^n with compact support [Tao (2011, ch. 1)], constructed such that its integral over \mathbb{R}^n equates to one. More formally,

$$\xi(x) \in \mathcal{C}_0^\infty(\mathbb{R}^n) \quad \text{and} \quad \int_{\mathbb{R}^n} \xi(w) dw = 1 .$$

Defining

$$\xi_\varepsilon(x) := \frac{1}{\varepsilon^n} \xi\left(\frac{x}{\varepsilon}\right) , \quad (3.30)$$

a function $f(\cdot)$ can be smoothened through $\xi_\varepsilon(\cdot)$ via a *convolution*

$$f_\varepsilon(x) := \xi_\varepsilon * f(x) := \int_{\mathbb{R}^n} \xi_\varepsilon(x - y) f(y) dy , \quad (3.31)$$

where $f_\varepsilon(\cdot)$ denotes the mollified version of $f(\cdot)$. Smoothing formulas according to Eq. (3.31) have some valuable mathematical properties.

Theorem 1 [Ziemer (1989)] *Suppose $\xi(\cdot)$ is a smoothing kernel and let $f_\varepsilon(\cdot)$ be given by Eq. (3.31).*

- (i) *If $f \in \mathcal{L}_{\text{loc}}^1(\mathbb{R}^n)$, then for every $\varepsilon > 0$, $f_\varepsilon(\cdot) \in \mathcal{C}^\infty(\mathbb{R}^n)$ and $\mathcal{D}^v f_\varepsilon = \mathcal{D}^v(\xi_\varepsilon * f) = (\mathcal{D}^v \xi_\varepsilon) * f$ for each multi-index v .*
- (ii) *If $f(\cdot) \in \mathcal{L}^p(\mathbb{R}^n)$, $1 \leq p < \infty$, then for all $x \in \mathbb{R}^n$, $f_\varepsilon(\cdot) \in \mathcal{L}^p(\mathbb{R}^n)$, $\|f_\varepsilon(x)\|_{\mathcal{L}^p} \leq \|f(x)\|_{\mathcal{L}^p}$, and $\lim_{\varepsilon \rightarrow 0} \|f_\varepsilon(x) - f(x)\|_{\mathcal{L}^p} = 0$.*

$\mathcal{C}^\infty(\mathbb{R}^n)$ contains all infinitely and continuously differentiable functions on \mathbb{R}^n , whereas $\mathcal{L}_{\text{loc}}^1(\mathbb{R}^n)$ and $\mathcal{L}^p(\mathbb{R}^n)$ denote the *Lebesgue spaces* of the locally integrable and, respectively, the p -th power integrable functions on \mathbb{R}^n [see also Tao (2011, ch. 1)]. An $\mathcal{L}^p(\cdot)$ function space comprises all measurable functions with a finite \mathcal{L}^p norm. Symbolized by $\|\cdot\|_{\mathcal{L}^p}$, the \mathcal{L}^p norm on \mathbb{R}^n is defined as

$$\|f\|_{\mathcal{L}^p} := \left(\int_{\mathbb{R}^n} |f(x)|^p dx \right)^{\frac{1}{p}} .$$

Despite having unbounded support, *derivative-of-Gaussian* (DoG) filters correspond to the most common class of *separable (partial) derivative filters* [see also Canny (1986), Mather and Koch (2010, ch. 7)]. As the name implies, DoGs are based on Gaussian kernels and thus of the form

$$\xi_\sigma(x) = \frac{1}{(2\pi\sigma^2)^{\frac{n}{2}}} \exp \left\{ -\frac{1}{2\sigma^2} \|x\|^2 \right\} . \quad (3.32)$$

Sketch: Edge Detection via DoG filters

We assume that we are given an image of a textured scene such as the brick wall in Fig. 3.3. The image matrix denoted by Z is presentable as a function $Z(\cdot)$ assigning one pixel value $z_w \in \mathbb{R}$ to each image coordinate $w = (w_1, w_2)^T \in W \subseteq \mathbb{R}^2$, i.e. $Z : w \mapsto Z(w) = z_w$. For identifiability reasons, we normalize the range of $Z(\cdot)$ to the $[0,1]$ -interval.



Figure 3.3: Original image of a brick wall.

The following two to three transformation stages can be executed to turn an image into a probability map via DoG filtering.

- (0.) Depending on the pixel value distribution in the image plane, it may be convenient to start with some preliminary image adjustment. *Median filters*, for example, are frequently used to even out outliers. They replace each measured pixel value by the median of all pixel values in its immediate surrounding and thus yield a mapping $w \mapsto \text{med}(z_{\sim w})$ where “ $\sim w$ ” symbolizes the neighborhood of w .

Fig. 3.4 results from the application of a median filter to Fig. 3.3. Here, each local neighborhood has been delimited by a circular disc of radius $r = 2$ pixels.

1. The actual smoothing is performed by estimating the partial derivatives of the original or preprocessed image function. As indicated above, it is convenient to use smoothing kernels such as the Gaussian kernel in Eq. (3.32). Referring to Eq. (3.27) – Eq. (3.29) as well as to Eq. (3.31) and item (i) of the subsequent theorem, an appropriately implemented Gaussian derivative filter yields a mapping

$$w \mapsto \sqrt{\left(\frac{\partial Z}{\partial w_1}\right)^2 + \left(\frac{\partial Z}{\partial w_2}\right)^2} = \sqrt{\mathcal{D}_1^2 Z(w) + \mathcal{D}_2^2 Z(w)} ,$$

where

$$\mathcal{D}_i Z(w) \approx (D_i \xi_\sigma) * Z(w) = -\frac{1}{2\pi\sigma^2} \sum_{y \sim w} z_y(w_i - y_i) \exp\left\{-\frac{1}{2\sigma^2} \|w - y\|^2\right\} .$$



Figure 3.4: Brick wall image (cf. Fig. 3.3) after median filtering. The range of the local neighborhoods has been determined by a circular disc of radius $r = 2$ pixels.

In this formal representation, “ $y \sim w$ ” indicates that y is a neighbor of w . A DoG filtering routine turns Fig. 3.4 into Fig. 3.5. Again, the radii of the neighborhoods have been set to $r = 2$ pixels. The standard deviation σ has been defined as $\sigma = \sqrt{\frac{r-1}{4}} = 1$.



Figure 3.5: Brick wall image (cf. Fig. 3.3) after median and DoG filtering. The range of the local neighborhoods has been determined by a circular disc of radius $r = 2$ pixels, and the standard deviation in the Gaussian kernels equates to $\sigma = \sqrt{\frac{r-1}{4}} = 1$.

2. To learn a probability map from a smoothened image, suitable distance transformations need to be carried out. It makes sense to first distinguish between fore- and background pixels, where the foreground pixels represent the inner parts and the background pixels the boundary regions of the local symmetry elements contained in the image. This differentiation can be performed based on a carefully determined threshold value t_z which turns the smoothened image into a binary map. Afterwards, the shortest path between each single image location and the background pixels can be computed. A normalization of the resulting distances to values between

0 and 1 allows to directly interpret them as probability measures. Each probability measure thus indicates how likely it is that the corresponding image coordinate is the center of a symmetry element and not part of the boundary domain.

Fig. 3.6 shows a probability map resulting from a distance transformation of Fig. 3.5, where $t_z = 0.02$.




Figure 3.6: Distance transformation of a brick wall image (cf. Fig. 3.3) after median and DoG filtering. The radius of the local neighborhoods has been set to $r = 2$ pixels, the standard deviation in the Gaussian kernels to $\sigma = \sqrt{\frac{r-1}{4}} = 1$, and the binarization threshold to $t_z = 0.02$.

The step sequence in 0. – 2. can flexibly be extended, elaborated, reversed or replaced by other smoothing and distance transforming techniques. Distances between local histograms, for instance, are frequently used to detect similar features and edges in a given image, or to compare two images with each other. In the pattern recognition community, the earth-mover’s distance [Pele and Werman (2009)] is among the most common tools for comparing color histograms.

In Chapter 6, we regard texture as a realization of a latent locally scaled point process, meaning that each texture element is associated with one point. For the point-texture assignment and hence for the estimation of the latent point process, we take preprocessed images such as the probability map in Fig. 3.6 as a basis.

Before establishing a connection between image analysis and spatial statistics, we provide a general introduction to the theory of locally scaled point processes.

4 Locally Scaled Spatial Point Processes

Chapter 2 is concerned with a comprehensive presentation of the most widely recognized classes of spatial point process models. It differentiates between processes with a homogeneous point intensity and processes with heterogeneously distributed points. Furthermore, patterns with independently distributed points are distinguished from patterns showing point interactions in the form of pairwise repulsions or clustering effects. What is missing in the introductory chapter, however, is the conjunction of inhomogeneity assumptions with interaction constraints. Research in this particular statistical field appears challenging, but indispensable due to its practical relevance.

Hahn et al. (2003) provide a point process approach allowing to model so-called *locally scaled spatial point patterns*. In contrast to formerly developed procedures inducing heterogeneity into a Markov model – such as the algorithms discussed by Ogata and Tanemura (1986), Baddeley et al. (2000), or Jensen and Nielsen (2001) –, locally scaled point process models due to Hahn et al. ensure that the variation in interaction is adapted to the variation in intensity. That is, local geometry remains unaffected by inhomogeneity, which makes regions with a high point intensity look like scaled versions of regions where the points are less densely packed. Locally scaled point processes, in fact, correspond to Markov processes that appear homogeneous up to a scale factor. A larger scale factor results in a lower point intensity and stricter interaction constraints, whereas the opposite effect occurs when the scale factor takes smaller values.

It is not astonishing that the development of modeling approaches to locally scaled spatial point patterns is of high relevance in terms of real-world applications. Thinking of forestry, for instance, the number of trees per unit area inter alia depends on the sizes of the trees and their root networks, and vice versa [see e.g. Shimatani and Kubota (2004), Eckel et al. (2009)]. Therefore, arbitrarily big trees cannot stand arbitrarily close to each other. Imagining a “regular forest” viewed from a bird’s-eye perspective, regions with broader tree silhouettes look like zoom-ins of regions where the trees are narrower.

In what follows, we first give a definition of inhomogeneous point processes by location-dependent scaling, particularly focusing on processes of Strauss type (see Section 4.1). The characteristics of a proper scaling function are discussed in Section 4.2. Building on these definitions, we motivate and introduce an efficient Bayesian inference framework for locally scaled spatial point processes in Section 4.3. As this framework is based on the enhancement of the Metropolis-Hastings (M-H) algorithm by an auxiliary variable scheme according to Alg. 2.7

(see Section 2.5), we devote Section 4.3 to the proposition of two appropriate techniques for generating the necessary auxiliary samples. Finally, Section 4.4 is concerned with a simulation study validating the performance of our approach. We also discuss the advantages of the Bayesian modeling framework, point out potential room for development and improvement, and sketch an alternative frequentist approach proposed by Prokešová et al. (2006).

4.1 Locally Scaled Point Process Models

As mentioned above, a locally scaled spatial point process is a point process with a varying point intensity and a conformably changing degree of point interaction. The location-dependent scaling properties thus make the process locally behave like the scaled version of a homogeneous template process. In terms of implementation, location-dependent scaling impacts on a homogeneous point process are obtained from a local scaling of all volume measures contained in its density. Therefore, the Markov property (cf. Eq. (2.8)) remains unaffected, and, locally confined, the Papangelou conditional intensity (cf. Eq. (2.21)) is proportional to that of the homogeneous template pattern [see also Hahn et al. (2003)]. The following definitions and descriptions are excerpted from the work of Hahn et al. (2003) as well as from the related discussions in Prokešová et al. (2006).

A prerequisite for a profound understanding of the motivation and the technical principles of locally scaled point processes consists in the familiarity with the concept of *scale invariance*. We assume a given a measure $\kappa(\cdot)$ on \mathbb{R}^d with σ -algebra \mathcal{B} . For any set $B \in \mathcal{B}$ associated with a scale factor c , a transformation of $\kappa(B)$ subject to c yields

$$\kappa_c(B) = \kappa(c^{-1}B) .$$

In this connection, a scale invariant function $h(\cdot|\boldsymbol{\kappa})$ with $\boldsymbol{\kappa} = (\kappa^{(1)}(\cdot), \kappa^{(2)}(\cdot), \dots)^\top$ can be characterized as follows.

Definition 4.1 A real-valued measurable function $h(\mathbf{X}|\boldsymbol{\kappa})$ on the sample space Ω is called *scale invariant* if, for all realizations \mathbf{x} of the random variable \mathbf{X} ,

$$h(c\mathbf{x}|\boldsymbol{\kappa}_c) = h(\mathbf{x}|\boldsymbol{\kappa}) ,$$

where $\boldsymbol{\kappa}_c = (\kappa_c^{(1)}(\cdot), \kappa_c^{(2)}(\cdot), \dots)^\top$ and $c > 0$.

The homogeneous point processes introduced in Section 2.2 – 2.3 have densities proportional to a scale invariant function $h(\cdot|\boldsymbol{\kappa})$. In this context, $\boldsymbol{\kappa}$ is defined as the set of the d^* -dimensional volume measures $\boldsymbol{\nu} = (\nu^0, \dots, \nu^d)^\top$ in \mathbb{R}^d , also known as *Hausdorff volume measures* of dimension d^* [see also Hausdorff (1918), Berger (2009, ch. 9)].

In line with the notation and definitions from Section 2, we consider a finite point process \mathbf{X} on a set $W \subset \mathbb{R}^2$. Note that, just as above, extensions to higher dimensions are possible though not discussed here [Hahn et al. (2003)]. We assume an inhomogeneous point configuration which we model by a *location-dependent scaling function* $c : \mathbb{R}^2 \mapsto \mathbb{R}_+$. Thus, $c(\cdot)$ does not necessarily take constant values in W . It has to be carefully selected or designed (see Section 4.2), such that it is possible to act on the d^* -dimensional volume measures ν^{d^*} in \mathbb{R}^2 , $d^* \in \{0, 1, 2\}$. We consider *locally scaled d^* -dimensional volumes* of the form

$$\nu_c^{d^*}(W) = \int_W c(y)^{-d^*} \nu^{d^*}(dy) \quad (4.1)$$

for all Borel sets $W \subseteq \mathbb{R}^2$ and $d^* \in \{1, 2\}$. Here, $\nu^1(\cdot)$ is the distance measure associated with the Euclidean norm, while $\nu^2(\cdot)$ is a rescaling of the usual two-dimensional Lebesgue measure. The third volume measure, $\nu^0(\cdot) = n(\cdot)$, is the counting measure and thus unaffected by the scaling function $c(\cdot)$. To ensure that $\nu_c^{d^*}(W) < \infty$ holds for $\nu^{d^*}(W) < \infty$, we require that $c(\cdot)$ has finite upper and lower bounds.

As commented on in view of Def. 4.1, homogeneous point process densities are equivalent to scale invariant functions $h(\cdot|\boldsymbol{\nu})$, up to some normalizing constant. We assume a homogeneous point process \mathbf{X} in $W \subseteq \mathbb{R}^2$ and a density $f_{\mathbf{X}}(\cdot)$ with respect to the unit rate Poisson process (cf. Def. 2.9). For each realization $\mathbf{x} \subset W$, we can evaluate

$$f_{\mathbf{X}}(\mathbf{x}) \propto h(\mathbf{x}, \boldsymbol{\nu}) . \quad (4.2)$$

Definition 4.2 A *locally scaled point process* \mathbf{X}_c in $W^{(c)} \subseteq \mathbb{R}^2$ with homogeneous template model $f_{\mathbf{X}}(\cdot)$ due to Eq. (4.2) has a density

$$f_{\mathbf{X}_c}^{(c)}(\mathbf{x}) \propto h(\mathbf{x}|\boldsymbol{\nu}_c)$$

relative to a Poisson process with the scaled volume measure $\nu_c^2(\cdot)$ as its intensity. To guarantee numerical accessibility, $h(\cdot|\boldsymbol{\nu}_c)$ must be integrable subject to the Poisson process with intensity $\nu_c^2(\cdot)$, and $c : \mathbb{R}^2 \mapsto \mathbb{R}_+$ must be (Borel-)measurable.

A common assumption is that $W^{(c)} = W$, allowing the conclusion that the density of the locally scaled point process \mathbf{X}_c with respect to the unit rate Poisson process conforms to

$$\begin{aligned} f_{\mathbf{X}_c}(\mathbf{x}) &= \exp \left\{ - \int_W [c(y)^{-2} - 1] \nu^2(dy) \right\} \prod_{x \in \mathbf{x}} c(x)^{-2} f_{\mathbf{X}_c}^{(c)}(\mathbf{x}) \\ &\propto \prod_{x \in \mathbf{x}} c(x)^{-2} f_{\mathbf{X}_c}^{(c)}(\mathbf{x}) . \end{aligned} \quad (4.3)$$

Considering a parametric setting where $\boldsymbol{\eta}$ denotes the vector of parameters contained in the scaling function, i.e. $c(\cdot) = c_{\boldsymbol{\eta}}(\cdot)$, and $\boldsymbol{\theta}$ comprises all model

parameters, we can rewrite Eq. (4.3) as

$$f(\mathbf{x}|\boldsymbol{\theta}) := f_{\mathbf{X}_{c_\eta}}(\mathbf{x}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta})^{-1} \prod_{\mathbf{x} \in \mathbf{x}} c_\eta(x)^{-2} f^{(c)}(\mathbf{x}|\boldsymbol{\theta}), \quad (4.4)$$

where $f^{(c)}(\mathbf{x}|\boldsymbol{\theta}) := f_{\mathbf{X}_{c_\eta}}^{(c)}(\mathbf{x}|\boldsymbol{\theta})$ and $\mathcal{Z}(\cdot)^{-1}$ is the normalizing constant formally defined according to Eq. (2.10).

In this chapter, we concentrate on *locally scaled Strauss models*. The locally scaled version of a homogeneous Strauss process density on $W \subseteq \mathbb{R}^2$ [Strauss (1975)] is of the form

$$f^{(c)}(\mathbf{x}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta})^{-1} \beta^{n(\mathbf{x})} \gamma^{s_{c,R}(\mathbf{x})},$$

where

$$s_{c,R}(\mathbf{x}) = \sum_{\{u,v\} \subseteq \mathbf{x}}^{\neq} \mathbb{1}[\nu_c^1([u,v]) \leq R] \quad (4.5)$$

counts the number of distinct point pairs in \mathbf{x} no further than R apart, based on the scaled distance measure $\nu_c^1(\cdot)$. Thus, $\nu_c^1([u,v])$ represents the scaled length of the line segment $[u,v]$ between u and v . We see that, in terms of notation, a locally scaled Strauss model differs from a homogeneous Strauss model according to Eq. (2.11) only in the additional scaling constraints imposed through $c(\cdot)$ and η . Relative to the unit rate Poisson process and hence following Eq. (4.4), a Strauss model with local scaling attributes takes the form

$$f(\mathbf{x}|\boldsymbol{\theta}) = \mathcal{Z}(\boldsymbol{\theta})^{-1} \left[\prod_{\mathbf{x} \in \mathbf{x}} c_\eta(x)^{-2} \right] \beta^{n(\mathbf{x})} \gamma^{s_{c,R}(\mathbf{x})}. \quad (4.6)$$

As discussed in Section 2.3, the generic Strauss model has three parameters, an intensity parameter $\beta > 0$, an interaction parameter $0 \leq \gamma \leq 1$, and an interaction radius $R \geq 0$ [Strauss (1975), Kelly and Ripley (1976)]. The interaction parameter γ controls the amount of repulsion in the model, which ranges from $\gamma = 1$ indicating no repulsion and hence resulting in a completely random Poisson pattern, to a hardcore model for $\gamma = 0$. As a quick reminder, a visual interpretation of the effect of γ can be gathered from Fig. 2.3.

In the following section, we first describe how a *proper scaling function* is constructed. Thereafter, we discuss the statistical characteristics of the exponential scaling function in more detail.

4.2 Proper Scaling Functions

As before, we act on a two-dimensional observation plane $W \subseteq \mathbb{R}^2$ in which we assume a locally scaled point process $\mathbf{X} \subset W$, with realizations denoted by \mathbf{x} . We focus on parametric scaling effects imposed through a scaling function $c_\eta(\cdot)$.

In general, determining $c_{\boldsymbol{\eta}}(\cdot)$ such that it is proper, and hence statistically and numerically valid, is not trivial. According to Hahn et al. (2003) and Prokešová et al. (2006), two important conditions must be satisfied.

On the one hand, $c_{\boldsymbol{\eta}}(\cdot)$ has to be presentable in a well-defined complete form. We follow Prokešová et al. and suggest, for identifiability, a normalization fulfilling

$$\int_W c_{\boldsymbol{\eta}}(y)^{-2} \nu^2(dy) = \nu^2(W) . \quad (4.7)$$

On the other hand, all scaled volume measures going into the model considered must be computable. A Strauss model due to Eq. (4.5), for instance, requires the calculation of scaled pairwise point distances. This means that, referring to Eq. (4.1),

$$v_c^1([u, v]) = \int_{[u, v]} c_{\boldsymbol{\eta}}^{-1}(y) \nu^1(dy)$$

needs to be available in closed form for all pairs $\{u, v\} \in \boldsymbol{x}$. Applying the coarea formula [Krantz and Parks (2008, ch. 5)], an alternative representation of $\nu_c^1(\cdot)$ which often proves easier to handle amounts to

$$\nu_c^1([u, v]) = \nu^1([u, v]) \int_0^1 c_{\boldsymbol{\eta}}^{-1}(u + t(v - u)) dt . \quad (4.8)$$

A particularly appealing class of scaling functions is the class of the *exponential scaling functions*,

$$c_{\boldsymbol{\eta}}(u) = \alpha(\boldsymbol{\eta}) \exp\{\boldsymbol{\eta}^\top g(u)\} , \quad (4.9)$$

for some $g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, where $\alpha(\cdot)$ is a normalizing factor. Prokešová et al. (2006) consider the special case where $g(\cdot)$ is the identity function, i.e. $g(u) = u$. Fig. 4.1 illustrates the effect of γ in an exponentially scaled Strauss model (cf. Eq. (4.6)) where $c_{\boldsymbol{\eta}}(u) \propto \exp\{\gamma u\}$ and $\gamma > 0$, which yields a decreasing point intensity along the horizontal axis. The locally scaled distances are here given by

$$\nu_c^1([u, v]) = \nu^1([u, v]) \frac{c_{\boldsymbol{\eta}}(u)^{-1} - c_{\boldsymbol{\eta}}(v)^{-1}}{\boldsymbol{\eta}^\top (v - u)} ,$$

for $u \neq v$, which can easily be calculated.

It follows from Eq. (4.7) that the normalizing constant $\alpha(\cdot)$ depends on the observation window W . In a rectangular window $W = [w_l, w_r] \times [w_b, w_t]$ with left, right, top and bottom margins, w_l , w_r , w_b and w_t ,

$$\alpha(\boldsymbol{\eta}) = \frac{1}{\sqrt{|W|}} \left(\frac{\exp\{-2\eta_1 w_l\} - \exp\{-2\eta_1 w_r\}}{2\eta_1} \right)^{\frac{1}{2}} \left(\frac{\exp\{-2\eta_2 w_b\} - \exp\{-2\eta_2 w_t\}}{2\eta_2} \right)^{\frac{1}{2}}, \quad (4.10)$$

where $|W| = (w_t - w_b)(w_r - w_l)$, $\boldsymbol{\eta} \geq 0$, and $\frac{0}{0} := 1$. For $W = [0, 1] \times [0, 1]$, in particular, Eq. (4.9) simplifies to such an extent that the entire scaling term

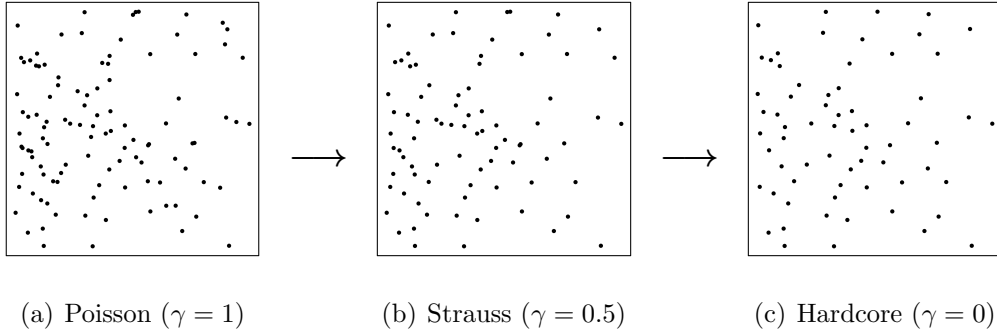



Figure 4.1: Simulations in $W = [0, 1] \times [0, 1]$ from locally scaled Strauss models with intensity $\beta = 100$, interaction range $R = 0.05$, different interaction weights γ , and an exponential scaling effect $\eta_1 = 1$ ($\eta_2 = 0$) according to Eq. (4.11). The point process realizations in (b) – (c) have been generated via dominated coupling from the past (see Section 4.3), yielding a progressive thinning of the template Poisson process in (a).

$c_{\boldsymbol{\eta}}(\cdot)$ takes the compact form


$$c_{\boldsymbol{\eta}}(u) = \left(\frac{1 - \exp\{-2\eta_1\}}{2\eta_1} \right)^{\frac{1}{2}} \left(\frac{1 - \exp\{-2\eta_2\}}{2\eta_2} \right)^{\frac{1}{2}} \exp\{\boldsymbol{\eta}^\top u\}. \quad (4.11)$$

Assuming that $\eta_{\bullet} > 0$ with \bullet standing for either element in $\{1, 2\}$, the respective exponential scaling effect is a denser point pattern close to the origin with increasing sparsity as the distance from the origin towards the direction represented by \bullet increases. The opposite effect is observable if $\eta_{\bullet} < 0$. If $\eta_{\bullet} = 0$, the point pattern is not scaled along \bullet , which is why the respective fraction in Eq. (4.9) is set to 1. Fig. 4.2 exemplifies the impact of different exponential scaling constraints on the point pattern itself as well as on the pairwise point distances.

We introduce two further proper specifications of $c_{\boldsymbol{\eta}}(\cdot)$ in the course of this work, one that imposes step-wise scaling constraints on a circular observation window (see Section 5.2), and another one that geometrically describes camera projections and thus induces a scaling of perspective (see Section 6.4). The next section is concerned with the proposition of a Bayesian inference concept for locally scaled spatial point processes. It is mainly based on the exchangeability M-H framework by Murray et al. (2012) and hence requires the implementation of a well-considered algorithm that generates auxiliary point data in an efficient and precise manner. A corresponding journal article will be submitted soon [Didden et al. (2015)].



(i) $\boldsymbol{\eta}^\top = (0, 0)$ (ii) $\boldsymbol{\eta}^\top = (0.5, 0.5)$ (iii) $\boldsymbol{\eta}^\top = (0.5, -1)$ (iv) $\boldsymbol{\eta}^\top = (-1, 1)$
 (a) Exponentially scaled distances from the point coordinates $(0.5, 0.5)^\top$.



(i) $\boldsymbol{\eta}^\top = (0, 0)$ (ii) $\boldsymbol{\eta}^\top = (0.5, 0.5)$ (iii) $\boldsymbol{\eta}^\top = (0.5, -1)$ (iv) $\boldsymbol{\eta}^\top = (-1, 1)$
 (b) Point process realizations corresponding to (a).

Figure 4.2: Exponential scaling effects of varying strength in $W = [0, 1] \times [0, 1]$. (a) exemplifies exponentially scaled distances under four different specifications of $\boldsymbol{\eta}$, whereby increasing darkness indicates decreasing distance. (b) shows four corresponding point process realizations sampled from a Strauss model with parameters $\beta = 100$, $\gamma = 0.5$, $R = 0.05$. The sampling has been performed via dominated coupling from the past (see Section 4.3).

4.3 Bayesian Inference Framework

As discussed in Section 2.3 and Section 2.5, interaction point process models suffer from the numerical intractability of their normalizing constant, which makes efficient and precise parameter estimation a challenging task. In the literature, several inference techniques are provided to handle the problem of an incomplete data model. The pseudo likelihood approach by Besag (1974), which we explain in Section 2.5 of this thesis (cf. Eq. (2.22)), is one of the most widely-used and common examples [see also Besag et al. (1982)]. It is the groundwork of diverse subsequent research studies [see e.g. Huang and Ogata (1999, 2002) and references therein]. A comprehensive toolbox with further frequentist solutions has been developed in the past decades. It provides several well-established inference strategies, such as approximate maximum likelihood and Monte Carlo likelihood techniques [see Ogata and Tanemura (1981, 1984), Geyer (1999)]. Walker (2011) proposes a latent variable procedure to overcome the necessity

to solve the integral that determines the normalizing constant. The so-called Gaussian process density sampler introduced by Murray et al. (2008) allows for Bayesian density estimation in a non-parametric manner [see also Adams et al. (2009)]. Focusing on point processes with location dependent scaling characteristics and an incomplete density according to Eq. (4.4), suitable frequentist inference concepts can be obtained from Ogata and Tanemura (1986), Nielsen and Jensen (2004), and Prokešová et al. (2006).

Estimators resulting from the maximization of a tractable adjustment of the likelihood, however, are not necessarily unbiased, which is one of the major motivations behind our proposition of an alternative Bayesian framework. Moreover, only few Bayesian approaches have been established so far, e.g. a computationally intensive Markov chain Monte Carlo (MCMC) routine based on importance sampling [Bognar (2005)], or a Markov approximation for hardcore Gibbs models [Rajala and Penttinen (2012)].

We go back to the locally scaled Strauss model on $W \subset \mathbb{R}^2$ introduced by means of Eq. (4.6). Given a suitable prior $p(\boldsymbol{\theta})$ for $\boldsymbol{\theta} = \{\beta, \gamma, R, \boldsymbol{\eta}\} \in \Theta$, the joint posterior under this model conforms to

$$p(\boldsymbol{\theta}|\mathbf{x}) = \mathcal{C}(\mathbf{x})^{-1} \times \mathcal{Z}(\boldsymbol{\theta})^{-1} \left[\prod_{\mathbf{x} \in \mathbf{x}} c_{\boldsymbol{\eta}}(\mathbf{x})^{-2} \right] \beta^{n(\mathbf{x})} \gamma^{s_{c,R}(\mathbf{x})} \times p(\boldsymbol{\theta}) ,$$

where

$$s_{c,R}(\mathbf{x}) = \sum_{\substack{\neq \\ \{u,v\} \subseteq \mathbf{x}}} \mathbb{1}[\nu_c^1([u,v]) \leq R] ,$$

and both $\mathcal{C}(\mathbf{x}) = \int_{\Theta} f(\mathbf{x}|\boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}$ and $\mathcal{Z}(\boldsymbol{\theta}) = \int_W f(\mathbf{x}|\boldsymbol{\theta}) d\mathbf{x}$ are not available in closed form (cf. Eq. (2.10) and Eq. (2.25)). Consequently, $p(\cdot|\mathbf{x})$ is doubly intractable (cf. Eq. (2.26)). To deal with this issue, we follow our sketch of the exchange algorithm by Murray et al. (2012) in Alg. 2.7. Unless the context otherwise requires, we start from the simplifying assumption that the parameters are uncorrelated in the prior, meaning that

$$p(\boldsymbol{\theta}) = \prod_{\theta \in \boldsymbol{\theta}} p_{\theta}(\theta) .$$

For each parameter $\theta \in \boldsymbol{\theta}$, we define an appropriate transition density $q_{\theta}(\cdot|\theta)$ from which we can directly draw proposals θ' . We denote the modified set of parameters by $\boldsymbol{\theta}' = \left\{ \{\boldsymbol{\theta} \setminus \{\theta\}\} \cup \{\theta'\} \right\}$. More precisely, we e.g. simulate $\theta' = \beta'$ from $q_{\beta}(\cdot|\beta)$ and write $\boldsymbol{\theta}' = \{\beta', \gamma, R, \boldsymbol{\eta}\}$. Along with each parameter proposal, an auxiliary point process realization \mathbf{w} is drawn from $f(\mathbf{x}|\boldsymbol{\theta}')$. The M-H ratio with auxiliary variable extension due to Eq. (2.27) therefore takes the general form

$$R_{\text{eMH}} = \frac{\mathcal{C}(\mathbf{x})^{-1} \mathcal{Z}(\boldsymbol{\theta}')^{-1} f^*(\mathbf{x}|\boldsymbol{\theta}') p_{\theta}(\theta')}{\mathcal{C}(\mathbf{x})^{-1} \mathcal{Z}(\boldsymbol{\theta})^{-1} f^*(\mathbf{x}|\boldsymbol{\theta}) p_{\theta}(\theta)} \frac{q_{\theta}(\theta|\theta')}{q_{\theta}(\theta'|\theta)} \frac{\mathcal{Z}(\boldsymbol{\theta})^{-1} f^*(\mathbf{w}|\boldsymbol{\theta})}{\mathcal{Z}(\boldsymbol{\theta}')^{-1} f^*(\mathbf{w}|\boldsymbol{\theta}')} ,$$

for $\theta \in \{\beta, \gamma, R, \boldsymbol{\eta}\}$. Reverting to the exemplary case where $\theta = \beta$, we obtain

$$\begin{aligned} R_{\text{eMH}} &= \frac{\left[\prod_{x \in \mathbf{x}} c_{\boldsymbol{\eta}}(x)^{-2} \right] \beta'^{n(\mathbf{x})} \gamma^{s_{c,R}(\mathbf{x})} p_{\beta}(\beta')}{\left[\prod_{x \in \mathbf{x}} c_{\boldsymbol{\eta}}(x)^{-2} \right] \beta^{n(\mathbf{x})} \gamma^{s_{c,R}(\mathbf{x})} p_{\beta}(\beta)} \frac{q_{\beta}(\beta|\beta')}{q_{\beta}(\beta'|\beta)} \frac{\left[\prod_{w \in \mathbf{w}} c_{\boldsymbol{\eta}}(w)^{-2} \right] \beta^{n(\mathbf{w})} \gamma^{s_{c,R}(\mathbf{w})}}{\left[\prod_{w \in \mathbf{w}} c_{\boldsymbol{\eta}}(w)^{-2} \right] \beta'^{n(\mathbf{w})} \gamma^{s_{c,R}(\mathbf{w})}} \\ &= \frac{\beta'^{n(\mathbf{x})}}{\beta^{n(\mathbf{x})}} \frac{p_{\beta}(\beta')}{p_{\beta}(\beta)} \frac{q_{\beta}(\beta|\beta')}{q_{\beta}(\beta'|\beta)} \frac{\beta^{n(\mathbf{w})}}{\beta'^{n(\mathbf{w})}}. \end{aligned}$$

The determination of the prior and transition densities usually appears as a minor matter. Nevertheless, clever decisions reduce computing time and increase efficiency. Therefore, it is important to act on distributional conditions that are easily tractable, not too restrictive and compatible with the domain of the parameters. Generating the auxiliary variables \mathbf{w} , in contrast, proves to be a major challenge. For this reason, we devote the following section to the discussion of a suitable perfect as well as an alternative MCMC-based simulation technique.

As mentioned in the introductory sections, we first and foremost use an appropriate implementation of the *dominated coupling from the past (CFTP)* algorithm by Berthelsen and Møller (2002, 2003) for drawing realizations from interaction point process distributions with or without local scaling attributes. Depending on the strength of the intensity, interaction and scaling effects, this perfect simulation routine might require a high computational effort and thus be very time-consuming. Therefore, we also discuss an alternative MCMC method, the *approximate double M-H sampler* by Liang (2010). As the name indicates, this sampling technique introduces additional inaccuracy into the entire exchangeability framework, which has to be regarded critically.

Dominated Coupling from the Past (CFTP)

Based on a conjunction of path sampling and perfect simulation, CFTP yields exact drawings from pairwise interaction point processes. Berthelsen and Møller (2002, 2003) introduce this method in different versions. To simulate from a locally scaled Strauss model, we choose a dominated CFTP sampler that is based on the construction of three types of auxiliary processes relative to the unit rate Poisson process, a starting backward birth-death process as well as an upper and a lower forward process. Because of the complex and complicated nature of the proposed path sampling procedure, the subsequent paragraphs first describe its step sequence in a very general manner and then provide a more precise sketch of the algorithm. For the sake of completeness, the key stages are once again summarized in Alg. 4.1.

The starting backward process serves as a template without interaction penalty. Depending on the target model, its states are either homogeneous or inhomogeneous point process realizations. In the latter case, it is important that the

heterogeneity is imposed through a location-dependent scaling factor or a function ensuring local stability. The upper and lower forward processes step-wisely impose interaction restrictions on the independent point locations resulting from the backward birth-death process. As soon as both limiting processes converge, they merge together and return realizations from the target interaction model.

For a more detailed formal description of the dominated CFTP framework with upper and lower limiting processes, we denote the backward states of the homogeneous or inhomogeneous Poisson template by $D^{(t)}$, $t = 0, \dots, -T$ and the target interaction process by \mathbf{X} . The sequences of the upper and lower processes are denoted by $U^{(t)}$ and $L^{(t)}$, where $t \in \{-T, \dots, 0\}$. As explained above, $D^{(t)}$ is developed based on a birth-death step sequence. The courses of $U^{(t)}$ and $L^{(t)}$ reversely depend on the point in- and decreases in $D^{(t)}$ as well as on the specification of the interaction constraints in \mathbf{X} . When both sequences coincide and merge to one single chain, every state of this chain can be regarded as a simulation from \mathbf{X} , as already stated.

