


# Characterising shape patterns with applications to biomedical image analysis



A thesis submitted for the degree of  
Doctor of Philosophy

Amelia Gontar  
B. Sc. (Hons)

College of Science and Engineering  
Flinders University

9 November 2018



# Contents

<b>List of figures</b>	<b>vii</b>
<b>List of tables</b>	<b>xi</b>
<b>Summary</b>	<b>xiii</b>
<b>Publications arising from the study</b>	<b>xv</b>
Refereed journal papers . . . . .	xv
Full-length refereed conference papers . . . . .	xv
Conference abstracts . . . . .	xvi
<b>Declaration</b>	<b>xvii</b>
<b>Acknowledgements</b>	<b>xix</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Previous work and motivation . . . . .	1
1.2 Objectives of the thesis . . . . .	3
1.3 Overview of the data sets . . . . .	4
1.4 Overview of the thesis . . . . .	6
<b>2 Background and literature review</b>	<b>9</b>
2.1 Shape analysis methods . . . . .	9
2.2 Texton-based methods . . . . .	12

---

2.2.1	Filter banks . . . . .	14
2.2.2	$N \times N$ neighbourhoods . . . . .	16
2.2.3	Oriented thickness textons . . . . .	17
2.2.4	Example from the literature . . . . .	19
2.3	Clustering methods . . . . .	19
2.3.1	$K$ -means clustering . . . . .	20
2.3.2	Gaussian mixture models . . . . .	20
2.3.3	Example from the literature . . . . .	21
2.4	Classification methods . . . . .	22
2.4.1	Linear discriminant analysis . . . . .	22
2.4.2	The Fisher classifier . . . . .	23
2.4.3	Example from the literature . . . . .	24
2.5	Biomedical background . . . . .	24
2.5.1	Yeast colony morphology . . . . .	24
2.5.2	Microstructure of cancellous bone . . . . .	27
2.5.3	Shape of marbling in beef . . . . .	29
2.6	Regular versus irregular shape patterns . . . . .	31
<b>3</b>	<b>Clustered shape primitives</b>	<b>35</b>
3.1	Theoretical framework . . . . .	35
3.1.1	Oriented thickness measures . . . . .	39
3.1.2	Features derived from best-fitting ellipsoids . . . . .	40
3.2	Significance and contribution . . . . .	43
<b>4</b>	<b>Experimental details</b>	<b>45</b>
4.1	Data sets . . . . .	46
4.1.1	Yeast colonies . . . . .	46
4.1.2	Cancellous bone . . . . .	48
4.1.3	Marbling in beef . . . . .	51
4.2	Spatial indices . . . . .	52



---

4.3	Implementation of oriented thickness measures . . . . .	56
4.3.1	Yeast colonies . . . . .	56
	Classification by growth stage . . . . .	57
	Rotation invariance . . . . .	59
	Classification by strain and nutrient conditions . . . . .	61
4.3.2	Cancellous bone . . . . .	64
	Shape and density information, unseparated . . . . .	65
	Shape information only . . . . .	66
	Shape and density information, separated . . . . .	66
	Density information only . . . . .	66
4.3.3	Marbling in beef . . . . .	67
4.4	Implementation of features derived from best-fitting ellipsoids . .	68
<b>5</b>	<b>Results and analysis</b>	<b>71</b>
5.1	Oriented thickness measures . . . . .	71
5.1.1	Yeast colonies . . . . .	72
	Classification by growth stage . . . . .	72
	Rotation invariance . . . . .	74
	Classification by strain and nutrient conditions . . . . .	75
5.1.2	Cancellous bone . . . . .	85
5.1.3	Marbling in beef . . . . .	86
5.2	Features derived from best-fitting ellipsoids . . . . .	93
5.2.1	Yeast colonies . . . . .	93
5.2.2	Cancellous bone . . . . .	97
5.2.3	Comparison with previous work . . . . .	99
5.2.4	Summary . . . . .	101
<b>6</b>	<b>Discussion</b>	<b>103</b>
6.1	Discussion of the yeast colonies . . . . .	103
6.2	Discussion of the cancellous bone . . . . .	105

---

6.3	Discussion of marbling in beef . . . . .	106
6.4	Discussion of the overall study . . . . .	107
6.5	Summary of the contributions of the thesis . . . . .	111
<b>7</b>	<b>Conclusions and future work</b>	<b>113</b>
7.1	Conclusions . . . . .	113
7.2	Future work . . . . .	114
	<b>References</b>	<b>117</b>

# List of figures

2.1	Microarchitecture of human bone. . . . .	11
2.2	The LM filter bank. . . . .	15
2.3	The Schmid filter bank. . . . .	16
2.4	Gabor filter bank. . . . .	17
2.5	Computation of one texture primitive $\mathbf{v}_p$ at point $p$ (red) as in Martin and Bottema, illustrated using an example in two dimensions. . . . .	18
2.6	Section of cancellous bone. . . . .	28
2.7	(a) Dual-energy X-ray absorptiometry image of the proximal femur. (b) Marbling present in one striploin. . . . .	33
3.1	Workflow for the theoretical framework for generating CSPs, illustrating computation of shape primitives, representation space, CSPs, CSP maps, and histograms of CSP occurrences. . . . .	38
3.2	Process of fitting an ellipse in two dimensions. . . . .	40
3.3	The difference between the length measurements $\ell_i$ , $i = 1, 2$ , and $L_i$ , $i = 1, 2$ . . . . .	41
3.4	Computation of two shape primitives $\mathbf{v}_{p,1}$ , $\mathbf{v}_{p,2}$ at points $p_1$ (red) and $p_2$ (blue), respectively, as in the two-dimensional example given in Figure 2.5. . . . .	43
4.1	Example images from the three groups of yeast colonies. . . . .	47

---

4.2	AWRI 796 colony number five grown in 50 $\mu\text{M}$ ammonium sulphate.	48
4.3	Examples of cancellous bone from rat tibiae taken eight weeks after the beginning of the study. . . . .	49
4.4	Splitting a block from the ovx group into four sub-blocks. . . . .	50
4.5	One slice taken from one striploin in the marbling data set. . . . .	52
4.6	The computation of a shape primitive for $D = 4$ . . . . .	57
4.7	Recording oriented thickness measures in an identical image, rotated. . . . .	60
4.8	Rotationally invariant method for computing shape primitives using oriented thickness measures. . . . .	61
4.9	A7-50 colony number five imaged at $t = 233$ hours. . . . .	69
5.1	The shape patterns that gave a perfect classification of A7-50 and AR-50 by strain. . . . .	78
5.2	The shape patterns that gave a perfect classification of A7-50 and A7-500 by nutrient concentration. . . . .	81
5.3	The shape patterns that gave a perfect classification when classifying A7-50, A7-500 and AR-50 by both strain and nutrient concentration. . . . .	83
5.4	One example colony from each of the three groups in the classification by strain and nutrient concentration. . . . .	84
5.5	Marbling proportion plotted as a function of the occurrence of CSP labelled number 10. . . . .	88
5.6	The CSPs identified as important for predicting marbling proportion, each represented using a histogram. . . . .	90
5.7	The CSPs identified as important for predicting marbling proportion, each represented by a drawing in three dimensions. . . . .	91

5.8	Examples of one slice of striploin taken from steers with (a) minimum (MP = 0.090, represented by slice 13 taken from steer number nine) and (b) maximum (MP = 0.286, represented by slice 13 taken from steer number 14) marbling proportions. . . .	92
5.9	An ellipse fit to a yeast colony. . . . .	94
5.10	An ellipse fit to a portion of a yeast colony. . . . .	95
6.1	Three regular shapes and corresponding shape primitives computed at different sampled points $p$ . . . . .	110



# List of tables

4.1	Details of the three yeast groups used in this study. . . . .	46
4.2	Details of the resolution of the yeast colony images. . . . .	47
4.3	Numerical details of the classification process for each data set. .	70
5.1	Classification accuracy as a function of $D$ , for the A7-50 strain with each of the eight observation times considered to be one experimental group. . . . .	73
5.2	The number of times (out of 30) that each of the spatial indices were chosen as one of the three best features for classification, as a function of $D$ . . . . .	73
5.3	Classification accuracy as a function of $D$ , for the A7-50 strain with each of the eight observation times considered to be one experimental group and oriented thickness measures modified to be rotationally invariant (Equation 4.8). . . . .	74
5.4	The mean and standard deviation ( $n = 30$ ) as a function of the number of clusters $K$ per class and number of angles $D$ chosen to compute the CSPs when classifying A7-50 and AR-50 by strain.	76
5.5	The accuracy score as a function of the number of clusters $K$ per class and number of angles $D$ used to compute the CSPs when performing LOOCV for the classification of A7-50 and AR-50 by strain. . . . .	76