For a locally scaled Strauss model on $W \subset \mathbb{R}^2$ with parameter set $\boldsymbol{\theta} = \{\beta, \gamma, R, \boldsymbol{\eta}\}$, intensity $\beta(\cdot) := \beta c_{\boldsymbol{\eta}}(\cdot)^{-2}$ and density $f(\cdot | \boldsymbol{\theta})$ due to Eq. (4.6), we implement the described dominated CFTP routine as follows:

0. We preliminarily generate an inhomogeneous point process realization $D^{(0)}$ by location-dependent thinning of a homogeneous template with intensity $\beta^{(*)} = \sup_W \{\beta(\cdot)\}$ (see also Section 2.4 and references therein). $D^{(0)}$ consists of $n(D^{(0)})$ points.
1. Based on $D^{(0)}$ and the scaling function considered, we generate a birth-death chain of Poisson point process realizations $D^{(t)}$, $t \in \{1, \dots, -T\}$, backwards in time. Therefor, we first compute the expected point intensity with respect to the observation window W , $\mu(W) = \beta |W|$, and then execute the subsequent commands:

For $t \in \{-1, \dots, -T\}$ and based on the initial point configuration $D^{(0)}$,

- (1) with probability $\frac{\mu(W)}{\mu(W) + n(D^{(t+1)})}$, generate a new point $x_b^{(t)} \sim \beta(x)$ and set

$$D^{(t)} = \left\{ D^{(t+1)} \cup \{x_b^{(t)}\} \right\},$$

otherwise, delete a randomly selected point $x_d^{(t+1)}$ from $D^{(t+1)}$, such that

$$D^{(t)} = \left\{ D^{(t+1)} \setminus \{x_d^{(t+1)}\} \right\}.$$

- (2) draw a mark $m_{t+1} \sim U[0, 1]$.

2. The upper and lower processes $U^{(\cdot)}$ and $L^{(\cdot)}$ are first started at $j = 0$, then at $j = -1$, etc. The last possible starting time is $j = -T$. Developed in forward direction, sequences of the form $U^{(j)} = \{U_j^{(j)}, \dots, U_0^{(j)}\}$ and $L^{(j)} = \{L_j^{(j)}, \dots, L_0^{(j)}\}$ are thus constructed. The starting states are defined as $U_j^{(j)} = D^{(j)}$ and $L_j^{(j)} = \emptyset$, whereas the succeeding stages are derived from the birth-death template $D^{(\cdot)}$ as sketched in the following:

For $t = \{j + 1, \dots, 0\}$,

$$(1) \text{ if } D^{(t)} = \{D^{(t-1)} \setminus \{x_b^{(t)}\}\},$$

$$U_t^{(j)} = \{U_{t-1}^{(j)} \setminus \{x_b^{(t)}\}\} \text{ and } L_t^{(j)} = \{L_{t-1}^{(j)} \setminus \{x_b^{(t)}\}\},$$

$$\text{otherwise, if } D^{(t)} = \{D^{(t-1)} \cup \{x_d^{(t)}\}\},$$

$$U_t^{(j)} = \begin{cases} \{U_{t-1}^{(j)} \cup \{x_d^{(t)}\}\}, & \text{if } m_t \leq \prod_{x \in L_{t-1}^{(j)}} \gamma^{s_{c,R}}(\{x, x_d^{(t)}\}), \\ U_{t-1}^{(j)}, & \text{else.} \end{cases}$$

$$L_t^{(j)} = \begin{cases} \{L_{t-1}^{(j)} \cup \{x_d^{(t)}\}\}, & \text{if } m_t \leq \prod_{x \in U_{t-1}^{(j)}} \gamma^{s_{c,R}}(\{x, x_d^{(t)}\}), \\ L_{t-1}^{(j)}, & \text{else.} \end{cases}$$

[note that $s_{c,R}(\{x, x_d^{(t)}\})$ is defined according to Eq. (4.5)]

$$(2) \text{ if } U_t^{(j)} = L_t^{(j)}, \text{ stop! Regard } \mathbf{x} := L_t^{(j)} \text{ as a realization from } \mathbf{X}.$$

Concerning indexing and, in particular, the determination of the earliest time stamp $-T$, Berthelsen and Møller (2003) propose to proceed as follows:

$$(i) \ T_{\min} := \begin{cases} \inf\{-t : \{D^{(t)} \cap D^{(0)}\} \neq \emptyset, \{D^{(t-1)} \cap D^{(0)}\} = \emptyset\}, & \text{if } D^{(0)} \neq \emptyset, \\ 0, & \text{else.} \end{cases}$$

$$(ii) \ j_k = -2^{-k} T_{\min}, \quad k \in \{0, 1, 2, \dots\}$$

$$(iii) \ T = \inf\{-j_k : U_0^{(j_k)} = L_0^{(j_k)}\} \Rightarrow U_0^{(-T)} = L_0^{(-T)} = \mathbf{x} \sim f(\mathbf{X}|\boldsymbol{\theta}).$$

Based on this notational convention, Alg. 4.1 illustrates the structural design of the dominated CFTP algorithm in its general form.

Given a point pattern with a high packing density and strong interaction or pronounced scaling effects, dominated CFTP requires a lot of CPU time. Thanks to intensive research on alternative MCMC methods, however, approximate techniques have been elaborated that are almost as precise as exact sampling, but reduce time costs considerably and hence save CPU power.

Algorithm 4.1: Sketch: Dominated coupling from the past. For more details, see the comments above and the notational definitions in (i)–(iii).

Data: Intensity function $\beta(\cdot)$

Result: Point process realization \mathbf{x}

Generate $D^{(0)} \sim \beta(\cdot)$;

Set $k = 0$;

repeat

 set $k = k + 1$;

 develop backwards $D^{(j_{k-1}-1)}, \dots, D^{(j_k)}$;

for $t \in \{j_k, \dots, j_{k-1}\}$ **do**

if $\{D^{(t)} \setminus D^{(t-1)}\} \neq \emptyset$ **then**

 draw $m_t \sim U[0, 1]$;

end

end

 develop forwards $(U_{j_k}^{(j_k)}, L_{j_k}^{(j_k)}), \dots, (U_0^{(j_k)}, L_0^{(j_k)})$;

until $U_0^{(j_k)} = L_0^{(j_k)}$;

Set $-T := j_k$;

Return $\mathbf{x} := U_0^{(-T)}$ as a drawing from $f(\mathbf{x}|\boldsymbol{\theta})$;

Approximate Double M-H Sampler

Just to name a few examples of MCMC-based simulation under incomplete modeling assumptions, Liang et al. (2007), for instance, propose a so-called stochastic approximation Monte Carlo algorithm, whereas Jin and Liang (2012) suggest to embed this concept into a Bayesian framework. The double M-H approach by Liang (2010) connects the exchangeability principle to an approximate auxiliary sampling process that is based on sequences of suitably determined M-H transition kernels. Owing to the fact that we have already introduced an implementation of the exchange algorithm by Murray et al. (2012) to analyze repulsive point configurations, the double M-H sampler also building on the work by Murray et al. suggests itself as an alternative approximate method. The additional loss of accuracy, however, needs to be treated with caution. It mainly depends on the number of iterations of the double M-H sampler. Theoretically, infinitely many MCMC steps have to be conducted to guarantee convergence. A reduction in the number of iterations makes the optimization process less precise, even if it returns a Markov chain that appears reasonable and sufficiently

long. In general, every additional MCMC routine incorporated into an existing MCMC framework yields a decrease in precision and has to be judged carefully.

As we have provided a detailed description of the exchange algorithm in Section 2.5, we begin our sketch of the double M-H sampler where the problem of constructing and evaluating the auxiliary variable scheme comes up (see Alg. 2.7). That is, we explain how to generate a sample of point data, $\mathbf{w} \sim f(\cdot|\boldsymbol{\theta}')$, via the construction of a Markovian chain of states $\mathbf{w}^{(0)} \rightarrow \mathbf{w}^{(1)} \rightarrow \dots \rightarrow \mathbf{w}^{(K)} := \mathbf{w}$. Based on appropriately predefined transition densities $q_{\boldsymbol{\theta}'}(\mathbf{w}^{(k)}|\mathbf{w}^{(k-1)})$, $k \in \{1, \dots, K\}$, a series of conventional M-H steps yields

$$\begin{aligned} f_{\boldsymbol{\theta}'}^{(K)}(\mathbf{w}|\mathbf{w}^{(0)}) &= f_{\boldsymbol{\theta}'}^{(K)}(\mathbf{w}^{(K)}|\mathbf{w}^{(0)}) \\ &= q_{\boldsymbol{\theta}'}(\mathbf{w}^{(1)}|\mathbf{w}^{(0)}) \times \dots \times q_{\boldsymbol{\theta}'}(\mathbf{w}^{(K)}|\mathbf{w}^{(K-1)}) = f(\mathbf{w}|\boldsymbol{\theta}'). \end{aligned} \quad (4.12)$$

It is convenient to use uniform densities or flat truncated Gaussians as transitions densities. We propose to let the initial state $\mathbf{w}^{(0)} = \{w_1^{(0)}, \dots, w_K^{(0)}\}$ be similar to the observed point pattern or to the auxiliary data set associated with the latest parameter proposal accepted subject to Alg. 2.7. Since $\mathbf{w}^{(0)}$ consists of K elements, it determines the parameter K in Eq. (4.12).

For each point $w_k^{(k-1)}$, we first set $w_{k_0}^{(k)} := w_k^{(k-1)}$ and then successively propose L random moves $w_{k_l}^{(k)}|w_{k_{l-1}}^{(k)} \sim q(w_{k_l}^{(k)}|w_{k_{l-1}}^{(k)})$, $l \in \{1, \dots, L\}$, within the given observation window. In short, conditional on $\boldsymbol{\theta}'$, $L \times K$ M-H move steps transform $\mathbf{w}^{(0)}$ into $\mathbf{w}^{(K)} = \mathbf{w}$ via the intermediate stages

$$\begin{aligned} \mathbf{w}^{(k)} &= \left\{ \left\{ \mathbf{w}^{(k-1)} \setminus \{w_k^{(k-1)}\} \right\} \cup \{w_k^{(k)}\} \right\} : \\ &\quad \left. w_k^{(k)} = w_{k_L}^{(k)} \leftarrow w_{k_{L-1}}^{(k)} \leftarrow \dots \leftarrow w_{k_0}^{(k)} = w_k^{(k-1)} \right\}. \end{aligned}$$

The next section presents and discusses results from an analysis of simulated point patterns via a CFTP-based execution of the exchange algorithm. Our goal is to evaluate the overall performance of the proposed Bayesian inference framework before applying it to real databases. For comparison, we additionally analyze the synthetic data via a pseudo likelihood approach.

4.4 Simulation Study

We generate five independent point patterns with exponential scaling characteristics due to Eq. (4.9) from a locally scaled Strauss model as introduced by means of Eq. (4.6) (see Fig. 4.3). The true parameter configuration considered is $\{\beta, \gamma, r, \eta_1, \eta_2\} = \{100, 0.30, 0.05, 1.5, 0.75\}$.

Assuming that all parameters are unknown and starting with different initial estimates, we compute several Markovian chains using the exchange algorithm




Figure 4.3: Samples from an exponentially scaled Strauss model. The observation windows are of size $W = [0, 1] \times [0, 1]$, and the true model parameters $\{\beta, \gamma, R, \eta_1, \eta_2\}$ have been set to $\{100, 0.3, 0.05, 1.5, 0.75\}$.

(cf. Alg. 2.7) with auxiliary variable schemes constructed via dominated CFTP (cf. Alg. 4.1). As priors, we employ uniform densities on the intervals gray shaded in Fig. 4.4. The parameter proposals in the M-H step are grouped such that each subset, $\{\beta, \gamma\}$, $\{R\}$, and $\{\eta_1, \eta_2\}$, is being updated at a time. Every chain is of length 100.000. Regarding the first 50.000 states as burn-in samples and eliminating or reducing autocorrelation effects by accounting for every 100th realization only, we generate histograms visualizing the estimates of the full conditional posteriors $\hat{p}(\theta | \theta^{-\{\theta\}})$, for $\theta \in \{\{\beta, \gamma\}, \{R\}, \{\eta_1, \eta_2\}\}$ (see Fig. 4.4). Owing to the big lengths of the Markovian state sequences, differences in the shapes of the histograms are not visible when comparing results from different chains.

In consideration of the fact that we analyze a sample of size five only, it is not surprising that the modes and medians of some of the estimated full conditional posteriors eventually differ from the true effects. The small sample size makes unintentional random trends in the patterns notably affect the estimation of the posterior densities. Fig. 4.4 (i) provides clear evidence that the observed repulsion effect in the simulated data tends to be stronger than the true underlying value of γ .

Having shown that the CFTP-based exchange algorithm yields the expected results, we implement a frequentist inference strategy for comparison. We apply the two-stage optimizer by Prokešová et al. (2006, Ch. 4) to the exponentially scaled point process realizations illustrated in Fig. 4.3. Following Prokešová et al., we first assess the values of the scaling parameters $\boldsymbol{\eta}$ via convenient likelihood maximization [Aldrich (1997)]. This is possible, since the exponential scaling function is available in normalized form and therefore identifiable (cf. Eq. (4.9) – Eq. (4.11)). Conditioned on the resultant estimates $\hat{\boldsymbol{\eta}}$, a pseudo likelihood routine according to Eq. (2.22) is implemented for estimating the intensity parameter β and the weight of interaction γ [see also Besag (1974)]. The interaction radius R serves as a nuisance parameter which we optimize with respect to a fine grid of values on the interval $]0, 0.1]$.

From the first stage of the frequentist point process analysis, we obtain $\hat{\boldsymbol{\eta}} = \{1.46, 0.75\}$ as ML estimates of the scaling parameters. Conditioned on these

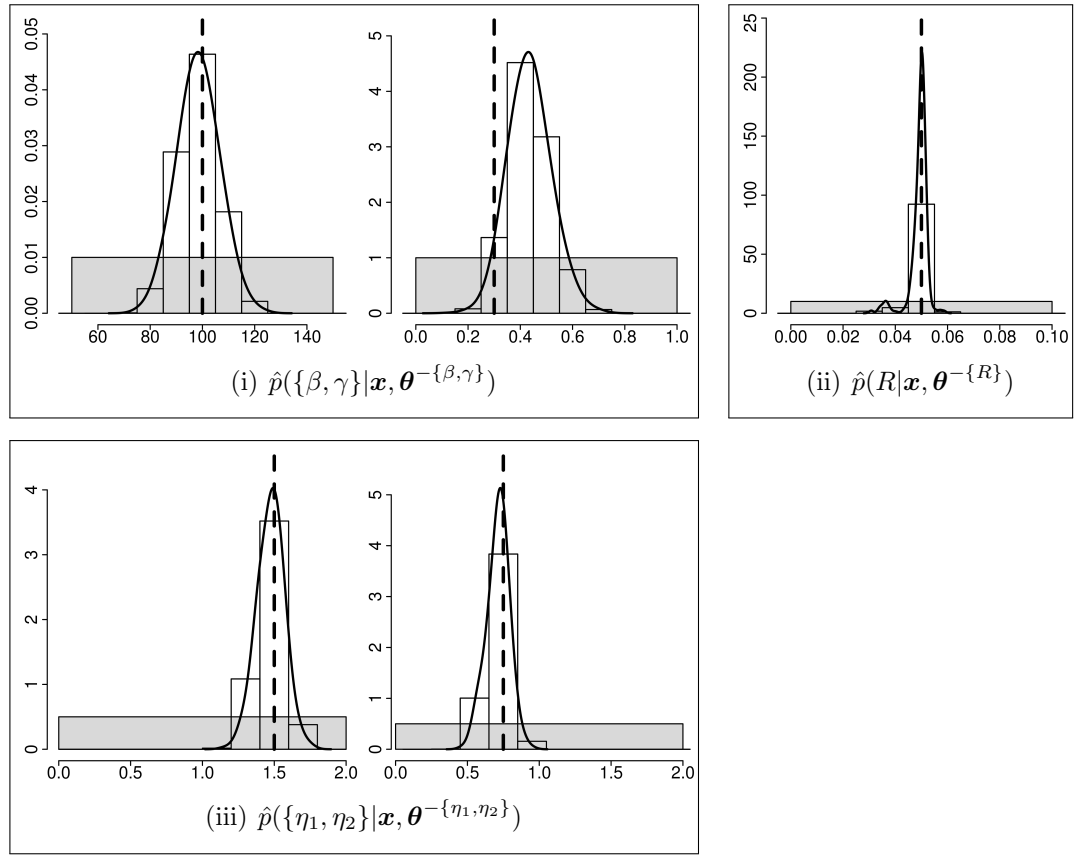


Figure 4.4: Full conditionals of the parameter subsets $\{\beta, \gamma\}$, $\{R\}$, and $\{\eta_1, \eta_2\}$. The dashed vertical lines mark the true values, and the horizontal gray bars the uniform parameter priors. The curves result from adjusting the full conditionals by a kernel density estimator [Silverman (1986, ch. 3 – 4)].




Figure 4.5: Optimization of the pseudo log-likelihood subject to the nuisance parameter R and the previously estimated scaling effects $\hat{\boldsymbol{\eta}} = \{1.46, 0.75\}$. The dash lines mark the true underlying value of R .

values, we calculate the maximum of the pseudo log-likelihood,

$$\hat{l}_P(\{\beta, \gamma\}|\{R, \hat{\eta}\}) := \max_{\{\beta, \gamma\}} \left\{ \log (L_P(\{\beta, \gamma\}|\{R, \hat{\eta}\})) \right\} ,$$

subject to every element of the grid considered for R , and plot the respective outcomes against R (see Fig. 4.5 (i)). Additionally, we graphically illustrate the ML estimates $\hat{\beta}$ and $\hat{\gamma}$ with respect to R (Fig. 4.5(ii) – (iii)),

$$\{\hat{\beta}, \hat{\gamma}\}|\{R, \hat{\eta}\} := \operatorname{argmax}_{\{\beta, \gamma\}} \left\{ \log (L_P(\{\beta, \gamma\}|\{R, \hat{\eta}\})) \right\}$$

The pseudo likelihood reaches its highest value under the parameter configuration $\theta_{ML} = \{\hat{\beta}, \hat{\gamma}, \hat{R}, \hat{\eta}_1, \hat{\eta}_2\} = \{99, 0.42, 0.0501, 1.46, 0.75\}$, whereas the estimated posterior resulting from our Bayesian approach takes its maximum under $\{98, 0.43, 0.0502, 1.49, 0.73\}$ and its mean under $\{98, 0.44, 0.0491, 1.47, 0.72\}$.

Summing up, we can state that, apart from a consistent data-driven overestimation of γ , both the frequentist approach by Prokešová et al. and our alternative Bayesian framework yield parameter estimates close to the true settings. There are only slight differences in the maximizers of the pseudo likelihood and the estimated posterior.

Since the CFTP-based exchange algorithm delivers the expected results, we now apply our Bayesian inference framework to real-world datasets where the ground truth is not known. We analyze cross-sections through maize stems in terms of their vascular bundle arrangement. The overall goal is a classification between two maize genotypes.

5 Analysis of Biological Image Data

Crop plants do not only serve as an important feeding stuff, but also as an essential basis for the generation of bio-ethanol and the development of organic production. After a series of mechanical and biological processes, raw material from the stems and leaves is turned into energy and fuel, respectively. Efficiency gains are achievable through an intelligent use and valorization of the agro-resources. Therefore, a well-founded understanding of the cellular plant structures is crucial and essential. We thank David Legland¹ and Marie-Françoise Devaux² for making data from different maize genotypes available to us.

Here, we examine cross-sections through two genotypes of maize stems in order to detect similarities and divergences in the spatial distributions of their vascular bundles (see Fig. 5.1). The positions of the bundles and their packing density have a strong impact on the water conductivity in the plants. Particularly for farmers and agricultural scientists, a sound answer to the question whether distinct maize genotypes differ in their water conductivity is of great importance and interest. Our work is similar to the research project by Legland et al. (2014). The authors propose a general framework for describing the cellular structures in the maize stems. It comprises preprocessing steps identifying the stem contours and the positions of the vascular bundles, a data normalization method, and the description of the bundle density and its variation by means of intensity maps.




Figure 5.1: Cellular structures in cross-sections through maize stems, where each small white patch represents one vascular bundle.

We explain below how the spatial arrangement of the vascular bundles in a maize stem can be modeled via a Bayesian locally scaled point process approach

¹INRA & AgroParisTech, *UMR 782 Food Process Engineering and Microbiology*, Thiverval-Grignon; INRA & AgroParisTech, *UMR 1318 Institut Jean-Pierre Bourgin*, Versailles

²INRA, UR 1268 *Biopolymers, Interactions and Assemblies*, Nantes

(cf. Section 4.1 – 4.3). To begin with, we normalize the given data appropriately (see Section 5.1) and introduce a feasible new scaling function (see Section 5.2). For the purpose of efficiency, we then propose a two-stage inference procedure (see Section 5.3). We visualize and discuss our results in Section 5.4 which concludes with a comprehensive and critical look at our Bayesian modeling framework. A journal article on Bayesian inference in locally scaled point processes, including the maize data problem, is in preparation [Didden et al. (2015)]

5.1 Data Preparation

As mentioned above, we concentrate on sets of maize data from two different genotypes, genotype I and genotype II. Our objective is to check whether the locally scaled Strauss model defined in Eq. (4.6) proves to be capable of detecting differences in the relative amount, the locations and the pairwise interactions of the vascular bundles. Per genotype, we are given three representative data realizations, that is, information from cross-sections through three different maize stems.

Thanks to the image preprocessing effort made by Legland et al. (2014), two data files describing each cross-section are at our disposal. One file contains the actual locations of the vascular bundles in the form of a point pattern $\mathbf{x}^{(0)} = \{x_1^{(0)}, \dots, x_n^{(0)}\}$ and the other file a fine grid of 200 to 300 boundary points manually determined and here denoted by $\mathbf{b}^{(0)} = \{b_1^{(0)}, \dots, b_{n_b}^{(0)}\}$ (see Fig. 5.2 (a) – (b) (i)). The data are arranged such that $b_c = (0, 0)^\top$ is the center of gravity.

We project the point coordinates, that is, the positions of the vascular bundles, into circular observation windows $D_{\mathcal{R}_I}(0)$ and $D_{\mathcal{R}_{II}}(0)$ of radii $\mathcal{R}_I = 1$ and $\mathcal{R}_{II} = 0.69$, respectively (see Fig. 5.2 (a) – (b) (ii)). The proportion between the radii reflects the proportion between the average diameters of the stems of both genotypes. Concerning the transformation process, we initially normalize each boundary point $b_i^{(0)}$, $i \in \{1, \dots, n_b\}$, by its distance from b_c which we denote by d_i . Thereafter, each inner point coordinate $x_j^{(0)}$, $j \in \{1, \dots, n\}$, is scaled by dividing the distance to its related boundary coordinate $b_i^{(0)}$ by d_i . We call an outer point $b_i^{(0)}$ the related boundary of $x_j^{(0)}$ if the acute angle between $x_j^{(0)}$, b_c and $b_i^{(0)}$ is smaller than the acute angle between $x_j^{(0)}$, b_c and any other outer point. This turns $\mathbf{b}^{(0)}$ into \mathbf{b} and $\mathbf{x}^{(0)}$ into \mathbf{x} , where \mathbf{b} and \mathbf{x} denote the normalized data. As already stated, both the original and the normalized point patterns are exemplified in Fig. 5.2.

At first view, we notice differences in the sizes of the stems. Moreover, we see that the number of the vascular bundles correlates positively with the stem diameters. The question comes up whether, irrespective of those absolute measures, the genotypes are classifiable subject to the relative spatial arrangements of their vascular bundles.

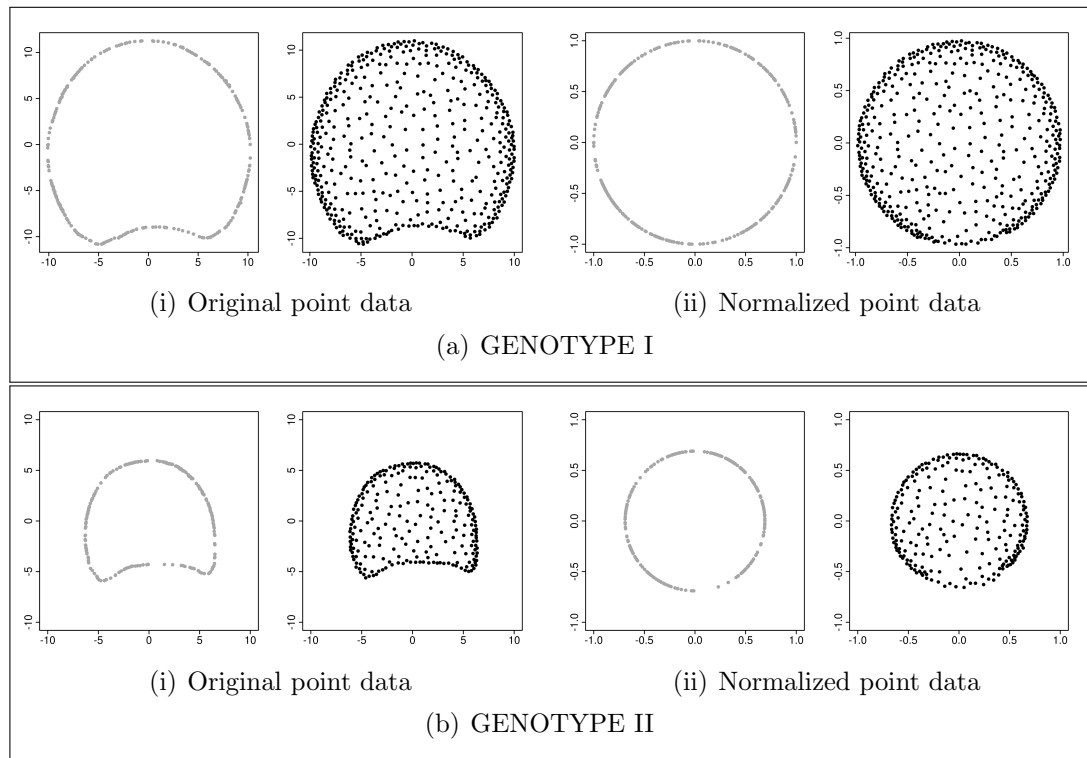


Figure 5.2: Cross-sections through representative maize stems before (i) and after (ii) a circular normalization. The gray points [l] serve as auxiliary boundary points, and the black points [r] show the observed (i) and normalized (ii) bundle locations.

Looking at Fig. 5.1 – 5.2 once again, it becomes obvious that the packing density of the vascular bundles is higher in the outer than in the inner parts of the stems. A closer inspection of the bundle arrangement allows for the assumption that within both the inner and the outer stem sections, the bundles are homogeneously distributed. For this reason, we have established a specific step scaling function which we introduce and explain in the next section (see in particular Eq. (5.1)). In accordance with our normalized maize data, it is defined on a circular two-dimensional observation plane.

5.2 Step-wise Scaling Constraints

We denote the circular observation window containing the point process data \mathbf{x} by $W = D_{\mathcal{R}}(b_c) \subset \mathbb{R}^2$, where \mathcal{R} stands for the disc radius and b_c , again, for the center of gravity. Without loss of generality, we start from the assumption that

$$W = D_{\mathcal{R}}(0) = \{u \in \mathbb{R}^2 : \|u\| \leq \mathcal{R}\} ,$$

where $\|\cdot\|$ denotes the Euclidean norm. The *step scaling function* is defined as

$$c_{\boldsymbol{\eta}}(u) = \alpha(\boldsymbol{\eta}) \left(\mathbb{1}[\|u\| \leq \eta_1] + \eta_2^{-\frac{1}{2}} \mathbb{1}[\eta_1 \leq \|u\| \leq \mathcal{R}] \right), \quad (5.1)$$

with normalizing constant

$$\alpha(\boldsymbol{\eta}) = \frac{1}{\mathcal{R}^2} (\eta_1^2 + \eta_2(\mathcal{R}^2 - \eta_1^2))^{\frac{1}{2}}. \quad (5.2)$$

Considering point data normalized to the unit disc $W = D_1(0)$, Eq. (5.1) – Eq. (5.2) simplify to

$$c_{\boldsymbol{\eta}}(u) = (\eta_1^2 + \eta_2(1 - \eta_1^2))^{\frac{1}{2}} \left(\mathbb{1}[\|u\| \leq \eta_1] + \eta_2^{-\frac{1}{2}} \mathbb{1}[\eta_1 \leq \|u\| \leq 1] \right). \quad (5.3)$$

Here, the step scaling function acts on two levels spatially delimited by η_1 and \mathcal{R} , the disc $D_{\eta_1}(0)$ for some $0 < \eta_1 < \mathcal{R}$ and the set $\{D_{\mathcal{R}}(0) \setminus D_{\eta_1}(0)\}$ which we abbreviate by $D_{\mathcal{R} \setminus \eta_1}(0)$. For the sake of a better overview, all our explanations are based on just two levels. Accounting for additional stages proves straightforward, though much more cumbersome. The second scaling parameter, $\eta_2 \geq 0$, controls the difference in the scaling between the two levels considered. Fig. 5.3 (a) shows a step-wisely scaled point pattern simulated from an inhomogeneous Strauss model.

To scale the distances between pairs of points $\{u, v\}$ in a pattern with two gradation levels, three cases must be taken into account. Both u and v may be located in the inner disc $D_{\eta_1}(0)$, or both points may lie in the outer circular surrounding $D_{\mathcal{R} \setminus \eta_1}(0)$, or, otherwise, one point may be contained in the inner and the other one in the outer region. If both points are located in the outer surface area, the scaling of their distance depends on whether their connecting line $[u, v]$ crosses the inner disc or not. Incorporating the auxiliary measure $\Delta := \nu^1([u, v])^{-1} \|u\| \|v\|$ and applying principles of triangular geometry [Berger (2009, ch. 10)], point distances step-wisely scaled by means of Eq. (4.8) fulfill

$$\frac{\nu_c^1([u, v])}{\nu^1([u, v])} = \begin{cases} \alpha(\boldsymbol{\eta})^{-1}, & \{u, v\} \in D_{\eta_1}(0), \\ \sqrt{\eta_2} \alpha(\boldsymbol{\eta})^{-1}, & \{u, v\} \in \{W \setminus D_{\eta_1}(0)\}, \Delta \geq \eta_1, \\ \tau^* \alpha(\boldsymbol{\eta})^{-1} + (1 - \tau^*) \sqrt{\eta_2} \alpha(\boldsymbol{\eta})^{-1}, & \{u, v\} \in D_{\mathcal{R} \setminus \eta_1}(0), \Delta < \eta_1, \\ \tau^{**} \alpha(\boldsymbol{\eta})^{-1} + (1 - \tau^{**}) \sqrt{\eta_2} \alpha(\boldsymbol{\eta})^{-1}, & u \in D_{\eta_1}(0), v \in D_{\mathcal{R} \setminus \eta_1}(0). \end{cases} \quad (5.4)$$

Here, $\tau^* = \nu^1([u, v])^{-1}(\tau_1 + \tau_2 - 1)$ and $\tau^{**} = \nu^1([u, v])^{-1}\tau_1$, with τ_1 and τ_2 solving the equations

$$\tau_1^{-1}(\tau_1^2 + \|v\|^2 - \eta_1^2) = \nu^1([u, v])^{-1}(\nu^1([u, v])^2 + \|v\|^2 - \|u\|^2),$$

$$\tau_2^{-1}(\tau_2^2 + \|u\|^2 - \eta_1^2) = \nu^1([u, v])^{-1}(\nu^1([u, v])^2 + \|u\|^2 - \|v\|^2),$$


such that $\tau^*, \tau^{**} \in [0, 1]$. That is, τ_1 and τ_2 are given by the values

$$\tau_1 = \frac{\|u-v\|^2 + \|v\|^2 - \|u\|^2 \pm \sqrt{(\|u-v\| + \|v\| - \|u\|)^2 - 4\|u-v\|^2(\|v\|^2 - \eta_1^2)}}{2\|u-v\|},$$

and

$$\tau_2 = \frac{\|u-v\|^2 + \|u\|^2 - \|v\|^2 \pm \sqrt{(\|u-v\| + \|u\| - \|v\|)^2 - 4\|u-v\|^2(\|u\|^2 - \eta_1^2)}}{2\|u-v\|}.$$

A visual demonstration of a step-wise scaling effect is provided in Fig. 5.3



(a) Step scaled distances from the point coordinates $(0,0)^\top$, $(0.3,0.3)^\top$, and $(0.6,0.6)^\top$ (b) Realization

Figure 5.3: Step scaling effects in a point pattern on $W = D_1(0)$. (a) exemplifies step-wisely scaled distances from three different points in W , where $\eta_1 = 0.75$ (dotted circles) and $\eta_2 = 7$. Increasing darkness indicates decreasing distance. (b) shows a point process realizations sampled from a Strauss model with parameters $\beta = 100$, $\gamma = 0.5$, $R = 0.05$, and scaling effects as in (a). The sampling has been performed via dominated coupling from the past (see Section 4.3).

With regard to the maize data, it seems possible to assess the value of η_1 in a grid-based and model-free manner. For this, we define a set of η_1 -proposals by multiplying the disc radii \mathcal{R}_I (genotype I) and \mathcal{R}_{II} (genotype II) by a sequence of 20 values ranging from 0.8 to 1.0. In each respectively resulting inner disc, we first calculate the pairwise point distances and then search for their minimum among each of the genotypes. The left plot in Fig. 5.4 (i) shows that the minimum pairwise point distance increases notably if $\frac{\eta_1}{\mathcal{R}_I} = \eta_1$ is decreased from 0.98 to 0.97, from 0.95 to 0.94, from 0.93 to 0.92, and from 0.85 to 0.84. Regarding genotype II, we observe the most striking changes where $\frac{\eta_1}{\mathcal{R}_{II}} = \frac{\eta_1}{0.69}$ is diminished from 0.95 to 0.94 and from 0.90 to 0.89 (see Fig. 5.5 (i)). Since the number of points in $D_{\eta_1 := 0.97\mathcal{R}_I}(0)$ and in $D_{\eta_1 := 0.94\mathcal{R}_{II}}(0)$ is close to the total amount of points in $D_{\mathcal{R}_I}(0)$ and $D_{\mathcal{R}_{II}}(0)$, we consider these settings inappropriate. We confirm our decision by means of Fig. 5.4 – 5.5 (ii) which show the decomposition of representative point patterns with respect to the potential specifications of η_1 worked out in (i). Looking at Fig. 5.4 (ii), a cut-off at $0.84\mathcal{R}_I$ also appears unsuitable. We consequently consider $\eta_1 := 0.92$ or $\eta_1 := 0.94$ for genotype I, and take $\eta_1 := 0.89 \cdot 0.69 = 0.6141$ as a basis for the analysis of genotype II.

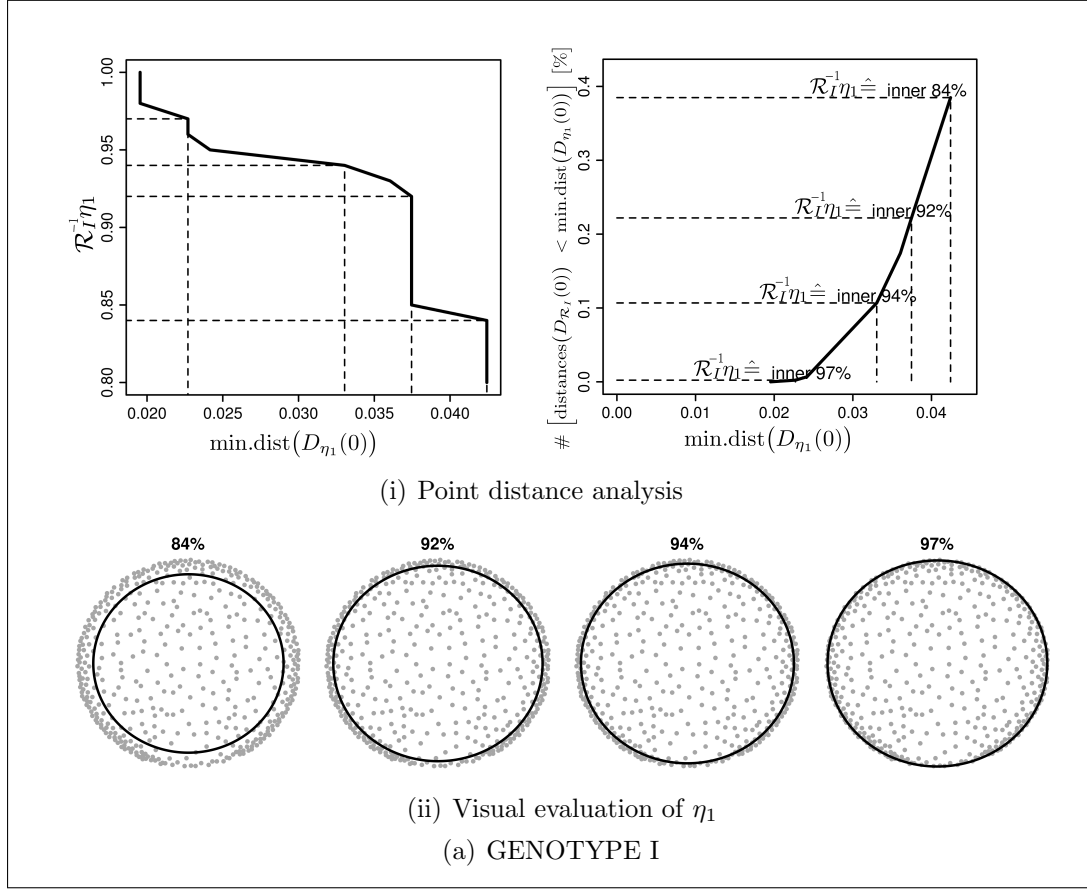


Figure 5.4: Determination of the inner stem radius η_1 as a percentage of the total radius $\mathcal{R}_I = 1$. Potential specifications of η_1 established based on the minimum pairwise point distance in $D_{\eta_1}(0)$ are illustrated in (i) which also shows the according percentage of the distances in $D_{\mathcal{R}_I}(0)$ that are smaller than the minimum distance in $D_{\eta_1}(0)$. In (ii), the potential inner radii (black) are plotted on top of a representative normalized point pattern (gray).

We regard the proposed specifications of η_1 as fixed and continue with the analysis of the remaining model parameters $\boldsymbol{\theta}^{-\{\eta_1\}} = \{\beta, \gamma, R, \eta_2\}$. Aiming at analyzing the preprocessed maize data in a Bayesian manner using the exchangeability method provided in Section 4.3, we incorporate the step scaling function from Eq. (5.1) in a Strauss model due to Eq. (4.6). The step-wise scaling attributes motivate us to develop a two-stage algorithm which, indeed, proves efficient and time-saving. In the following section, the two inference stages are described in detail, and all relevant intermediate results are graphically and verbally presented.

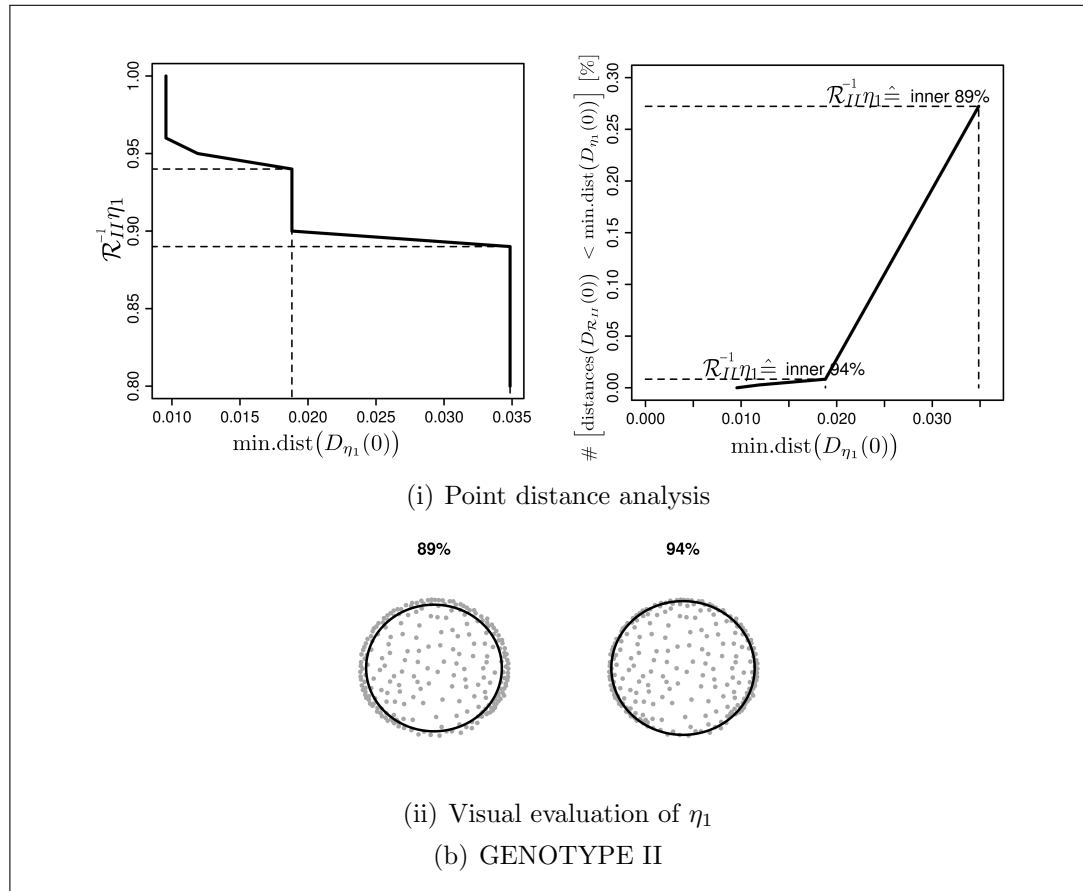


Figure 5.5: Determination of the inner stem radius η_1 as a percentage of the total radius $\mathcal{R}_{II} = 0.69$. Potential specifications of η_1 established based on the minimum pairwise point distance in $D_{\eta_1}(0)$ are illustrated in (i) which also shows the according percentage of the distances in $D_{\mathcal{R}_{II}}(0)$ that are smaller than the minimum distance in $D_{\eta_1}(0)$. In (ii), the potential inner radii (black) are plotted on top of a representative normalized point pattern (gray).

5.3 Two-Stage Inference Procedure

An inhomogeneous Strauss model with a step-wisely scaled point density contains five unknown parameters, the intensity parameter β , the weight of interaction γ , the interaction radius R , the radius η_1 of the inner disc and the scaling factor η_2 (cf. Eq. (5.1) – Eq. (5.3) and Fig. 5.3). We recall that Fig. 5.2 shows the original and the normalized locations of the vascular bundles in two maize stems. As the points representing the bundles obviously repulse each other, their arrangement suggests that γ is close to 0. Moreover, as discussed above, their packing density looks substantially lower in the inner than in the outer stem sections, which makes us expect that $\eta_2 \gg 1$. Referring to our discussion in Section 4.3, we assume that the pronounced interaction and scaling effects

in the point data cause long CPU run-times if we implement an exchange algorithm (cf. Alg. 2.7) with an auxiliary variable scheme exactly generated via dominated CFTP (cf. Alg. 4.1).

We propose a two-stage inference procedure, the structure and implementation of which is based on the intuitive meaning of the step scaling function as well as on its formal statistical definition. We denote the inner points with coordinates in $D_{\eta_1}(0)$ by $\mathbf{x}^{(1)} := \{x \in \mathbf{x} : \|x\| < \eta_1\}$ and the remaining points located in $D_{\mathcal{R} \setminus \eta_1}(0)$ by $\mathbf{x}^{(2)} := \{x \in \mathbf{x} : \eta_1 \leq \|x\| < \mathcal{R}\}$, where \mathcal{R} stands for the normalized stem radius, as introduced earlier. This notation allows us to decompose the joint posterior of $\boldsymbol{\theta} = \{\beta, \gamma, R, \boldsymbol{\eta}\}$ (cf. Eq. (2.24)) into

$$\begin{aligned} p(\boldsymbol{\theta}|\mathbf{x}) &= p(\boldsymbol{\theta}|\mathbf{x}^{(1)}, \mathbf{x}^{(2)}) = \frac{f(\mathbf{x}^{(1)}|\boldsymbol{\theta}) p(\boldsymbol{\theta}) f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}, \boldsymbol{\theta})}{\mathcal{C}(\mathbf{x}^{(1)}, \mathbf{x}^{(2)})} \\ &\propto p(\boldsymbol{\theta}|\mathbf{x}^{(1)}) \times f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}, \boldsymbol{\theta}), \end{aligned} \quad (5.5)$$

where

$$\begin{aligned} f(\mathbf{x}^{(1)}|\boldsymbol{\theta}) &= \mathcal{Z}^{(1)}(\boldsymbol{\theta})^{-1} \prod_{x \in \mathbf{x}^{(1)}} c_{\boldsymbol{\eta}}(x)^{-2} \beta^{n(\mathbf{x}^{(1)})} \gamma^{s_{c,R}(\mathbf{x}^{(1)})} \\ &= \mathcal{Z}^{(1)}(\boldsymbol{\theta})^{-1} \alpha(\boldsymbol{\eta})^{-2n(\mathbf{x}^{(1)})} \beta^{n(\mathbf{x}^{(1)})} \gamma^{s_{c,R}(\mathbf{x}^{(1)})} \\ &= \mathcal{Z}^{(1)}(\boldsymbol{\theta})^{-1} \beta^{*n(\mathbf{x}^{(1)})} \gamma^{s_{R^*}(\mathbf{x}^{(1)})} \end{aligned} \quad (5.6)$$

with

$$\begin{aligned} s_{R^*}(\mathbf{x}^{(1)}) &:= \sum_{\substack{\neq \\ \{u,v\} \subseteq \mathbf{x}^{(1)}}} \mathbb{1}[\nu^1([u,v]) \leq R^*] \\ &= \sum_{\substack{\neq \\ \{u,v\} \subseteq \mathbf{x}^{(1)}}} \mathbb{1}[\alpha(\boldsymbol{\eta})^{-1} \nu^1([u,v]) \leq R] = s_{c,R}(\mathbf{x}^{(1)}), \end{aligned}$$

and

$$f(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}, \boldsymbol{\theta}) = \mathcal{Z}^{(2)}(\boldsymbol{\theta})^{-1} \prod_{x \in \mathbf{x}^{(2)}} c_{\boldsymbol{\eta}}(x)^{-2} \beta^{n(\mathbf{x}^{(2)})} \gamma^{s_{c,R}(\mathbf{x}^{(2)}|\mathbf{x}^{(1)})} \quad (5.7)$$

$$\text{with } s_{c,R}(\mathbf{x}^{(2)}|\mathbf{x}^{(1)}) := \sum_{\substack{u \in \mathbf{x}^{(2)}, \\ v \in \{\mathbf{x}^{(1)} \cup \mathbf{x}^{(2)}\}}} \mathbb{1}[\nu_c^1([u,v]) \leq R].$$

$\mathcal{Z}^{(1)}(\cdot)^{-1}$ and $\mathcal{Z}^{(2)}(\cdot)^{-1}$ denote the normalizing constants with respect to $D_{\eta_1}(0)$ and $D_{\mathcal{R} \setminus \eta_1}(0)$ (cf. Eq. (2.10)). If we look at $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$ separately, each sub-pattern appears homogeneous, as already noted in Section 5.1. However, both partitions depend on each other due to the pairwise point distances between them.

Due to the factorization in Eq. (5.5), we first analyze $\mathbf{x}^{(1)}$ to obtain posterior information on the three unknown parameters in Eq. (5.6), $\{\beta, \gamma, R\}$. Afterwards, we add $\mathbf{x}^{(2)}$ as supplementary data and update the estimates of the

posterior distributions accordingly. This requires the consideration of η_2 as a fourth unknown parameter, since it appears as a scaling factor in Eq. (5.7). Although the set of parameters associated with $\mathbf{x}^{(1)}$ is hence a subset of the parameters describing $\mathbf{x}^{(2)}$ conditional on what has been learned from $\mathbf{x}^{(1)}$, a sequential analysis of the two datasets is statistically valid if the prior of η_2 does not depend on the priors of the remaining parameters. Here, we use uniform prior densities on suitable intervals, as specified later.

First Stage: Assessing $p(\boldsymbol{\theta}|\mathbf{x}^{(1)})$

To estimate

$$p(\boldsymbol{\theta}|\mathbf{x}^{(1)}) \propto f(\mathbf{x}^{(1)}|\boldsymbol{\theta}) \times p(\boldsymbol{\theta}) ,$$

we let $p(\boldsymbol{\theta})$ be the product of $p(\beta)$, $p(\gamma)$, and $p(R)$ which we define as independent uniform densities on the intervals $[50, 250]$, $[0, 1]$, and $[0.02, 0.08]$. Since the vascular bundles appear to be uniformly distributed in the inner parts of the stems and we hence assume $\mathbf{x}^{(1)}$ to be homogeneous, the scale factor η_2 is not being analyzed in the first stage of the inference procedure.

Concerning the technical issues, we implement the exchange algorithm by Murray et al. (2012), as sketched in Alg. 2.7. In each cycle, the parameters are updated one after the other. To collect the auxiliary variables needed for the calculation of the exchange ratio, we follow our outline of the dominated CFTP procedure by Berthelsen and Møller (2002, 2003), summarized in Alg. 4.1. Since $\mathbf{x}^{(1)}$ can be modeled under homogeneity assumptions (cf. Eq. (5.6)), dominated CFTP works much faster here than in an immediate application to the entire location-dependently scaled configuration $\mathbf{x} = \{\mathbf{x}^{(1)} \cup \mathbf{x}^{(2)}\}$. A side advantage is that the function `rStrauss()` from the R-package ‘`spatstat`’ can be used, which provides a CFTP routine for simulating from homogeneous Strauss models [Baddeley and Turner (2005)]. The results, however, refer to transformed versions of the actual parameters and hence need to be retransformed for consistency purposes (cf. Eq. (5.6)).

We determine that $\beta'|\beta \sim N_0^{300}(\beta, 2.5)$, $\gamma'|\gamma \sim N_0^1(\gamma, 0.025)$, and $R'|R \sim N_0^{0.15}(R, 0.00125)$, meaning that truncated Gaussian densities serve as M-H transition kernels $q_\theta(\theta'|\theta)$ for $\theta \in \boldsymbol{\theta}^{-\{\eta\}}$ (cf. Alg. 2.2). Almost every chain of parameter states clearly converges after less than 100.000 iterations. After 250.000 iterations, we stop our simulations, exclude the first 125.000 outcomes as burn-in results, and extract posterior information from every 100th triple $\hat{\boldsymbol{\theta}} = \{\hat{\beta}, \hat{\gamma}, \hat{R}\}$ of MCMC states. Fig. 5.6 illustrates the estimates of the full conditional posteriors.

Since we observe little variability in $\hat{p}(R|\mathbf{x}^{(1)}, \beta, \gamma, \eta_1 = 0.94)$ (genotype I) and $\hat{p}(R|\mathbf{x}^{(1)}, \beta, \gamma, \eta_1 = 0.61)$ (genotype II), it appears convenient and computationally efficient to decide for $\eta_1 = 0.94$ (genotype I) and to condition our further analysis on the *maximum a posteriori* (MAP) estimates of R , $\hat{R}_{\text{MAP}} = 0.0422$

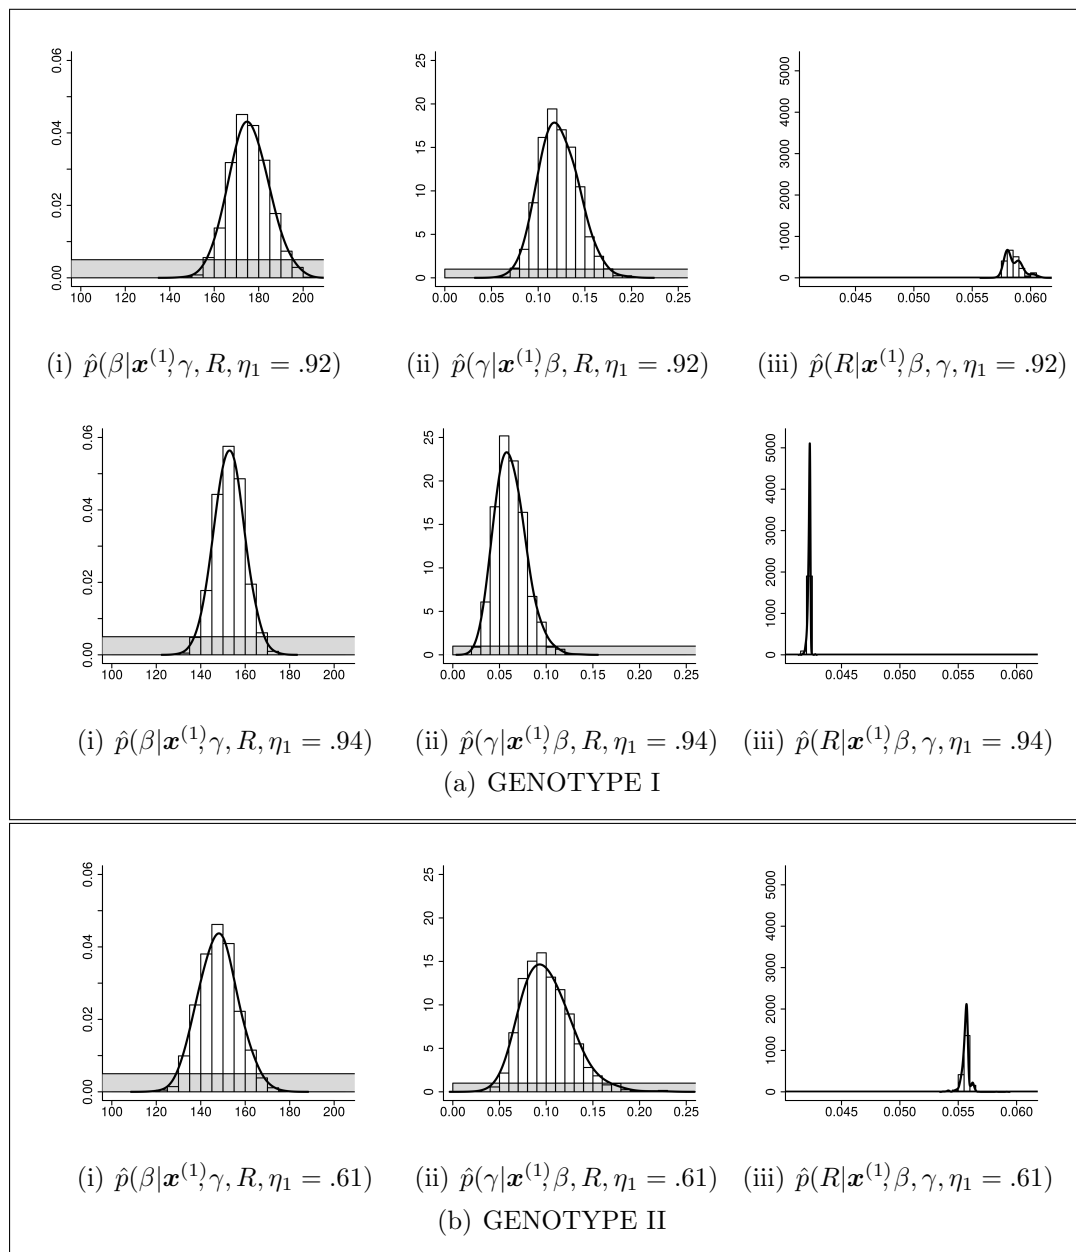


Figure 5.6: Full conditional posteriors describing $p(\boldsymbol{\theta}|\mathbf{x}^{(1)})$, and gray bars marking the parameter priors. The curves result from a kernel density adjustment [Silverman (1986, ch. 3 – 4)].

(genotype I) and $\hat{R}_{\text{MAP}} = 0.0557$ (genotype II). The MAP estimates correspond to the modes of the posterior distributions and hence to the parameter combination that maximizes the joint posterior. From the decision to fix R , we expect shorter burn-in phases of the Markovian chains approaching $p(\boldsymbol{\theta}|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$.

To avoid a rash conclusion, however, we first compare the competing models by using the *probability integral transform (PIT)* recently proposed by Thorarindottir (2013). PIT calibration diagnostics are based on rank statistics

which allow to detect structural differences between the observed data and reference point patterns simulated from the estimated posterior models. In a first step, we subdivide the circular observation window into 64 disjoint cells of equal size and generate 1000 reference point patterns. In a cell-wise fashion, we calculate the rank of the number of observed points among the number of reference points and normalize the outcome to an $[0, 1]$ -interval. The closer the resulting rank distribution is to standard uniformity the better the model fits the given data. Fig. 5.7 shows that both the histogram where $\eta_1 = 0.92$ and the




Figure 5.7: Results from PIT calibration diagnostics visualized through normalized rank distributions.

histogram where $\eta_1 = 0.94$ are biased towards higher ranks. For $\eta_1 = 0.92$, we furthermore observe a clear underrepresentation of the lower ranks, which reinforces our preference for $\eta_1 = 0.94$.

Second Stage: Assessing $p(\boldsymbol{\theta}|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$

In the second Bayesian inference stage, we estimate $p(\boldsymbol{\theta}|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$. For this purpose, we set $p(\boldsymbol{\theta}|\mathbf{x}^{(1)}) = \hat{p}(\boldsymbol{\theta}^{-\{\eta_2\}}|\mathbf{x}^{(1)}) \times p(\eta_2)$ in Eq. 5.5, where $\hat{p}(\boldsymbol{\theta}^{-\{\eta_2\}}|\mathbf{x}^{(1)})$ conforms to the posterior estimated before and $p(\eta_2)$ is defined as a uniform prior on $[2, 10]$. Based on the knowledge gained from the first stage of the analysis, we assume that $\boldsymbol{\theta}^{-\{\eta_2\}}|\mathbf{x}^{(1)}, \mathbf{x}^{(2)}$ takes values on the discrete domain of the MCMC states $\{\hat{\beta}, \hat{\gamma}\}|\mathbf{x}^{(1)}$ computed previously. It furthermore depends on R and η_1 , both serving as constants.

Iteratively, we upgrade the full conditionals $p(\{\beta, \gamma\}|\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \boldsymbol{\theta}^{-\{\beta, \gamma\}})$ and $p(\eta_2|\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \boldsymbol{\theta}^{-\{\eta_2\}})$ via the exchangeability framework in Alg. 2.7, with exact and approximate auxiliary sampling steps. We let $\hat{p}(\boldsymbol{\theta}^{-\{\eta_2\}}|\mathbf{x}^{(1)})$ determine the discrete prior and transition density of $\{\beta, \gamma\}$. The auxiliary variable scheme incorporated into the M-H ratio for evaluating new $\{\beta, \gamma\}$ proposals is generated via the dominated CFTP routine in Alg. 4.1 [Berthelsen and Møller (2003)].

For modeling the transitions of the scaling effects, we assume that $\eta'_2|\eta_2 \sim N_2^{10}(\eta_2, 0.25)$. To shorten the computer runtime, we draw the auxiliary realizations in an approximate manner, following Liang (2010). As pointed out in Section 4.3, Liang proposes to sample from distributions not presentable in complete form via sufficiently long MCMC chains. Here, we let the auxiliary variable associated with the latest parameter proposal accepted serve initial state $\mathbf{w}^{(0)}$. We assume that it consists of K points, i.e. $\mathbf{w}^{(0)} = \{w_1^{(0)}, \dots, w_K^{(0)}\}$. To determine $f(\mathbf{w}|\boldsymbol{\theta}) = f^{(K)}(\mathbf{w}^{(K)}|\mathbf{w}^{(0)})$ by means of Eq. (4.12), we first define $w_{k_0}^{(k)} := w_k^{(k-1)}$, for $k \in \{1, \dots, K\}$, and then successively propose and evaluate 10 random moves

$$w_{k_l}^{(k)}|w_{k_{l-1}}^{(k)} \sim q(w_{k_l}^{(k)}|w_{k_{l-1}}^{(k)}) := N(w_{k_{l-1}}^{(k)}, 0.125), \quad l \in \{1, \dots, 10\},$$

within the circular observation windows $D_{\mathcal{R}_I}(0)$ and $D_{\mathcal{R}_{II}}(0)$, respectively. Uniform densities on $D_{\mathcal{R}_I}(0)$ and $D_{\mathcal{R}_{II}}(0)$ are considered as marginal priors for the point locations. Conditional on $\boldsymbol{\theta}' = \{\beta, \gamma, R, \eta_1, \eta'_2\}$, $10 \times K$ M-H move steps thus transform $\mathbf{w}^{(0)}$ into $\mathbf{w}^{(K)} = \mathbf{w}$ via the intermediate stages

$$\begin{aligned} \mathbf{w}^{(k)} &= \left\{ \left\{ \mathbf{w}^{(k-1)} \setminus \{w_k^{(k-1)}\} \right\} \cup \{w_k^{(k)}\} \right\} : \\ &\quad \left. w_k^{(k)} = w_{k_{10}}^{(k)} \leftarrow w_{k_9}^{(k)} \leftarrow \dots \leftarrow w_{k_1}^{(k)} \leftarrow w_{k_0}^{(k)} = w_k^{(k-1)} \right\}. \end{aligned}$$

Liang (2010) argues that suggesting one move per point is sufficient from a theoretical standpoint, and confirms this statement by means of several data examples. We nonetheless prefer to increase the number of shift proposals to 10. In a simulation study as well as based on a subset of the maize data, we have demonstrated that no striking differences between the embedding of dominated CFTP in the exchange algorithm and the use of transition kernels according to Liang can be recognized in view of the estimated posteriors. The respective results are not shown in this work.

Having fixed η_1 in a mainly data-driven manner and R based on the negligible variations in $\hat{p}(R|\mathbf{x}^{(1)}, \boldsymbol{\theta}^{-\{R\}})$, the second and final stage of our approach returns exact and approximate drawings from $p(\{\beta, \gamma\}|\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \boldsymbol{\theta}^{-\{\beta, \gamma\}})$ and $p(\eta_2|\boldsymbol{\theta}^{-\{\eta_2\}})$. These samples describe the joint target distribution $\hat{p}(\boldsymbol{\theta}|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$. Fig. 5.8 illustrates the estimates of the full conditional posteriors.

Our two-stage approach proves to be very efficient in its application to the genotype-I samples. In the second stage, the Markovian chain reaches its equilibrium state after less than 5000 iterations. In total, we execute 25.000 iterations and regard the first 7.500 iterations as burn-in phase. To avoid autocorrelation effects, we use every 50th chain state for the description of the point process posterior. 350 estimated parameter configurations therefore determine $\hat{p}(\boldsymbol{\theta}|\mathbf{x})$. Genotype II requires about ten times as many iterations of the exchangeability process as genotype I. We thus compute a sequence of 250.000 parameter

realizations, remove the first 125.000 outcomes and base our further analysis on 350 parameter states sampled in regular intervals from the remaining 75.000 outcomes.

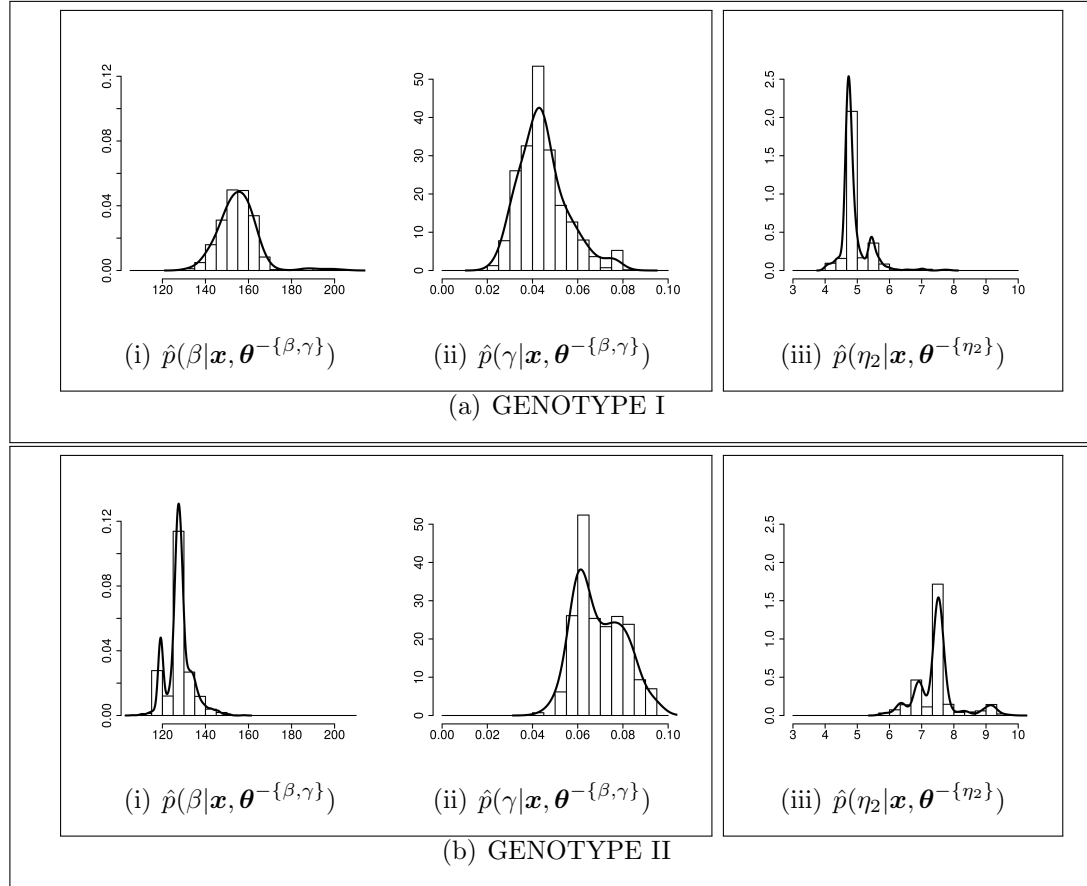


Figure 5.8: Full conditionals describing $p(\theta|\mathbf{x})$. The curves result from a kernel density adjustment [Silverman (1986, ch. 3 – 4)].

Fig. 5.8 visualizes the full conditional posterior distributions of the collected chain states. Conditional on the complete type-I datasets, the Markovian chain states of β tend towards higher and the states of γ towards smaller values than conditional on the type-II data, as already concluded from Fig. 5.8 (i) – (ii). However, for genotype I, the estimated intensity effect appears to be higher than before, while the estimated weight of interaction takes smaller values. Regarding genotype II, the estimated interaction penalty has also become stricter, though here, the intensity effect has been adjusted downwards. This leads to the conclusion that the packing density of the inner type-II bundles is higher when estimated based on the inner point data only than when estimated conditional on all point locations. There is, in other words, an evident discrepancy between the prior knowledge on the vascular bundle distribution obtained from the first stage of the analysis and the distributional properties of the vascular bundles

in the whole stems, which reduces speed of convergence in the second inference stage.

For both data series, the scaling effect η_2 between the bundle arrangement in the inner and in the outer stem segments has precisely been assessed in that there is little variation in its estimates, and the corresponding density curves show one clear maximum peak (see Fig. 5.8 (iii)). However, η_2 appears to be substantially higher in the type-II stems than in the plants of genotype I.

The following section is devoted to a comprehensive evaluation of our final results, a critical look at the proposed Bayesian inference scheme, and an outlook on possible further developments.

5.4 Conclusion and Outlook

We have analyzed two series of cross-sections through maize stems differing in their genotypes. Apart from the obvious fact that the stems representing genotype I are wider than those representing genotype II, no structural differences between the two series have been observable at first view. Apart from the original cross-sectional image data, we have been given a point process realization where each point represents one vascular bundle. To analyze these point process data, we have developed a locally scaled Strauss model (cf. Eq. (4.6)) with a properly defined step scaling function (cf. Eq. (5.1) – Eq. (5.3)). Inference has been based on the exchangeability framework presented in Section 4.3 (cf. Alg. 4.1). The proposed model comprises five unknown parameters, an intensity parameter β , two interaction parameters γ and R , and two scaling parameters η_1 and η_2 .

In a preprocessing stage, we have projected the point data to circular discs of fixed radii reflecting the proportion between the stem diameters of the two genotypes (cf. Fig. 5.2). The scaling parameter η_1 marking the boundary between the denser packed outer stem regions and the inner parts of the stems has been assessed in a deterministic manner (cf. Fig. 5.4 – 5.5), whereas suitable priors have been chosen for the remaining parameters.

Looking at both parts of the stems separately, each sub-pattern appears homogeneous and seems to follow an unscaled Strauss process distribution. We have argued that, regarding genotype I, the inner 94% of the stem area is less densely packed with vascular bundles than the outer 6%, whereas with regard to genotype II, a proportion of 89% versus 11% seems reasonable. The inner stem radius η_1 has accordingly been specified. Therefore, our first conclusion is that the two genotypes differ in the relative width of their sparsely packed inner stem regions.

We have discussed the statistical appropriateness of a two-stage approach where the first step corresponds to the estimation of the posterior distribution

of β , γ and R with respect to the inner stem regions, and the second step yields a description of the posterior of all parameters, conditional on the complete dataset as well as on the results from the previous step.

	GENOTYPE I ($\eta_1 := 0.94$)				GENOTYPE II ($\eta_1 := 0.61$)		
	Mode	Mean	Variance		Mode	Mean	Variance
β	152.93	152.81	4.32e+01		148.24	147.92	7.17e+01
γ	0.0575	0.0616	2.58e-04		0.093	0.1007	6.68e-04
R	0.0422	0.0422	1.20e-08		0.0557	0.0556	1.17e-07

Table 5.1: Results from the first stage of the analysis.

We first summarize the outcomes from the first inference stage in Tab. 5.1. As discussed in the preceding section, the bundle intensity parameterized through β proves to be higher in the plants of genotype I than in the type-II plants (cf. Fig. 5.8 (i)). At the same time, the interactions between the vascular bundles tend to be stronger in the stems of type I than in those of type II. The very small values of γ , particularly in view of genotype I, indicate that the modeled point patterns are similar to realizations of a hardcore Strauss process. Since there is almost no variation in the estimates of the interaction radii R (cf. Fig. 5.6 (iii)), we have embedded their MAP estimates $\hat{R}_{\text{MAP}} = 0.0422$ (genotype I) and $\hat{R}_{\text{MAP}} = 0.0557$ (genotype II) as constants in the second part of the analysis.

	GENOTYPE I ($R := 0.0422$, $\eta_1 := 0.94$)				GENOTYPE II ($R := 0.0557$, $\eta_1 := 0.61$)		
	Mode	Mean	Variance		Mode	Mean	Variance
β	154.50	155.19	8.48e+01		127.75	129.84	2.12e+01
γ	0.0433	0.0447	1.12e-04		0.0608	0.0694	1.14e-04
η_2	4.6747	4.9004	1.90e-01		7.4752	7.4644	4.15e-01

Table 5.2: Results from the second stage of the analysis.

Tab. 5.2 sums up what we have obtained from the second Bayesian modeling stage. Regarding the final results for the full data sets, the density curves estimated for β and γ have slightly to moderately been shifted. The modes and the means of γ have been adjusted downwards, but only for genotype I, we accordingly observe an increase in the estimates of β . For genotype II, the MAP estimate and the posteriori mean of β take considerably smaller values than in Tab. 5.1. This means that a smaller point intensity is now penalized by stronger interaction constraints. The discrepancy between the packing density estimated based on the reduced homogeneous and the packing density assessed for complete inhomogeneous type-II data explains the comparatively low speed

of convergence of Alg. 4.1. In general, the packing density in the inner stem segments is substantially lower than the packing density in the outer parts of the stems. It turns out that the scale factor η_2 associated with the bundle arrangement in the type-II maize plants is remarkably higher than the scaling effect associated with genotype I (see Fig. 5.8 (iii)).

A graphical summary of all results from our two-stage Bayesian modeling approach is given in Fig. 5.9. which does not only allow for direct comparisons between the two genotypes, but also provides a compact overview of the outcomes from the first and the second part of the analysis.

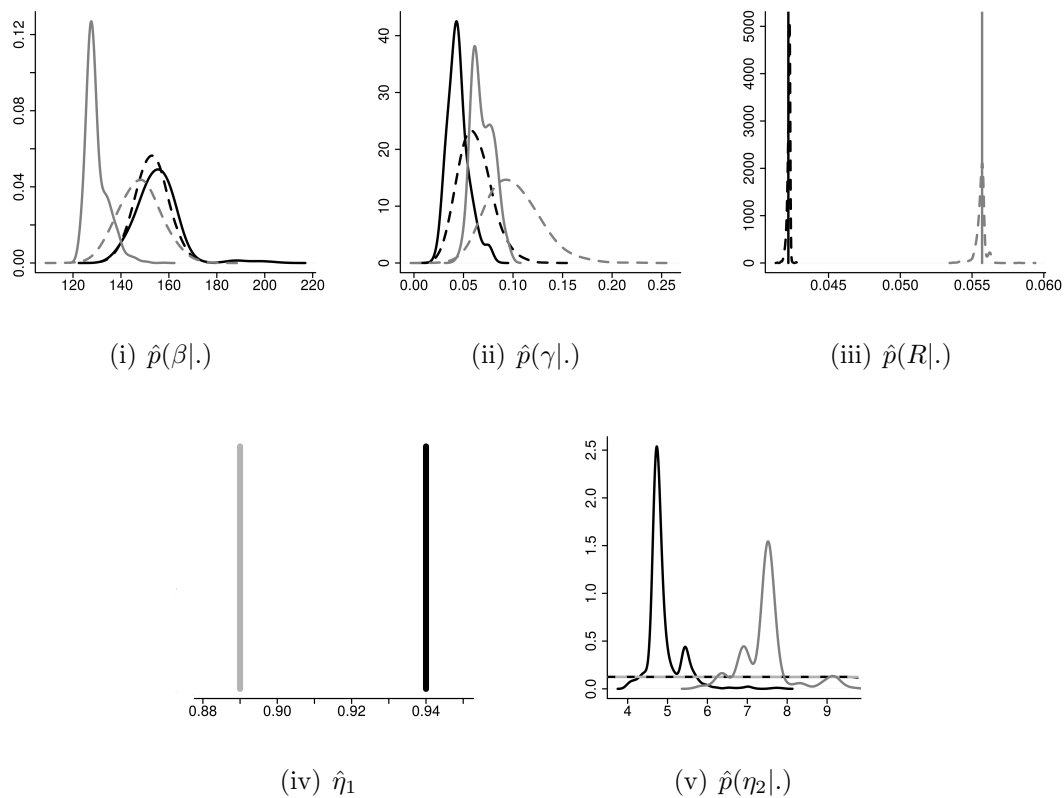


Figure 5.9: Full conditionals describing $p(\theta|\mathbf{x}^{(1)})$ and $p(\theta|\mathbf{x}^{(1)}, \mathbf{x}^{(2)})$. The black (genotype I) and gray (genotype II) curves result from a kernel density adjustment [Silverman (1986, ch. 3 – 4)]. The adjusted density curves from both the first and the second part of the analysis are illustrated, which allows for immediate comparisons. The dashed lines correspond to the curves in Fig. 5.6 and the solid lines to the curves in Fig. 5.8.