5.6	The number of times (out of 20 runs) each of the spatial indices was chosen as one of the best features when classifying A7-50 and AR-50 by strain and performing LOOCV. . . . .	77
5.7	The mean and standard deviation ( $n = 30$ ) as a function of the number of clusters $K$ per class and number of angles $D$ used to compute the CSPs when classifying A7-50 and A7-500 by nutrient concentration. . . . .	79
5.8	The accuracy score obtained using LOOCV when classifying A7-50 and A7-500 by nutrient concentration, as a function of the number of clusters $K$ per class and number of angles $D$ used to compute the CSPs. . . . .	79
5.9	The number of times (out of 19 runs) each of the spatial indices was chosen as one of the best features for classification when performing LOOCV for the classification of A7-50 and A7-500 by nutrient concentration. . . . .	80
5.10	The mean and standard deviation ( $n = 30$ ) of the classification accuracy as a function of the number of clusters $K$ per class and number of angles $D$ used to compute the CSPs for the three-group classification problem using A7-50, A7-500, AR-50. . . . .	82
5.11	The accuracy score obtained using LOOCV for the three-group classification problem as a function of the number of clusters $K$ per class and number of angles $D$ used to compute the CSPs using A7-50, A7-500, and AR-50. . . . .	82
5.12	The number of times (out of 29 runs) each of the spatial indices was chosen as one of the best features for classification, using LOOCV for the three-group classification problem using A7-50, A7-500 and AR-50. . . . .	83
5.13	Performance as a function of the classification process on cancellous bone data. . . . .	86



---

5.14	The coefficient of determination $R^2$ as a function of CSP label. . . . .	87
5.15	The best coefficient of determination achieved using linear regression, and the corresponding best combination of features as a function of the number of features used to fit the regression model. . . . .	87
5.16	The values of the CSPs labelled 2, 3, 4, 8, and 10. . . . .	88
5.17	The proportion of ellipses fit to all 80 images whose centroids fell outside the colony as a function of $r$ (in $\mu\text{m}$ ) when ellipses were fit to yeast colony data. . . . .	93
5.18	Best classification accuracy as a function of $r$ ( $\mu\text{m}$ ) for yeast colony data. . . . .	96
5.19	The proportion of ellipsoids fit to all 90 sub-blocks whose centroids fall outside the bone as a function of $r$ (in $\mu\text{m}$ ), for cancellous bone data. . . . .	97
5.20	Best classification accuracy as a function of $r$ ( $\mu\text{m}$ ) for cancellous bone data. . . . .	99



# Summary

The development of shape analysis methods is an important field of study, especially in the context of biomedical image analysis. An understanding of shape patterns provides information on morphology, function, growth, abnormalities, and so on. Various shape analysis methods – including elliptic Fourier analysis, multiple resolution skeletons, landmark methods and statistical shape models – have been covered extensively in the literature. However, these methods have been developed to analyse regular shape patterns and do not extend well to objects exhibiting irregular shape patterns, such as those common in biomedical settings.

In this thesis, a method for shape analysis is developed specifically for cases where the shape patterns are highly irregular. The idea of clustered shape primitives is introduced, in which local shape patterns are captured and clustered to form representative shape patterns occurring commonly throughout an object or group of objects. Histograms of the occurrences of the clustered shape primitives are used to characterise the shape patterns, and may be used for classification or for regression to predict a given property of the data.

The method is demonstrated and explored using three example data sets: *Saccharomyces cerevisiae* yeast colonies exhibiting pseudohyphal growth, cancellous bone in rat tibiae, and marbling in beef. Each of these types of objects exhibit highly irregular shape patterns that cannot be adequately described using existing shape analysis methods, for example because the consistent placement of landmark points between samples is difficult or impossible.

---

The most significant contribution of clustered shape primitives is that complex shape patterns are learned automatically. There is no need to define any features or landmark points at the outset of the study, which is an improvement on previous methods wherever irregular shape patterns are present. Another advantage of the method proposed here is that the number of features can be kept small, since the most important features are learned and selected automatically. This avoids the need to define a very large number of complex geometric features and ensures that important features are not overlooked – these are examples of problems that may arise when a large list of features is defined at the outset of the study.

Classification methods based on clustered shape primitives achieved competitive results wherever only binary data was available and the shape patterns exhibited were complex. The method could have an impact on shape analysis in biomedical images, and the current study forms a basis for future work in this direction.

# Publications arising from the study

## Refereed journal papers

A. Gontar, H. Tronnolone, B. J. Binder, and M. J. Bottema. Characterising shape patterns using features derived from best-fitting ellipsoids. *Pattern Recognition*, 83:365–374, 2018.

A. Gontar, M. J. Bottema, B. J. Binder, and H. Tronnolone. Characterizing the shape patterns of dimorphic yeast pseudohyphae. *Royal Society Open Science*, 5:180820, 2018.

## Full-length refereed conference papers

A. Gontar, S. Williams, and M. J. Bottema. Characterising 3D structure of cancellous bone. In *Proceedings of the International Conference on Digital Image Computing: Techniques and Applications*, pages 365–371, Gold Coast, Australia, 2016.

## Conference abstracts

M. J. Bottema and A. Gontar. Characterising complex shape in biomedical images. In *39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, Jeju Island, Korea, 2017.

# Declaration

I certify that this thesis does not incorporate without acknowledgment any material previously submitted for a degree or diploma in any university; and that to the best of my knowledge and belief it does not contain any material previously published or written by another person except where due reference is made in the text.

Amelia Gontar





# Acknowledgements

First and foremost, I would like to thank Murk Bottema for his support, advice, understanding and friendship throughout my candidature. His advice regarding the research project was invaluable, but he has also supported me with opportunities to attend conferences, to teach, and to immerse myself in the world of academia during my candidature. Murk has always gone above and beyond what is expected of him as a supervisor, often being available for advice in the evenings, on weekends, and even while on holiday. (I was tempted to write a quip along the lines of “I’d like to thank my supervisor, but I can’t bring myself to do it” here, but Murk really does deserve my sincere thanks.)

I would like to thank Simon Williams for his advice and friendship, and for giving me the opportunity to teach a topic very early on in my candidature. After learning the ropes of lecturing, it was a very rewarding experience and I’m grateful for the opportunity.

For his invaluable advice about writing, for standing in for my classes on days of illness, and for providing encouragement and support, I’d like to thank Michael Haythorpe.

Special thanks to Ben Binder and Hayden Tronnolone from the University of Adelaide, for approaching me with some interesting ideas and opportunities for collaboration. Discussions with Ben and Hayden re-ignited my interest in my project at a difficult time during my candidature – I am very thankful for their eagerness to work with me and share their ideas. In addition, their advice and encouragement was invaluable.

I would like to thank my mum, Irena, for her constant support and (not-so-)occasional hearty meal and glass of wine; and my dad, Andrzej, for his unwavering support despite being a huge distance away geographically. To my second family (the Rodgers) – thank you for adopting me as one of your own and for your continued support and encouragement.

To everyone at Croydon Kings F. C. – thank you for a brilliant few years, and for always giving me a chance to have fun, laugh, and generally unwind from the stresses of my academic life. To all of my friends – Sarah F., Alice F., Alice O., Sylvia, Sarah C., Ben, Chris, Alice J., Hosna, Bec, Stu – thank you for your support and interest in my endeavours.

I would like to acknowledge the contribution of the Australian Government Research Training Program (RTP) Scholarship.

Last, but certainly not least, my warmest thanks go to Niall – for his unwavering support in all of my endeavours, for always being there for a laugh, a cry, or a whinge, and for making the last few years of my life amazing.

# Chapter 1

## Introduction

In this chapter, the thesis is introduced and placed in context. In Section 1.1, the context and motivation for the current study is discussed. Section 1.2 outlines the specific objectives of the thesis. The methods developed throughout the thesis are evaluated on three diverse data sets, and these are discussed in Section 1.3. Finally, Section 1.4 provides an overview of the thesis, including the major contributions of the study.