We generally conclude that the spatial bundle arrangement in cross-sections through maize stems can accurately be described by an inhomogeneous Strauss process distribution (cf. Eq. (4.6)) where the inhomogeneity is modeled through a step scaling function according to Eq. (5.1) – Eq. (5.3). The two-stage inference scheme suggested in Section 5.3 enables us to classify two genotypes of

maize plants by the packing density and the mutual dependencies of their vascular bundles. The two genotypes differ most strikingly in the estimates of the interaction radius R and the scaling factor η_2 . There is only minor variation in the Markovian chain states of these two parameters. Although the density curves approximating the estimated full conditionals of β and γ overlap to some extent, the corresponding modes and means clearly differ from each other.

Not only in its application to the cross-sectional maize data, our implementation of the exchange algorithm by Murray et al. (2012) proves flexible, efficient and promising. It may be used as a tool for Bayesian point process analysis whenever the spatial point arrangement is assumed to follow location-dependent scaling and interaction constraints that can be modeled by a proper scaling function. The determination of a numerically feasible and identifiable scaling function, however, is not trivial at all. Hahn et al. (2003) and Prokešová et al. (2006) propose an exponential scaling function which is discussed in Chapter 4 of this thesis. We have developed two further scaling functions, the step scaling function discussed and used above and a scaling function that allows to deduce information on the camera orientation towards a textured three-dimensional scene from a given two-dimensional image. The latter is introduced in Section 6.4 of the following chapter.

6 Shape from Texture using Locally Scaled Point Processes

Given a three-dimensional (3D) scene projected onto a two-dimensional (2D) plane, geometric properties such as camera positioning and angle are normally hard to comprehend and difficult to reproduce, particularly, if only one single image is available. In case that objects of identical or similar 3D shape appear in different parts of the image, local 2D characteristics of their arrangement, surfaces and sizes may facilitate a global geometric description of the scene. It is evident that prior information on the original 3D shapes of the objects provides a major advantage. In this chapter, we introduce a locally scaled point process framework for estimating *shape from texture* [see also Didden et al. (2013)]. Given one 2D image of a textured 3D scene, we turn the texture arrangement into a point process realization from which we infer scaling parameters describing the orientation of the camera towards the scene.

Referring to related projects, Section 6.1 introduces and motivates the research question. We propose a first intuitive point process approach in Section 6.2. Our actual inference scheme comprising two modeling stages is described in Section 6.3 – 6.4, where Section 6.3 is concerned with the question of how to learn a point process realization from an image preprocessed according to Section 3.2, and Section 6.4 presents a locally scaled point process model to immediately estimate the camera angle. For evaluation purposes, several data examples are discussed in Section 6.5. We conclude Chapter 6 with a brief discussion and an outlook on potential further developments (see Section 6.6).

6.1 Motivation

Natural images contain a variety of perceptual information enabling the viewer to infer the 3D shapes of objects and surfaces [Tuceryan and Jain (1998)]. Stevens (1980) observes that surface geometry mainly has three effects on the appearance of texture in images: foreshortening and scaling of texture elements, and a change in their density. In his seminal work, Gibson (1950) proposes the *slant*, the angle between a normal to the surface and a normal to the image plane, as a measure for surface orientation. Stevens amends this by introducing the *tilt*, the angle between the projection of the surface normal onto the image plane and a fixed coordinate axis in the plane. In our work, we directly infer the surface normal from a single image taken under standard perspective projection.

Statistical procedures for estimating surface orientation often make strong

assumptions on the regularity of texture. Witkin (1981) assumes observed edge directions provide the necessary information, while Blostein and Ahuja (1989) focus on circular texture elements with uniform intensity. Blake and Marinos (1990) consider the bias of the orientation of line elements isotropically oriented on a plane in 3D space, along with a computational approach related to Kanatani's texture moments [Kanatani (1989)]. The bias is here defined as the error induced by the orientation of the plane under *orthographic projection*. An orthographic projection, also known as orthogonal projection, is the mapping of a 3D object onto a plane in 2D space such that all projection rays are orthogonal to the projection plane.

Malik and Rosenholtz (1997) locally estimate texture distortion in terms of an affine transformation of adjacent image patches. The strong homogeneity assumption underlying this approach has been relaxed by Clerc and Mallat (2002), to a condition that is difficult to verify in practice. Forsyth (2006) eliminates assumptions on the non-local structure of textures, e.g. on perfect homogeneity, altogether and aims to estimate shape from the deformation of individual texture elements. Loh and Hartley (2005) criticize prior work due to the restrictive constraints related to homogeneity, isotropy, stationarity or orthographic projection, and claim to devise a shape-from-texture approach in the most general form. Their work, however, also relies on estimating the deformation of single texture elements, similar to Forsyth (2006).

We propose a general framework for inferring shape from *near-regular texture* by applying the locally scaled point process model of Hahn et al. (2003). Texture is termed near-regular if all texture elements are of identical or similar shape and spatially arranged in a clear and consistent manner. According to Liu et al. (2009, ch. 4), near-regularity can be regarded as strongly related to approximate symmetry.

Our approach briefly sketched in Fig. 6.1 enables the simultaneous representation of local variability and global regularity in the spatial arrangement of texture elements which are thought of as a marked point process. We preprocess the image [1.] to obtain a probability map [2.] representing an unnormalized intensity estimate for the underlying point process. Subsequently, we infer the point locations [3.] and learn the parameters of a locally scaled point process model [4.] to obtain a compact description of the 3D image attributes [5.].

As our main goal is not the detection of individual texture elements but the extraction of 3D information, we omit the exact modeling of each single texture element. Thus, our sole assumption regarding texture element shape is approximate convexity which offers considerable flexibility.

We start from the definition of a plane Π in 3D as

$$\Pi = \{\mathcal{X} \in \mathbb{R}^3 : \langle \boldsymbol{\delta}, \mathcal{X} \rangle + d = 0\} , \quad (6.1)$$

with $\|\boldsymbol{\delta}\| = 1$, which corresponds to a rewriting of Eq. (3.20). That is, d denotes

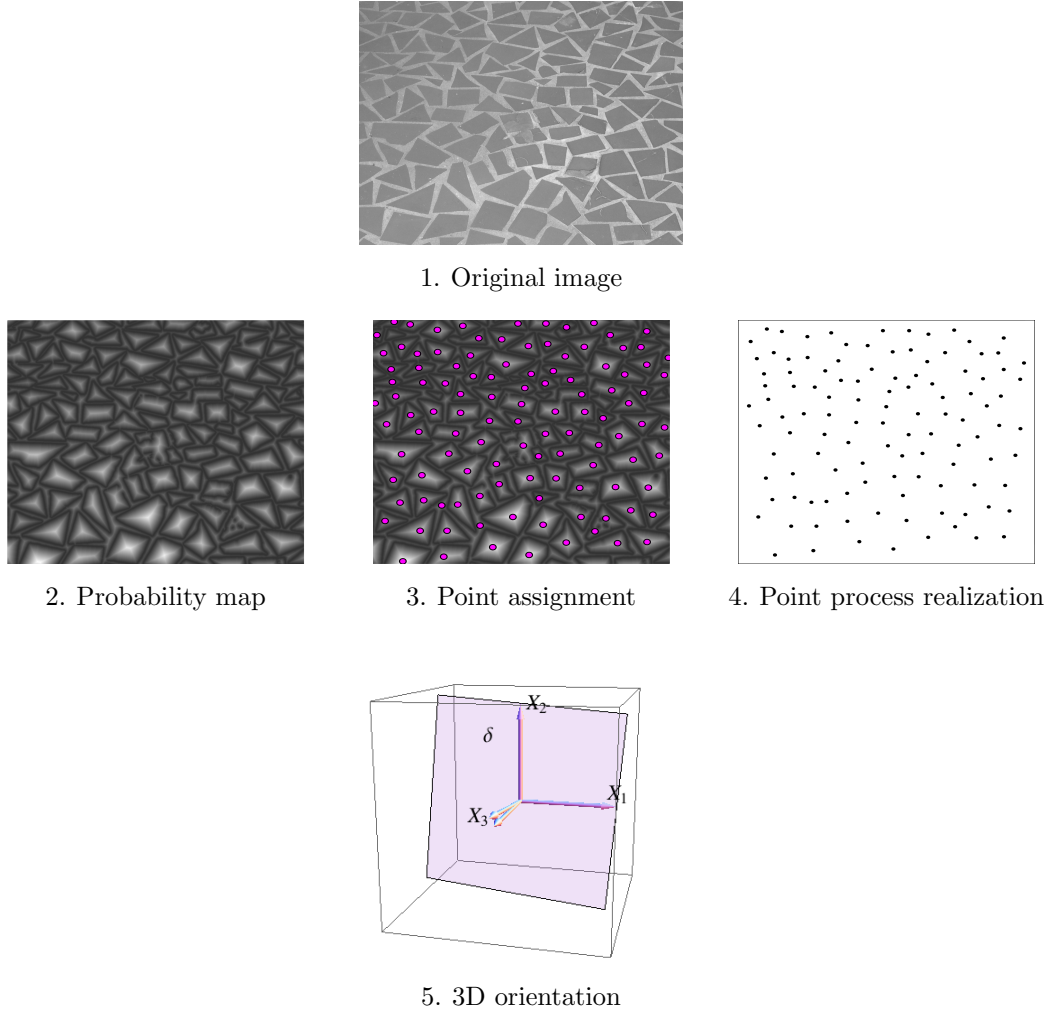


Figure 6.1: Sketch of the sequence of operations proposed for inferring shape from texture via a locally scaled point process.

the distance between Π and the origin. We assume the unknown *unit normal* δ to be oriented from Π towards the camera, forming obtuse angles $\langle \delta, \mathcal{X} \rangle < 0$ with projection rays $\lambda \vec{\mathcal{X}} = \lambda \mathcal{X}$, $\lambda \in \mathbb{R}^+$. Following the notational principles introduced in Chapter 3, we let $\mathcal{X} = (\mathcal{X}_1, \mathcal{X}_2, \mathcal{X}_3)^\top$ stand for the world and $x = (x_1, x_2)^\top$ for the image coordinates. The image domain is denoted by $W \subset \mathbb{R}^2$, and we assume the image to be scaled to have fixed area, $\nu^2(W) = a$, where $\nu^2(\cdot)$ is the 2D volume measure explained in Section 4.1.

We consider the basic pinhole camera [Hartley and Zisserman (2000, ch. 6)], and among the internal parameters, we only look at the focal length $f > 0$ which depends on the field of view. Fig. 6.2 illustrates our camera model. We also refer to Fig. 3.1. Following Eq. 3.1, we identify image points and rays of the projective plane \mathbb{P}^2 through the homogeneous coordinates

$$\tilde{x} = (x_1, x_2, -f)^\top. \quad (6.2)$$




Figure 6.2: 3D camera model relative to the coordinate axes \mathcal{X}_{a1} , \mathcal{X}_{a2} , and \mathcal{X}_{a3} . The camera with focal length f is oriented towards the negative \mathcal{X}_3 -halfspace. The scaled visible image domain is $W = [-\frac{a}{2}, \frac{a}{2}] \times [-\frac{1}{2}, \frac{1}{2}]$. Given the field of view in terms of an angle ρ_c , we have $f = \frac{a/2}{\tan(\rho_c/2)}$.

An image point \tilde{x} given by Eq. (6.2) meets Π in $\lambda\tilde{x}$ with

$$\lambda = -\frac{d}{\langle \boldsymbol{\delta}, \tilde{x} \rangle}, \quad \lambda > 0. \quad (6.3)$$

It follows that a point \mathcal{X} in Π is related to the image point \tilde{x} through

$$\mathcal{X} = \mathcal{X}(x_1, x_2) = -\frac{d}{\langle \boldsymbol{\delta}, \tilde{x} \rangle} \tilde{x}. \quad (6.4)$$

A *homogeneous texture* covering Π induces an *inhomogeneous texture* on the 2D image plane W , with density given by the *surface element*

$$\begin{aligned} \partial\mathcal{X} &= \left\| \frac{\partial\mathcal{X}}{\partial x_1} \times \frac{\partial\mathcal{X}}{\partial x_2} \right\| \nu^2(dx) \\ &= -\frac{d^2 f}{\langle \boldsymbol{\delta}, \tilde{x} \rangle^3} \nu^2(dx). \end{aligned} \quad (6.5)$$

Taking, for instance, the *fronto-parallel plane* defined through $\boldsymbol{\delta} = (0, 0, 1)^\top$ results by Eq. (6.2) merely in the constant scale factor $(d/f)^2$, that is, in the homogeneous density $(d/f)^2 \nu^2(dx)$. However, for arbitrary orientation $\boldsymbol{\delta}$, this factor depends on $x = (x_1, x_2)^\top$. Eq. (6.5) then quantifies perspective foreshortening and inhomogeneity of the texture, respectively, as illustrated in Fig. 6.3. Therefore, Eq. (6.5) mathematically represents the visually apparent *texture gradient*.

In addition to the preliminary background on camera projection provided in Section 3.1, Eq. (6.1) – Eq. (6.5) and Fig. 6.1 – 6.2 introduce the notational and technical principles our shape-from-texture modeling framework is based on.

In the following sections, we approach the problem of how to transform a probability map generated based on DoG filtering and distance transforming




Figure 6.3: Mappings of regular homogeneous point patterns in 3D space onto a 2D plane $W = [-\frac{1}{2}, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]$. The simulations are based on the parameters $d = 20$ and $\rho_c = 54^\circ$ ($f = 0.98$).

techniques (cf. Section 3.2) into an estimate of its underlying *latent point process* from which we extract geometric information on the orientation of the camera towards the 3D scene. As possible solutions, we provide two modeling concepts differing in their complexity, step sequences and emphases (see Section 6.2 and Section 6.3 – 6.4).

The first algorithm is based on a Gibbs modeling approach comprehending one first-order term and two terms of interaction (see Section 6.2). We initially assume that the texture arrangement, and thus the latent Gibbs process, follows exponential scaling constraints. The probability map is treated as a covariate. A straightforward grid-based optimization routine returns both a point process realization and the estimates of the unknown model parameters. Afterwards, the exponential scaling parameters are related to a spherical coordinate representation determining the orientation of the camera towards the 3D scene.

Our second inference concept provides a more flexible and sophisticated approach to the estimation of shape from texture [see also Didden et al. (2013)]. It starts with two alternative algorithms to detect the latent points in a model-free manner (see Section 6.3). To analyze the resulting pattern, an inhomogeneous Strauss process density is taken as a basis, where the inhomogeneity is directly induced by a scaling of perspective through spherical coordinates (see Section 6.4). We here estimate the scaling parameters via a maximum composite likelihood routine.

The second approach can be seen as an advancement of the first procedure in that it immediately estimates spherical coordinates describing the direction of the textured 3D plane relative to the camera. However, it does not include the estimation of exponential scaling parameters and requires the implementation of separate algorithms for the point assignment (cf. Section 6.3) and for the

estimation of the model parameters (cf. Section 6.4). Which method to apply therefore depends on the overall goal of the analysis.

6.2 Intuitive Preliminary Framework

As outlined above, we regard texture as a realization of a latent spatial point process where each point represents the symmetry center of exactly one near-regular texture element. Due to our principal notation, we let \mathbf{X} denote the point process, $\mathbf{x} = \{x_1, \dots, x_n\}$ one of its realizations and $W \subset \mathbb{R}^2$ the observation window. For clarity of exposition, we assume W to be bounded and scaled such that $W = [0, a] \times [0, 1]$.

We consider a finite Gibbs model with an unknown number of points and a set of parameters $\boldsymbol{\theta}$, as defined in Eq. (2.13). Including the previously computed probability map Z as a covariate in the model equation (cf. Section 3.2), the Gibbs density takes the form

$$f(\mathbf{x}|\boldsymbol{\theta}, Z) = \mathcal{Z}(\boldsymbol{\theta}, Z)^{-1} \exp \{ -U(\mathbf{x}|\boldsymbol{\theta}, Z) \} . \quad (6.6)$$

$U(\cdot|\boldsymbol{\theta}, Z)$ corresponds to the total Gibbs energy associated with the point pattern \mathbf{x} . As discussed in Section 2.3, the normalizing constant of a Gibbs model, here represented by $\mathcal{Z}(\cdot)^{-1}$, is usually intractable (cf. Eq. (2.10)) and hence requires the set-up of either Monte Carlo sampling or approximate inference algorithms [see e.g. Møller and Waagepetersen (2003, ch. 7 – 9)]. Such methods tend to be difficult to implement in practice.

In our setting, the process \mathbf{x} is a latent variable, that is, both \mathbf{x} and the parameter vector $\boldsymbol{\theta}$ in Eq. (6.6) are unknown. The goal of the analysis is to obtain the most likely pattern under the probability map Z , such that the points are at least a distance $R > 0$ apart and with no two points inside the same texture element. Mathematically expressed, we therefore define $U(\cdot|\boldsymbol{\theta}, Z)$ as

$$U(\mathbf{x}|\boldsymbol{\theta}, Z) := - \sum_{u \in \mathbf{x}} \underbrace{\phi_1(u|Z)}_{:=Z(u)} + \sum_{\substack{\{u,v\} \subseteq \mathbf{x} \\ u \neq v}} \phi_2(\{u,v\}|Z, m_1) + \sum_{\substack{\{u,v\} \subseteq \mathbf{x} \\ u \neq v}} \phi_3(\{u,v\}|\boldsymbol{\theta}) . \quad (6.7)$$

Here, $Z(u)$ denotes the value of the probability map Z at location $u \in W$, and $\phi_2(\cdot|Z, m_1)$ and $\phi_3(\cdot|\boldsymbol{\theta})$ are non-negative functions that model the interactions between all pairs of distinct points, $\{u, v\} \subset W$.

The interaction function $\phi_2(\cdot|Z, m_1)$ assigns zero density to point patterns with two points inside the same texture element. As stressed in Section 6.1, we assume that all elements are approximately convex in shape. We define two points $u \in \mathbf{x}$ and $v \in \mathbf{x}$ to lie inside the same texture element if the probability

map takes solely high values on the entire line $[u, v]$ between the two points. That is,

$$\phi_2(\{u, v\}|Z, m_1) = \begin{cases} 0, & \text{if } \min_{w \in [u, v]} Z(w) \geq \frac{\min\{Z(u), Z(v)\}}{m_1}, \\ \infty, & \text{otherwise,} \end{cases} \quad (6.8)$$

for some positive constant $m_1 > 1$ which is assumed fixed. The second interaction function $\phi_3(\cdot|\boldsymbol{\theta})$ ensures that a hardcore rule is not violated with

$$\phi_3(\{u, v\}|\boldsymbol{\theta}) = \begin{cases} 0, & \text{if } \nu_c^1([u, v]) \leq R, \\ \infty, & \text{otherwise.} \end{cases} \quad (6.9)$$

In conformity with the previous chapters, $R > 0$ denotes the interaction radius or range, and $\nu_c^1(\cdot)$ measures the locally scaled (Euclidean) distance between any two points in \boldsymbol{x} .

In Section 4.1 – 4.2, we have described by means of Eq. (4.1) and Eq. (4.8) how to determine an inhomogeneous distance function according to Hahn et al. (2003) and Prokešová et al. (2006). We recall that the authors propose to induce inhomogeneity through a normalized exponential scaling function $c_{\boldsymbol{\eta}}(\cdot)$ which is formally presented in Eq. (4.9), along with some theoretical and technical explanations. Here, we assume that each point $u \in \boldsymbol{x}$ associated with a texture element is exponentially scaled through

$$c_{\boldsymbol{\eta}}(u) = \alpha(\boldsymbol{\eta}) \exp(\boldsymbol{\eta}^\top u),$$

where $\alpha(\cdot)$ is the normalizing constant and $\boldsymbol{\eta} = (\eta_1, \eta_2)^\top \in \mathbb{R}^2$, according to our previous definitions. The Gibbs model in Eq. (6.6) thus has a total of three unknown parameters $\boldsymbol{\theta} = \{R, \eta_1, \eta_2\}$. Concerning the form of $\alpha(\cdot)$, we refer to Eq. (4.10) in Section 4.2. Since we consider an observation window of dimension $W = [0, a] \times [0, 1]$, it follows that

$$\alpha(\boldsymbol{\eta}) = \frac{1}{\sqrt{a}} \left(\frac{1 - \exp\{-2a\eta_1\}}{2\eta_1} \right)^{\frac{1}{2}} \left(\frac{\exp\{1 - \exp\{-2\eta_2\}\}}{2\eta_2} \right)^{\frac{1}{2}},$$

for $\boldsymbol{\eta} \geq 0$ and $\frac{0}{0} := 1$. As also discussed in Section 4.2, the resulting scaled pairwise distances are of the form

$$\nu_c^1([u, v]) = \nu^1([u, v]) \frac{c_{\boldsymbol{\eta}}(u)^{-1} - c_{\boldsymbol{\eta}}(v)^{-1}}{\boldsymbol{\eta}^\top (v - u)},$$

for any subset of points, $\{u, v\} \subset W$. Exponential scaling effects of different strength and orientation are visualized in Fig. 6.4 as well as earlier in Fig. 4.2.

In our setting, both the point process \boldsymbol{x} and the parameters $\boldsymbol{\theta} = \{R, \eta_1, \eta_2\}$ are unknown. To simultaneously infer the most likely point pattern $\hat{\boldsymbol{x}}$ and the associated parameters $\hat{\boldsymbol{\theta}}$ for a given probability map Z , we apply the approximate grid-based optimization loop in Alg. 6.1. The thresholds m_1 (cf. Eq. (6.8))




Figure 6.4: Examples of distances from the point $(0, 0)^\top$ in an observation window $W = [-\frac{1}{2}, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]$ under exponential scaling assumptions due to Eq. (4.9). Darker shades of gray indicate smaller distances.

and m_2 have to be set by hand, as discussed below. The same holds for the search grids.

Alg. 6.1 returns a valid pair $\{\hat{\mathbf{x}}, \hat{\boldsymbol{\theta}}\}$ in the sense that $f(\hat{\mathbf{x}}|\hat{\boldsymbol{\theta}}, Z) > 0$, while it is not guaranteed that it finds the global optimum. In particular, different results may be obtained due to the user-defined search grids and threshold values. Therefore, we execute the algorithm under several initial conditions and choose the pair $\{\hat{\mathbf{x}}, \hat{\boldsymbol{\theta}}\}$ for which the total energy U in Eq. (6.7) is minimized. Since Alg. 6.1 has been developed under the consideration of a high packing density of the point pattern, the estimate \hat{R} for the interaction parameter is expected to be close to the minimum inter-points distance of the configuration $\hat{\mathbf{x}}$. This assumes that the texture elements are fairly densely packed in the plane.

To explain the next stage of the proposed modeling procedure, we refer to the introductory definitions given at the end of Section 6.1. Eq. (6.5) shows how to relate the density of an inhomogeneous texture on a 2D image plane to the homogeneous texture covering the original 3D plane Π . This density, defined through the surface element

$$\partial\mathcal{X} = -\frac{d^2 f}{\langle \boldsymbol{\delta}, \tilde{\mathbf{x}} \rangle^3} \nu^2(dx) ,$$

can be connected to the inhomogeneous intensity $c_{\boldsymbol{\eta}}(\cdot)^{-2}$ determined by the exponential scaling function in Eq. (4.9). For this, we first extend $\partial\mathcal{X}$ itself to a proper scaling function (cf. Section 4.2). Referring to the area-preserving method suggested by Prokešová et al. (2006, see also Eq. (4.7)), we yield the normalization constant $\alpha(\cdot)$ by solving

$$|W| = a = \int_W \alpha(\boldsymbol{\delta}, d, f)^{-2} \partial\mathcal{X} ,$$

Algorithm 6.1: Optimization algorithm to infer a point process realization and its parameters from a probability map.

Data: Image matrix $Z(\cdot)$, observation window W , threshold values m_1 and m_2

Result: Point process estimate $\hat{\mathbf{x}}$, parameter estimates \hat{R} and $\hat{\boldsymbol{\eta}}$

Initialize $\hat{R} = \hat{R}^{(0)}$, such that $\hat{R}^{(0)}$ is small;

Initialize $\hat{\boldsymbol{\eta}} = \hat{\boldsymbol{\eta}}^{(0)} = (\hat{\eta}_1^{(0)}, \hat{\eta}_2^{(0)})^\top$, such that $\hat{\eta}_1^{(0)} + \hat{\eta}_2^{(0)}$ is high;

Initialize $\hat{\mathbf{x}} = \emptyset$ and set $Z^* = Z$;

```

while  $\max_{w \in W} \{Z^*(w)\} > m_2$  do
    update  $\hat{\mathbf{x}} = \{\hat{\mathbf{x}} \cup \{x\}\}$ , where  $x = \operatorname{argmax}_{w \in W} \{Z^*(w)\}$ ;
    for  $w \in W$  do
        if  $\nu_c^1([w, x] \leq \hat{R})$  then
            set  $Z^*(w) = 0$ ;
        end
    end
    if  $\exists \{u, v\} \subseteq \hat{\mathbf{x}} : \phi_2(\{u, v\} | Z, m_1) = \infty$  then
        if  $\hat{\eta}_1 = \hat{\eta}_2 = 0$  then
            increase  $\hat{R}$ ;
            reset  $\hat{\boldsymbol{\eta}} = \hat{\boldsymbol{\eta}}^{(0)}$ ;
        end
        else
            decrease  $\hat{\boldsymbol{\eta}}$ ;
        end
        reset  $\hat{\mathbf{x}} = \emptyset$  and  $Z^* = Z$ ;
    end
end

```

which results in

$$\alpha(\boldsymbol{\delta}, d, f)^{-2} = \frac{(a\delta_1 - 2f\delta_3 - \delta_2)(a\delta_1 - 2f\delta_3 + \delta_2)(a\delta_1 + 2f\delta_3 - \delta_2)(a\delta_1 + 2f\delta_3 + \delta_2)}{16d^2 f^2 \delta_3}. \quad (6.10)$$

As a byproduct, the unknown plane parameter d cancels. This parameter sets the absolute scale and cannot be inferred from a single image. Thus, given the focal length f , it remains to estimate the orientation $\boldsymbol{\delta}$. To this end, we choose

spherical coordinates

$$\boldsymbol{\delta} = \boldsymbol{\delta}(\rho_1, \rho_2) = (\sin \rho_1 \cos \rho_2, \sin \rho_1 \sin \rho_2, \cos \rho_1)^\top, \quad \rho_1 \in [0, \pi^*], \rho_2 \in [0, 2\pi], \quad (6.11)$$

with restriction of the range of ρ_1 through π^* due to the last condition in Eq. (6.1). We numerically solve the problem

$$(\hat{\rho}_1, \hat{\rho}_2)^\top = \min_{\substack{0 \leq \rho_1 \leq \pi^* \\ 0 \leq \rho_2 \leq 2\pi}} S(\rho_1, \rho_2), \quad (6.12)$$

$$\text{where } S(\rho_1, \rho_2) = \sum_{\mathbf{x}^* \in \mathbf{x}^*} \left(\frac{1}{c_{\hat{\boldsymbol{\eta}}}(\mathbf{x}^*)^2} - \frac{d^2 f}{\alpha(\boldsymbol{\delta}(\rho_1, \rho_2), d, f)^2 |\langle \boldsymbol{\delta}(\rho_1, \rho_2), \tilde{\mathbf{x}}^* \rangle|^3} \right)^2.$$

Here, $\mathbf{x}^* \subset W$ denotes a set of regularly sampled reference points on the visible image domain W , which in turn determines by Eq. (6.2) the set $\tilde{\mathbf{x}}^*$ of homogeneous coordinates $\tilde{\mathbf{x}}^* = (x_1^*, x_2^*, -f)^\top$.

We apply our entire framework to a (960×1280) -image of a brick wall. Fig. 6.5 (a) – (b) shows the selected scene and the associated probability map computed as discussed in Section 3.2. Based on Alg. 6.1, we determine the underlying latent point process and illustrate one of its realizations in Fig. 6.5 (c).

We assume that $R \geq 0.05$ and that $\boldsymbol{\eta} \in ([0, 2], [0, 2])^\top$. Our search starts on a rough grid which is iteratively refined. To decrease $\boldsymbol{\eta}$ due to the instructions in Alg. 6.1, we gradually decrease $\eta_1 + \eta_2$. The constants m_1 and m_2 are both set to 4. Results appear to be marginally impacted by alternative choices of the initial settings. Therefore, we execute our algorithm 25 times and select the output with the smallest energy due to Eq. (6.7).

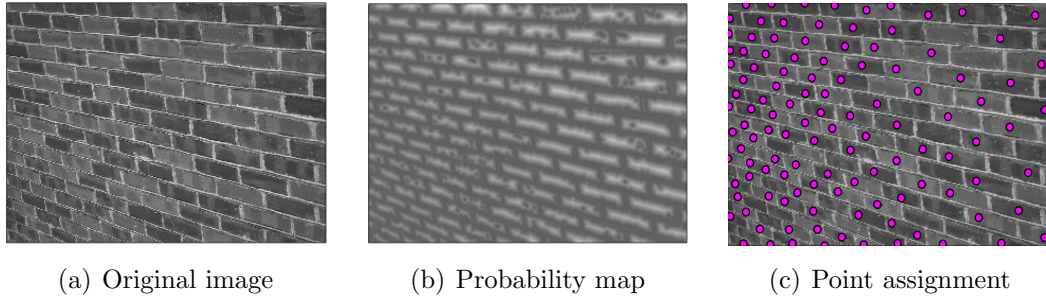


Figure 6.5: Learning a point process realization from the image of a brick wall by means of Section 3.2 and Alg. 6.1. The 2D brick arrangement appears to be exponentially scaled in both x- and y-direction. The according parameter estimates are $\hat{R} = 0.088$ and $\hat{\boldsymbol{\eta}} = (0.80, 0.09)^\top$.

Based on the estimate $\hat{\boldsymbol{\eta}}$ and Eq. (6.10) – Eq. (6.12), we illustrate the process of determining the orientation $\hat{\boldsymbol{\delta}}$ for the brick wall scenario shown in Fig. 6.5 (a). The only assumption made is a *standard wide-angle* value, $\rho_c = 54^\circ$, determining the field of view (cf. Fig. 6.2, caption). Fig. 6.6 shows that the parametrization

of our exponential scaling function fits quite well to the geometric counterpart emerging from Eq. (6.5). Fig. 6.7, finally, indicates that the optimization problem in Eq. (6.12) is remarkably well-behaved despite non-convexity. We refer to the figure captions for further details and discussion.




Figure 6.6: Model validation I: First term (a), second term (b), and the squared residual (c) under the sum of the nonlinear least-squares problem in Eq. (6.12) as functions of $x \in W$, evaluated for the numerically determined minimizer $(\hat{\rho}_1, \hat{\rho}_2)^\top$. The maximal residual value of 0.1 indicates reasonable accuracy.

The question comes up whether we can directly incorporate geometric scaling constraints into a point process modeling framework, without taking the intermediate step of estimating the parameters of an exponential scaling function. In what follows, we propose to infer shape from texture via a locally scaled Strauss model, where the inhomogeneity is imposed through a specifically developed *perspective scaling function*. This inference framework, introduced by Didden et al. (2013), requires the latent points to be localized in the image plane beforehand. Therefore, the following section is devoted to the introduction of two algorithmic strategies to learn a point pattern from a probability map computed according to Section 3.2.

6.3 Estimation of the Latent Point Process

In contrast to the shape-from-texture approach introduced above, we now estimate the latent point process realization in a model-free manner. Again, we first apply an image preprocessing strategy according to what has been sketched in Section 3.2, so that we obtain a probability map $Z = \{Z(w) : w \in W^{(0)}, 0 \leq Z(w) \leq 1\}$. $Z(\cdot)$ represents the spatial arrangement of the texture elements on the original image plane which we here denote by $W^{(0)}$, and which is assumed to be rectangular. Recalling that the value of the probability map in $w \in W^{(0)}$, $Z(w)$, indicates how likely it is that $w = (w_1, w_2)^\top$ is the symmetry center of




Figure 6.7: Model validation II and estimation of δ : (a) Level lines of the objective $S(\rho_1, \rho_2)$ (cf. Eq. (6.12)) within a relevant region of the parameter space. The white region for values of ρ_1 approaching $\pi/2$ (i.e. $\cos \rho_1 \rightarrow 0$) corresponds to large values of S . The plot reveals the non-convexity of the objective S , but also the existence of a single minimizer $(\hat{\rho}_1, \hat{\rho}_2)$ with a large basin of attraction. There is a second minimizer (not shown) corresponding to the sign reversal $-\hat{\delta} = -\delta(\hat{\rho}_1, \hat{\rho}_2)$ that can be ignored due to the last condition in Eq. (6.1). (b) The orientation $\hat{\delta} = \delta(\hat{\rho}_1, \hat{\rho}_2)$ finally inferred from the brick image shown in Fig. 6.5.

a texture element, we access the latent point process based on the information in Z by searching for local maxima in Z . For this purpose, we provide two different approaches, one that identifies the maxima in a rather focused manner and another one that acts on the entire observation window.

The first procedure starts from the assumption of an unknown number of symmetry elements. Based on appropriately specified threshold values, it searches locally for maxima in Z and then ensures that the maxima are delimited from each other by boundary segments. The second approach requires the number of points associated with the symmetry centers to be fixed in advance. Globally oriented, it uses a weighted distance measure to localize one point after the other in Z , and to estimate a *Voronoi tessellation* from the resulting point pattern.

Neighborhood-Based Point Detection

As stated right above, the algorithm introduced in this section starts with a local search for maxima in Z . We therefore set $W_u := [u_1 - k_1, u_1 + k_1] \times [u_2 - k_1, u_2 + k_1]$ for all $u \in W^{(0)}$ and some $k_1 > 0$, and compute

$$\Phi := \{u \in W^{(0)} : W_u \subset W^{(0)}, Z(u) = \max_{w \in W_u} \{Z(w)\}\} . \quad (6.13)$$

We then define a neighborhood relation “ \sim ” on Φ which says that $u \sim v$ if

$$\min_{w \in [u, v]} \{Z(w)\} \geq k_2 \max \{Z(u), Z(v)\}, \quad \text{for } \{u, v\} \subseteq \Phi, \quad (6.14)$$

where $[u, v]$ denotes the line from u to v , and k_2 is a constant with $0 < k_2 < 1$. Φ can be rewritten as a union of n_0 disjoint neighborhood components, $\Phi = \bigcup_{i=1}^{n_0} C_i$, such that each $u \in C_i$ is a point with at least one neighboring point in $\{C_i \setminus \{u\}\}$. Under the assumption that the texture elements are close to convex sets, two points u and v in Φ are neighbors if and only if they likely fall within the same texture element. Hence, we estimate a realization of the latent point process \mathbf{X}_0 in the observation window $W^{(0)}$ as

$$\hat{\mathbf{x}}_0 := \{x_1, \dots, x_{n_0} : Z(x_i) = \max_{u \in C_i} \{Z(u)\}\}. \quad (6.15)$$

To avoid boundary effects, we afterwards eliminate all elements of $\hat{\mathbf{x}}_0$ that are not located in $W := [w_{l_0} + k_1, w_{r_0} - k_1] \times [w_{b_0} + k_1, w_{t_0} - k_1]$, where w_{l_0} , w_{r_0} , w_{b_0} and w_{t_0} denote the left, right, bottom and top margins of $W^{(0)}$. We obtain

$$\hat{\mathbf{x}} = \{x : x \in \{\hat{\mathbf{x}}_0 \cap W\}\}. \quad (6.16)$$

The entire step sequence proposed is summarized in Alg. 6.2, and Section 6.5 provides some illustrative examples (see e.g. Fig. 6.12 (ii)).

The next section introduces another similar algorithm (cf. Alg. 6.3) which does not require any threshold values, but a predefined number of points.

Point Detection using Voronoi Tessellations

As before, we assume a probability map Z in an observation window $W \subset \mathbb{R}^2$, and again, we let Z describe the image of a textured 3D scene in that $Z(w)$ is the probability of $w \in W$ representing the symmetry center of a texture element. Regarding the set of all symmetry centers as a realization $\mathbf{x} = \{x_1, \dots, x_n\}$ of a latent point process, where n is assumed to be known, we propose an optimization procedure that estimates \mathbf{x} and immediately transforms W into a Voronoi tessellation comprising n Voronoi cells.

In general terms, conditional on a point pattern $\mathbf{x} = \{x_1, \dots, x_n\} \subset W$, a Voronoi tessellation on W corresponds to a partition of W into n cells, $\{W_1, \dots, W_n\}$, such that

$$W_i := \{w \in W : \|w - x_i\| \leq \|w - x_j\|, \forall j \neq i, x_i \in \mathbf{x}, x_j \in \mathbf{x}\}. \quad (6.17)$$

It follows that $W = \left\{ \bigcup_{i=1}^n W_i \right\}$ and that all cells are of convex shape. The tessellation can be seen as a mapping $w \mapsto V(w)$, where $V(w) \in \{1, \dots, n\}$ is the set of the cell labels. If $V(w) = i$, then $x_i \in \mathbf{x}$ is the closest point to w , and both x_i and w lie in W_i . For more information on spatial tessellations and Voronoi diagrams, see Okabe et al. (2009).

Algorithm 6.2: Neighborhood-based point process estimation.

Data: Probability map $Z(\cdot)$, observation windows $W^{(0)}$ and $W \subseteq W^{(0)}$, threshold values k_1 and k_2 (goes into “ \sim ”)

Result: Point process estimate $\hat{\mathbf{x}}$

Set $W_u := [u_1 - k_1, u_1 + k_1] \times [u_2 - k_1, u_2 + k_1]$, $\forall u \in W^{(0)}$;

Compute $\Phi := \{u \in W^{(0)} : W_u \subset W^{(0)}, Z(u) = \max_{w \in W_u} \{Z(w)\}\}$;

Initialize $n_0 = 0$;

Set $i = 1$;

while $\{\mathbf{u} \subseteq \Phi : n(\mathbf{u}) > 1, u_i \sim u_j, \forall i \neq j\} \neq \emptyset$ **do**

take one $C_i := \{u \in \mathbf{u} : \mathbf{u} \subseteq \Phi, n(\mathbf{u}) > 1, u_j \sim u_k, \forall j \neq k\}$ from Φ ;

set $\Phi = \{\Phi \setminus C_i\}$;

set $n_0 = n_0 + 1$;

set $i = i + 1$;

end

Compute $\hat{\mathbf{x}}_0 := \{x_1, \dots, x_{n_0} : Z(x_i) = \max_{u \in C_i} Z(u)\}$;

Return $\hat{\mathbf{x}} := \{x : x \in \{\hat{\mathbf{x}}_0 \cap W\}\}$;

Alg. 6.3 alternates between the detection of new points in W and the refinement of the estimated Voronoi tessellation. Conditional on the pixel value distribution in Z , it iteratively searches for the shortest weighted distance between each image coordinate and the set of already identified points. The pixel associated with the longest minimum distance becomes a new member of the point process estimate $\hat{\mathbf{x}}$, provided that less than n points have so far been selected.

We propose to employ an *asymmetrically weighted distance* measure of the form

$$\delta_{\mathbf{x}}(u, w|Z) := \|u - w\| Z(w) \left(1 - \min_{v \in [u, w]} \{Z(v)\}\right), \quad (6.18)$$

where $u \in \mathbf{x}$ and $w \in \{W \setminus \mathbf{x}\}$. We recall that, due to our principal notation, $[u, w]$ is defined as the line segment connecting u with v . Since only $Z(w)$ goes into $\delta_{\mathbf{x}}(u, w|Z)$, but not $Z(u)$, Eq. (6.18) measures pairwise distances in an asymmetric manner.

The weighting factors in Eq. (6.18) depend on $Z(w)$ as well as on the minimum value $Z(\cdot)$ takes on $[u, w]$. As a consequence, the Voronoi tessellation obtained conditional on the points in $\hat{\mathbf{x}}$ is driven by three measures, by the Euclidean

Algorithm 6.3: Point process estimation and Voronoi tessellation.**Data:** Probability map $Z(\cdot)$, observation window W , number of cells n **Result:** Point process estimate $\hat{\mathbf{x}}$, Voronoi tessellation $\hat{V}(\cdot)$ Set $V^{(1)}(w) = 1, \forall w \in W$;Find $x_1 = \operatorname{argmax}_{w \in W} \{Z(w)\}$ and set $\mathbf{x}^{(1)} = \{x_1\}$;Set $i = 2$;**while** $i \leq n$ **do** find $x_i := \operatorname{argmax}_{w \in \{W \setminus \mathbf{x}^{(i-1)}\}} \left\{ \min_{u \in \mathbf{x}^{(i-1)}} \{ \delta_{\mathbf{x}}(u, w|Z) \} \right\}$; set $\mathbf{x}^{(i)} = \{\mathbf{x}^{(i-1)} \cup \{x_i\}\}$; set $V^{(i)} = V^{(i-1)}$; **foreach** $w \in W$ **do** **if** $\|w - x_i\| < \|w - x_{V^{(i)}(w)}\|$ **then** $V^{(i)}(w) = i$; **end** **end****end**Return $\hat{\mathbf{x}} = \mathbf{x}^{(n)}$;Return $\hat{V}(\cdot) = V^{(n)}(\cdot)$;

distances between all pairs of points which should be as high as possible, by the probability values in $\hat{\mathbf{x}}$ which should also be as high as possible, and by $\min_{w \in [x_i, x_j]} \{Z(w)\}$ which should be as small as possible for $i \neq j$, to ensure that different Voronoi cells represent different symmetry elements.

Depending on the spatial gray value distribution in Z , a fine-tuning of the factors determining the distance measure $\delta_{\mathbf{x}}(\cdot|Z)$ in Eq. (6.18) may appear suitable. For instance, we can use

$$\delta_{\mathbf{x}}(u, v|Z) := \|u - v\| Z(v)^{l_1} \left(1 - \min_{w \in [u, v]} \{Z(w)^{l_2}\}\right),$$

where $l_1 > 0$ and $l_2 \geq 0$.

Fig. 6.8 illustrates results obtained from an application of our Voronoi tessellation method to a real-word data example. Since we currently exclusively deal with images where the texture elements are clearly arranged and easy to count, we have not yet established a mechanism that automatically fixes n .

Due to the fact that two different distance measures are considered for the estimation of the latent point process and for the determination of the Voronoi

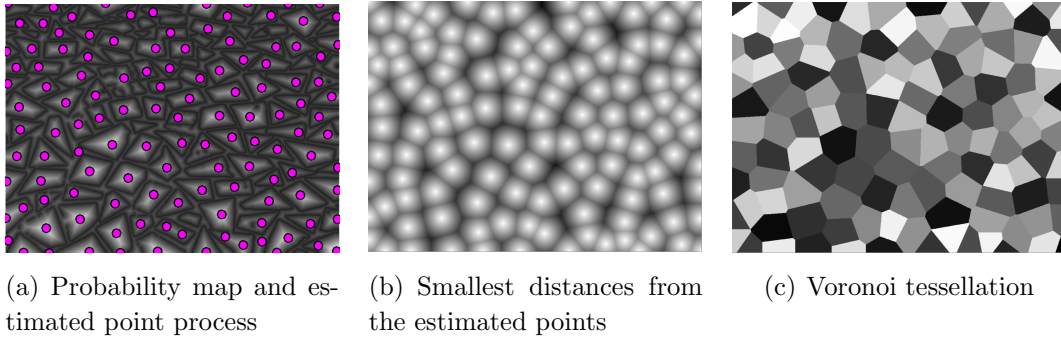


Figure 6.8: Voronoi tessellation framework. (a) shows the given probability map and an estimate of the associated latent point process, and (b) illustrates the pixel-wise distance to the closest estimated point. Regarding each point as the center of a Voronoi cell, the resulting Voronoi tessellation is visualized in (c); darker cells indicate an earlier date of selection. Boundary effects are not eliminated.

cells, Alg. 6.3 still leaves room for improvement. Furthermore, the question arises how to tackle boundary effects. If the estimation of the latent point process is of main interest, we suggest to first compute $\hat{\mathbf{x}}$ based on Alg. 6.3 and to subsequently eliminate those components of $\hat{\mathbf{x}}$ that are part of the outermost border of W . Edge effects in the Voronoi tessellation are more difficult to handle. Within the scope of this project, we ignore these effects and refer to Kenkel et al. (1989) for more detailed discussions.

Having localized the latent point pattern associated with the image of a textured plane in 3D space, the modeling framework presented in Section 6.4 delivers parameter estimates describing the orientation of the camera towards the scene.

6.4 Model Construction and Inference

We recall that in Section 2.1, we have described a point process as a random counting measure $N(\cdot)$, where $N(B)$ is the number of events in a Borel subset B of the relevant state space, in our context the image domain W . Due to Eq. (2.2), the intensity measure of the point process is given by $\beta(B) = E(N(B))$, and the associated intensity function is

$$\beta(x) = \lim_{\nu^2(dx) \rightarrow 0} \frac{(EN(dx))}{\nu^2(dx)}.$$

For a homogeneous point process, it holds that $\beta(x) = \beta$ for some $\beta > 0$, while for an inhomogeneous point process where the inhomogeneity stems from local scaling constraints according to Hahn et al. (2003), we obtain

$$\beta(x) = \beta c_{\boldsymbol{\eta}}(x)^{-2}.$$

As explained by means of Eq. (4.6) – Eq. (4.8), a spatial scaling function $c_{\boldsymbol{\eta}} : \mathbb{R}^2 \rightarrow \mathbb{R}_+$ acts as a local deformation in that it locally affects distances and areas. Due to Prokešová et al. (2006) and Section 4.2, it is statistically proper if it is identifiable, e.g. by virtue of the normalization in Eq. (4.7), and if it ensures that the scaled pairwise point distances can be computed in an exact manner, e.g. based on Eq. (4.8).

Referring to Eq. (6.5), the surface element $\boldsymbol{\partial}\mathcal{X} = \frac{d^2 f}{\langle \boldsymbol{\delta}, \tilde{x} \rangle^3} \nu^2(dx)$ describes the density of a heterogeneous texture on a 2D image plane emerging from the camera projection of a homogeneously textured plane in 3D space. We have shown that $\boldsymbol{\partial}\mathcal{X}$ can be extended to an identifiable scaling function via the area-preserving normalization proposed by Prokešová et al. (cf. Eq. (6.10)). For $W = [w_l, w_r] \times [w_b, w_t]$, we obtain

$$c_{\boldsymbol{\delta}}(x) = \frac{\alpha(\boldsymbol{\delta}, d, f) \sqrt{|\langle \boldsymbol{\delta}, \tilde{x} \rangle|^3}}{d\sqrt{f}}, \quad (6.19)$$

with

$$\begin{aligned} \alpha(\boldsymbol{\delta}, d, f) &= d \sqrt{\frac{f}{2}} \left(-(w_l + w_r)\delta_1 - (w_b + w_t)\delta_2 + f\delta_3 \right)^{\frac{1}{2}} \\ &\times (w_l\delta_1 + w_b\delta_2 - f\delta_3)^{-\frac{1}{2}} \\ &\times (w_l\delta_1 + w_t\delta_2 - f\delta_3)^{-\frac{1}{2}} \\ &\times (w_r\delta_1 + w_b\delta_2 - f\delta_3)^{-\frac{1}{2}} \\ &\times (w_r\delta_1 + w_t\delta_2 - f\delta_3)^{-\frac{1}{2}}, \end{aligned} \quad (6.20)$$

and $\tilde{x} = (x_1, x_2, -f)^\top$ according to Eq. (6.2). As discussed earlier, a convenient side effect of this formal definition is the canceling of the unknown plane parameter d . We call $c_{\boldsymbol{\delta}}(\cdot)$ a *perspective scaling function*. In conformity with our previous shape-from-texture approach, we have spherical coordinates

$$\boldsymbol{\delta} = \boldsymbol{\delta}(\eta_1, \eta_2) = (\sin \eta_1 \cos \eta_2, \sin \eta_1 \sin \eta_2, \cos \eta_1)^\top, \quad (6.21)$$

with $\eta_1 \in [0, \pi^*]$ and $\eta_2 \in [0, 2\pi]$. The upper limit π^* restricting the range of the scaling parameter η_1 ensures that $\langle \boldsymbol{\delta}, \tilde{x} \rangle < 0$ and therefore depends on the focal length f as well as on the size and location of the observation window W . Apart from a change in the notation, i.e. $\rho_1 \mapsto \eta_1$ and $\rho_2 \mapsto \eta_2$, this parameterization corresponds exactly to the parameterization used in Section 6.2 (cf. Eq. (6.11)). We intentionally replace ρ_1 and ρ_2 by η_1 and η_2 to put emphasis on the concept of immediately incorporating angles as scaling effects $\boldsymbol{\eta}$ into the point process model.

Under perspective scaling assumptions, the inhomogeneous intensity $\beta(\cdot)$ becomes

$$\beta(x) = \beta \frac{d^2 f}{\alpha(\boldsymbol{\delta}(\eta_1, \eta_2), d, f)^2 |\langle \boldsymbol{\delta}(\eta_1, \eta_2), \tilde{x} \rangle|^3}. \quad (6.22)$$

The perspective scaling function is computationally tractable and, as under the exponential scaling constraints discussed in Section 4.2, the accordingly scaled distance function is available in closed form,

$$\nu_c^1([u, v]) = \nu^1([u, v]) \frac{1}{\alpha(\boldsymbol{\delta}, d, f)} \left| \frac{2d\sqrt{f}}{\langle \boldsymbol{\delta}, \tilde{u} - \tilde{v} \rangle} \left(\frac{1}{\langle \boldsymbol{\delta}, -\tilde{u} \rangle^{\frac{1}{2}}} - \frac{1}{\langle \boldsymbol{\delta}, \tilde{v} \rangle^{\frac{1}{2}}} \right) \right| ,$$

provided that $\langle \boldsymbol{\delta}, \tilde{u} \rangle < 0$ and $\langle \boldsymbol{\delta}, \tilde{v} \rangle < 0$ is fulfilled for all pairs of distinct points, $\{u, v\} \subset W$. This compact representation can easily be obtained by applying the coarea formula from Eq. (4.8) [see also Krantz and Parks (2008, ch. 5)].

Visual examples of scaled distances are given in Fig. 6.9. When compared to Fig. 6.4, Fig. 6.9 clarifies that the perspective scaling constraints result in similar distance transformations as the exponential scaling, while also providing a coherent description of the perspective foreshortening.




Figure 6.9: Examples of distances from the point $(0, 0)^\top$ in $W = [-1/2, 1/2] \times [-1/2, 1/2]$ under a scaling of perspective due to Eq. (6.22). The internal parameters correspond to the settings in Fig. 6.3. Darker shades of gray indicate smaller distances.

For a given image, we assume that the focal length f is known. It remains to estimate the parameters β , η_1 and η_2 of the intensity function in Eq. (6.22) based on the estimated point pattern $\hat{\mathbf{x}}$ (cf. Eq. (6.16)). The desired 3D image information, the slant and the tilt of the surface, may then be characterized by the scaling parameter estimates $\hat{\eta}_1$ and $\hat{\eta}_2$. We propose to perform the parameter estimation by maximizing the composite likelihood given by Eq. (2.23) [Lindsay (1988)], which takes the form

$$L_C(\boldsymbol{\theta}) = L_C(\beta, \boldsymbol{\eta}) = \exp \{ -\beta \nu^2(W) \} \beta^n \prod_{x \in \hat{\mathbf{x}}} c_{\boldsymbol{\eta}}(x)^{-2} . \quad (6.23)$$

The maximum composite likelihood estimate for β is $\hat{\beta} = \frac{n}{\nu^2(W)}$. For the remaining two parameters which are the parameters of interest in our setting, we

maximize

$$\log(L_C(\boldsymbol{\eta}|\hat{\beta})) = n \log\left(\frac{n}{\nu^2(W)} - 1\right) + \sum_{x \in \hat{\mathbf{x}}} \log(c_{\boldsymbol{\eta}}(x)^{-2}). \quad (6.24)$$

Eq. (6.24) is equivalent to the *profile composite log-likelihood*, i.e. the logarithm of the right-hand side of Eq. (6.23) with β replaced by its estimate $\hat{\beta}$. The profile (composite/pseudo) log-likelihood is defined as the logarithm of the (composite/pseudo) likelihood of a subset of parameters $\boldsymbol{\theta}^{(2)} \subset \boldsymbol{\theta}$, conditional on the maximum likelihood estimates $\hat{\boldsymbol{\theta}}^{(1)}$ of the remaining parameters $\boldsymbol{\theta}^{(1)} = \{\boldsymbol{\theta} \setminus \boldsymbol{\theta}^{(2)}\}$ [see also Murphy and Van der Vaart (2000)]. If $\hat{\mathbf{x}}$ is a realization of a Poisson process, the estimates of β , η_1 and η_2 are identical with the maximum likelihood estimates (cf. Section 2.5).

In applications to synthetic as well as real-world image data, we examine the performance of the composite likelihood approach under perspective scaling assumptions. We prefer this framework to the preliminarily introduced strategy from Section 6.2, since it saves us the additional estimation of interaction and exponential scaling parameters and the implementation of a computationally more expensive grid-based algorithm. Concerning the generation of the point process realizations, both Alg. 6.2 and Alg. 6.3 return almost identical results (compare e.g. Fig. 6.8 (a) and Fig. 6.12 (b, middle)). Here, we use Alg. 6.2 without any objective reason. The most striking outcomes are described and discussed in the following section.

6.5 Case Studies

First, we present the results of a simulation study. We analyze sets of 3D point coordinates that have been sampled from either a perfectly regular pattern or a homogeneous Poisson process and subsequently been projected onto the 2D plane $W = [-1/2, 1/2] \times [-1/2, 1/2]$ (see Fig. 6.3 and Fig. 6.10). We estimate the scaling parameters associated with the synthetic patterns by maximizing the composite log-likelihood in Eq. (6.24). The true parameter values and the corresponding estimates are given in Tab. 6.1. While the estimation procedure is able to reconstruct the true values with a reasonable accuracy, the results are slightly better for the regular than for the random patterns. These outcomes are representative for several further such examples not shown here. We conclude that our inference framework allows us to identify the scaling parameters of the perspective scaling function irrespective of the second-order structure of the point process. That is, for quantifying the scaling effects, a model not accounting for interaction is sufficient even if the points are repulsive as in Fig. 6.3.

In what follows, we evaluate and discuss the point detection procedure proposed in Eq. (6.13) – Eq. (6.16) prior to the estimation of the actual scaling




Figure 6.10: Simulated Poisson point patterns with 3D shape determined by the outer normal in the subfigure captions. The internal parameters correspond to the settings in Fig. 6.3.

Pattern type	$(\eta_1, \eta_2)^\top$	$(\hat{\eta}_1, \hat{\eta}_2)^\top$
Regular	$(45^\circ, 0^\circ)^\top$	$(45.5^\circ, 0.0^\circ)^\top$
Poisson	$(45^\circ, 0^\circ)^\top$	$(46.2^\circ, 0.7^\circ)^\top$
Regular	$(30^\circ, 45^\circ)^\top$	$(29.9^\circ, 45.7^\circ)^\top$
Poisson	$(30^\circ, 45^\circ)^\top$	$(26.2^\circ, 45.5^\circ)^\top$

Table 6.1: True angles and composite likelihood estimates of the surface normals of the simulated point patterns in Fig. 6.3 and Fig. 6.10. Regular pattern type refers to the images in Fig. 6.3 and Poisson type to the images in Fig. 6.10.

parameters. To gain an impression of how concavity in the shapes of the texture elements biases the estimation of the point process and the model parameters, we generate images of size 1800×1800 pixels with varying proportions and arrangements of non-convex shapes (see Fig. 6.11). The true scaling effects are $\eta_1 = 20^\circ$ and $\eta_2 = 25^\circ$, and the overall packing density is the same in each image. Given the image in Fig. 6.11 (a) containing convex shapes only, all texture elements are correctly detected if the threshold k_1 determining the width and height of the local neighborhoods takes a value between 65 and 75 pixels. For the scenes in Fig. 6.11 (b) – (e) which are partly covered by concave shapes, $k_1 = 65$ seems more appropriate than $k_1 = 75$. Since all images are binary, they may directly be interpreted as probability maps, and the threshold k_2 in Eq. (6.14) hence becomes irrelevant. To avoid boundary effects, we do not consider the outer 75 pixels as potential point locations.

Fig. 6.11 shows the textured scenes and the estimated point process realizations for $k_1 = 65$. The scaling parameters estimated from the point process




Figure 6.11: Point process estimation in scenes with different amounts of non-convex texture elements. The point detection via Alg. 6.2 is based on $k_1 = 65$, the focal length corresponds to $f = 1.17$ ($\rho_c = 54^\circ$), and the true angles are $\eta_1 = 20^\circ$ and $\eta_2 = 25^\circ$.

realizations in Fig. 6.11 are illustrated in Tab. 6.2. We see that a grouping of the non-convex shapes increases the point detection and modeling inaccuracy in comparison to a random distribution of the non-convex among the convex texture elements.

For the analysis of real natural scenes, we apply our methodological framework to the set of tiling and brick images shown in Fig. 6.12. The given images are of size 1280×960 pixels, but during the preprocessing, they are downsized to 1166×846 pixels in order to eliminate boundary effects in the point detection. To be precise, we cut off bounding boxes of width 115 pixels from the original scenes, since we consider for the estimation of the latent point process \mathbf{X} that $k_1 \in \{35, 45, \dots, 105, 115\}$. For each value on this grid, we generate a point process realization $\hat{\mathbf{x}}$ and subsequently estimate its scaling parameters, i.e. the angles η_1 and η_2 determining the unit normal vector $\boldsymbol{\delta}$ of the original plane in 3D space.

	# points	$(\hat{\eta}_1, \hat{\eta}_2)^\top$	concave shapes [%]	arrangement
(a)	84	$(20.32^\circ, 23.66^\circ)^\top$	0	—
(b)	86	$(19.21^\circ, 31.00^\circ)^\top$	20	random
(c)	85	$(13.66^\circ, 20.89^\circ)^\top$	33	random
(d)	86	$(16.34^\circ, 39.20^\circ)^\top$	33	one cluster
(e)	79	$(14.31^\circ, 38.36^\circ)^\top$	33	four clusters

Table 6.2: Composite likelihood estimates of the surface normals of the simulated point patterns in Fig. 6.11. Random means that the concave elements are randomly distributed among the convex elements, whereas in a clustered arrangement, the concave shapes appear in groups. The number of texture elements in the visible image planes is equal to 89, and the true scaling parameters are $\eta_1 = 20^\circ$ and $\eta_2 = 25^\circ$.

The point detection is very robust in the selection of the threshold value k_2 . Threshold values from 0.15 to 0.5 have limited effects on the results which are somewhat more sensitive to changes in the neighborhood size k_1 . For the tiling images, neighborhood dimensions from 55×55 to 95×95 pixels result in similar point patterns and hence in similar scaling parameter estimates, while for the bricks image, slightly smaller neighborhoods seem to be needed. The estimated components of $\boldsymbol{\delta}$ are graphically evaluated by means of Fig. 6.13. For the first tiling image in (a) and especially for the brick scene in (c), we obtain a considerable k_1 -dependent variation in the estimates of $\boldsymbol{\delta}$, particularly in $\hat{\delta}_1$ and $\hat{\delta}_2$. Regarding the analysis of the second tiling scene in (b), the specification of k_1 has a lower impact.

We afterwards test our modeling assumptions by means of log-likelihood ratios. That is, for each value specified for k_1 and the resulting parameter estimates $\hat{\beta}$ and $\hat{\boldsymbol{\eta}}$, we sample 100.000 point process realizations from the respective locally scaled Poisson model, e.g. by following the step sequence in Alg. 2.4. We then calculate the ratios between the log-likelihood computed for the point-texture assignment and the log-likelihoods of the simulated point data. This is statistically valid, since the composite likelihood of a Strauss process formally corresponds to the density of an inhomogeneous Poisson process (cf. Eq. 2.23 and discussions).

Fig. 6.14 (i) shows that most of the obtained results are smaller than one. Although the observation window has already been downsized in the preprocessing stage, an additional reduction of its width makes the ratios take values close to one (see Fig. 6.14 (ii)). We therefore conclude that the boundary regions of the given images are still susceptible to inconsistencies in the point assignment, which may be due effects of light and shadow on the gray-value distributions.




Figure 6.12: Estimation of shape from texture: Original natural scene (left), a realization of the latent point process plotted on top of the associated probability map (middle), and the estimated 3D surface orientation towards the camera (right). The point detection via Alg. 6.2 is based on $k_1 = 75$ in (a) – (b) and $k_1 = 55$ in (c), and the field of view is assumed to be driven by a standard wide angle setting, i.e. $\rho_c = 54^\circ$.

Based on the number of correctly detected points which we here count manually, we decide for neighborhoods of size 75×75 pixels for the tiling scenes and 55×55 pixels for the bricks scene, with a threshold of $k_2 = 0.25$ in all cases. The probability maps and the resulting point patterns are shown in the middle column of Fig. 6.12.




Figure 6.13: Estimation of the unit normal δ with respect to different values of the neighborhood size k_1 .

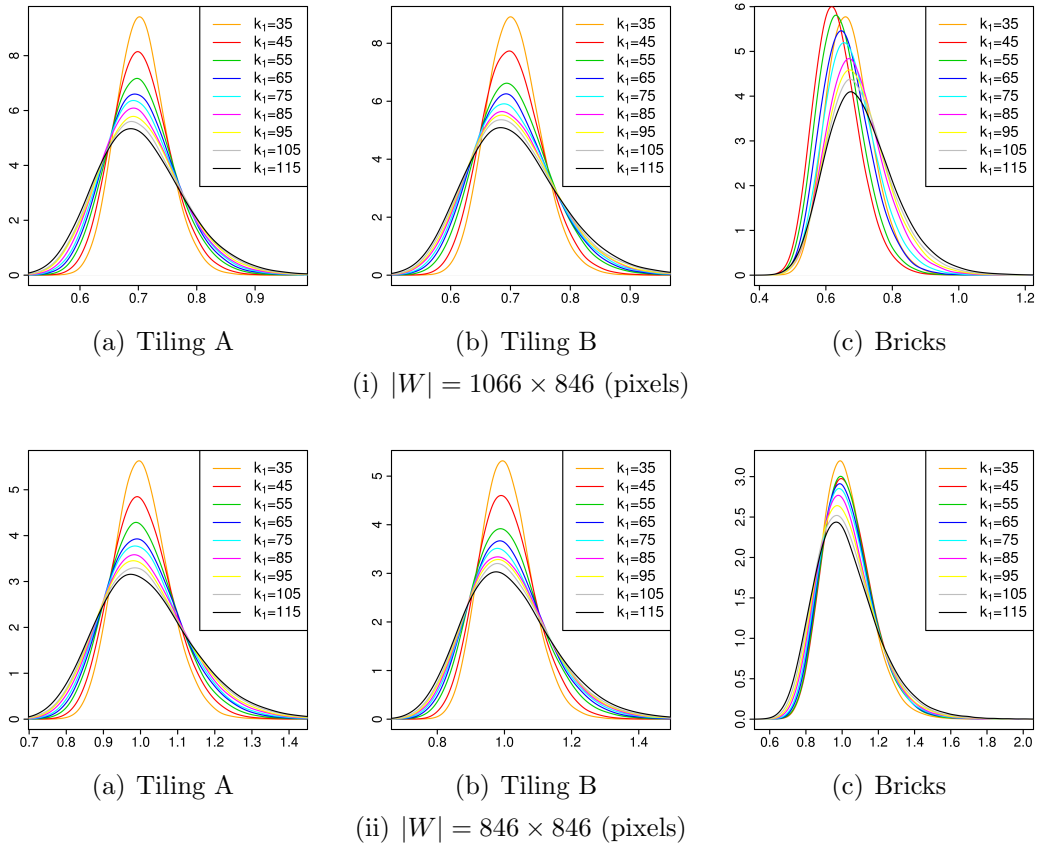


Figure 6.14: Log-likelihood ratios with respect to different values of k_1 . The numerator is defined as the log-likelihood of the estimated point configuration, whereas the denominator contains the log-likelihood of a simulated reference point process. In total, 100.000 reference patterns are considered per value of k_1 .

For deriving information on camera positioning and angle from the estimated point configurations, we project the points onto observation windows $W := [-0.69, 0.69] \times [-0.50, 0.50]$. We further assume that the field of view corresponds to a standard wide angle setting of $\rho_c = 54^\circ$ and hence take $f = \frac{0.69}{\tan(\rho_c/2)} = 1.35$ as a basis, the same settings as in the simulation examples above. The resulting scaling parameter estimates are listed in Table 6.3, and the 3D orientation of the camera towards the textures is illustrated in the right column of Fig. 6.12.

Texture	$(\hat{\eta}_1, \hat{\eta}_2)^\top$
(a) Tiling A	$(22.1^\circ, 94.7^\circ)^\top$
(b) Tiling B	$(12.2^\circ, 66.7^\circ)^\top$
(c) Bricks	$(36.0^\circ, 44.1^\circ)^\top$

Table 6.3: Estimated perspective scaling effects for the natural scenes in Fig. 6.12.

6.6 Outlook

Building on the recently developed locally scaled point processes [Hahn et al. (2003)], Chapter 6 introduces a framework for inferring 3D information from the 2D image of a textured scene. The modeling concepts discussed in Section 6.2 – 6.4 are quite flexible regarding assumptions on the texture composition in that they only require the texture elements to be close to convex in shape. Useful information related to surface orientation can thus successfully be extracted from a suitable probability map representation of the image. The pragmatic inference procedures allow to quickly determine a realization of the latent point process and to assess its scaling effects in a logically consistent manner.

The newly suggested scaling function in Eq. (6.19) – Eq. (6.20) quantifies perspective foreshortening and the resulting inhomogeneity of the texture. It is statistically proper in the sense that it is well-defined and that it allows for an exact calculation of locally scaled distance and volume measures.

The separation of image preprocessing on the one hand, and point detection and parameter estimation on the other hand offers great flexibility. We believe that the locally scaled point process framework can be applied in more general settings to analyze point patterns in images, for instance, as a new additional inference step in the texture detection algorithms discussed by Lafarge et al. (2010). Due to the low computational budget of our approaches, it also seems feasible to combine them with image segmentation where 3D information is needed for several segments within an image, each of which may be covered with a different type of texture.

There are further considerable avenues for development. One area for future advancement is to build a large hierarchical framework on our current inference steps, so that the image preprocessing, the point detection, and the parameter estimation are joined in an iterative fashion. Besides increasing compactness, such a coupling of the preprocessing and estimation stages may enhance process reliability. A fully Bayesian inference procedure along the lines of the Gibbs modeling strategy by Rajala and Penttinen (2012) appears to be a reasonable alternative to our frequentist approach. Finally, an extension of our framework to estimating smooth non-planar surfaces is conceivable, based on locally estimating the orientation of tangent planes. We sum up that plenty of room is left for embellishing our current inference concept.

Point process models have previously been used in image analysis applications where the goal is the detection of texture elements [see e.g. Lafarge et al. (2010)]. The respective approaches usually apply a marked point process framework, with marks describing the texture elements. Such set-ups rely on a good geometric description of individual texture elements, limiting the class of feasible textures. The next chapter deals with marked point process approaches to the analysis of textured 3D scenes. In particular, we provide a sketch of how a Bayesian marked point process model may be applied to estimate shape from texture (see Section 7.3).

7 Analysis of Textured 3D Scenes using Marked Point Processes

Texture learning and texture synthesis, the process of algorithmically constructing a large digital image from a smaller image sample, are two important tasks in computer vision. An integrated part of these procedures is the learning and modeling of three dimensional (3D) geometric attributes associated with a two dimensional (2D) textured image with depth effect. We apply a spatial statistics approach to this problem and develop a Bayesian marked point process model to extract structural and geometric attributes from such images. For this, we regard a textured scene as a realization of a marked point process, where the marks correspond to distinctive patches of texture appropriately scaled and rotated. Based on adequate data and prior distributional assumptions, posterior realizations of the marked point process are drawn by using a Gibbs sampler with an incorporated birth-death-move Metropolis Hastings (M-H) step. Our model set-up is particularly suitable for applications to 3D scenes with near-regular textures, such as brick walls (cf. Fig. 7.9).

This chapter starts with the discussion of related projects, particularly stressing the marked point process concept by Lafarge et al. (2010), which is presented in detail in Section 7.2. Orienting our research towards that methodological framework, we have developed a general strategy of how to estimate shape from texture via a marked point process approach. We introduce this framework in Section 7.3, discuss some first results in Section 7.4, and give an outlook on possible follow-ups in Section 7.5.

7.1 Motivation

In the past decades, marked point process models have frequently been used to analyze images of textured scenes. Descombes and Zerubia (2008), for instance, have developed a marked point process approach that allows to extract and describe the outline of rectangular buildings in dense urban areas. A similar framework has been suggested by Tournaire et al. (1965) for modeling road markings with a special focus on the dashed lines. Lacoste et al. (2005) model a more general class of line networks in remotely sensed images via a marked point process approach. This class includes, among other things, hydrographic networks. Medical data samples have been analyzed by Sun et al. (2007). The authors propose a two-steps modeling scheme which automatically detects vascular trees on angiograms. To detect and count pink flamingos, the so-called

Phoenicoptera Rosea, in aerial photographs of their colonies, Descamps et al. (2008) employ a multiple birth-death process. A similar problem has been tackled by Perrin et al. (2005) who propose a method for estimating the packing density of trees based on aerial images of the tree crowns.

For the purpose of analyzing crowds of people, Ge and Collins (2010) present a Bayesian marked point process model that allows to count individuals in noisy scenes. The authors regard the human shapes as similar in that they are describable as geometric transformations of each other. In contrast to Ge and Collins (2010), Lafarge et al. (2010) present a marked point process approach that is not only capable of counting geometric objects, but also allows to extract structures of interest from the given scene or to replace texture by a spatial arrangement of geometric template objects. Lafarge et al. take a codebook of areal and linear descriptors, including line objects, rectangular shapes and circles, as a basis to geometrically describe the textured 3D scene. Via a sophisticated Gibbs modeling approach, the authors localize these descriptors in the given image, which allows them to count objects of similar structure, to extract line networks and buildings from aerial photographs, or to represent texture elements by simpler geometric shapes. The following section is concerned with a detailed sketch of the modeling approach by Lafarge et al..

7.2 Geometric Feature Extraction

As stated above, Lafarge et al. (2010) present a flexible and widely applicable marked point process approach that tackles different questions associated with the geometric analysis of textured scenes. This general framework often proves to be less precise than other more specific techniques, such as those referred to in Section 7.1. Its versatileness, however, turns out to be a great advantage.

In what follows, we sketch the marked point process procedure proposed by Lafarge et al. (2010). We first explain the authors' definition of a mark, then continue with the description of their Gibbs modeling framework, and briefly discuss the proposed optimization approach in the end of this section.

Marks

We recall that in Chapter 2, we have denoted a spatial marked point process in a product space $A = (W \times M)$ as

$$\mathbf{Y} = \{(X, m_X) : X \in \mathbf{X}, m_X \in M\},$$

where \mathbf{X} is the unmarked point process in $W \subset \mathbb{R}^2$, and each element $X \in \mathbf{X}$ is assigned with a mark m_X from the mark space M .

For modeling textured scenes in an efficient and flexible manner, Lafarge et al. suggest a marked point process approach, where the marks are taken from a

mark library containing the seven geometric structures shown in Fig. 7.1. The corresponding parameter domains are given in the figure caption. Consequently, the entire mark space M can be decomposed into a union of seven subsets,

$$M = \bigcup_{q=1}^7 M_q, \quad (7.1)$$

which vary in the type and number of their parameters. Point processes with such a decomposable mark domain are commonly termed *multi-marked* (see also Fig. 2.1 (c) in Section 2). The geometrical features in Fig. 7.1 can be assigned to two classes, the class of the *linear descriptors* comprising the line objects in (a), and the class of the *areal descriptors* represented by the circular and rectangular shapes in (b) – (c).

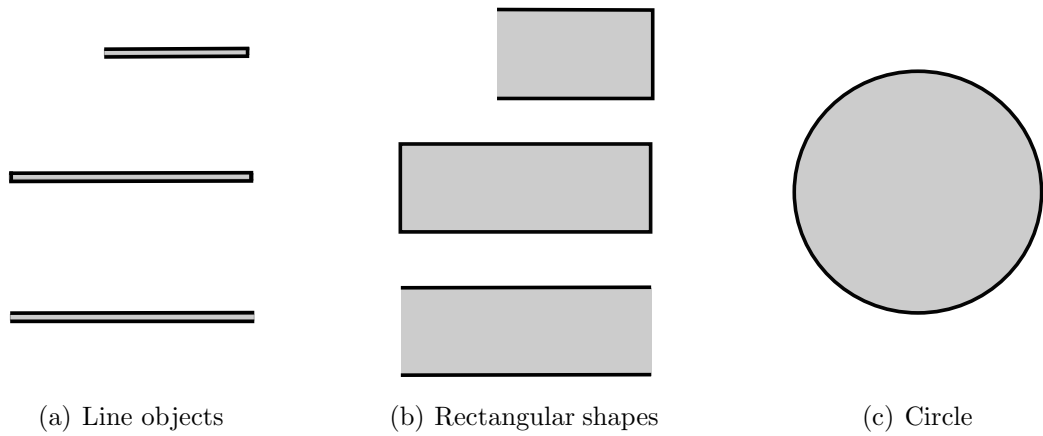


Figure 7.1: Finite library of geometric marks. With regard to the linear structures in (a), definition domains of the form $[\theta^{\min}, \theta^{\max}] \times [0, \pi]$ are used for lines and line segments, whereas $[\theta^{\min}, \theta^{\max}] \times [0, 2\pi]$ is considered for line ends. Concerning the shapes in (b), rectangles and bands are defined on $[\theta_1^{\min}, \theta_1^{\max}] \times [\theta_2^{\min}, \theta_2^{\max}] \times [0, \pi]$, while the parameter domain of a band end conforms to $[\theta_1^{\min}, \theta_1^{\max}] \times [\theta_2^{\min}, \theta_2^{\max}] \times [0, 2\pi]$. For circles (c), the parameter range is given by $[\theta_r^{\min}, \theta_r^{\max}]$.