### 1.1 Previous work and motivation

This thesis is about characterising irregular shapes but the methods have their roots in image texture analysis. In an early review paper on texture analysis, Haralick defined texture as the spatial distribution of grey tones [49]. Since then, many studies have appeared on texture analysis. In particular, texture analysis in biomedical images is an important and well-developed area of research [28, 65, 79, 99, 131]. A typical objective is to determine if there are textural differences between images corresponding to different experimental groups, and thus whether texture can be used to automatically distinguish between groups. Many examples of such questions arise in mammography, for instance whether or not there are textural differences between images of a benign and

a malignant mass [78, 84, 103], or between images corresponding to different densities of breast tissue [44, 74, 111, 112]; or whether texture can be used to predict the risk of a woman developing breast cancer in the future [75, 76, 77].

Since 1999, when the notion of a texton was formalised by Malik *et al.* [90], textons have been widely used for texture analysis in biomedical images [38, 65, 74, 75, 77, 79, 111]. First, a texture primitive is recorded at each pixel to capture information about the local texture at that pixel. There are many ways in which the local texture may be captured, for example by recording the pixel intensities inside the neighbourhood about the central pixel or by recording filter responses. Regardless of the method chosen for capturing local texture information, the key is that certain texture patterns are expected to occur commonly throughout a particular group of images. A clustering algorithm is applied to the collection of texture primitives, called the representation space, and the centres of the clusters are called textons. Thus, textons describe texture patterns that are common throughout the collection of images. Each pixel is mapped to the label of the texton closest to its texture primitive, and thus the image is replaced by a texton map. Finally, histograms of the frequency of texton occurrences are computed for each image. Such histograms represent the texture content of each image and may be used as input into a classifier. In cases where experimental groups exhibit unique texture patterns, histograms of texton occurrences should discriminate well between the experimental groups and have been shown to do so in a wide variety of contexts as cited above.

Despite an extensive body of literature on texture analysis, including some theoretical studies on the possible variations of texton implementations [72, 90, 133, 134], little has been done on extending the idea of textons to binary data for shape analysis. Since the key idea in texton-based characterisation methods is the clustering step, extracting local shape patterns and then finding commonly occurring shape patterns through clustering is theoretically possible. In fact, in a preliminary study by Martin and Bottema, the notion of textons

was widened and applied to characterising cancellous bone in rats [91]. In that study, the authors explored whether there are differences in the microstructure of cancellous bone obtained from the tibiae of normal, diseased, and treated rats. The results of the study showed that a shape analysis method based on an extension of textons could be applied to binary data to detect differences in shape between experimental groups.

However, this idea remains underdeveloped, with no other studies on the topic (to the author's knowledge) appearing in the literature. For this reason, the focus of this thesis is to develop an extension of textons that may be used for shape analysis, and to explore this idea in detail and in a variety of contexts. Perhaps the biggest difficulty lies in choosing a sensible method for extracting local shape patterns, analogous to extracting textural information by considering an image patch or filter responses. Specific research questions are as follows.

1. How does the choice of method for local shape extraction affect classification performance?
2. Is one variation of the method always better regardless of the type of data, or must the implementation of the method be carefully considered depending on the structure of the data?

## 1.2 Objectives of the thesis

Throughout the thesis, a method for shape analysis is developed, based on an extension of well-known texture analysis methods.

First, the idea of a shape primitive is introduced, in which the local shape at a point is recorded. This may be done in a variety of ways, and some of these possibilities are explored. Then, the shape primitives are clustered, and the centres of the clusters are taken to be clustered shape primitives (CSPs) – these represent shape patterns that occur commonly throughout the objects being analysed. Finally, the frequency of the occurrence of each CSP is viewed

as one shape descriptor for each object. These shape descriptors, or features, may be used in a classifier to distinguish between experimental groups, or in a regression as a predictor of a particular characteristic of the data set.

As mentioned earlier, there are choices to be made in computing CSPs, in particular the way in which the local shape at a point is captured most effectively. In this thesis, the idea of oriented thickness measures is explored, which involves measuring the thickness of the object at a pre-specified number of directions and passing through a central point. Alternatively, a set of best-fitting ellipsoids may be fit to the object, with features based on major and minor axes of the ellipsoids. Here, both size and orientation parameters could be included. Many of the choices considered may depend on the characteristics of the data, for instance the dimension, or whether or not the data is expected to have inherent rotation invariance (for instance, does an organism grow with bias in a certain direction, or is there a direction of loading for bone?). As such, the idea of CSPs is implemented and developed on three diverse binary data sets. These are described in the following section.

### 1.3 Overview of the data sets

The first data set comprises binary images of the yeast *Saccharomyces cerevisiae*, with colonies of three different strains imaged at various times after initiation of growth. “On” pixels correspond to the presence of yeast cells and “off” pixels correspond to the absence of yeast cells in that patch [6]. *S. cerevisiae* may grow either by the budding of single cells, or as filaments of unseparated cells called pseudohyphae. The pseudohyphal mode of growth is known to be triggered by nitrogen deprivation [6], and many studies have been conducted on the genetic control and pathways of this type of growth [23, 43, 80, 97, 98, 106, 107, 113]. However, little has been done towards quantifying the spatial patterns in colonies exhibiting pseudohyphal growth, especially quantifying the transition from uniform to pseudohyphal growth [6]. This is important because

such work could lead to the automatic classification of yeast species and stages of growth.

The second data set comprises micro-computed tomography ( $\mu$ -CT) scans of cancellous bone obtained from the tibiae of rats in three groups: normal, oestrogen deprived, and oestrogen deprived and treated with bisphosphonates [37, 91]. The data is three-dimensional, with “on” voxels corresponding to bone and “off” voxels corresponding to surrounding soft tissue or air. This data set is important for studying the microstructure of cancellous bone, which may lead to a better understanding of the effect of diseases such as osteoporosis [61, 115, 123]. In particular, ascertaining whether treatment can restore the microstructure of the bone to normal is of interest, since this could lead to a better understanding of treatment methods and the process of bone remodelling. Bone data has traditionally been described using a set of standard parameters, which are based on two-dimensional histomorphometry [108]. Recently,  $\mu$ -CT has allowed for a three-dimensional view of objects such as cancellous bone. The idea of using CSPs could be a sensible approach towards quantifying this shape information in three dimensions because all available information can be utilised, as opposed to considering only individual slices in two dimensions.

The third data set contains images of the intramuscular fat, or marbling, appearing in striploins obtained from Australian Angus steers on a control diet and a vitamin A supplemented diet [69, 127]. Different consumer markets worldwide prefer different amounts of marbling in beef [34, 60, 102], and various studies have shown that diets low in vitamin A tend to increase marbling [69, 102, 127]. However, little is known about the shape of marbling. For example, as the amount of marbling increases, does the shape of the marbling change or does the marbling simply become thicker? Answering such questions could lead to a more thorough understanding of the biological processes that govern the formation of marbling. CSPs may be used to characterise the shape patterns in marbling. However, the data set presents a few problems in the sense that each

striplon has been cut into slices, and each slice has been imaged separately. The slices can easily be aligned to create a three-dimensional reconstruction of the marbling, but this results in the distance between the slices being much larger than the resolution of the pixels within each slice. Although CSPs may still be used, the choices made (for example, the angles at which to make oriented thickness measurements, or whether or not finding a set of best-fitting ellipsoids is sensible) must be considered carefully.

A characteristic common to all of the data sets is that the shape patterns are highly irregular, in the sense that they are not adequately described by standard parameters that may be defined at the outset of the study. For example, although human cancellous bone has previously been described by the rod-and-plate model [82, 129], this does not work well in describing the microstructure of cancellous bone from rat tibiae because the direction of loading is not as pronounced as in humans. Various authors have described cancellous bone in rats using structural parameters [10, 12, 41, 110], but these do not appear to adequately characterise the microstructure of the bone. For example, the studies cited reach conclusions that are inconsistent with each other. Although yeast colonies have been classified using pre-defined parameters such as fractal dimension and the area of the colony [119], such approaches require a large number of features to be defined at the outset of the study to avoid important features being overlooked. In addition, in that study, only six features were identified as important, suggesting that a majority of standard parameters are redundant as descriptive features.