Lafarge et al. characterize the texture elements by compositions of overlapping or directly adjacent shapes taken from the mark library as well as by their positioning in the image. As mentioned above, the object positions are seen as realization \mathbf{x} of an unmarked point process \mathbf{X} . In what follows, we explain the Gibbs point process framework developed by the authors for modeling such a multi-marked setting.

Model

Lafarge et al. (2010) propose a Gibbs model due to Eq. (2.14) for the reason that it allows to approach object dependencies in a flexible manner. Consequently,

important interaction structures can be accounted for, while the less relevant ones may be eliminated based on appropriately formulated restrictions. We recall that the Gibbs modeling framework proposed in Section 6.4 has been motivated by exactly the same technical and practical considerations.

In their work, Lafarge et al. proceed from the assumption that the number of objects is unknown, which means that the marked point process associated with the object configuration in the image consists of a random number of components. Similarly to the decomposition of the mark space M into the subspaces M_q where $q \in \{1, \dots, 7\}$ (cf. Eq. (7.1)), the authors subdivide the domain of the marked point process \mathbf{Y} (cf. Eq. (2.4)),

$$O_{\mathbf{Y}} = \{ \mathbf{y} = \{(x, m_x) : x \in \mathbf{x}, m_x \in M\} \subseteq A : n(\mathbf{x}_B) < \infty, \forall B \subseteq W \} ,$$

into a union of sub-domains $O_{\mathbf{Y}}^{(k)}$, $k \in \mathbb{N}$. Each sub-domain supports a fixed amount of geometric objects, that is, a fixed number of linear, rectangular and circular shapes (cf. Fig. 7.1). The probability distribution on the entire domain $O_{\mathbf{Y}}$ is therefore transformed into a mixture of the distributions on $O_{\mathbf{Y}}^{(k)}$, each assigned with a Gibbs density of the form

$$f^{(k)}(\mathbf{y}) \propto \exp \{ -U^{(k)}(\mathbf{y}) \} = \exp \{ -U_C^{(k)}(\mathbf{y}) - U_R^{(k)}(\mathbf{y}) \} , \quad (7.2)$$

where \mathbf{y} is a marked point process realization and $\boldsymbol{\theta}$ the set of model parameters. $U^{(k)}(\cdot)$ is the total Gibbs energy as explained in Section 2.3 by means of Eq. (2.13). It is defined by Lafarge et al. as the sum of a so-called *data coherence term* and a *regularization constraint*.

The data coherence term sums up the local energy measures associated with the marked points in \mathbf{y} . We thus have

$$U_C^{(k)}(\mathbf{y}) = \sum_{y \in \mathbf{y}} u_C^{(k)}(y|\boldsymbol{\theta}) .$$

To ensure that $u_C^{(k)}(\cdot|\boldsymbol{\theta})$ is appropriately specified, that is, to ensure that it reflects the coherence between y and the given image data in a satisfying manner, several conditions need to be fulfilled. First, the area of the selected object type must be accounted for, so that neither areal shapes are preferred to linear descriptors, nor vice versa. The second requirement is the selection of attractive features, which means that attractive features must be assigned with a negative local energy. Finally, as a prerequisite for using diffusion dynamics in the respective optimization stage, Lafarge et al. point out the necessity of a differentiable data coherence term. They suggest to apply

$$u_C^{(k)}(y|\boldsymbol{\theta}) = \begin{cases} \sqrt{\frac{\sigma_{\text{in}}^2(y) + \sigma_{\text{out}}^2(y) + \epsilon}{s(\mu_{\text{in}}(y) - \mu_{\text{out}}(y))^2}} - \theta_{\text{attr}} , & \text{if } \mu_{\text{in}}(y) \geq \mu_{\text{out}}(y) , \\ \infty , & \text{else} , \end{cases} \quad (7.3)$$

where $\mu_{\text{in}}(\cdot)$ and $\sigma_{\text{in}}(\cdot)$ denote the mean pixel intensity and the standard deviation inside the geometric object, and $\mu_{\text{out}}(\cdot)$ and $\sigma_{\text{out}}(\cdot)$ measure the mean intensity and standard deviation among the surrounding pixels. The authors suggest to set the width of the outer border to two pixels. \mathcal{S} corresponds to the total area covered by the object and its surrounding environment. The relational operator appearing in the first case of Eq. (7.3) serves a control tool. If the user decides for $\mu_{\text{in}} > \mu_{\text{out}}$, bright objects are preferred to dark shapes, whereas $\mu_{\text{in}} < \mu_{\text{out}}$ results in the opposite effect. Neutrality is achieved by setting $\mu_{\text{in}}(y) \neq \mu_{\text{out}}(y)$. The infinitesimally small positive auxiliary variable ϵ guarantees differentiability, and the threshold θ_{attr} representing an unknown model parameter supports the selection of attractive features and controls the sensitiveness of the data fitting process.

Besides the locally evaluated data coherence term $U_C^{(k)}(\cdot)$, Lafarge et al. include a regularization constraint in their Gibbs modeling approach (cf. Eq. (7.2)). This constraint imposed through $U_R^{(k)}(\cdot)$ accounts for prior assumptions on the layout of the geometrical features in the image. It penalizes overlaps and therefore controls the extent of interaction between marked points. An appropriate implementation of $U_R^{(k)}(\cdot)$ reduces object dependencies to the essential ones and thus provides a general model of the non-overlapping geometric shapes. To achieve such a target, meaningful connections between the detected geometrical features need to be established. The authors propose a regularization term of the form

$$U_R^{(k)}(\mathbf{y}) = \sum_{\substack{u \neq v \\ \{y_u, y_v\} \subseteq \mathbf{y}}} \left[\exp \left\{ \zeta_{\text{pen}} \varphi(\{y_u, y_v\}) \right\} - 1 \right], \quad (7.4)$$

where $y_u = (u, m_u)$ and $y_v = (v, m_v)$ denote any two distinct marked points, and $\varphi(\{y_u, y_v\}) \in [0, 1]$ measures their relative area of intersection. This area is then penalized by the weight ζ_{pen} . Lafarge et al. suggest to set ζ_{pen} to a high value, e.g. $\zeta_{\text{pen}} = 100$, such that large overlaps become heavily weighted. Minor overlaps, in contrast, have a negligible impact on the penalization of the total Gibbs energy in Eq. (7.2).

If the interaction constraint in Eq. (7.4) turns out to be too general in an application to a given image, Lafarge et al. consider two possible extensions of $U_R^{(k)}(\cdot)$, both allowing to handle object overlaps in a more precise manner. The first extension adjusts the penalization of the overlap of neighboring objects. Instead of subtracting 1 from the exponential term in Eq. (7.4), a new parameter ζ_{attr} is introduced. It balances between object attraction and repulsion. We thus have

$$U_R^{(k)}(\mathbf{y}) = \sum_{\substack{u \neq v \\ \{y_u, y_v\} \subseteq \mathbf{y}}} \left[\exp \left\{ \zeta_{\text{pen}} \varphi(\{y_u, y_v\}) \right\} - \zeta_{\text{attr}} \right]. \quad (7.5)$$

By setting $\zeta_{\text{attr}} = \exp\{0.1 \zeta_{\text{pen}}\}$, connected objects with a relative area of intersection of up to 10% are preferred to object configurations where the relative

overlap is greater than 10%.

Apart from a regularization of interaction in terms of object overlaps, certain application examples require a control tool that penalizes the difference in orientation between adjacent objects. To achieve the desired mutual alignment, Lafarge et al. suggest the implementation of

$$U_R^{(k)}(\mathbf{y}) = \sum_{\substack{u \neq v \\ \{y_u, y_v\} \subseteq \mathbf{y}}} \left[\exp \left\{ \zeta_{\text{pen}} \varphi(\{y_u, y_v\}) \right\} - 1 \right] + \zeta_{\text{al}} \sum_{\substack{u \neq v \\ \{y_u, y_v\} \subseteq \mathbf{y}}} Al(\{y_u, y_v\}) , \quad (7.6)$$

where $0 \leq Al(\{y_u, y_v\}) \leq 1$ evaluates the mutual alignment of the marked points y_u and y_v . By setting $Al(\{y_u, y_v\})$ to a fix value in $[0, 1]$, the proportion of the selected rotation-invariant objects can be influenced by the user. The balance between the alignment and the actual interaction criterion is adjustable through the parameter ζ_{al} .

If necessary, both Eq. (7.5) and Eq. (7.6) can be applied simultaneously. A simple combination of both terms yields

$$U_R^{(k)}(\mathbf{y}) = \sum_{\substack{u \neq v \\ \{y_u, y_v\} \subseteq \mathbf{y}}} \left[\exp \left\{ \zeta_{\text{pen}} \varphi(\{y_u, y_v\}) \right\} - \zeta_{\text{attr}} \right] + \zeta_{\text{al}} \sum_{\substack{u \neq v \\ \{y_u, y_v\} \subseteq \mathbf{y}}} Al(\{y_u, y_v\}) ,$$

which still corresponds to a numerically feasible extension of Eq. (7.4).

The Gibbs modeling approach developed by Lafarge et al. to extract geometrical features from images of textured 3D scenes comprises several random parameters represented by $\boldsymbol{\theta}$. First, the total amount of marked points is assumed to be unknown. The proportion between the different linear and areal descriptors taken from the mark library is treated as a random variable, too, and each single feature is characterized by a set of unknown mark parameters as given in the caption of Fig. 7.1. Finally, the attraction threshold θ_{attr} appearing in the data coherence term is contained in $\boldsymbol{\theta}$. Such a complex set-up and, in particular, the object-dependent number of mark parameters requires a well-elaborate inference framework which allows to jointly select multiple objects and guarantees fast a computer runtime.

Inference

We recall that we have cited the articles by Descamps et al. (2008), Lacoste et al. (2005), Descombes and Zerubia (2008), Perrin et al. (2005), Sun et al. (2007), and Tournaire et al. (1965) in Section 7.1. In all these works, the implementation of *jump diffusion processes* [Grenander and Miller (1994)] has proven to be a suitable inference strategy. Jump diffusion processes couple Markov chain Monte Carlo (MCMC) sampling with appropriately specified and evaluated Langevin equations. A variety of image analysis problems can thus be tackled.

Lafarge et al. (2010) also include a jump diffusion routine in their Gibbs modeling framework. The authors implement an algorithm that alternates between jumps from $O_{\mathbf{y}}^{(k)}$ to $O_{\mathbf{y}'}^{(k')}$ where $k, k' \in \{1, 2, \dots\}$, and a parameter fine-tuning process. The jump steps are conducted via a birth-death-move Metropolis-Hastings routine [Geyer and Møller (1994)] which is similar to the M-H-sampler in Alg. 2.6. We recall that a detailed theoretical description of the M-H sampling concept is provided in Section 2.4.

Given the current state \mathbf{y} of the marked point process, a transition to the state \mathbf{y}' here implicates the addition of a new geometric object from the mark library, the removal of an existing mark, or the replacement of one object by another differently shaped object. Due to the basic principle of M-H simulation (cf. Eq. (2.17)), the acceptance probability of the change considered conforms to

$$P_{\text{acc}}(\mathbf{y}'|\mathbf{y}) = \min \left\{ 1, \frac{q(\mathbf{y}|\mathbf{y}')}{q(\mathbf{y}'|\mathbf{y})} \exp \left\{ \frac{U^{(k')}(\mathbf{y}') - U^{(k)}(\mathbf{y})}{T} \right\} \right\},$$

where the so-called *relaxation temperature* T can be seen as a tool that regulates the entire optimization process. T decreases in the number of computing steps of the jump diffusion algorithm. The smaller it gets the more importance is attached to the diffusions compared to the jumps. Lafarge et al. determine T based on *simulated annealing*, according to the definitions by Van Laarhoven and Aarts (1987) and Salamon et al. (2002). Given a complex function with many unknowns, simulated annealing returns an approximate solution that is close to the global optimum. Therefore, sufficiently many time stages, i.e. computational steps, need to be passed through. Independently of the initial state, the outcome with the smallest possible energy is reached after the predetermined amount of time. For more information on the meaning and determination of T in the context of jump diffusion processes, we make reference to Grenander and Miller (1994) and Lafarge et al. (2010).

It remains to shortly describe what is meant by *Langevin equations/diffusions* and *diffusion dynamics/equations/processes*, respectively. Langevin equations, named after their inventor Paul Langevin, conform to stochastic differential equations that characterize the time evolution of a certain subset of variables in a given function. In the strict sense, Langevin diffusions are driven by Brownian motions [Langevin (1908)]. However, in the past century, the term Langevin/diffusion dynamics/equations/processes has been broadened considerably. Its meaning and implementation in the context of jump diffusion optimization has comprehensively been explained by Geman and Hwang (1986).

Lafarge et al. combine their birth-death-move sampler with diffusion equations to significantly accelerate convergence. That is, between each two M-H steps, a stochastic diffusion routine is applied to Eq. 7.2, whereby the current state \mathbf{y} of the multi-marked point process serves as the initial configuration.

The diffusion process then acts on the continuous subspace \mathbf{y} is living on.

We currently work on the development of a marked point process approach to the estimation of shape from texture, similarly to the procedures introduced in Chapter 6. This alternative concept is related to the work by Lafarge et al. (2010) in that it detects template objects from a mark library in the image and allows to describe their local 2D geometric properties. Our future target is to learn global 3D geometries, such as camera positioning and angle, from the arrangement and attributes of the local features.

7.3 Shape from Texture via Marked Point Processes

Our model synthesizes approved computer vision methodology with recent developments in point process research. For a better understanding of the modeling concept, we particularly refer to the preliminaries provided in Chapter 2 – 3.1 of this thesis as well as to Foley et al. (1997, ch. 5) and Møller and Waagepetersen (2003, ch. 6 – 7). Given an image of a textured 3D scene, our overall goal is to estimate the global geometric attributes of the scene via an appropriately implemented Bayesian hierarchical inference mechanism. In contrast to the methods described in Chapter 6, this mechanism does not require any image preprocessing. Before explaining its technical details, we briefly outline its fundamental conception in an informal and intuitive manner.

We regard the texture arrangement in the given image as a realization of an inhomogeneous marked point process with independent and unpredictable marks (see Def. 2.4 – 2.5). A small set of representative texture elements taken from a training image without depth effect serves as object library. Although our objective is not the reconstruction of the image of interest, we act on the assumption that it can fully be recomposed by the sparse selection of characteristic features contained in the library. However, since there is an unspecified but obvious dept effect in the image of interest, a geometric adjustment of the features is necessary for localizing them on the image plane. Our proposed procedure is thus similar to a puzzling process, where every puzzle piece may be chosen several times, each time requiring an individual Euclidean similarity transformation. The closer the pieces are located to the vanishing point, for instance, the more strongly they need to be scaled down. At the same time, their packing density increases. Our aim is to infer shape from texture based on the gradual change in the packing density as well as in the sizes and orientation of the geometric objects, i.e. the “puzzle pieces”. We may then also be able to determine the continuation of the scene beyond the image borders.

Fig. 7.2 shows a training image (a), a randomly constructed test image consisting of two geometrically transformed parts of the training image (b), and a

suitable real-world test scenario (c). In Section 7.4, we discuss first results from the analysis of the images in (b) and (c). As library of template objects, a set of geometrical features from the planar test scenario in (a) is being used. The

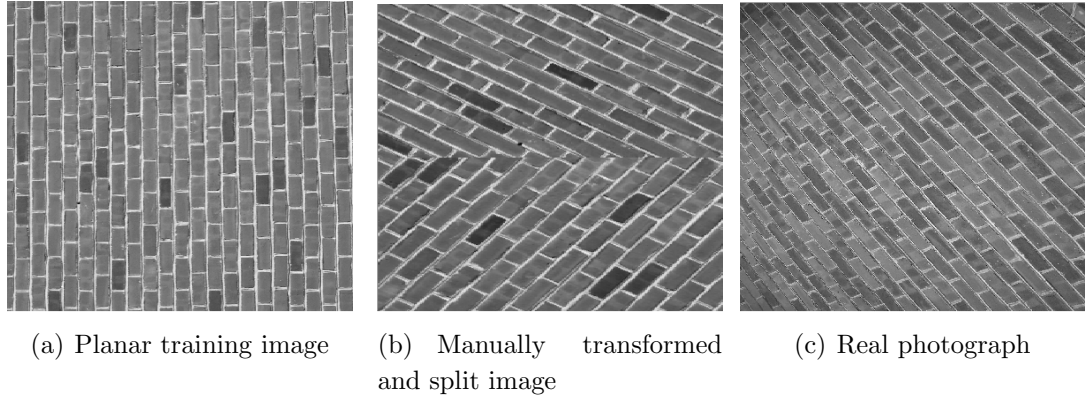


Figure 7.2: Images of brick walls considered for dictionary learning and model evaluation purposes.

next sections are devoted to a detailed description of our marked point process model and the presentation of a suitable inference concept.

Marks

In the context of this research project, marks correspond to Euclidean similarity transformations of representative patches of texture with no 3D structure, as explained above. We therefore first define a suitable library consisting of rectangular image sections cut out from a training scenario which shows the relevant texture, but no depth, warping or rotational effect.

The library may be constructed in an automatic way, e.g. through an appropriate implementation of the K-SVD approach by Aharon et al. (2006). Coupling convenient singular value decomposition [Eckart and Young (1936)] with a generalized k-means clustering algorithm [Hartigan and Wong (1979)], K-SVD is capable of extracting the most informative features from a large image matrix.

At present, we prefer to randomly select representative sections from the training image. In compliance with the work of Aharon et al. (2006), we call these sections *atoms* and the library a *dictionary*. To facilitate the marking process, we normalize the pixel values in each atom, such that they have zero mean. We define marks as Euclidean similarity transformations of dictionary atoms. The entire mark space M thus contains the labels of the atoms and the domains of the geometric parameters $\boldsymbol{\eta}$, i.e. the scaling parameters η_1 and η_2 as well as the rotation parameter η_ρ . Concerning the scaling effects, η_1 shrinks or stretches a dictionary atom in horizontal and η_2 in vertical direction. Rotations are always carried out with respect to the center of the atom. To keep the set of the geometric parameters as small as possible, we do not consider shifts at

this point. Shifts are indirectly involved in the modeling, as move steps of the M-H inference framework proposed further below. Fig. 7.3 illustrates results from the scaling and rotation of a small rectangular cut-out of Fig. 7.2 (c).




Figure 7.3: Euclidean similarity transformations (b) – (c) of a dictionary atom (a) cut out from a training image without depth effect (cf. Fig. 7.2 (a)).

We denote the dictionary by $\mathbf{a} = \{a_1, \dots, a_K\}$. Assuming that the k -th untransformed dictionary atom a_k takes values $a(w^{(0)}|k)$ on $w^{(0)} \in W^{(0)}$, where $w^{(0)} = (w_1^{(0)}, w_2^{(0)})^\top$, a scaling of the atom yields

$$\begin{pmatrix} w_1 \\ w_2 \end{pmatrix} = \begin{pmatrix} \eta_1 & 0 \\ 0 & \eta_2 \end{pmatrix} \begin{pmatrix} w_1^{(0)} \\ w_2^{(0)} \end{pmatrix}.$$

That is, through the effects of η_1 and η_2 , each pixel value $a(\cdot|k)$ on $w^{(0)} \in W^{(0)}$ is mapped to a new location $w \in W$ where $w = (w_1, w_2)^\top$. The orientation of the dictionary atom can be influenced by a coordinate mapping of the form

$$\begin{pmatrix} w_1 \\ w_2 \end{pmatrix} = \begin{pmatrix} \cos \eta_\rho & -\sin \eta_\rho \\ \sin \eta_\rho & \cos \eta_\rho \end{pmatrix} \begin{pmatrix} w_1^{(0)} \\ w_2^{(0)} \end{pmatrix},$$

which corresponds to Eq. 3.10 in Section 3.1. The degree of rotation is therefore determined by η_ρ . For a simultaneous change of both the surface area and the orientation of the dictionary atom, we implement

$$\begin{pmatrix} w_1 \\ w_2 \end{pmatrix} = \begin{pmatrix} \eta_1 \cos \eta_\rho & -\eta_2 \sin \eta_\rho \\ \eta_1 \sin \eta_\rho & \eta_2 \cos \eta_\rho \end{pmatrix} \begin{pmatrix} w_1^{(0)} \\ w_2^{(0)} \end{pmatrix}. \quad (7.7)$$

We write $a(w^{(0)}|k) = a(w|k, \boldsymbol{\eta})$. Fig. 7.4 gives an example of how a concrete set of parameter values affects the geometric properties of a quadratic object. For more technical details and additional information, see Foley et al. (1997, ch. 5).

Before marking an image location with a scaled and rotated dictionary atom, a color matching between the mark and the respective image section should be carried out. Recalling that the pixel values in each dictionary atom have zero mean, this can easily be performed by adding the mean pixel intensity in the image section to the pixel values of the mark. In case that marks partly cover

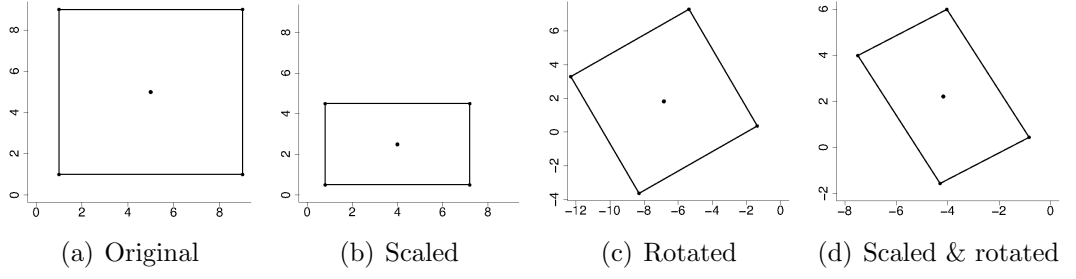


Figure 7.4: Euclidean similarity transformations of a square using $\eta_1 = 0.8$ and $\eta_2 = 0.5$ in (b) and (d), and $\eta_\rho = \frac{2}{3}\pi$ in (c) and (d).

each other, we suggest an averaging at the respective locations which we label by $(*)$. Assuming that J marks overlap in w and assuming that these J marks are independent and unpredictable due to Def. 2.4 – 2.5, we propose to compute

$$a^{(*)}(w|k^{(*)}, \boldsymbol{\eta}^{(*)}) = \frac{1}{J} \sum_{j=1}^J a(w|k^{(j)}, \boldsymbol{\eta}^{(j)}) \quad (7.8)$$

$$\text{and} \quad \sigma_a^2{}^{(*)}(w|k^{(*)}, \boldsymbol{\eta}^{(*)}) = \frac{1}{J^2} \sum_{j=1}^J \sigma_a^2(w|k^{(j)}, \boldsymbol{\eta}^{(j)}) .$$

In doing so, $a(w|k^{(j)}, \boldsymbol{\eta}^{(j)})$ corresponds to the pixel value of the $k^{(j)}$ -th geometrically transformed and color-adjusted dictionary atom at location w . Its variance is denoted by $\sigma_a^2(w|k^{(j)}, \boldsymbol{\eta}^{(j)})$. For the sake of simplicity, we replace $\sigma_a^2(\cdot|k^{(j)}, \boldsymbol{\eta}^{(j)})$ by a constant variance σ_a^2 whenever appropriate. The second line of Eq. (7.8) then becomes

$$\sigma_a^2{}^{(*)}(w|k^{(*)}, \boldsymbol{\eta}^{(*)}) = \frac{1}{J} \sigma_a^2 . \quad (7.9)$$

Having explained how marks are defined and obtained in the context of this research project, the following section is concerned with the introduction of our actual modeling framework.

Model

As briefly outlined in the beginning of Section 7.3, our objective is the implementation of a Bayesian hierarchical inference mechanism that allows us to extract geometric knowledge from the image of a textured 3D scene. Therefore, several distributional definitions need to be established. On the one hand, a model describing the pixel value distribution in the given image has to be determined. On the other hand, meaningful prior assumptions have to be made on the spatial distribution of the marked points and on the marks themselves. Hyper priors may additionally be specified.

Our hierarchical point process model is assumed to act on a rectangular grid of pixels defining the observation window W . We denote the given image matrix by $Z(\cdot)$ and the pixel value in $w \in W$ by $Z(w) = z_w$. This project is confined to the analysis of gray-scale images, which means that $Z(\cdot)$ takes values on a continuous interval. Therefore, a Gaussian model appears to be suitable for describing the image data. The parameters of the model, i.e. mean and variance, are assumed to depend on the current state \mathbf{y} of the marked point process which we define as $\mathbf{y} = \{(x, m_x) : x \in \mathbf{x}, m_x \in M\}$, in conformity with our previous notation (see also Eq. (2.4)). Depending on whether $w \in W$ is currently covered by multiple objects, exactly one object, or no object at all, it lies in the subset $W_a^{(*)}$, in $\{W_a \setminus W_a^{(*)}\}$, or in $W_{\bar{a}}$, where \bar{a} indicates the absence of a coverage and $\{W_{\bar{a}} \cup W_a\} = W$. In accordance with this decomposition as well as with Eq. (7.8) – Eq. (7.9), we define the expected pixel value in w as $a^{(*)}(w|\mathbf{x}, k_w^{(*)}, \boldsymbol{\eta}_{k_w}^{(*)})$, $a(w|\mathbf{x}, k_w, \boldsymbol{\eta}_{k_w})$, or $\mu_{\bar{a}}$, and the respective variance as $\sigma_a^{2(*)}(w|\mathbf{x}, k_w^{(*)}, \boldsymbol{\eta}_{k_w}^{(*)})$, σ_a^2 , or $\sigma_{\bar{a}}^2$. As before, \mathbf{x} denotes a realization of the non-marked point process \mathbf{X} and k_w the label(s) of the dictionary atom(s) covering w . In summary, the mean and variance measures depend on the current state of the point process \mathbf{x} marked by dictionary atoms that have been taken from $\mathbf{a} = \{a_1, \dots, a_K\}$ and geometrically transformed through individual scaling and rotation parameters $\boldsymbol{\eta} = \{\eta_1, \eta_2, \eta_\rho\}$.

The partition of the observation window into non-covered “background sections” and “foreground sections” covered by at least one geometrically transformed dictionary atom implicates that the Gaussian data model can be decomposed into

$$\begin{aligned}
 \prod_{w \in W} \underbrace{f(Z|\mathbf{y})}_{f(z_w|\mathbf{x}, k_w, \boldsymbol{\eta}_{k_w})} &\propto \left(\frac{1}{\sigma_{\bar{a}}^2}\right)^{\frac{|W_{\bar{a}}|}{2}} \exp \left\{ - \sum_{w \in W_{\bar{a}}} \frac{(z_w - \mu_{\bar{a}})^2}{2\sigma_{\bar{a}}^2} \right\} \\
 &\times \prod_{w \in \{W_a \setminus W_a^{(*)}\}} \left(\frac{1}{\sigma_a^2}\right)^{\frac{1}{2}} \exp \left\{ - \frac{(z_w - a(w|\mathbf{x}, k_w, \boldsymbol{\eta}_{k_w}))^2}{2\sigma_a^2} \right\} \\
 &\times \prod_{w \in W_a^{(*)}} \left(\frac{1}{\sigma_a^{2(*)}(w|\mathbf{x}, k_w^{(*)}, \boldsymbol{\eta}_{k_w}^{(*)})}\right)^{\frac{1}{2}} \exp \left\{ - \frac{(z_w - a^{(*)}(w|\mathbf{x}, k_w^{(*)}, \boldsymbol{\eta}_{k_w}^{(*)}))^2}{2\sigma_a^{2(*)}(w|\mathbf{x}, k_w^{(*)}, \boldsymbol{\eta}_{k_w}^{(*)})} \right\},
 \end{aligned} \tag{7.10}$$

given the current state of the marked point process, \mathbf{y} . Here, $|W_{\bar{a}}|$ equates to the number of pixels in $W_{\bar{a}}$. We recall that proposals of how to compute the mean and variance terms are given in Eq. (7.8) – Eq. (7.9).

Since the Gaussian likelihood in Eq. (7.10) depends, inter alia, on the current estimate \mathbf{x} of the latent point process \mathbf{X} , our Bayesian modeling framework requires an appropriately specified point process prior. As discussed in Section 2.1,

this thesis deals with simple point processes, meaning that we do not consider situations in which two or more than two points may be assigned to exactly the same location. Here, we furthermore assume a random number of points that are independently and homogeneously distributed on the discrete grid of coordinates determining W . We decide for a homogeneous Poisson process prior according to Def. 2.8 and Eq. (2.5), and thus have

$$f(\mathbf{x}|\beta) = \exp\{-\beta|W|\} \beta^{n(\mathbf{x})}$$

for a given realization \mathbf{x} of the latent point process \mathbf{X} . The intensity parameter β serves as a hyper parameter. Depending on the existing knowledge, it may either be fixed or assigned with a suitable *hyper prior* $p(\beta)$.

Concerning the drawing of the dictionary atoms from $\mathbf{a} = \{a_1, \dots, a_K\}$, we a priori assume that

$$k^{(i)} \sim U\{1, 2, \dots, K\}, \quad \forall i \in \{1, \dots, n(\mathbf{x})\},$$

which means that each atom is selected with probability $\frac{1}{K}$. Similarly, we propose to take non- or little-informative priors as a basis for the scaling and rotation parameters η_1 , η_2 and η_ρ . Here, we use uniform probability distributions on discrete intervals.

Finally, adequate priors need to be specified for the mean and variance parameters in Eq. (7.10). While the expected values in W_a are determined by the pixel value distribution inside the selected and geometrically transformed dictionary atoms and hence by the current realizations of \mathbf{x} , k and $\boldsymbol{\eta}$, we sample $\mu_{\bar{a}}$ from a specific parameter prior. The same holds for the unknown variances σ_a^2 and $\sigma_{\bar{a}}^2$. For computational reasons, we suggest a normal-inverse- χ^2 or a normal-inverse- Γ model as described by O'Hagan et al. (2004, ch. 11) and thus have e.g.

$$\begin{aligned} (\mu_{\bar{a}}, \sigma_{\bar{a}}^2) &\sim N-inv-\chi^2(\mu_{0_{\bar{a}}}, \kappa_{0_{\bar{a}}}, \alpha_{0_{\bar{a}}}, \vartheta_{0_{\bar{a}}}), & \sigma_{\bar{a}}^2 &\sim inv-\chi^2(\alpha_{0_{\bar{a}}}, \vartheta_{0_{\bar{a}}}), \\ \mu_{\bar{a}}|\sigma_{\bar{a}}^2 &\sim N(\mu_{0_{\bar{a}}}, \frac{\sigma_{\bar{a}}^2}{\kappa_{0_{\bar{a}}}}), \end{aligned}$$

$$\text{and} \quad \sigma_a^2 \sim inv-\chi^2(\alpha_{0_a}, \vartheta_{0_a}).$$

It is convenient to specify a Gaussian prior for the mean parameter and inverse- Γ or inverse- χ^2 priors for the variance parameters, since these densities correspond to conjugate densities of the Gaussian data model. Such a set-up has the advantage that the posteriors take the same well-defined forms as the priors, as explained in Section 2.5. Given σ_a^2 , $\sigma_a^{2(*)}(\cdot|\mathbf{x}, k^{(*)}, \boldsymbol{\eta}_k^{(*)})$ can be calculated due to Eq. (7.9).

From our prior and data distributional assumptions, we derive a posterior density of the form

$$p(\mathbf{y}|Z) = \frac{f(Z|\mathbf{y}) p(\mathbf{y})}{\mathcal{C}(Z)}, \quad (7.11)$$

where

$$p(\mathbf{y}) = p(\mathbf{x}, \boldsymbol{\theta}) = f(\mathbf{x}|\beta) p(\beta) p(\sigma_a^2) p(\mu_{\bar{a}}|\sigma_a^2) p(\sigma_a^2) \prod_{i=1}^{n(\mathbf{x})} p(\boldsymbol{\eta}^{(i)}) p(k^{(i)}) , \quad (7.12)$$

according to Eq. (2.24). $\mathcal{C}(\cdot)^{-1}$ is the normalizing constant of the posterior model. Since it is numerically intractable, Eq. (7.11) is not available in complete form and hence only accessible via a simulation framework (cf. Eq. 2.25). In the following, we propose a birth-death-move M-H inference scheme which is similar to the step sequence in Alg. 2.2 and the sampler suggested by Lafarge et al. (2010) (cf. Section 7.2).

Inference

To estimate realizations \mathbf{y} of the latent marked point process \mathbf{Y} given the data matrix Z , we iteratively match dictionary atoms with the data by placing geometrically transformed versions of them in W , or by modifying or deleting existing ones. As mentioned right above, a birth-death-move M-H sampler appears appropriate for this purpose. We first decompose the joint posterior $p(\mathbf{y}|Z)$ from Eq. (7.11) into the conditional posteriors

$$\begin{aligned} \text{(A)} \quad & p(\mathbf{x}, \boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}} | Z, \mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2) \quad \text{and} \\ \text{(B)} \quad & p(\sigma_a^2 | Z, \mathbf{x}, \boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}}), \quad p(\sigma_{\bar{a}}^2 | Z, \mathbf{x}, \boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}}), \\ & p(\mu_{\bar{a}} | Z, \mathbf{x}, \boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}}) . \end{aligned}$$

To sample from the full conditional marked point process posterior in (A), we use an M-H sampler conforming to Alg. 2.2 [see also Møller and Waagepetersen (2003, ch. 7)]. This reversible jump framework comprises birth steps where the addition of a new marked point is proposed, death steps where one randomly selected marked point may be removed, and move steps that allow to modify a marked point. We suggest three types of modifications, the replacement of the dictionary atom while retaining its geometric parameters, the change of the geometric mark parameters while retaining the label of the atom, and the shifting of the marked point. As usual, we suggest to take uniform or truncated Gaussian densities as a basis for the transition proposals.

After each birth, death or move step, a new sample of mean and variance parameters has to be drawn from the updated full conditional posteriors. Choosing a normal-inverse- χ^2 prior model for the mean and the variances, the posteriors in (B) correspond to Gaussian and inverse- χ^2 densities with posterior parameters conforming to weighted averages of the prior parameters and statistics from the data. A Gibbs sampling routine similar to Alg. 2.1 with reversible jump M-H steps according to Alg. 2.6 allows to iteratively simulate from the

Algorithm 7.1: Bayesian inference scheme to estimate the posterior of a latent marked point process under the assumption of a Gaussian data model. $|W_{\bar{a}}|$ and $|W_a|$ denote the current number of pixels in $W_{\bar{a}}$ and W_a , $\bar{Z}_{W_{\bar{a}}}$ and \bar{Z}_{W_a} are the respective mean pixel intensities, and $S(Z_{W_{\bar{a}}})$ and $S(Z_{W_a})$ the empirical standard deviations.

Data: Image matrix Z , prior parameters $\{\alpha_{0_{\bar{a}}}, \vartheta_{0_{\bar{a}}}, \mu_{0_{\bar{a}}}, \kappa_{0_{\bar{a}}}, \alpha_{0_a}, \vartheta_{0_a}\}$, dictionary $\mathbf{a} = \{a_1, \dots, a_K\}$, prior intervals $\{\eta_1^{\min}, \dots, \eta_1^{\max}\}$, $\{\eta_2^{\min}, \dots, \eta_2^{\max}\}$, $\{\eta_\rho^{\min}, \dots, \eta_\rho^{\max}\}$, burn-in threshold T_0

Result: Point process realizations $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$,
 mark parameters $\boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(T_0)}}, \dots, \boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(T)}}$,
 mean and variance parameters $\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(T_0)}, \dots, \{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(T)}$

Set $t = 1$;

while $t \leq T$ **do**

determine $\mathbf{x}^{(t)}$ and $\boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(t)}}$ according to Alg. 2.2;

compute $\mu_{p_{\bar{a}}} = \left(\frac{\kappa_{0_{\bar{a}}}}{\kappa_{0_{\bar{a}}} + |W_{\bar{a}}|}\right) \mu_{0_{\bar{a}}} + \left(\frac{|W_{\bar{a}}|}{\kappa_{0_{\bar{a}}} + |W_{\bar{a}}|}\right) \bar{Z}_{W_{\bar{a}}}$;

$\kappa_{p_{\bar{a}}} = \kappa_{0_{\bar{a}}} + |W_{\bar{a}}|$;

$\alpha_{p_{\bar{a}}} = \alpha_{0_{\bar{a}}} + |W_{\bar{a}}|$;

$\alpha_{p_{\bar{a}}} \vartheta_{p_{\bar{a}}} = \alpha_{0_{\bar{a}}} \beta_{0_{\bar{a}}} + (|W_{\bar{a}}| - 1) S(Z_{W_{\bar{a}}})^2 + \frac{\kappa_{0_{\bar{a}}}}{\kappa_{0_{\bar{a}}} + |W_{\bar{a}}|} (\bar{Z}_{W_{\bar{a}}} - \mu_{0_{\bar{a}}})^2$;

compute $\alpha_{p_a} = \alpha_{0_a} + |W_a|$;

$\alpha_{p_a} \vartheta_{p_a} = \alpha_{0_a} \beta_{0_a} + (|W_a| - 1) S(Z_{W_a})^2 + \frac{\kappa_{0_a}}{\kappa_{0_a} + |W_a|} (\bar{Z}_{W_a} - \mu_{0_a})^2$;

draw $\sigma_a^2 \sim \text{inv}\chi^2(\alpha_{p_a}, \vartheta_{p_a})$;

$\sigma_{\bar{a}}^2 \sim \text{inv}\chi^2(\alpha_{p_{\bar{a}}}, \vartheta_{p_{\bar{a}}})$;

$\mu_{\bar{a}} | \sigma_{\bar{a}}^2 \sim N\left(\mu_{p_{\bar{a}}}, \frac{\sigma_{p_{\bar{a}}}^2}{\kappa_{p_{\bar{a}}}}\right)$;

set $\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(t)} = \{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}$;

end

Return $\mathbf{x}^{(T_0)}, \dots, \mathbf{x}^{(T)}$;

Return $\boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(T_0)}}, \dots, \boldsymbol{\theta}^{-\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(T)}}$;

Return $\{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(T_0)}, \dots, \{\mu_{\bar{a}}, \sigma_{\bar{a}}^2, \sigma_a^2\}^{(T)}$;

whole posterior model. Alg. 7.1 illustrates an appropriate implementation of this inference concept. To evaluate the practical application of our Bayesian modeling approach, we take the manually constructed split image in Fig. 6.5 (b) and the real-world scenario in Fig. 6.5 (c) as a basis.