## 1.4 Overview of the thesis

The major contribution of the thesis is the development of CSPs, including the development and exploration of various ways to capture shape patterns at a local level. The study demonstrates that CSPs are an effective means of characterising shape patterns in biomedical data, which are often highly irregular. In



particular, important shape patterns are learned automatically from the data, avoiding the need to define a list of features at the outset of the study. This is especially important in cases where defining a set of shape parameters is difficult due to the irregularity of the shape patterns present, as is the case with the data sets considered here. The key is that excessively large lists of pre-defined features are avoided, and important features are not overlooked. As a result, CSP-based methods have the potential to be effective in a wide range of applications.

The outline of this thesis is as follows. Chapter 2 comprises a detailed literature review. Classical shape analysis methods, including active shape models and active appearance models, are discussed. Classical texture-based methods used for texture analysis in greyscale images also reviewed. The methods for clustering and classification used throughout this thesis are reviewed. A detailed background of yeast colony morphology, structure of cancellous bone, and shape of marbling in beef is presented. Since the shape patterns analysed throughout the study may be considered highly irregular, this chapter also includes a brief discussion on what is meant by regular and irregular shapes.

In Chapter 3, the CSPs are formally introduced and developed. Two methods for capturing the local shape at a point are presented, namely oriented thickness measures and features derived from best-fitting ellipsoids. The CSPs are then implemented on the three data sets. The experimental details, including a detailed description of the data sets and a discussion of the choices made in implementing CSP-based classification and regression methods, are given in Chapter 4. The results of these experiments are summarised and analysed in Chapter 5.

Chapter 6 comprises a discussion of the contribution of CSPs towards characterising shape patterns in biomedical data. In particular, the two methods for extracting shape features – oriented thickness measures and best-fitting ellipsoids – are compared on each of the data sets individually. Finally, the study is

discussed as a whole, including a discussion of the advantages and disadvantages of CSP-based methods and the cases in which the method is most efficient.

Chapter 7 comprises a summary of the conclusions that may be drawn from this study. Some shortcomings of the current work and ideas for future work are briefly described here.

## Chapter 2

# Background and literature review

This chapter provides a detailed literature review of existing shape analysis methods (Section 2.1) and texture-based methods for texture analysis (Section 2.2). The clustering methods used in this thesis, including  $K$ -means clustering and Gaussian mixture models, are reviewed in Section 2.3. The classification methods used throughout, including linear discriminant analysis and the Fisher classifier, are reviewed in Section 2.4. The biomedical background for the three data sets is discussed in Section 2.5 along with relevant results obtained in previous studies. Since the shapes of objects in these data sets may be considered highly irregular, the notions of regular and irregular shape patterns are loosely defined in Section 2.6.

### 2.1 Shape analysis methods

Shape analysis is vital in many biological applications, providing information on growth mechanisms and assisting in the detection of a variety of medical conditions [117]. In a review of shape analysis techniques, Loncaric defined the shape of an object as a binary image representing the extent of the object,

and grouped shape analysis methods based on whether the method considers the boundary of the object or the interior [85]. Extensive work has been done in both of these areas, and a comprehensive review has been conducted by Pavlidis [109].

Many studies have appeared in which the shapes of two-dimensional objects are characterised in a biological context. Elliptic Fourier analysis has been used to describe the shape of otoliths in fish [64], leaves for plant species identification [11, 95], mussel shells [58], and the human mandible [14]. Multiple resolution skeletons have been used to characterise the geometry of non-rigid objects, such as pseudopods on white blood cells [29]. Landmark methods have been used for describing the shape of the human frontal bone [7] and Old World Talpidae (mole) skulls [118]. The latter studies used projections of three-dimensional computed tomography (CT) scans onto a two-dimensional plane, which means that not all available information was utilised.

Cootes *et al.* extended landmark methods to the idea of point distribution models [20], leading to statistical shape models (SSMs). The central idea of SSMs is to extract the mean shape and several modes of variation from a collection of training shapes using statistical methods [55]. SSMs encompass both active shape models (ASMs) and active appearance models (AAMs).

ASMs are similar to active contour models (“Snakes”) [62], and for this reason are sometimes called “Smart Snakes” [20], in the sense that data from the image is used to iteratively deform the shape model in order to fit the model to the data. However, the key step in ASMs is that constraints for allowable deformations are learned from training data [57]. ASMs have been used to model the shape of resistors, hands [19, 20], the left ventricle in echocardiograms [18, 20], prostate in magnetic resonance (MR) images, ventricles in brain MR images [18], and car brake components [21]. ASMs have also been used for classification, for example to classify plant seeds, and to recognise faces and handwritten postcodes [70]. Hill *et al.* extended ASMs to three dimensions




Figure 2.1: Microarchitecture of human bone. Left: healthy bone. Right: osteoporotic bone. The bone does not exhibit any significant morphological features that could be used as landmark points for consistent labelling between samples. Image created by Gtiroufflet and licenced under CC BY-SA 3.0.

using contours, and used these to segment MR head images [57]. Lu *et al.* used a combination of ASMs, AAMs, and texture analysis to estimate bone fracture risk from dual-energy X-ray absorptiometry images of the proximal femur [87].

An advantage of ASMs is that they can be implemented on either binary or greyscale images. Only binary images are necessary for the model to be implemented, but if greyscale data is available, then information about pixel intensities can be incorporated into the model [18, 21, 70]. AAMs may be viewed as the subsequent extension to ASMs, and combine a model of shape variation with a model of texture variation [17]. AAMs have been used to model the appearance of faces [17], and for segmentation of three-dimensional cardiac MR images [2]. AAMs are designed to incorporate both shape and textural information to model the “appearance” of an object.

Although modelling shape using SSMs is an automated process, the landmark points need to be manually located in the training images in a consistent

manner. In particular, landmark points placed on one sample must correspond to equivalent landmarks on the other samples [20]. For example, in the study by Lu *et al.*, the authors had to identify significant morphological features of the femur including the femoral neck and greater trochanter for each sample, and landmark points had to be placed at these points consistently between samples [87]. SSMs may be successfully implemented to model the shapes of objects for which finding such landmark points is possible, but are not appropriate for objects that are highly irregular in the sense that there are no recognisable landmarks for comparing samples (Figure 2.1).

## 2.2 Texton-based methods

Texton-based methods are useful for characterising image texture. While texture does not play a direct role in this thesis, the idea of textons leads to the formulation of the methods for shape analysis that constitute the main contribution of this study.

Wherever greyscale images are available, texture analysis may be used for detection [28, 73, 79, 84, 103, 104], classification [38, 44, 65, 78, 99, 101, 111, 112, 132], risk estimation [74, 75, 76, 77], *etc.* in various biomedical contexts.

In general, an image is viewed as a map  $X : \mathbb{R}^2 \rightarrow \mathbb{R}$ . Often, a data set comprising a collection of images may be separated into some number of groups. For example, the data set could comprise regions of interest (ROIs) taken from mammograms. For a detection problem, the ROIs could be split into two groups: those corresponding to normal tissue, and those corresponding to a mass. For a classification problem, the ROIs usually correspond to masses that have been detected previously, with the two groups labelled “benign” or “malignant”. Another example of a classification problem in biomedical image analysis is assessing the risk of breast cancer based on whole mammograms, in particular incorporating information about breast cancer history and breast density (Sections 2.2.4, 2.3.3, and 2.4.3). Texton-based methods have widely

been used to characterise texture in such groups of images, both in a biomedical context [38, 44, 74, 75, 76, 111, 112] and in other images [72, 90, 125, 133, 134].

The general framework for texture analysis using the idea of textons is as follows. The local texture at a point  $p \in X$  is represented by a vector

$$\mathbf{v}_p = (v_{p,1}, v_{p,2}, \dots, v_{p,L}), \quad (2.1)$$

called the texture primitive at  $p$ . The way in which the local texture at  $p$  is determined is a question that has been studied extensively (Sections 2.2.1–2.2.3).

The collection of all texture primitives in a data group is called the representation space for that group. Clusters in the representation space indicate common local texture patterns in the corresponding group and are called textons (Algorithm 1).

---

**Algorithm 1** Computation of textons

---

**Input:** collection of pairs  $(X_i, G_i)$ ,  $i = 1, 2, \dots, I$ ,  
 $G_i$  chosen from group labels  $\{G_g : g = 1, 2, \dots, m\}$ ,  
number  $K$  of textons per class  
**Output:**  $M = Km$  cluster centres (textons)

```

initialise array  $T$ 
for  $g = 1$  to  $m$  do
  initialise array  $F_g$ 
  for all  $X_i$  in group  $G_g$  do
    for all  $p$  in  $X_i$  do
      compute  $\mathbf{v}_p$  (Equation 2.1, Sections 2.2.1–2.2.3)
      append  $\mathbf{v}_p$  to  $F_g$ 
    end for
  end for
   $T_k^g$ ,  $k = 1, 2, \dots, K \leftarrow$  clustering on  $F_g$  with  $K$  clusters
end for
append  $T_k^g$  to  $T$ 
return  $T$ 

```

---

The texton map associated with an image  $X$  is obtained by replacing every point  $p \in X$  by the label of the texton closest to  $\mathbf{v}_p$ . The histogram of texton label occurrences for  $X$  (Algorithm 2) represents the texture content of  $X$ . The

frequency of the occurrence of each texton may then be viewed as a feature for classifying  $X$ . Thus, texture patterns are quantified and patterns occurring more commonly in one group over another are used to distinguish between groups.

---

**Algorithm 2** Computation of histograms of texton occurrences

---

**Input:** collection of images  $X_i$ ,  $i = 1, 2, \dots, I$ ,  
collection of  $M = Km$  textons (Algorithm 1)

**Output:** collection of histograms  $h_i$ ,  $i = 1, 2, \dots, I$

```

for  $i = 1$  to  $I$  do
  initialise  $\hat{h}_i$ 
  for  $p$  in  $X_i$  do
    compute  $v_p$  (Equation 2.1, Sections 2.2.1–2.2.3)
    find index  $j$  of texton closest to  $v_p$ 
    replace  $p$  with index  $j$ 
  end for
  for  $j = 1$  to  $M$  do
    count number of times index  $j$  occurred in  $X_i$ 
    append number of occurrences to  $\hat{h}_i$ 
  end for
   $h_i \leftarrow$  set sum of  $\hat{h}_i$  to equal one
end for
return all  $h_i$ ,  $i = 1, 2, \dots, I$ 

```

---

### 2.2.1 Filter banks

The majority of classical texture analysis studies have used the filter bank version of textons [22, 72, 90, 111, 112, 125, 134]. In the filter bank version, the texture primitives are given by Equation 2.1 with

$$v_{p,l} = h_l * X(p), \quad l = 1, 2, \dots, L,$$

where  $h_l$  is the  $l$ -th filter in the filter bank,  $X(p)$  is the value of the image at the point  $p$ , and  $*$  denotes convolution. Different choices may be made for the filter bank. A popular filter bank for two-dimensional images is the Leung-Malik (LM) set [72, 90, 134]. The LM set consists of 48 filters, and so up to 48 filter responses may be obtained at a single pixel. This results in texture primitives






Figure 2.2: The LM filter bank. These 48 filters were designed to detect the presence of bars, edges and spots of various sizes in the image. Image is publicly available from the Visual Geometry Group, Department of Engineering Science, University of Oxford.

of dimension  $L = 48$ . The set of filters contains the first and second derivatives of Gaussians at six orientations and three scales, eight Laplacian of Gaussian filters, and four Gaussian filters (Figure 2.2). The filters are designed to pick out bars, edges, and spots of various sizes.

A potential drawback of the LM set is that it is not rotationally invariant, that is, two rotated images with identical textures do not give the same filter responses. In other words, the same texture pattern occurring at two different orientations is not recognised as the same pattern. Depending on the context, one may wish to use a rotationally invariant filter bank such as the Schmid set [125] (Figure 2.3). The Schmid set consists of 13 rotationally invariant filters, resulting in 13-dimensional texture primitives. Note that, in the LM filter bank one may choose to record only the maximum filter response across all orientations, giving texture primitives of length eight (one for each orientation). These responses are rotationally invariant. Whether or not using a rotationally invariant filter bank is always advantageous is an open question, and in many cases the choice of filter bank appears to be *ad hoc*.

Gabor filters [27] have been used extensively to recognise oriented features, including texture, in many diverse applications. A few examples include the recognition of handwritten characters [128], face recognition [138], and more recently secure cloud-based identification [48]. The filter at a point  $(x, y)$  is




Figure 2.3: The Schmid filter bank. All of the filters are rotationally invariant. Image is publicly available from the Visual Geometry Group, Department of Engineering Science, University of Oxford.

given by

$$h(x, y, \lambda, \phi, \sigma_x, \sigma_y) = \frac{1}{2\pi\sigma_x\sigma_y} \exp\left(-\frac{1}{2}\left(\frac{x'^2}{\sigma_x^2} + \frac{y'^2}{\sigma_y^2}\right)\right) \exp\left(i\frac{2\pi x'^2}{\lambda}\right),$$

where  $x' = x \cos \phi + y \sin \phi$ ,  $y' = y \cos \phi - x \sin \phi$ , and  $\sigma_x$  and  $\sigma_y$  are the standard deviations of the Gaussians in the  $x$  and  $y$  directions, respectively. The wavelength and orientation of the filter are denoted by  $\lambda$  and  $\phi$ , respectively. A Gabor filter bank is usually generated using several scales and orientations [116]. For instance, generating filters at five scales and eight orientations generates a filter bank of 40 filters (Figure 2.4). If this collection of Gabor filters is used to generate textons, the texture primitives are of length  $L = 40$ . Gabor filter banks are used to detect lines and edges.

### 2.2.2 $N \times N$ neighbourhoods

In the  $N \times N$  neighbourhood version of textons [44, 73, 74, 133], the texture primitives are given by Equation 2.1 with

$$v_{p,l} = x_l, \quad l = 1, 2, \dots, N^2,$$

where  $x_l$  is the image intensity value at pixel number  $l$  within the  $N \times N$  neighbourhood centred at  $p$ . These texture primitives may be thought of as




Figure 2.4: Gabor filter bank. These filters were generated using five scales and eight orientations. Reprinted from Haghighat *et al.* [48], with permission from Elsevier.

the output of applying  $N^2$  filters, with each filter being an  $N \times N$  patch comprising  $N^2 - 1$  zeros and a single entry of one [75]. Thus, textons based on  $N \times N$  neighbourhoods are a specific example of textons based on filter banks. The  $N \times N$  neighbourhood version of textons is computationally less expensive than the general filter bank version, but is not rotationally invariant.

Varma and Zisserman showed that the  $N \times N$  neighbourhood method can outperform filter bank methods [133]. The authors also showed that there is no significant difference between including the central pixel and leaving it out, and there is very little difference between choosing  $N = 3, 5$ , or  $7$ . The  $N \times N$  neighbourhood method has been used to classify whole mammograms into four breast density groups [44].

### 2.2.3 Oriented thickness textons

Oriented thickness textons were first introduced by Martin and Bottema, where local shape vectors were computed and clustered in order to characterise the structure of cancellous bone in rat tibiae [91]. The shape at a point  $p$  within the




Figure 2.5: Computation of one texture primitive  $\mathbf{v}_p$  at point  $p$  (red) as in Martin and Bottema [91], illustrated using an example in two dimensions. Here, the lengths of line segments (red lines) through  $p$  lying entirely within the bone (shaded) were measured at four orientations. Thus, the texture primitive at  $p$  is given by  $\mathbf{v}_p = (v_{p,1}, v_{p,2}, v_{p,3}, v_{p,4})$ , where the lengths are measured in directions  $0$ ,  $\pi/4$ ,  $\pi/2$ , and  $3\pi/4$ , respectively.

bone was characterised by the lengths of line segments through  $p$  lying entirely within the bone at various orientations. These lengths, called oriented thickness measures, were used as components in a local shape representation vector and clusters of these vectors were viewed as textons. The data consisted of a three-dimensional binary array with “on” voxels representing bone and “off” voxels representing soft tissue or air. Specifically, the texture primitives were given by Equation 2.1, where  $v_{p,l}$  was the length of the longest line segment of on pixels through  $p$  in direction  $\theta_l$  (Figure 2.5). These could be thought of as “shape” primitives rather than texture primitives.

In that preliminary study, the clustered shape representation vectors were viewed as “textons”. The study demonstrated that recording local shape patterns, clustering these patterns, and using the occurrence of the cluster centres as shape descriptors was a viable method for distinguishing between experimental groups exhibiting different shape patterns. This was applied in a context where only binary data was available.

### 2.2.4 Example from the literature

Recently, Li *et al.* [75] considered 320 images from the Digital Database of Screening Mammography [54]. The images were evenly selected from four categories describing mammographic appearance, and there were two experimental groups: high risk and low risk for breast cancer. High risk images were defined as images of the normal breast of women identified to have a mass (benign or malignant) in the breast at screening time, whereas low risk mammograms were randomly selected from women found to have no masses during screening. The images were randomly and evenly split into training and testing groups comprising 160 images each. The objective was to classify mammograms as either high or low risk [75].