7.4 Case Studies

First, we use Alg. 7.1 for the analysis of the split brick image in Fig. 6.5 (b) which comprises two similarity transformations of the planar training scene in Fig. 6.5 (a). The geometrically modified scenes are arranged one below the other. While the upper image section labeled by I results from a geometric transformation through the parameter combination $\boldsymbol{\eta}_I = \{\eta_1, \eta_2, \eta_\rho\}_I = \{0.8, 0.6, -0.2\pi\}$, the lower part corresponds to a transformation through $\boldsymbol{\eta}_{II} = \{0.7, 0.7, 0.3\pi\}$.

We take a section from the training scene as a basis for the dictionary learning process. From this section containing 300×300 pixels, we cut out two rectangular patches of size 75×75 pixels which we define as our dictionary atoms (see Fig. 7.5 (a)). A priori, we assume that each atom becomes selected with a probability of 50%. Concerning the priors of the geometric parameters, we take discrete uniform densities on sufficiently fine grids as a basis. We a priori expect that $0.5 \leq \eta_1 \leq 1.0$, $0.5 \leq \eta_2 \leq 1.0$ and $-\frac{1}{2}\pi \leq \eta_\rho \leq \frac{1}{2}\pi$, and base the transition proposals on the same uniform distributional conditions.

Starting with an empty set of marked points, we execute 200.000 iterations of Alg 7.1, i.e. $T = 200.000$. We eliminate the first 125.000 results as burn-in states and base our further analysis on every 750th realization contained in the remaining chain, in order to avoid autocorrelation effects. The extracted samples describe the posterior distribution of the latent marked point process given the visible image data.




Figure 7.5: Modeling of a latent marked point process. (a) shows two quadratic patches cut out from a planar training section and used as dictionary atoms, (b) corresponds to the manually constructed scenario of interest, and (c) visualizes one realization of the marked point process.

Fig. 7.5 (c) visualizes one representative estimate of the marked point process. We see that this sample from the posterior looks very similar to the original scene in Fig. 7.5 (b). On closer inspection, however, the long edges of the bricks are better fitted than the short edges. We thus suspect that here, our algorithm estimates the scaling effect in horizontal direction, η_1 , more precisely than the scaling effect in vertical direction, η_2 . Fig. 7.6 confirms this assumption.

Furthermore, we observe a high level of accuracy and little variation in the estimates of η_ρ .

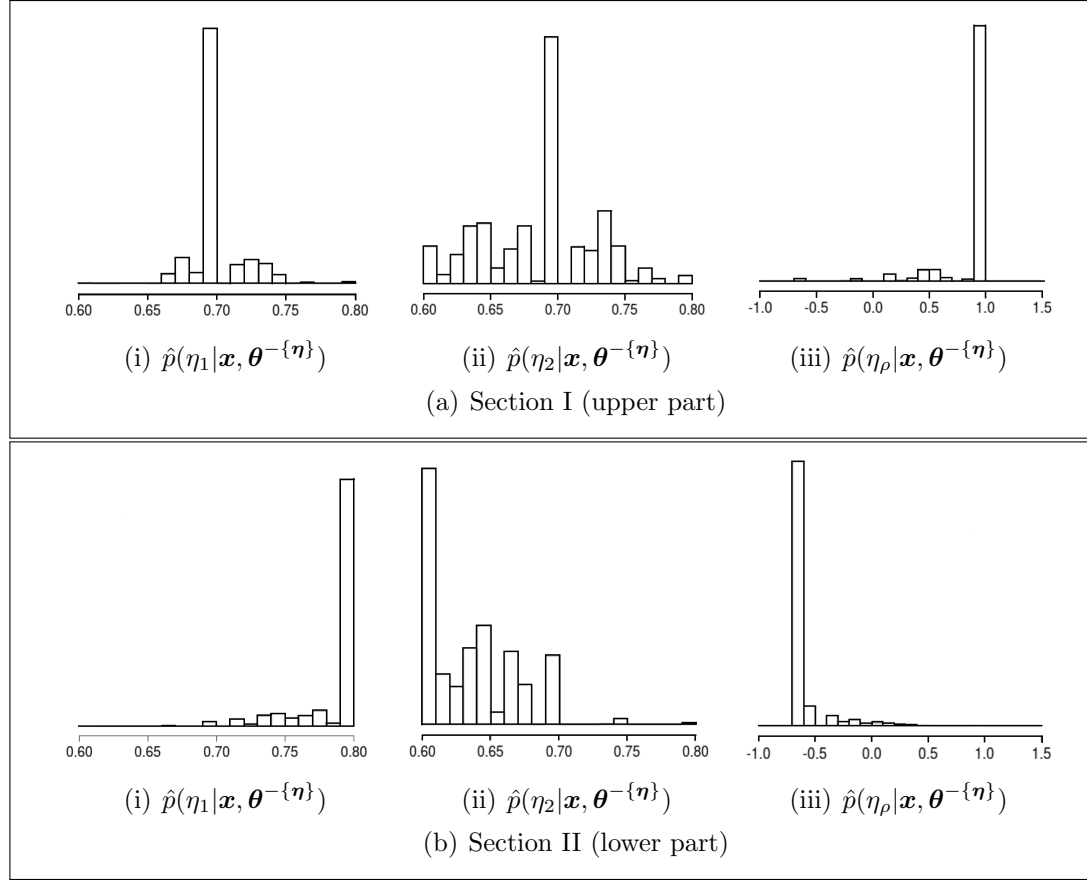


Figure 7.6: Estimates of the full conditional posteriors of the geometric parameters, subdivided into results with respect to the upper (a) and results with respect to the lower (b) section of Fig. 7.5.

While frequentist modeling approaches yield maximum-likelihood (ML) estimates, posterior results from a Bayesian inference procedure allow for the computation of the maximum a posteriori (MAP) estimates, as already explained in Section 4.4 and Section 5.3. MAP estimates are equal to the modes of the posterior. Since our prior and transition distributional assumptions ensure that the realizations of $\boldsymbol{\eta}$ take values on discrete bounded intervals, it is sufficient to search for the most frequent parameter combination among the states of the Markovian chain. This search is carried out pixel-wisely, which means that, for each $w \in W$ covered by a dictionary atom in at least 80% of all marked point process realizations, we identify the most frequently accepted set of atom-transforming parameters $\hat{\boldsymbol{\eta}}_{\text{MAP}}$. The results for all three geometric parameters are shown in Fig. 7.7.

We see that, apart from some gaps in the boundary region between the upper and lower image section, almost every pixel is steadily covered by marks.




Figure 7.7: Pixel-wise MAP estimates of the geometric parameters $\boldsymbol{\eta} = \{\eta_1, \eta_2, \eta_\rho\}$. White patches indicate that in less than 80% of the marked point process realizations, the respective coordinates are covered by a geometrically transformed dictionary atom.

According to the conclusions drawn from Fig. 7.10, the pixel-wise visualization of the MAP estimates points out that the true rotation effects in the upper and the lower image sections are almost perfectly reproduced. The MAP estimates of the scaling parameters in horizontal direction suffer from slight variations, whereas the estimation of the scaling effects along the vertical axis turns out to be grossly inaccurate. As stated above, this means that the long brick edges are better identified than the small edges.

To examine the speed of convergence of Alg. 7.1, Fig. 7.8 illustrates the increase in the amount of points (a), the gradient of the log-likelihood given the parameter estimates (b), and the improvement in the goodness-of-fit (c). The goodness-of-fit is measured through the *mean squared error* (MSE), that is, through the sum of the squared differences between the true pixel values and the pixel intensity under the current state of the marked point process.

Fig. 7.8 (b) confirms that our Bayesian hierarchical modeling framework performs well in an application to Fig. 7.5 (b). The increase in the log-likelihood given the current estimate of the point and mark process is considerable and reaches its equilibrium state after about 100.000 iterations. The same holds for the decrease in the MSE measuring the similarity between the given image and the pixel value distributions stemming from the marked point process realizations. Regarding the development of the point intensity in Fig. 7.8 (a), however,




Figure 7.8: Convergence of the algorithm and overall model performance. The horizontal dashed line in (b) illustrates the value of the log-likelihood calculated under the assumption of an empty set of marked points.

an equilibrium has not been reached. We therefore suspect that our algorithm tends to identify redundant points. By including strict interaction constraints, this problem may be controlled or tackled.

Having shown that Alg. 7.1 allows to detect local geometries in a manually transformed brick image, we now analyze a non-manipulated scenario of size 1050×1000 pixels, where the geometric ground truth is unknown (see Fig. 7.9 (b)). For logical reasons, we regard the previously used dictionary as too restrictive. Instead of keeping it unchanged, we thus enhance its size, while reducing the dimension of the single atoms. Fig. 7.9 (a) shows seven small rectangular patches cut out from the training section and defined to serve as dictionary atoms.

We make the same prior and transition distributional assumptions as before. In contrast to above, however, we now expect that each of the three geometric parameters takes values between 0.3 and 1.0. Since we are not interested in perfectly reconstructing the original image, but in determining its local scale and rotation parameters, we subdivide it into 12 partitions of size 350×250 pixels. Parallel programming therefore yields a gain in computer efficiency and time.

We apply Alg. 7.1 to the 12 sections and set the total number of iterations per section to 200,000. Again, the first 125,000 outcomes are regarded as burn-in states, and every 750th element of the remaining chain is taken as a basis for describing the posterior of the marked point process. Fig. 7.9 (c) exemplifies one randomly selected realization of the latent marked point process. The closer a brick is located to the camera, the bigger it appears in the 2D image. Since the proposed algorithm turns out to recognize large objects better than objects of small size, some of the 12 sections are densely packed with marked points, whereas others contain considerable gaps.

We hypothesize that the estimated rotation effects are the same in each part




Figure 7.9: Estimation of a marked point process realization. (a) shows seven quadratic patches cut out from a planar training section and used as dictionary atoms, (b) corresponds to the image of interest, and (c) visualizes one realization of the marked point process. For reasons of clarity, the estimated points themselves are not highlighted.

of the image, but expect clear differences in the full conditional posteriors of the scaling parameters. To investigate these assumptions, we inspect two parts of the point process realization in Fig. 7.9 (c) more closely, which we labeled by I and II. Since the distance from the camera to part I of the original 3D plane has obviously been larger than its distance to part II, the bricks in section I of the given 2D image are smaller than the bricks in section II. The question comes up whether Alg. 7.1 allows us to clearly quantify this difference.

Fig. 7.10 visualizes the full conditional posteriors of the parameter set $\boldsymbol{\eta}$, separately computed based on the spatial pixel value distribution in image section I and in image section II. As expected, the full conditional posteriors of the rotation parameters are of almost identical shape, whereas the histograms illustrating the posteriors of the scaling effects are quite dissimilar. Although there is much variability in the estimation of η_1 and, in particular, of η_2 , a difference in the modes and means is clearly visible. It turns out that

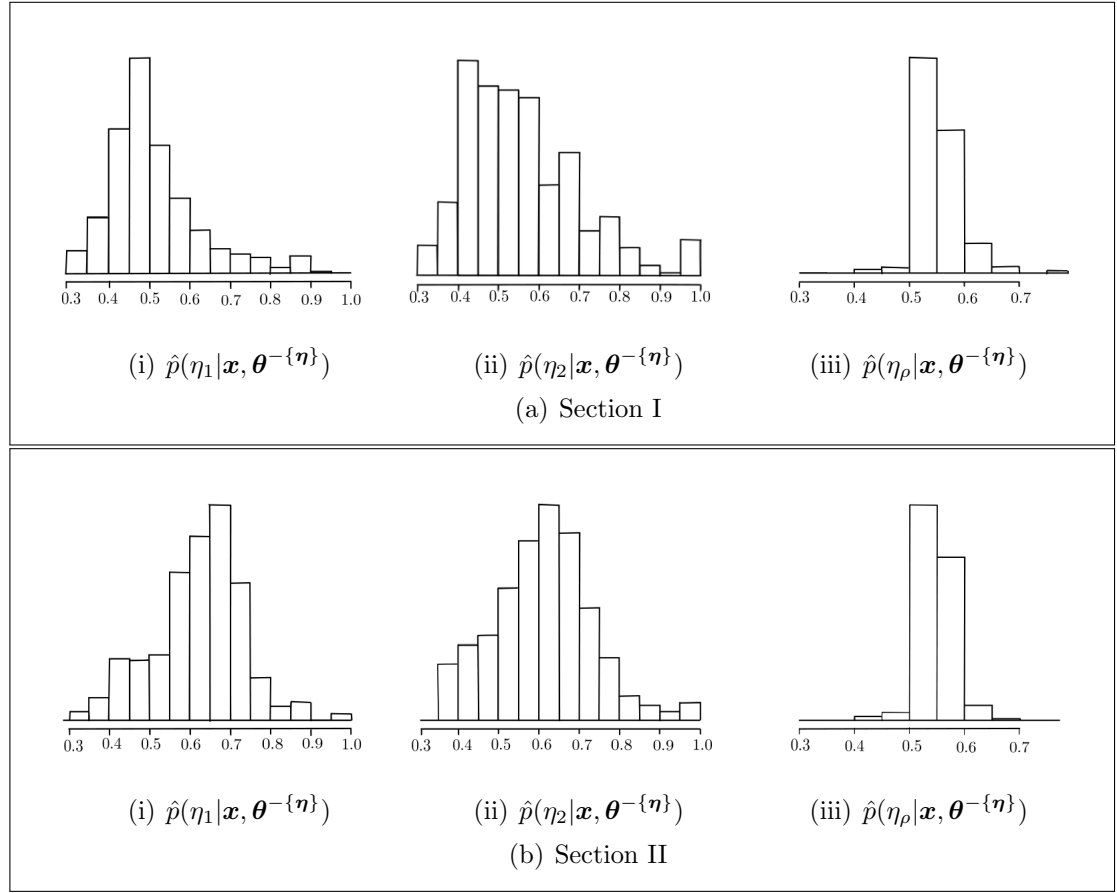


Figure 7.10: Empirical distributions of the geometric posterior estimates for both parts of the split image.

$\hat{\boldsymbol{\eta}}_{I_{\text{MAP}}} = \{0.46, 0.44, 0.53\}$ is the most frequently accepted parameter combination in I, while $\hat{\boldsymbol{\eta}}_{II_{\text{MAP}}} = \{0.66, 0.63, 0.53\}$ is the MAP estimate for section II.

7.5 Outlook

We have started to develop a Bayesian marked point process model that assesses 3D geometries in 2D images of textured surfaces. For the sake of simplicity, we have used images with pixel values measurable on a continuous gray-level scale. As data distribution, a Gaussian likelihood has thus proven to be appropriate (cf. Eq. (7.10)). A priori, we have assumed that the latent point process follows a homogeneous spatial Poisson process distribution (cf. Eq. (2.6)).

Each mark has been specified by two groups of parameters, the label of the marking object and its geometric properties. Marking objects have been defined as distinctive patches of texture without depth effect, stored as atoms in a dictionary. As potential geometric deformations of these objects, we have considered Euclidean similarity transformations, i.e. shrinking, stretching and

rotations (cf. Eq. (7.7)). Shifts have not been accounted for at this points. All model parameters have been assigned with conjugate or uniform priors.

In terms of inference, we have developed Alg. 7.1 which matches appropriately scaled and rotated dictionary atoms with the given image. The reversible jump MCMC algorithm alternates between simple Gibbs sampling stages (cf. Alg. 2.1) and a birth-death-move M-H routine (cf. Alg. 2.6). From the posterior distributional properties of the model parameters, we have drawn conclusion about the local geometries in the image. Our further target is to describe the orientation of the camera towards the 3D space based on the marked point process realizations, especially focusing on the estimated local 2D geometries. Images of brick walls with both known and unknown geometric characteristics have been used to evaluate the entire modeling framework (cf. Fig. 6.5).

Our Bayesian marked point process approach delivers promising results in that it detects and describes local 2D geometries in images of textured 3D scenes. However, there is still room for improvement and further refinement. On the one hand, our modeling routine tends to strongly overestimate the number of latent points, which makes it difficult to deduce 3D information from the spatial point distribution. This problem may be approached by the inclusion of interaction constraints (cf. Section 2.3). On the other hand, an inhomogeneous point process prior with exponential or perspective scaling attributes as suggested in Section 6.4 may be more appropriate than a homogeneous prior model, though the increase in the model complexity should be evaluated critically. If we decide to include location-dependent scaling constraints in our point process prior, similar constraints should also be embedded in the priors of the geometric parameters.

Summing up, we assume that an enhancement of our Bayesian hierarchical point process model by appropriately specified interaction and scaling functions may yield a more accurate geometric description of the local image features. In the optimal case, it even allows to directly infer shape from texture.

8 Discussion

In this concluding section, a summary of our research projects and results is given. We discuss the overall performance of the modeling approaches newly introduced in Chapter 4 – 7, point out their advantages and disadvantages, and reveal potential for future advancement and improvement. Thereby, special emphasis is given to evaluating the effort made to couple spatial statistics with image analysis.

Our first contribution is a Bayesian inference framework for analyzing inhomogeneous spatial point processes, where the heterogeneity stems from a scaling function that locally adjusts the arrangement and relative amount of the points (see Chapter 4). Intensity and interaction become equally scaled in the sense that, in different regions of the observation window, the point packing density varies only by a scale factor [Hahn et al. (2003)]. In general, locally scaled point process models are not presentable in a complete form, as a normalizing constant is missing in the likelihood. A posterior model for the unknown intensity, interaction and scaling parameters thus implies two unknown terms, the normalizing constant of the likelihood which depends on the unknown parameters and that of posterior itself which depends on the data. Concentrating on locally scaled Strauss models, an exchangeability framework according to Murray et al. (2012) proves to be suitable for dealing with the double-intractability. Since it requires the implementation of an auxiliary variable scheme where the auxiliary variables correspond to samples from the likelihood, methods for drawing realizations from an incomplete point process density need to be considered. We recommend to generate perfect simulations via dominated coupling from the past (CFTP) [Berthelsen and Møller (2003)].

In an application to simulated point patterns, it turns out that our Bayesian inference framework is capable of accurately reflecting the ground truth. The maximum a posterior estimates are similar to the maximum likelihood estimates resulting from an application of the pseudo likelihood approach suggested by Prokešová et al. (2006).

As every Markov chain Monte Carlo routine, the exchange algorithm only returns sound results if a sufficiently high number of iterations has been executed and the chain has for certain reached its equilibrium state. The incorporated CFTP sampler requiring the implementation of multiple birth-death processes may be computationally expensive. If the point intensity in a given pattern is high and the scaling and interaction effects are strong, that is, if the pattern is close to regular and densely packed with points, CPU run-times tend to be excessively long. In such cases, approximate MCMC-based samplers may be

more appropriate than a CFTP algorithm, despite the loss in accuracy.

Subsequent to the theoretically oriented Chapter 4, we demonstrate by means of selected research studies how expedient and relevant locally scaled point processes turn out to be in applied sciences (see Chapter 5 – 7).

In Chapter 5, we use our newly proposed CFTP-based exchangeability framework to approach an agricultural research problem. We analyze two series of cross-sections through the stems of maize plants. Each series corresponds to one genotype. Regarding the spatial distribution of the vascular bundles in the stems as a realization of an inhomogeneous point process, our objective is the detection of genotype-specific characteristics in the packing densities. In other words, we are interested in the derivation of classification rules from the arrangement of the vascular bundles in the stems.

With the aid of adequate preprocessing mechanisms, the bundles are replaced by points and projected onto circular observation windows that reflect the proportion between the stem diameters. The diameters measured for the type-II plants are on average 0.69 times the radii obtained for the type-I plants. Furthermore, it becomes immediately apparent that the cross-sectional area of the maize stems can be partitioned into two sections, an outer section where the bundles are very densely packed and an inner section where their density is considerably lower. Assuming homogeneity inside both sections, a demarcation line between them can be determined in a model-free manner.

We have developed a step scaling function which is proper in that it is well-defined and numerically tractable with regard to the calculation of the relevant locally scaled volume and distance measures. Under minor formal adjustments, it may also be applied to rectangular data spaces and point processes with more than two gradations in the density. For efficiency purposes, we propose a two-stage inference procedure to preliminarily examine the inner stem areas via a Bayesian homogeneous Strauss model and to consider the gained posterior information as prior distributional knowledge for the modeling of the complete datasets. The full databases are analyzed by means of the exchange algorithm, where the auxiliary variables are partly generated via perfect simulation and partly via an approximate MCMC sampler [Liang (2010)]. In terms of convergence, it turns out that our two-phases framework is more suitable for modeling the type-I than for describing the type-II data. This is also reflected in the fact that, with regard to genotype II, the outcomes obtained from the first stage of the analysis differ considerably from the final results.

Our most striking conclusion concerns the estimated scaling factor quantifying the difference in the packing density between the inner and the outer stem areas, which is considerably larger for genotype II than for genotype I. The overall bundle intensity tends to be higher in the plants of type I than in the plants of

type II. Concerning the interaction parameters, a concise conclusion cannot be drawn. It seems that the interaction range between the type-I vascular bundles is smaller than the interaction range in the type-II stems, whereas the type-II data appear to feature a less strict weighting of interaction.

Summing up, our newly developed Bayesian inference framework has enabled us to classify two genotypes of maize plants in terms of their vascular bundle arrangement. Nevertheless, it needs to be stressed that only three cross-sectional images have been used per genotype. In order to draw reliable and trustworthy conclusions, the sample size needs to be substantially increased. Moreover, it remains to interpret the results from an agricultural point of view and to extend the analysis to more than two genotypes.

Our most seminal contributions in terms of interdisciplinary research are discussed in Chapter 6 and Chapter 7, where we couple image processing and scene analysis with statistical methods. Assuming that we are given one single image of a textured surface in three-dimensional (3D) space, our goal is a global geometric description of the 3D scene, especially in view of camera orientation.

Chapter 6 is inspired by the idea to associate texture with a hidden homogeneous point process and each texture element with exactly one latent point. Projected onto a two-dimensional (2D) image plane, the arrangement of the texture elements appears inhomogeneous due to the angle of the camera. For this reason, we have established a multi-stage inference framework by means of which an inhomogeneous point process realization can be learned from the image of a textured plane in 3D space, and geometrically be analyzed in terms of camera orientation, under appropriately determined location-dependent scaling assumptions. To facilitate the estimation of the point process realization, we turn the original image into a probability map. The higher the probability value of a certain pixel is, the more likely it represents the symmetry center of a texture element.

Our first approach focuses on the development of a Gibbs model, where the Gibbs energy includes hardcore interaction and exponential scaling constraints. A greedy algorithm simultaneously identifies points and estimates the interaction radius as well as the exponential scaling parameters of the model. Afterwards, the exponential scaling function is replaced by another well-defined function quantifying perspective foreshortening. Through least squares fitting, the estimated exponential scaling effects are substituted by spherical coordinates describing the orientation of the camera towards the textured plane. Having successfully demonstrated that this substitution yields negligibly small residuals, we have worked on the development of a locally-scaled point process model where the inhomogeneity is directly described through a perspective scaling function.

It can easily be shown that the newly suggested perspective scaling function is proper in that it is not only identifiable, but also applicable to locally scale pairwise point distances. Therefore, we can immediately incorporate it into a Strauss modeling framework. Concerning the estimation of the latent point process, we now propose to learn a point pattern from the given probability map prior to the analysis of the parameters and in a model-free manner. For this purpose, we provide two algorithms of a similar nature. Both algorithms yield a reasonably accurate point-texture assignment, provided that the texture elements are close to convex in shape. Despite the second-order structure of the Strauss process, we estimate its scaling and intensity parameters based on the composite likelihood principle.

Synthetic as well as real-world image data are taken as a basis for the evaluation of the proposed modeling routine which proves to be efficient and accurate. We show that, given a scene with convexly and concavely shaped texture elements, a violation of the convexity condition causes minor biases in the estimation of the camera orientation if only few non-convex shapes are evenly distributed among densely packed convex texture elements.

Although the separation of image preprocessing on the one hand, and point identification and parameter estimation on the other hand yields a flexible multi-stage approach, a merging of the individual steps into a large hierarchical overall framework may enhance process compactness and reliability. Another further research objective may consist in the development of a similar fully Bayesian inference concept for estimating shape from texture. Depending on the selected type of data model, this concept may include a CFTP-based exchange algorithm for estimating the posterior distributions of the unknown parameters.

So far, we have exclusively dealt with images of textured planes in 3D space, but not with images of arbitrary uneven surfaces. Room for advancement is thus left in view of an extension and generalization of our current framework. A possible starting point may be the consideration of smooth non-planar surfaces. There is also potential for further development in terms of image segmentation if the scene projected onto the image plane is composed of differently textured sub-regions. Since the proposed inference stages are computationally inexpensive, upgrading them to a reasonable and manageable extent should be unproblematic in terms of CPU runtime.

While the point process framework introduced in Chapter 6 immediately learns a point process realization from the image of a textured scene and then infers 3D geometric knowledge from just the local scaling properties of the resulting point pattern, the alternative modeling strategy presented in Chapter 7 is intended to estimate shape from texture based on the local 2D geometric deformations of the texture elements after their projection onto the image plane.

Therefore, a marked point process framework is proposed, where the marks corresponds to Euclidean similarity transformations of representative patches of texture learned from a reference image without depth effect. We believe that knowledge on camera positioning and angle can be derived from the spatial distribution of these patches in the image of interest as well as from their local 2D geometries. As a byproduct, systematic as well as random irregularities in the texturing of the 3D scene are expected to be detectable.

So far, we randomly select representative template patches from the training image and store them as atoms in a dictionary. We use a Bayesian hierarchical marked point process model to deduce posterior information from a weighted mixture of the data distribution and appropriately specified parameter priors. As data distribution, that is, as model for the spatial pixel value distribution in the image, we consider a mixture of Gaussians, whereas barely informative conjugate or non-informative priors are specified for the unknown model parameters. A Gibbs sampling framework with an embedded reversible jump birth-death-move Metropolis-Hastings (M-H) routine is proposed for estimating the parameter posteriors.

This project has not yet been completed. Having developed a Bayesian inference scheme that successfully estimates the local 2D deformations of texture after a projection onto an image plane, it remains to learn 3D information from the estimated local geometries, such as parameters describing the orientation of the camera towards the textured scene. A modification of the current inference framework by imposing location-dependent scaling constraints seems appropriate in this context.

Despite its incompleteness, the Bayesian marked point process method has two advantages over the locally scaled point process toolbox presented in Chapter 6. On the one hand, it does not require any image preprocessing, and on the other hand, the texture elements may be of any shape from a theoretical viewpoint. From a practical viewpoint, though, it is questionable how sensitive the approach is to irregularities in the contour and coloring of the texture elements. We believe that the sensitiveness particularly depends on the quality of the dictionary learning process. A general mechanism automatically determining the dictionary size and atoms would be helpful in any case.

Summing up in brief, we have shown that locally scaled point processes are a comprehensive and flexibly applicable modeling class for spatial point patterns. Although an unknown constant in the likelihood makes inference complicated, elaborate procedures, such as our newly developed CFTP-based exchange algorithm, have been made available to overcome this problem. This interdisciplinary work has successfully established a connection between point process theory and challenging problems related to the analysis of images. In the context

of a biological research study, we have successfully demonstrated that a suitably specified locally scaled Strauss model is capable of detecting genotype-specific characteristics in the distribution of the vascular bundles in maize plants. A further toolbox of point process procedures has specifically been established for the estimation of shape from texture. It thus couples spatial statistics with computer vision and pattern analysis. We believe that the range of open and unsolved problems approachable by efficient and elaborate point process methods is endless, and that it is worth making effort to formulate appropriate research questions.

The measure of greatness in a scientific idea is the extent to which it stimulates thought and opens up new lines of research.

– Paul A. M. Dirac (August 8th, 1902 – October 20th, 1984)

Bibliography

- Adams, R. P., I. Murray, and D. J. C. MacKay (2009). Nonparametric Bayesian Density Modelling with Gaussian Processes. In *arXiv preprint arXiv:0912.4896*.
- Aharon, M., M. Elad, and A. Bruckstein (2006). K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation. *IEEE T. Signal Proces.* 54(11), 4311–4322.
- Aldrich, J. (1997). R. A. Fisher and the Making of Maximum Likelihood 1912 – 1922. *Statist. Sci.* 12, 3, 133–220.
- Aloimonos, J. (1988). Shape from Texture. *Biol. Cybern.* 58, 345–360.
- Baddeley, A., P. Gregori, J. Mateu, R. Stoica, and D. Stoyan (Eds.) (2006). *Case Studies in Spatial Point Process Modelling*, Volume 185 of *Lecture Notes in Statistics*. Springer.
- Baddeley, A. and R. Turner (2005). Spatstat: An R Package for Analyzing Spatial Point Patterns. *J. Stat. Soft.* 12, 6, 1–42.
- Baddeley, A., R. Turner, J. Møller, and M. Hazelton (2005). Residual Analysis for Spatial Point Processes. *J. R. Statist. Soc. B* 67, 5, 617–666.
- Baddeley, A. J., J. Møller, and R. Waagepetersen (2000). Non- and Semiparametric Estimation of Interaction in Inhomogeneous Point Patterns. *Statist. Neerlandica* 54, 329–350.
- Becker, F., S. Petra, and C. Schnörr (2014). Optical Flow. In O. Scherzer (Ed.), *Handbook of Mathematical Methods in Imaging* (2nd ed.). Springer.
- Bedford, T. and J. Van den Berg (1997). A Remark on the Van Lieshout and Baddeley J-Function for Point Processes. *Adv. Appl. Prob.*, 19–25.
- Berger, M. (2009). *Geometry I*. Springer.
- Bernardo, J. M. and A. F. M. Smith (2009). *Bayesian Theory*, Volume 405. John Wiley & Sons.
- Berthelsen, K. K. and J. Møller (2002). A Primer on Perfect Simulation for Spatial Point Processes. *Bull. Braz. Math. Soc.* 33, 3, 351–367.
- Berthelsen, K. K. and J. Møller (2003). Likelihood and Non-Parametric Bayesian MCMC Inference for Spatial Point Processes based on Perfect Simulation and Path Sampling. *Scand. J. Stat.* 30, 549–564.

- Besag, J. (1974). Spatial Interaction and the Statistical Analysis of Lattice Systems. *J. R. Statist. Soc. B* 36, 2, 192–236.
- Besag, J. (1977). Some Methods of Statistical Analysis for Spatial Data. *Bull. Int. Stat. Inst.* 42, 77–92.
- Besag, J. and P. Clifford (1989). Generalized Monte Carlo Significance Tests. *Biometrika* 76(4), 633–642.
- Besag, J. and P. J. Diggle (1977). Simple Monte Carlo Tests for Spatial Pattern. *Appl. Stat.*, 327–333.
- Besag, J., R. Milne, and S. Zachary (1982). Point Process Limits of Lattice Processes. *J. Appl. Probab.* 19, 210–216.
- Blake, A. and C. Marinos (1990). Shape from Texture: Estimation, Isotropy and Moments. *Artif. Intellig.* 45, 323–380.
- Blostein, D. and N. Ahuja (1989). Shape from Texture: Integrating Texture-Element Extraction and Surface Estimation. *IEEE Trans. Patt. Anal. Mach. Intell. PAMI-11*, 1233–1251.
- Bognar, M. A. (2005). Bayesian Inference for Spatially Inhomogeneous Pairwise Interacting Point Processes. *Comput. Stat. Data An.* 49, 1–18.
- Brix, A. and J. Møller (2001). Space-Time Multi Type Log Gaussian Cox Processes with a View to Modelling Weeds. *Scand. J. Stat.* 28(3), 471–488.
- Brooks, S. P. and B. J. T. Morgan (1995). Optimization using Simulated Annealing. *The Statistician*, 241–257.
- Cam, L. L. (1986). The Central Limit Theorem around 1935. *Stat. Sci.* 1, 78–91.
- Canny, J. (1986). A Computational Approach to Edge Detection. *IEEE Trans. Patt. Anal. Mach. Intell.* (6), 679–698.
- Casella, G. and E. I. George (1992). Explaining the Gibbs Sampler. *Am. Stat.* 46, 3, 167–174.
- Clerc, M. and S. Mallat (2002). The Texture Gradient Equation for Recovering Shape from Texture. *IEEE Trans. Patt. Anal. Mach. Intell.* 24(4), 536–549.
- Cressie, N. A. C. (1993). *Statistics for Spatial Data* (Revised ed.). Wiley and Sons, New York.
- Daley, D. and D. Vere-Jones (2003). *An Introduction to the Theory of Point Processes, Volume I: Elementary Theory and Methods*. Springer.

- Daley, D. and D. Vere-Jones (2007). *An Intorduction to the Theory of Point Processes, Volume II: General Theory and Structure*. Springer.
- Dawid, P. (1984). Present Position and Potential Developments: Some Personal Views: Statistical Theory: The Prequential Approach. *J. R. Statist. Soc. A* 147, 2, 278 – 292.
- Descamps, S., X. Descombes, A. Béchet, and J. Zerubia (2008). Automatic Flamingo Detection using a Multiple Birth and Death Process. In *ICASSP 2008. IEEE Int. Conf. Acoust. Spee.*, pp. 1113–1116.
- Descombes, M. O. X. and J. Zerubia (2008). A Marked Point Process of Rectangles and Segments for Automatic Analysis of Digital Elevation Models. *IEEE Trans. Patt. Anal. Mach. Intell.* 30(1), 105–119.
- Didden, E.-M., T. L. Thorarinsdottir, A. Lenkoski, and D. Legland (in Preparation, 2014/2015). Efficient Framework for Bayesian Inference in Locally Scaled Spatial Point Processes.
- Didden, E.-M., T. L. Thorarinsdottir, A. Lenkoski, and C. Schnörr (2013). Shape from Texture using Locally Scaled Point Processes. *arXiv preprint arXiv:1311.7401*.
- Diebold, F. X., T. A. Gunther, and A. S. Tay (1998). Evaluating Density Forecasts with Applications to Financial Risk Management. *Int. Econ. Rev.* 39, 863–883.
- Dwass, M. (1957). Modified Randomization Tests for Nonparametric Hypotheses. *Ann. Math. Stat.*, 181–187.
- Eberhard, D. A. J., J. D. Zechar, and S. Wiemer (2012). A Prospective Earthquake Forecast Experiment in the Western Pacific. *Geophys. J. Int.* 190, 1579–1592.
- Eckart, C. and G. Young (1936). The Approximation of One Matrix by Another of Lower Rank. *Psychometrika* 1(3), 211–218.
- Eckel, S., F. Fleischer, P. Grabarnik, M. Kazda, A. Särkkä, and V. Schmidt (2009). Modelling Tree Roots in Mixed Forest Stands by Inhomogeneous Marked Gibbs Point Processes. *Biom. J.* 51, 3, 522–539.
- Edelman, A. J. (2012). Positive Interactions between Desert Granivores: Localized Facilitation of Harvester Ants by Kangaroo Rats. *PLoS ONE* 7, 2, e30914.
- Faugeras, O. and Q.-T. Luong (2001). *The Geometry of Multiple Images*. MIT Press.

- Fernández, R., P. A. Ferrari, and N. L. Garcia (2002). Perfect Simulation for Interacting Point Processes, Loss Networks and Ising Models. *Stoch. Process. Appl.* 102, 63–88.
- Foley, J. D., A. Van Dam, S. K. Feiner, and J. F. Hughes (1997). *Computer Graphics: Principle and Practice* (2nd ed.). Addison-Wesley Longman Publishing Co.
- Forsyth, D. (2006). Shape from Texture without Boundaries. *Int. J. Comp. Vision* 67(1), 71–91.
- Ge, W. and R. Collins (2010). Crowd Density Analysis with Marked Point Processes. *IEEE Signal Proc. Mag.* 27, 5, 107–123.
- Gelfand, A. E., P. J. Diggle, M. Fuentes, and P. Guttorp (Eds.) (2010). *Handbook of Spatial Statistics*. Chapman & Hall/CRC.
- Geman, S. and C.-R. Hwang (1986). Diffusions for Global Optimization. *SIAM J. Control Optim.* 24(5), 1031–1043.
- Geyer, C. J. (1999). Likelihood Inference for Spatial Point Processes. *Stochastic Geometry: Likelihood and Computation* 80, 79–140.
- Geyer, C. J. and J. Møller (1994). Simulation Procedures and Likelihood Inference for Spatial Point Processes. *Scand. J. Stat.* 21, 4, 359–373.
- Gibson, J. J. (1950). *The Perception of the Visual World*. Boston, MA: Houghton Mifflin.
- Gneiting, T., F. Balabdaoui, and A. E. Raftery (2007). Probabilistic Forecasts, Calibration and Sharpness. *J. R. Statist. Soc. B* 69, 2, 243–268.
- Green, P. J. (1995). Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination. *Biometrika* 82, 4, 711–32.
- Grenander, U. and M. Miller (1994). Representation of Knowledge in Complex Systems. *J. R. Statist. Soc.* 56, 4, 1–33.
- Hahn, U., E. B. V. Jensen, M.-C. Van Lieshout, and L. S. Nielsen (2003). Inhomogenous Spatial Point Processes by Location-Dependent Scaling. *Adv. Appl. Prob. (SGSA)* 35, 319–336.
- Hartigan, J. A. and M. A. Wong (1979). Algorithm AS 136: A K-Means Clustering Algorithm. *Appl. Stat.*, 100–108.
- Hartley, R. and A. Zisserman (2000). *Multiple View Geometry in Computer Vision*. Cambridge Univ. Press.