Textons were computed as described in Section 2.2 using  $N \times N$  neighbourhoods (Section 2.2.2), with  $N = 3, 5,$  and  $7$  and the central pixel left out, and using Gabor filters (Section 2.2.1) with wavelength  $\lambda = 20$  and orientations  $\phi = k\pi/10, k = 1, 2, \dots, 10$ . Second-order textons were computed using the  $N \times N$  neighbourhood method on the texton maps (Section 2.2), where the label for each pixel  $p$  was treated as an “intensity” value [75]. The texton maps and histograms of second-order texton occurrences were then computed as described in Section 2.2. Similarly, third-order textons were computed using the  $N \times N$  neighbourhood method (Section 2.2.2) on the second-order texton maps [75].

## 2.3 Clustering methods

Methods used for clustering a collection of vectors are reviewed here. This is not intended to be a comprehensive list of clustering methods, but is a list of methods used throughout this study. These methods are presented here in order to justify choices made later in the thesis (Sections 4.3 and 4.4). A detailed presentation of clustering methods in general may be found in Hastie *et al.* [51],

and code may be found on the associated website [50].

### 2.3.1 $K$ -means clustering

Consider a set  $E_L$  of  $L$ -dimensional unlabelled vectors. The purpose of  $K$ -means clustering is to separate  $E_L$  into  $K$  clusters and find the mean (centre) of each cluster.

One of the first theoretical studies of the  $K$ -means procedure was conducted by MacQueen in 1967 [88]. The author started with  $K$  clusters, each of which contained one point randomly picked from  $E_L$ . At each iteration, a new point was added from  $E_L$  and assigned to the cluster whose mean it was closest to. The cluster means were re-computed to take into account the new point. The algorithm was complete once the entire feature space was added. The “ $K$ -means” were the means of the  $K$  resulting clusters.

The  $K$ -means clustering algorithm has since evolved [33, 51]. Initially, a set of  $K$  centres is chosen from  $E_L$ . Then, for each centre  $C_k$ ,  $k = 1, \dots, K$ , the points that are closer to  $C_k$  than to any other centre are assigned to the cluster corresponding to this centre (in other words, a Voronoi diagram is constructed for the  $K$  centres). The mean of each cluster is re-computed. This procedure is iterated until it converges. Since the  $K$  centres are chosen randomly at the outset, the algorithm is non-deterministic.

The  $K$ -means algorithm has been experimentally shown to produce good results, for example in medical imaging [75], in image segmentation [90], and in texture classification [72, 134].

### 2.3.2 Gaussian mixture models

The purpose of Gaussian mixture models (GMMs) is to separate a set  $E_L$  of  $L$ -dimensional unlabelled vectors into  $K$  clusters, each of which is described by a Gaussian density. Each cluster has associated with it a mean and a covariance

matrix. The GMM is given by

$$f(x) = \sum_{k=1}^K \alpha_k \phi(x; \mu_k, \Sigma_k),$$

where each  $\alpha_k$  is a mixing proportion with  $\sum_k \alpha_k = 1$ , and  $\phi(x; \mu_k, \Sigma_k)$  is a Gaussian density with mean  $\mu_k$  and covariance matrix  $\Sigma_k$ . The multivariate Gaussian density is given by

$$\phi(x; \mu, \Sigma) = \frac{1}{(2\pi)^{\frac{L}{2}} |\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(x - \mu)' \Sigma^{-1} (x - \mu)\right),$$

where  $\mu$  is a vector of length  $L$  and  $\Sigma$  is the  $L \times L$  covariance matrix. The parameters  $\alpha_k, \mu_k, \Sigma_k$  are estimated using the expectation-maximisation (EM) algorithm. Initial choices are made for the estimates  $\hat{\alpha}_k, \hat{\mu}_k, \hat{\Sigma}_k$ . In the E-step, each observation in  $E_L$  is assigned a weight for each cluster, based on the probability of the corresponding Gaussian distributions. In the M-step, the estimates are re-computed, with each observation contributing to the weighted means and covariance matrices of each cluster [51].

The GMM algorithm may be thought of as a “soft” version of  $K$ -means clustering. If the Gaussian distributions are assumed to have a common covariance matrix with very small variance, then the means of the GMMs are analogous to the cluster centres obtained by  $K$ -means clustering. An advantage of GMMs is that observations are classified as belonging to a cluster using some probability measure, so there is an indication of how “strongly” an observation belongs to a cluster. Each cluster is defined by the mean and covariance.

### 2.3.3 Example from the literature

In the example from the Li *et al.* study (Section 2.2.4), cluster centres were found using  $K$ -means clustering (Section 2.3.1) with  $K = 5$  on the training images in each of the four categories describing mammographic appearance, resulting in 20 cluster centres (textons) [75]. This choice was made in order to avoid missing any

fundamental textures from different parenchymal patterns [75]. This procedure was repeated for first-, second-, and third-order textons, resulting in 60 textons in total.

## 2.4 Classification methods

The classification methods used throughout the thesis are discussed in this section. As with the clustering methods (Section 2.3), this is not a comprehensive list of existing classification methods, but comprises the methods used in the current study. The classification methods are presented generally here in order to justify choices made throughout the thesis (Sections 4.3 and 4.4). A thorough presentation of classification methods may be found in Hastie *et al.* [51], and examples of code may be found on the associated website [50].

### 2.4.1 Linear discriminant analysis

Consider a training set of  $M$ -dimensional vectors, each with known group labels. Training a classifier is equivalent to determining each class posterior

$$P(G | X) = P(G = g | X = x),$$

the probability that observation  $x$  falls into group  $g$ , for every observation in the training set. Let  $f_g(x)$  be the group-conditional density of the training set in group  $G = g$ . Let  $\pi_g$  be the prior probability of group  $g$ , that is, the probability of any observation falling into group  $g$  if there were no group labels. Set  $\sum_{g=1}^m \pi_g = 1$ , where  $m$  is the total number of groups. Then, the group posteriors are computed using the formula

$$P(G = g | X = x) = \frac{f_g(x)\pi_g}{\sum_{j=1}^m f_j(x)\pi_j}.$$

The group-conditional densities may be modelled by any choice of func-



tions  $f_g(x)$ . A classifier is called linear if these functions result in a linear decision boundary, for instance if each group density is modelled by a multivariate Gaussian

$$f_g(x) = \frac{1}{(2\pi)^{\frac{M}{2}} |\Sigma|^{\frac{1}{2}}} \exp\left(-\frac{1}{2}(x - \mu_g)' \Sigma^{-1} (x - \mu_g)\right)$$

with mean vector  $\mu_g$  and covariance matrix  $\Sigma$ , which is assumed to be common to all groups. Two groups  $g$  and  $j$  may be compared using the log-ratio

$$\log \frac{P(G = g | X = x)}{P(G = j | X = x)} = \log \frac{\pi_g}{\pi_j} - \frac{1}{2}(\mu_g + \mu_j)' \Sigma^{-1} (\mu_g - \mu_j) + x' \Sigma^{-1} (\mu_g - \mu_j). \quad (2.2)$$

The decision boundary between group  $g$  and group  $j$  is a hyperplane [51].

### 2.4.2 The Fisher classifier

The Fisher classifier was developed by Fisher in 1936 [39], and is a well-known example of a linear discriminant function. The Fisher decision criterion arises from Equation 2.2 when the observations are classified into two groups only: observation  $x$  falls into group  $g$  if and only if the log-ratio is non-negative, and into group  $j$  if and only if the log-ratio is negative. In most cases,  $\pi_g = \pi_j = 0.5$  (for instance, when the prior probability of an observation falling into either group  $g$  or  $j$  is unknown), in which case the decision rule reduces to observation  $x$  being assigned to

$$\begin{cases} \text{Group } g & \text{if } (\mu_g - \mu_j)' \Sigma^{-1} x \geq \frac{1}{2} (\mu_g - \mu_j)' \Sigma^{-1} (\mu_g + \mu_j), \\ \text{Group } j & \text{if } (\mu_g - \mu_j)' \Sigma^{-1} x < \frac{1}{2} (\mu_g - \mu_j)' \Sigma^{-1} (\mu_g + \mu_j). \end{cases}$$

In practice, the  $M$ -variate means  $\mu_g$  and  $\mu_j$  are estimated from the observations in the training set of images [51], and  $\Sigma$  is approximated by  $S_{\text{pooled}}$ , the pooled sample covariance matrix derived from the individual sample covariance matrices  $S_g$  and  $S_j$  of groups  $g$  and  $j$ , respectively [71].

### 2.4.3 Example from the literature

In the Li *et al.* example (Sections 2.2.4 and 2.3.3), among other methods, the Fisher classifier (Section 2.4.2) was trained and tested on features based on first-, second-, and third-order textons and combinations of these [75]. The best classification into high or low risk groups was achieved using the  $N \times N$  neighbourhood method (Section 2.2.2), with  $N = 3$ , and a combination of first- and second-order textons [75].

The Li *et al.* study provides an example of how texton-based (Section 2.2), clustering (Section 2.3), and classification (Section 2.4) methods may be used to approach a complex biomedical problem. Such methods may be useful in the context of developing computed-aided diagnostic tools [44, 75, 79, 111, 112].

## 2.5 Biomedical background

In this section, the background, previous results, and open questions related to the three data sets used throughout this thesis are discussed.

### 2.5.1 Yeast colony morphology

Many yeasts are considered useful for humans. For example, the yeast *Saccharomyces cerevisiae* is useful in producing bread, wine, and ale, and is often referred to as “bakers’ yeast” or “brewers’ yeast” [6]. This yeast is also used to produce some antibiotics and vaccines [106]. The yeast *Saccharomyces boulardii*, a subtype of *S. cerevisiae*, is used to prepare a probiotic for the prevention and treatment of various diarrhoeal disorders [35]. For this reason, the development of new superior strains and optimisation of growth behaviour of yeasts is of interest.

Since the 1980s, there has been growing evidence that *S. cerevisiae* should be considered as a possible opportunistic pathogen. In 1980, Kiehn *et al.* showed that *S. cerevisiae* had been isolated from the sputum in cancer patients, but

accounted for only 0.8% of total yeast isolates [66]. In 1984, Eng *et al.* gave evidence of *S. cerevisiae* infections in patients who had various underlying disorders and had undergone prior courses of antibiotics [36]. There has also been evidence that patients with Crohn's disease have an antibody to *S. cerevisiae*, although the yeast is unlikely to be the cause of Crohn's disease [26, 42, 89]. In a review of clinical reports in 1995, Hazen found that this yeast may be isolated from human mucosal sites, such as the gastrointestinal tract, and has been found in human blood [53]. Reynolds and Fink showed that *S. cerevisiae* can form a biofilm on surfaces, for example medical equipment, leading to resistance to microbial therapies [114]. More recently, Enache-Angoulvant and Hennequin conducted an analysis of 92 documented cases of proven invasive *S. cerevisiae* infection, with 15 cases (16.3%) occurring before 1990 and 76 cases (82.6%) occurring between 1990 and May 2005 [35]. They definitively argued that *Saccharomyces* organisms should be added to the list of fungal pathogens, and this argument has since been supported by other authors [107, 130]. In this case, inhibiting growth is of interest.

Yeasts such as *S. cerevisiae* may grow either by the budding of single cells or as multicellular filaments called pseudohyphae, which are chains of unseparated cells. Pseudohyphal growth is known to be triggered by nutrient (especially nitrogen) deprivation [43] and is thought to represent a scavenging response [24]. There are different strains of the yeast *S. cerevisiae*, with only some strains exhibiting pseudohyphal growth. McCusker *et al.* showed that there is a significant association between pseudohyphal growth and virulence [94]. Pseudohyphal growth in the yeast *Candida albicans* has also been linked to virulence [94], which suggests that studying pseudohyphal growth in general could lead to better identification of pathogenic yeast strains. Much is known about the genetic control and pathways of pseudohyphal growth [23, 43, 80, 97, 98, 106, 107, 113], but shortcomings remain. For instance, it is not routinely possible to distinguish *S. boulardii* from *S. cerevisiae* strains [35]. The tests used currently involve anal-

ysis of mitochondrial DNA, but these tests have not been fully evaluated [35].

Yeasts are commonly studied by growing a collection of cells on a solid medium, such as agar, and recording a two-dimensional top-down image of the resulting colony from which the morphology is analysed [6, 43, 86, 120]. Such studies typically require high-throughput assays, which produce large data sets, and a considerable variety of metrics have been employed to quantify the morphology exhibited in the resulting images. Filamentous and invasive growth has been measured by the relative size of the filamentous or invasive portion of the colony, or, for invasive growth only, by comparing the pixel intensity of pre- and post-washed colonies [120, 139, 140]. While studies on single-gene deletion alleles have used these metrics to understand the regulation of filamentous growth [120], they provide only limited information on the pattern and growth behaviour. Both experimental and simulated data have been quantified by the relative size of the squared colony perimeter to the area, and through the coefficient of variation of the colony boundary [13]. The fractal dimension has been shown to quantify the morphologies of bacterial colonies, which possess similar colony morphologies to yeast [40, 92].

Recently, Binder *et al.* began studying pseudohyphal growth in *S. cerevisiae* in terms of the two-dimensional spatial patterns appearing in the yeast colonies [6]. This work was based on spatial statistics and spatial metrics developed for cell biology images in general [1, 4, 5]. Ruusuvuori *et al.* also classified *S. cerevisiae* colonies into two phenotypes (smooth versus fluffy) by computing 427 intensity- and texture-based features and using logistic regression [119]. Of the features considered, only six were identified as important for classifying the colony morphology, with the most important being the fractal dimension, average entropy texture measure within the colony, and the area of the colony. Such approaches are able to provide a suitably accurate classification but rely on the specification of a large number of features, most of which may not be useful or are unrelated to yeast morphology, while other important features may

be overlooked.

The work developed in this thesis focusses on shape patterns being learned automatically from the data, which is especially important in contexts where the shape patterns are highly irregular – such as those exhibited in yeast pseudohyphae. A method that automatically analyses the complex shape patterns exhibited by yeast pseudohyphae may lead to better understanding of growth behaviour. This could, in turn, lead to a better understanding of how to optimise or inhibit growth as required.

### 2.5.2 Microstructure of cancellous bone

Osteoporosis is a skeletal disease characterised by low bone mass and abnormalities in the skeleton [61, 123]. The disease is associated with decreased production of oestrogen following the menopause in women and lower testosterone levels with advancing age in men, and leads to a rapid rise in the risk of fractures [25, 105]. Women are at a higher risk of developing osteoporosis than men because the decrease in oestrogen production in women is much more rapid than the decrease in testosterone levels in men [105]. White women have a one in six risk of sustaining a hip fracture in their lifetime [25], which increases to a one in three risk at 90 years of age [96]. In turn, the mortality and economic costs associated with fractures (especially hip fractures) are high. Although death cannot be directly attributed to hip fracture, 10–20% more women die than expected for their age within the first year following a hip fracture [16, 25]. Fractures in the USA cost an estimated 20 billion USD per year, with hip fractures accounting for more than one third of this total [25]. In England in 1999, hip fractures cost approximately 850 million GBP for the year [25]. In 2014 in Australia, an estimated 1.2 million people were affected by osteoporosis [105]. Osteoporosis is commonly treated using bisphosphonates, which are intended to reduce and eventually reverse bone loss [105].

Kleerekoper *et al.* showed that the biomechanical competence of cancellous




Figure 2.6: Section of cancellous bone. This block has been obtained from rat tibia using  $\mu$ -CT and is of approximate size  $0.52\text{ mm} \times 0.52\text{ mm} \times 0.52\text{ mm}$ . The rod and plate model attempts to characterise this structure in terms of vertical rods and horizontal plates only. Reprinted from Gontar *et al.* [45] © 2016 IEEE.

bone depends on the trabecular microstructure, in addition to the absolute amount of bone present [67]. The mechanical [81, 82, 83], morphometric [129] and topological [122] properties of human cancellous bone have been analysed in various studies based on the rod and plate model. However, the concept of bone “quality” is currently unmeasurable and warrants further study [115].

Studies have been conducted on cancellous bone obtained from rat tibiae in an effort to understand the effect of osteoporosis on bone microstructure more directly in experimental conditions [9, 10, 12, 37, 41, 110]. In such research, bone loss is experimentally induced in growing female rats through ovariectomy surgery. The induced bone loss is not strictly defined as osteoporosis, but the structural patterns mimic the disease and provide a means of studying the disease more directly. Such research is difficult because cancellous bone from rats has an irregular shape and the rod and plate model may be too simplistic to capture differences in structure patterns of cancellous bone in various disease states (Figure 2.6).