- Hastings, W. K. (1970). Monte Carlo Sampling Methods using Markov Chains and their Applications. *Biometrika* 57, 1, 97–109.
- Hausdorff, F. (1918). Dimension und Äußeres Maß. *Math. Ann.* 79, 157–179.
- Huang, F. and Y. Ogata (1999). Improvements of the Maximum Pseudo-Likelihood Estimators in Various Spatial Statistical Models. *J. Comput. Graph. Stat.* 8, 510–530.
- Huang, F. and Y. Ogata (2002). Generalized Pseudo-Likelihood Estimates for Markov Random Fields on Lattice. *Ann. I. Stat. Math.* 54, 1–18.
- Hughes, J. F., A. Van Dam, M. M., D. F. Sklar, J. D. Foley, S. K. Feiner, and K. Akeley (2014). *Computer Graphics: Principles and Practice*. Addison-Wesley.
- Illian, J., A. Penttinen, H. Stoyan, and D. Stoyan (2008). *Statistical Analysis and Modelling of Spatial Point Patterns*, Volume 70. John Wiley & Sons.
- Jähne, B. (1989). *Digitale Bildverarbeitung*, Volume 3. Springer.
- Jensen, E. B. V. and L. S. Nielsen (2001). Inhomogeneous Markov Processes by Transformation. *Bernoulli* 6, 761–782.
- Jin, I. H. and F. Liang (2012). Use of SAMC for Bayesian Analysis of Statistical Models with Intractable Normalizing Constants. *Comput. Stat. Data An.*
- Kanatani, K. (1989). Shape from Texture: General Principle. *Artif. Intell.* 38, 1–48.
- Kelly, F. P. and B. D. Ripley (1976). A Note on Strauss’s Model for Clustering. *Biometrika* 63(2), 357–360.
- Kendall, W. S. (1990). A Spatial Markov Property for Nearest-Neighbour Markov Point Processes. *J. Appl. Probab.* 27, 4, 767–778.
- Kendall, W. S. (1998). *Probability Towards 2000*, Chapter Perfect Simulation for the Area-Interaction Point Process, pp. 218 – 234. Springer.
- Kendall, W. S. and J. Møller (2000). Perfect Simulation using Dominating Processes on Ordered Spaces, with Application to Locally Stable Point Processes. *Adv. Appl. Prob.* 32, 844 – 865.
- Kenkel, N. C., J. A. Hoskins, and W. D. Hoskins (1989). Edge Effects in the Use of Area Polygons to Study Competition. *Ecology*, 272–274.

- Klaver, R. W., D. Backlund, P. E. Bertelt, M. G. Erickson, C. J. Knowles, P. R. Knowles, and M. C. Wimberley (2012). Spatial Analysis of Northern Goshawk Territories in the Black Hills, South Dakota. *The Condor* 114, 3, 532–543.
- Krantz, S. G. and H. R. Parks (2008). *Geometric Integration Theory*. Birkhäuser Boston.
- Lacoste, C., X. Descombes, and J. Zerubia (2005). Point Processes for Unsupervised Line Network Extraction in Remote Sensing. *IEEE Trans. Patt. Anal. Mach. Intell.* 27(10), 1568–1579.
- Lafarge, F., G. Gimel'Farb, and X. Descombes (2010). Geometric Feature Extraction by a Multi-Marked Point Process. *IEEE Trans. Patt. Anal. Mach. Intell.* 32(9), 1597–1609.
- Langevin, P. (1908). Sur la Théorie du Mouvement Brownien. *C. R. Acad. Sci. (Paris)* 146, 530–533.
- Legland, D., M. F. Devaux, and F. Guillon (2014). Statistical Mapping of Maize Bundle Intensity at the Stem Scale using Spatial Normalisation of Replicated Images. *PloS ONE* 9(3), e90673.
- Liang, F. (2010). Double Metropolis-Hastings Sampler for Spatial Models with Intractable Normalizing Constants. *J. Stat. Comput. Sim.* 80, 9, 1007–1022.
- Liang, F., C. Liu, and R. J. Carroll (2007). Stochastic Approximation in Monte Carlo Computation. *J. Am. Stat. Assoc.* 102 (477), 305–320.
- Lindsay, B. G. (1988). Composite Likelihood Methods. *Contemp. Math.* 80, 221–239.
- Liu, Y., H. Hel-Or, C. Kaplan, and L. Van Gool (2009). Computational Symmetry in Computer Vision and Computer Graphics. *Found. Trends Comp. Graphics and Vision* 5(1–2), 1–195.
- Loh, A. and R. Hartley (2005). Shape from Non-Homogeneous, Non-Stationary, Anisotropic, Perspective Texture. In *Proc. BMVC*, pp. 69–78.
- Malik, J. and R. Rosenholtz (1997). Computing Local Surface Orientation and Shape from Texture for Curved Surfaces. *Int. J. Comp. Vision* 23(2), 149–168.
- Mather, P. and M. Koch (2010). *Computer Processing of Remotely-Sensed Images: An Introduction*. John Wiley and Sons.

- Metropolis, N., A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller (1953). Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys* 21, 1087–1092.
- Meyn, S. P. and R. L. Tweedie (2009). *Markov Chains and Stochastic Stability*. Cambridge University Press.
- Møller, J., A. M. Pettitt, R. Reeves, and K. K. Berthelsen (2006). Miscellanea – an Efficient Markov Chain Monte Carlo Method for Distributions with Intractable Normalising Constants. *Biometrika* 93, 451–458.
- Møller, J. and M. Sørensen (1994). Statistical Analysis of a Spatial Birth-and-Death Process Model with a View to Modelling Linear Dune Fields. *Scand. J. Stat.* 21, 1, 1–19.
- Møller, J. and R. P. Waagepetersen (2003). *Statistical Inference and Simulation for Spatial Point Processes*. Chapman & Hall/CRC.
- Murphy, S. A. and A. W. Van der Vaart (2000). On Profile Likelihood. *J. Am. Statist. Assoc.* 95 (450), 449–465.
- Murray, I., Z. Ghahramani, and D. MacKay (2012). MCMC for Doubly-Intractable Distributions. *arXiv preprint arXiv:1206.6848*.
- Murray, I., D. J. C. MacKay, and R. P. Adams (2008). The Gaussian Process Density Sampler. *Adv. Neural Inf. Process. Syst.*, 9–16.
- Nielsen, L. S. and E. B. V. Jensen (2004). Statistical Inference for Transformation Inhomogeneous Point Processes. *Scand. J. Stat.* 31, 1, 131–142.
- Nunan, N., K. Wu, I. M. Young, J. W. Crawford, and K. Ritz (2003). Spatial Distribution of Bacterial Communities and their Relationships with the Micro-Architecture of Soil. *FEMS Microbiol. Ecol.* 44, 203–215.
- Ogata, Y. and M. Tanemura (1981). Estimation of Interaction Potentials of Spatial Point Patterns through the Maximum Likelihood Procedure. *Ann. Inst. Statist. Math.* 33, 315–338.
- Ogata, Y. and M. Tanemura (1984). Likelihood Analysis of Spatial Point Patterns. *J. R. Statist. Soc. B* 46, 496–518.
- Ogata, Y. and M. Tanemura (1986). Likelihood Estimation of Interaction Potentials and External Fields of Inhomogeneous Spatial Point Patterns. In *Pacific Statistical Congress*, pp. 150–154.
- O’Hagan, A., J. Forster, and M. G. Kendall (2004). *Bayesian Inference*. Arnold London.

- Okabe, A., B. Boots, K. Sugihara, and S. N. Chiu (2009). *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*, Volume 501. John Wiley & Sons.
- Papangelou, F. (1974). The Conditional Intensity of General Point Processes and an Application to Line Processes. *Z. Wahrscheinlichkeitstheorie verw. Geb., Springer* 28, 207–226.
- Pele, O. and W. Werman (2009). Fast and Robust Earth Mover’s Distances. In *Proc. Int. Conf. Comp. Vision (ICCV)*.
- Perrin, G., X. Descombes, and J. Zerubia (2005). Adaptive Simulated Annealing for Energy Minimization Problem in a Marked Point Process Application. In *EMMCVPR*, Lect. Notes Comput. Sc., pp. 3–17. Springer.
- Prokešová, M., U. Hahn, and E. B. V. Jensen (2006). Statistics for Locally Scaled Point Processes. In A. Baddeley, P. Gregori, J. Mateu, R. Stoica, and D. Stoyan (Eds.), *Case Studies in Spatial Point Process Modelling*, Volume 185 of *Lecture Notes in Statistics*, pp. 99–123. Springer.
- Rajala, T. and A. Penttinen (2012). Bayesian Analysis of a Gibbs Hard-Core Point Pattern Model with Varying Repulsion Range. *Comput. Stat. Data An.*.
- Ramírez Galarza, A. I. and J. Seade (2007). *Introduction to Classical Geometries*. Birkhäuser Architektur Basel.
- Ripley, B. D. (1976). The Second-Order Analysis of Stationary Point Processes. *J. Appl. Probab.*, 255–266.
- Ripley, B. D. (1977). Modelling Spatial Patterns. *J. R. Statist. Soc. B*, 172–212.
- Salamon, P., P. Sibani, and R. Frost (2002). *Facts, Conjectures, and Improvements for Simulated Annealing*. SIAM.
- Schmid, V. and T. Feilke (2012). Vorlesungsmaterial ‘Räumliche Statistik’, Wintersemester 2012/2013.
- Shimatani, K. and Y. Kubota (2004). Spatial Analysis for Continuously Changing Point Patterns along a Gradient and its Application to an *Abies Sachalinensis* Population. *Ecol. Model.* 180, 2, 359–369.
- Silverman, B. W. (1986). *Density Estimation for Statistics and Data Analysis*, Volume 26. CRC press.
- Stevens, K. A. (1980). Surface Perception from Local Analysis of Texture and Contour. Technical Report AI-TR 512, MIT Technical Report, Artificial Intelligence Laboratory.

- Stoyan, D. and H. Stoyan (2000). Improving Ratio Estimators of Second Order Point Process Characteristics. *Scand. J. Stat.* 27(4), 641–656.
- Strauss, D. J. (1975). A Model for Clustering. *Biometrika* 62(2), 467–475.
- Sun, K., N. Sang, and T. Zhang (2007). Marked Point Process for Vascular Tree Extraction on Angiogram. In A. Yuille, S.-C. Zhu, D. Cremers, and Y. Wang (Eds.), *EMMCVPR*, Volume 4679 of *Lect. Notes Comput. Sc.*, pp. 467–478. Springer.
- Tao, T. (2011). *An Introduction to Measure Theory*, Volume 126. Am. Math. Soc.
- Thorarinsdottir, T. L. (2013). Calibration Diagnostics for Point Process Models via the Probability Integral Transform. *Stat* 2, 150–158.
- Tournaire, O., N. Paparoditis, and F. Lafarge (1965). Rectangular Road Marking Detection with Marked Point Processes. *Photogrammetric Image Analysis* 36(3/W49A), 149–154.
- Tuceryan, M. and A. K. Jain (1998). Texture Analysis. In C. H. Chen, L. F. Pau, and P. S. P. Wang (Eds.), *Handbook of Pattern Recognition and Computer Vision(2nd Edition)*, Singapore, pp. 207–248. World Scientific.
- Van Laarhoven, P. J. and E. H. Aarts (1987). *Simulated Annealing*. Springer.
- Van Lieshout, N. M. N. and A. J. Baddeley (1996). A Nonparametric Measure of Spatial Interaction in Point Patterns. *Stat. Neerl.* 50(3), 344–361.
- Walker, S. G. (2011). Posterior Sampling when the Normalizing Constant is Unknown. *Commun. Stat. – Simul. C.* 40, 784–792.
- Waller, L. A. (2005). Bayesian Thinking in Spatial Statistics. *Handbook of Statistics* 25, 589–618.
- Warren, P. A. and P. Mamassian (2010). Recovery of Surface Pose from Texture Orientation Statistics under Perspective Projection. *Biol. Cybern.* 103, 199–212.
- Witkin, A. P. (1981). Recovering Surface Shape and Orientation from Texture. *Artif. Intellig.* 17, 17–45.
- Yamada, I. and J.-C. Thill (2004). Comparison of Planar and Network K -Functions in Traffic Accident Analysis. *Journal of Transport Geography* 12, 2, 149–158.
- Ziemer, W. P. (1989). *Weakly Differentiable Functions*. Springer.

Subject Index

- approximate double M-H sampler, 69
- areal descriptor, 123
- assignment field, 11, 45
 - assignment vector, 53
 - focus of expansion, 53
- assignment vector, 53
- asymmetrically weighted distance, 108
- autocorrelation, 27

- Bernoulli process, 19
- binomial point process, 19
 - homogeneous, 20

- calibrated coordinates, 48
- camera projection, 45
 - assignment field, 45
 - calibration matrix, 48
 - epipolar line, 50
 - epipoles, 52
 - essential matrix, 52
 - fundamental matrix, 51
 - pinhole camera, 47
 - projection matrix, 48
 - projective plane, 46
 - standard wide angle, 104
 - two-view geometry, 45
- CFTP, 69
- composite likelihood, 37
- convolution, 56
- coordinate representation
 - calibrated, 48
 - homogeneous, 47
 - inhomogeneous, 47
 - normalized, 48

- data coherence term, 124

- derivative-of-Gaussian (DoG) filter, 56
- detailed balance condition, 28
- dictionary, 129
 - atom, 129
- diffusion
 - dynamics, 127
 - equation, 127
 - Langevin, 127
 - process, 127
- DoG, 56
- dominated coupling from the past (CFTP), 35
- double intractability, 39

- epipolar line, 50
- epipoles, 52
- essential matrix, 52
- exchange algorithm, 40
- exponential family, 39
- exponential scaling, 11

- fix-n Gibbs process, 23
- focus of expansion, 53
- fronto-parallel plane, 98
- full conditionals, 28
- fundamental matrix, 51

- geostatistical marking, 17
- Gibbs energy, 25
 - data coherence term, 124
 - regularization constraint, 124
 - total, 25
- Gibbs model, 12
- Gibbs process, 10, 22
 - energy, 25
 - fix n, 23

- Gibbs sampler, 28
- gradient, 55
- hardcore model, 24
- Hastings ratio, 29
- Hessian matrix, 55
- homogeneous coordinates, 11, 47
- homography, 54
- hyper prior, 133
- inhomogeneous coordinates, 11, 47
- intensity, 15
 - conditional, 36
 - first order, 15
 - Papangelou, 36
 - second order, 15
- interaction point process, 23
 - isotropic, 23
 - marked, 25
 - pairwise interaction, 23
- jump diffusion process, 126
- Langevin
 - diffusions, 127
 - dynamics, 127
 - equation, 127
 - process, 127
- Laplacian, 55
- latent point process, 99
- latent variable, 100
- Lebesgue space, 56
 - \mathcal{L}^p norm, 56
- likelihood, 21
 - composite, 37
 - profile composite, 113
 - pseudo, 36
- linear descriptor, 123
- locally scaled point process, 63
- locally scaled Strauss model, 64
- location-dependent
 - scaling, 11, 63
 - scaling function, 63
- volume, 63
- M-H, 28
- marked point process, 11, 17
 - independent marks, 17
 - independent marks model, 17
 - mark density, 21
 - mark distribution, 21
 - multi-marked, 123
 - pairwise interaction, 25
 - Poisson, 21
 - random field model, 17
 - unpredictable marks, 17
- Markov chain, 27
- Markov chain Monte Carlo (MCMC),
 - 10, 26
 - Markov chain, 27
 - burn-in phase, 27
- Markov point process, 10, 22
- Markov property, 10, 22
- maximum a posteriori (MAP) estimate, 85, 137
- MCMC, 26
- mean squared error, 138
- median filter, 57
- Metropolis algorithm, 30
- Metropolis-Hastings (M-H)
 - algorithm, 28
 - birth ratio, 34
 - birth-death algorithm, 34
 - birth-death-move algorithm, 34
 - death ratio, 34
 - double M-H sampler, 69
 - Hastings ratio, 29
 - move algorithm, 33
 - move ratio, 34
- monomial, 54
- Monte Carlo integration, 27
- Monte Carlo sampling, 27
- multi-index, 54
- multi-marked point process, 18, 123

- near-regularity, 11, 96
- normalized coordinates, 48
- orthographic projection, 96
- partial derivative, 55
- perfect simulation, 35
- pinhole camera, 47
 - basic, 47
 - calibration matrix, 48
 - focal length, 47
- PIT, 86
- point process, 15
 - (weakly) isotropic, 16
 - (weakly) stationary, 16
 - homogeneous, 11
 - inhomogeneous, 11
 - latent, 99
 - locally scaled, 1
 - marked, 11, 17
 - multi-marked, 18, 123
 - strongly isotropic, 15
 - strongly stationary, 15
- Poisson point process, 10, 20
 - homogeneous, 20
 - inhomogeneous, 20
 - marked, 21
 - standard, 20
 - unit rate, 20
- posterior, 38
- prior, 38
 - conjugate, 39
 - hyper, 133
 - non-informative, 39
- probability integral transform (PIT), 86
- probability map, 11, 96
- profile composite likelihood, 113
- projective plane, 46
- pseudo likelihood, 36
- regularization constraint, 124
- relaxation temperature, 127
- rigid transformation, 48
 - Euclidean, 49
 - proper, 49
 - roto-translation, 49
 - special Euclidean group, 49
 - special orthogonal group, 49
- scale invariance, 62
- scale invariant function, 62
- scaling function, 11, 63
 - exponential, 65
 - perspective, 105
 - proper, 11, 64
 - step, 11, 80
- separable (partial) derivative filter, 56
- shape from texture, 12, 95
- simulated annealing, 127
- smoothing kernel, 55
- spatial point process, 10, 13
 - aggregated, 18
 - locally scaled, 61
 - random, 18
 - regular, 18
 - simple, 13
- special Euclidean group, 49
- special orthogonal group, 49
- spherical coordinates, 12
- standard wide angle, 104
- step scaling function, 11
- Strauss process, 11, 23
 - disc process, 25
 - hardcore, 24
 - homogeneous, 24
 - locally scaled, 11, 64
- surface element, 98
- surface geometry, 95
 - slant, 95
 - surface normal, 95
 - surface orientation, 95
 - tilt, 95
- surface orientation

- surface element, 98
- unit normal, 97
- texture, 95
 - approximately convex shapes, 96
 - distortion, 96
 - elements, 95
 - gradient, 98
 - homogeneous, 98
 - inhomogeneous, 98
 - near-regular, 96
 - texture arrangement, 95
- texture gradient, 98
- two-view geometry, 11, 45
- void probability, 20
- volume measure, 62
 - d^* -dimensional, 62
 - Hausdorff, 62
 - locally scaled, 63
- Voronoi cell, 107
- Voronoi tessellation, 106

Notation Index

A	product space
B	bounded set
B_{\bullet}	bounded set containing \bullet
$\mathcal{C}, \mathcal{C}()$	normalizing constant
$C(\bullet, \star)$	covariance of \bullet and \star
C_i	i -th neighborhood component
$D_{\star}(\bullet)$	disc of the radius \star with center point \bullet
$D_{\star \setminus \star \star}(\bullet)$	ring-shaped surface with outer and inner radii \star and $\star \star$, and center point \bullet
$D^{(\bullet)}, L^{(\bullet)}, U^{(\bullet)}$	state \bullet of an auxiliary process
$E(\bullet)$	expectation of \bullet
E	essential matrix
H	homography
F	fundamental matrix
I_{dd}	d -dimensional identity/unit matrix
K, K'	camera matrices
$L_C()$	composite likelihood
$L_P()$	pseudo likelihood
M	mark space
M_{\bullet}	subset of the mark space M
$N(\bullet)$	random number of points in \bullet
P, P'	projection matrices
$P(\bullet)$	probability of the event \bullet
$P(\bullet \star)$	conditional probability of \bullet given \star
$P_{\text{acc}}(\bullet \star)$	probability of accepting a transition from \star to \bullet
R	parameter describing the interaction radius/range
R_{MH}	Hastings ratio
R_{eMH}	Hastings ratio extended by an auxiliary variable scheme
R_b, R_d, R_m	birth ratio, death ratio, move ratio
R_t	t -th drawing from a uniform density on $[0, 1]$
$S(\bullet, \star)$	sum of squared residuals
$S()$	empirical standard deviation
T	relaxation temperature
T_{min}	threshold
$U(), U^{(\bullet)}()$	total Gibbs energy
$U_C^{(\bullet)}()$	data coherence term
$U_R^{(\bullet)}()$	regularization term
$U(\bullet \star)$	total Gibbs energy of \bullet given \star
$V(\bullet), V^0(\bullet)$	Voronoi cell of \bullet
W	2D observation window or image plane
W, W'	image (plane) before and after projection

W_a	subset of W covered by geometrically transformed dictionary atoms
$W_a^{(*)}$	subset of W covered by more than one geometrically transformed dictionary atom
$W_{\bar{a}}$	subset of W not covered by any geometrically transformed dictionary atom
$W^{(0)}$	original observation window or image plane
$\mathcal{X}_{a1}, \mathcal{X}_{a2}, \mathcal{X}_{a3}$	axes of a 3D coordinate system
$\tilde{\mathcal{X}}$	homogeneous coordinates of a 3D scene point
\mathcal{X}	scene point in \mathbb{R}^d , $d \geq 3$
$\vec{\mathcal{X}}$	vector pointing from the origin to \mathcal{X} or $\tilde{\mathcal{X}}$
X	random point
\mathbf{X}	random variable, e.g. a random point process
$\mathbf{X}_{-\{i\}} = \{\mathbf{X} \setminus \{X_i\}\}$	\mathbf{X} without its i -th component X_i
$\mathbf{X}_{\sim i}$	all components of \mathbf{X} located in the neighborhood of X_i
\mathbf{X}_B	random point process in B
$\mathbf{X}^{(0)}$	initial point process (with potential boundary effects)
Y	random marked point
\mathbf{Y}	random marked point process
Z	original image or probability map
Z^*	auxiliary image
$Z(\bullet)$	value of Z at location \bullet
\bar{Z}_\bullet	mean pixel intensity in the subset \bullet
$\mathcal{Z}, \mathcal{Z}()$	normalizing constant
$a()$	assignment vector
a	image border or area
\mathbf{a}	dictionary containing K atoms $\{a_1, \dots, a_K\}$
$a(\bullet \star)$	pixel value of a marking dictionary atom in \bullet given \star
$a^{(*)}(\bullet \star)$	average value of marking dictionary atoms in \bullet given \star
b_c	center of gravity
$\mathbf{b}^{(0)}$	original pattern of boundary points $\{b_1^{(0)}, b_2^{(0)}, \dots\}$
\mathbf{b}	normalized pattern of boundary points $\{b_1, b_2, \dots\}$
c	scale factor
$c(), c_\bullet()$	scaling function (parameterized through \bullet)
$d()$	probability of deleting a point in a thinning framework
d, d_\bullet	distance from the origin or from the center of gravity
d, d^*	dimensions
du, dv	circular discs surrounding u and v
e, e'	epipoles
f	discrete or continuous density
$f(\bullet)$	marginal density of \bullet
$f(\bullet \star)$	density of \bullet , given \star
$f^{(c)}()$	locally scaled density
$f_\bullet(), f_\bullet^{(c)}()$	(scaled) density with respect to \bullet

$f^{(*)}()$	unnormalized density
$g()$	function
$h(\bullet \star)$	scale invariant function
h	shift/translation vector
k_1, k_2	fixed parameters
$k^{(i)}$	label of the i -th selected dictionary atom
$k^{(*)}, k_{\bullet}^{(*)}$	mixture of atom labels (in \bullet)
k_w	label of the dictionary atom covering w
l, l'	epipolar lines
l_1, l_2	fixed parameters
$l_P()$	pseudo log-likelihood
m, m_{\bullet}	mark
\mathbf{m}	mark attributes
m_{\bullet}	marking of \bullet
m_1, m_2	fixed parameters
n	counting measure
n'	proposed number of points
$n(\bullet)$	observed number of points in \bullet
n_b	number of boundary points
$p(), p_{\bullet}()$	parameter prior (for \bullet)
$p(\bullet \star)$	parameter posterior of \bullet given \star
p_0	principal point of a 2D image plane
$p_b(), p_d()$	probability to propose a birth/death step
p_M	discrete or continuous mark density
$p_M()$	mark (prior) density
$q(\bullet \star), q_{\bullet}(\cdot \star)$	transition density (for \bullet), conditional on \star
$q_b(\bullet \star), q_d(\bullet \star)$	density of a birth/death transition conditional on \star
r	measured distance
$s_{\bullet}()$	interaction function conditional on \bullet
$s_{\bullet}(\star \star\star)$	interaction function conditional on \bullet and $\star\star$
$(s_1, s_2)^T$	vector shifting the origin of a 2D coordinate system
$t, t()$	auxiliary measure
t_z	binarization threshold
$u_C^{(\bullet)}(\bullet \star)$	coherence of \bullet with the data, given \star
u, v, w	points or image locations
w_l, w_r, w_b, w_t	left, right, bottom and top margins of W
$w_{l_0}, w_{r_0}, w_{b_0}, w_{t_0}$	left, right, bottom and top margins of $W^{(0)}$
$w^{(0)}$	original image location
w	image location
x_{a1}, x_{a2}	axes of a 2D coordinate system
x_e	focus of expansion
x, z	single point realizations
\vec{x}	vector pointing from the origin to x or \tilde{x}
\tilde{x}, \tilde{x}'	homogeneous coordinate representations of x and x'
x^*, \tilde{x}^*	single template or reference point

$x_{V^0(\bullet)}$	point contained in the same Voronoi cell as \bullet
\tilde{x}^c	calibrated/normalized coordinate representation of \tilde{x}
$\mathbf{x}, \mathbf{x}^{(\bullet)}$	realization of a random variable/point process
$\mathbf{x}_{-\{i\}} = \{\mathbf{x} \setminus \{x_i\}\}$	\mathbf{x} without its i -th component x_i
$\mathbf{x}_{\sim i}$	all components of \mathbf{x} located in the neighborhood of x_i
$\mathbf{x}^{(0)}$	initial point process realization
\mathbf{x}_B	point process realization in B
\mathbf{x}_B^*	realization of a homogeneous template process in B
$\mathbf{x}', x'_b, x'_d, x_b^{(0)}, x_d^{(0)}, x'_i$	proposals
y, y_u, y_v	single marked point realization
\mathbf{x}^*	reference point pattern
\mathbf{y}	realization of a marked point process
$\mathbf{y}_{-\{i\}} = \{\mathbf{y} \setminus \{y_i\}\}$	\mathbf{y} without its i -th component y_i
z_\bullet	pixel or probability value in \bullet
$z_{\sim \bullet}$	pixel or probability values in the neighborhood of \bullet
\mathbf{z}	realization of a random variable
$Al()$	function measuring the alignment of objects
$binomial()$	binomial point process distribution
$const$	constant term or value
$det()$	determinant
$i.i.d$	“independent and identically distributed”
$inv\chi^2()$	inverse- χ^2 distribution
$med()$	median
$N()$	Gaussian/normal distribution
N_\bullet^*	truncated Gaussian/normal distribution
$Pois()$	Poisson distribution
$Poisson()$	Poisson point process distribution
$U[...]$	uniform density distribution on a continuous interval
$U\{\dots\}$	uniform probability distribution on a discrete interval
$div()$	divergence
\mathcal{A}	σ -algebra
\mathcal{B}	Borel σ -algebra
$\mathcal{C}^\infty, \mathcal{C}^\infty()$	space of the continuously and infinitely differentiable functions
$\mathcal{C}_0^\infty, \mathcal{C}_0^\infty()$	space of the continuously and infinitely differentiable functions with compact support
Δ	auxiliary measure
\mathcal{D}_i	partial derivative operator: differentiation w.r.t. the i -th component of a vector
\mathcal{D}^{v_i}	partial derivative operator: derivative of order v_i
\mathcal{E}	epipolar 3D plane
Θ	parameter space
Λ, Λ_\bullet	rotation matrix

Λ^i	i -th column of the rotation matrix Λ
$\mathcal{L}_{\text{loc}}^1, \mathcal{L}_{\text{loc}}^1()$	Lebesgue space of the locally integrable functions
$\mathcal{L}^p, \mathcal{L}^p()$	Lebesgue space of the p -th power integrable functions
Π	plane in 3D space
$O_{\mathbf{X}}$	domain of the point process \mathbf{X}
$O_{\mathbf{Y}}$	domain of the marked point process \mathbf{Y}
$O_{\mathbf{Y}}^{(\bullet)}$	subspace of the domain $O_{\mathbf{Y}}$
Φ	process of local maxima
\mathcal{P}	probability measure
$\mathcal{R}, \mathcal{R}_{\bullet}$	radius of a circular data space
\mathcal{S}	joint area of an object and its immediate surrounding
Ω	sample space
$(O_{\mathbf{X}}, \sigma_O)$	measurable space
$(\Omega, \mathcal{A}, \mathcal{P})$	probability space
$(W \times M)$	product space of the point and the mark space, $M \& W$
$(\mathbb{R}^2 \times M)$	product space of the point and the mark space, $\mathbb{R}^2 \& M$
α	parameter
$\alpha_{0_a}, \alpha_{0_{\bar{a}}}, \alpha_{p_a}, \alpha_{p_{\bar{a}}}$	prior/posterior parameters in W_a and $W_{\bar{a}}$
$\alpha()$	normalizing constant
β	intensity parameter
$\beta^{(*)}$	template intensity parameter
$\beta()$	(first-order) intensity function
$\beta_2()$	second-order intensity function
$\beta_{\bullet}()$	intensity function w.r.t. \bullet
$\beta(\bullet \star)$	conditional intensity of \bullet given \star
γ	parameter describing the weighting of interaction
δ	unit vector, i.e. $\ \delta\ = 1$
$\tilde{\delta}$	adjustment of δ under homogeneous coordinates
$\delta_{\mathbf{x}}(\bullet, \star Z)$	weighted difference between $\bullet \in \mathbf{x}$ and \star , given Z
ϵ	infinitesimally small value
ε	smoothing parameter
ζ_{al}	weight of the mutual alignment of marked points
ζ_{attr}	variable balancing between attraction and repulsion
ζ_{pen}	variable penalizing object overlaps
η_1, η_2	scaling of the horizontal/vertical range
η_{ρ}	degree of rotation
$\boldsymbol{\eta}, \boldsymbol{\eta}_{\bullet}$	vector of geometrical parameters
$\eta_{\bullet}^{\min}, \eta_{\bullet}^{\max}$	limits of the range of η_{\bullet}
$\boldsymbol{\eta}^{(i)}$	geometric parameters of the i -th selected dictionary atom
$\boldsymbol{\eta}_{k_w}$	parameters of the dictionary atom k that covers w
$\boldsymbol{\eta}^{(*)}, \boldsymbol{\eta}_{\bullet}^{(*)}$	mixture of realizations of $\boldsymbol{\eta}$ (in \bullet)
θ	one parameter

θ_j	j -th parameter
$\vartheta_{0_a}, \vartheta_{0_{\bar{a}}}, \vartheta_{p_a}, \vartheta_{p_{\bar{a}}}$	prior(0)/posterior(p) parameters in W_a and $W_{\bar{a}}$
$\theta_{\bullet}^{\min}, \theta_{\bullet}^{\max}$	limits of the range of θ_{\bullet}
θ_{attr}	unknown attraction threshold
θ	set of all parameters
$\theta^{-\{\bullet\}}$	θ without \bullet
$\theta', \theta', \theta'_j$	parameter proposal(s)
κ	set of measures $\kappa^{(1)}, \kappa^{(2)}, \dots$
κ_c	set of scaled measures $\kappa_c^{(1)}, \kappa_c^{(2)}, \dots$
$\kappa_{0_{\bar{a}}}, \kappa_{p_{\bar{a}}}$	prior(0)/posterior(p) parameters in $W_{\bar{a}}$
λ	line-determining factor
$\mu()$	mean
$\mu_{\text{in}}, \mu_{\text{out}}$	mean pixel intensity inside/outside an object
$\mu_{\bar{a}}$	mean in $W_{\bar{a}}$
$\mu_{0_{\bar{a}}}, \mu_{p_{\bar{a}}}$	prior(0)/posterior(p) mean in $W_{\bar{a}}$
$\nu^d(\bullet)$	d -dimensional volume (Lebesgue, Hausdorff) measure on \bullet
$\nu_c^d(\bullet)$	scaled d -dimensional volume (Lebesgue, Hausdorff) measure on \bullet
ν, ν_c	set of (scaled) volume measures
ξ, ξ_{\bullet}	smoothing kernels (parameterized through \bullet)
$\xi(), \xi_{\bullet}()$	smoothing kernels (parameterized through \bullet)
π^*	threshold restricting the range of angles
ρ, ρ_{\bullet}	angle of rotation
σ^2	variance parameter
σ_{\bullet}	smallest σ -algebra subject to \bullet
$\sigma_{\text{in}}^2, \sigma_{\text{out}}^2$	variance of the pixel intensity inside/outside an object
$\sigma_a, \sigma_{\bar{a}}$	standard deviation in W_a and $W_{\bar{a}}$
$\sigma_a^2(\bullet \star)$	pixel variance in \bullet given \star if \bullet is covered by one dictionary atom
$\sigma_a^2(*) (\bullet \star)$	pixel variance in \bullet given \star if \bullet is covered by several dictionary atoms
$\tau^*, \tau^{**}, \tau_1, \tau_2$	intermediate variables
τ^2	variance parameter, variability measure
\mathbf{v}	multi-index (v_1, v_2, \dots)
$\phi_1()$	first-order interaction function
$\phi_2(), \phi_3()$	second-order interaction functions
$\varphi()$	interaction function
$\hat{\theta}_{\text{ML}}, \hat{\beta}_{\text{ML}}, \hat{\gamma}_{\text{ML}}, \dots$	maximum likelihood estimates
$\hat{\theta}_{\text{MAP}}, \hat{\beta}_{\text{MAP}}, \hat{\gamma}_{\text{MAP}}, \dots$	maximum a posteriori estimates
β^*, R^*	linear transformations of β and R
2D, 3D	“two-dimensional”, “three-dimensional”
\emptyset	empty set

$\mathbb{1}[\dots]$	indicator function: $\mathbb{1}[\dots] = 1$, if “...” is TRUE; $\mathbb{1}[\dots] = 0$, otherwise
\mathbb{N}	set of the natural numbers (positive integers)
\mathbb{P}^d	d-dimensional projective plane
\mathbb{R}^d	d-dimensional space of the real numbers
$\mathbb{R}_{\setminus\{0\}}$	real numbers without zero
$\mathbb{R}^{\bullet \times \star}$	matrix space of the real numbers
$SE(\bullet)$	special Euclidean group of dimension \bullet
$SO(\bullet)$	special orthogonal group of dimension \bullet
$(\dots)^\top$	(coordinate) vector
$\{\dots\}$	set
$[\bullet, \star]$	line segment between \bullet and \star
$\{\bullet \cup \star\}$	union of \bullet and \star
$\{\bullet \cap \star\}$	intersection of \bullet and \star
$\{\bullet \setminus \star\}$	elimination of \star from \bullet
$*$	convolution operator
\wedge	estimate
\sim	general neighborhood relation
\sim_\bullet	proximity relation evaluated at \bullet
$\star \sim \bullet$	all neighbors \star of \bullet
\propto	proportionality relation
\approx	similarity relation
\simeq	equivalence relation
$ \ $	area measure
$\ \ \ $	Euclidean distance or norm
$\ \ \ _{\mathcal{L}^p}$	\mathcal{L}^p norm
$\langle \ \rangle$	dot/scalar/inner product
\mapsto	mapping
$\partial, \partial_\bullet$	(partial) derivative operators
$\partial\mathcal{X}$	surface element
∇	gradient operator
∇^2	symbol for a Hessian matrix
Δ	symbol for a Laplacian matrix

$$0^0 \quad := \quad 1$$

$$\frac{0}{0} \quad := \quad 1$$

$$\mathcal{D}^{\boldsymbol{v}} \quad := \quad \mathcal{D}_1^{v_1} \mathcal{D}_2^{v_2} \dots = \frac{\partial^{|\boldsymbol{v}|}}{\partial x_1^{v_1} \dots \partial x_n^{v_n}} \, , \quad |\boldsymbol{v}| = \sum_{i=1}^n v_i$$

$$\bar{h}_{\bullet} \quad := \quad \frac{1}{\mathcal{X}_{\bullet}} \, h$$

$$[h]_{\times} \quad := \quad \left(\begin{array}{ccc} 0 & -h_3 & h_2 \\ h_3 & 0 & -h_1 \\ -h_2 & h_1 & 0 \end{array} \right)$$