Traditionally, the structure of cancellous bone has been defined in terms of a set of standard parameters based on two-dimensional histomorphometry [108].

Using these parameters as descriptors for cancellous bone in rats, studies have shown that a single injection of zoledronic acid (a bisphosphonate) inhibits changes in bone structure induced by ovariectomy [10], and that a course of zoledronic acid treatment over several weeks facilitates full reversal of bone loss [110]. In contrast, Campbell *et al.* showed that the ovariectomy surgery causes irreversible changes in the trabecular structure after only two weeks, which are reflected in changes in the connectivity [12]. These studies suggest that, while zoledronic acid treatment is beneficial and facilitates some recovery of the cancellous bone, the treatment does not return the microstructure to normal. Since the classic parameters are based on two-dimensional histomorphometry, they do not take full advantage of three-dimensional information that can now be obtained by  $\mu$ -CT scanning. The automated shape analysis method developed in this thesis takes advantage of the available data to a greater extent and have the potential to provide insight into the three-dimensional microstructure of cancellous bone.

### 2.5.3 Shape of marbling in beef

Intramuscular fat, or marbling, in beef cattle has an influence on the quality of the beef [47]. Consumers in different markets choose beef based on varying preferences and perceptions about marbling. For example, in the USA in the 1950s, leanness was the most desired factor [15, 126]. The Canadian market prioritises tenderness, followed by flavour [60]. The Australian market values both leanness and tenderness [34], whereas Japanese consumers prefer a higher degree of marbling [34, 102]. The New Zealand market prioritises leanness, tenderness, flavour, juiciness, in addition to a minimal amount of bone and excess fat [3]. Grunert *et al.* showed that European consumers have a negative perception of fat in beef, which they regard as a sign of bad quality [47].

In the scientific literature, the effect of marbling on the palatability of beef is also disputed. Many authors have argued that increased marbling has positive

effects on juiciness, tenderness, taste, and overall palatability [31, 32, 47, 52, 93, 124, 137]. Grunert *et al.* argued that the European consumers' expectation that low levels of fat lead to good palatability is dysfunctional [47]. On the other hand, Jeremiah showed that relationships between measures of marbling and palatability attributes were not strong [59]. The differences in the conclusions drawn by these authors could be attributed to differences in the cattle (for example, breed, age, and sample size) or the age of the beef [3] used in the studies. Furthermore, Barton argues that, although the classification and grading of beef should ideally reflect consumer requirements, this has not yet been achieved [3]. In order to produce beef of the quality desired by markets worldwide, understanding the factors that produce high levels of marbling is of interest.

Various authors have examined the effect of vitamin A on marbling, which may be described by intramuscular fat percentage (IMF%), a quantitative measure, or marbling score assigned by an expert grader, a qualitative measure. In an early study using Tajima steers of Japanese Black cattle, Oka *et al.* showed that qualitative marbling grade was highly correlated with IMF% [102]. While this result is not surprising, it highlights that in a discussion of marbling, the definitions of IMF%, marbling score, and "marbling" must be carefully considered. For the remainder of the thesis, the term IMF% will refer to the percentage of intramuscular fat in a given cut of beef, marbling score will refer to the qualitative score assigned by an expert grader, and "marbling" will refer to the intramuscular fat present in the beef viewed as an object in its own right.

Oka *et al.* showed that marbling score was negatively correlated with vitamin A levels in the serum in the blood, for cattle of ages less than 23 months [102]. The authors showed that, while a high energy diet increases fat in general (both intramuscular and subcutaneous fat), a low vitamin A diet may increase marbling only [102]. These results have been supported replicated in various studies using 12-month old Angus steers raised in Australia [69, 127]. Vita-



min A deficiency is thought to increase marbling because it is known that retinoic acid, a derivative of vitamin A, inhibits the terminal maturation of adipocytes [69, 102, 127, 135]. A deficiency of retinoic acid is therefore thought to cause higher levels of intramuscular fat being deposited.

Considering only IMF% or marbling score is simplistic in the sense that the spatial distribution and shape of the marbling is not taken into account. If IMF% or marbling score is used to describe the marbling, then only one value is assigned to a cut of beef and information about the microstructure of the marbling is lost. In a preliminary study on the shape of marbling, Bottema *et al.* considered the marbling in striploins taken from the Australian Angus steers and showed that most of the marbling in one striploin is concentrated in a single large object [8]. However, many questions about the shape of the marbling remain unanswered. For example, does increasing the amount of marbling change the shape patterns appearing in the marbling – do new branches form, or does the thickness of the fat simply increase? Are there differences in the marbling shape between animals on the control diet compared with the vitamin A supplemented diet? Answering such questions could lead to a better understanding of the mechanisms underlying marbling formation and the diet needed to produce the desired marbling in cattle.

## 2.6 Regular versus irregular shape patterns

The shape patterns appearing in the biomedical data in Section 2.5 may be considered highly irregular. To that end, in this section, regular and irregular shape patterns are loosely defined and the differences between these are discussed.

Shape patterns could be defined as regular if two criteria are satisfied: they can be described by geometric parameters, and identifiable features exist that can be used as landmark points to compare samples. For example, the geometry of the proximal femur may be described by geometric parameters such as hip axial length, femoral length, femoral neck width and neck-shaft angle [30, 87].

Furthermore, images of the proximal femur obtained from two different patients are similar in the sense that significant morphological features can be identified and used as landmark points. In particular, these landmark points can be placed consistently between samples (Section 2.1). For these reasons, the shape of the proximal femur is regular in some sense.

On the other hand, in this sense, the biomedical examples used in this study (Section 2.5) cannot be considered regular. Although the geometry of the yeast colonies could be described by certain parameters such as the length and number of pseudohyphae, these would fail to capture the complex fractal-like shape patterns occurring on the boundary of the pseudohyphae. Ruusuvaori *et al.* used geometric parameters including area, diameter, and length of the major and minor axes of an ellipse fitted to the shape of *S. cerevisiae* colonies, but these were combined with textural features to form a list of 427 descriptive features and most turned out to be unimportant for identifying the morphology of the colony [119]. This result suggests that the shape patterns exhibited by *S. cerevisiae* colonies are highly irregular and cannot be described by a set of standard geometric features or by placing landmark points.

Similarly, the quality of cancellous bone obtained from rat tibiae has been described by structural parameters such as bone volume ratio, bone surface ratio, trabecular thickness, trabecular separation, trabecular number, connectivity density, structural model index, and cortical thickness [9, 10, 12]. However, these assign one value to the entire bone and do not directly quantify the geometry or the shape patterns present in the bone. Furthermore, there are no easily identifiable morphological features that facilitate the direct comparison of samples – in human tibiae, the cancellous bone could be described by the rod and plate model [82, 129, 136], implying some degree of regularity in the structure of the bone, but the direction of loading in rats appears to be much less pronounced (Section 2.5.2). Thus, the shape patterns appearing in cancellous bone obtained from rat tibiae are considered highly irregular.




Figure 2.7: (a) Dual-energy X-ray absorptiometry image of the proximal femur. The image shows approximately equally spaced landmark points on the boundary of the femur. Reprinted from Lu *et al.* [87], with permission from Taylor & Francis. (b) Marbling present in one striploin. The placement of landmark points on samples of this type has not been attempted. Reprinted from Bottema *et al.* [8] and freely available from Scotland’s Rural College under the Freedom of Information Scotland Act (FOISA).

Finally, the scores assigned to describe marbling in beef, marbling score and IMF%, are qualitative and simplistic, respectively (Section 2.5.3), and do not adequately describe the shape of the marbling. In addition, the placement of landmark points at equivalent points on two different samples would be impossible in practice due to the lack of significant morphological features (Figure 2.7).

For these reasons, the shape of the biomedical examples considered in this study cannot adequately be described by a set of standard geometric parameters, or by placing landmark points and using SSMs. Hence, the focus of this thesis is to develop a method for describing highly irregular shape patterns by learning the important patterns directly and automatically from the data. Such a method would avoid the need to define a large list of geometric features at the outset of the study and to identify an appropriate set of landmark points on each training sample. This is important because, in the cases described here, defining landmark points consistently between samples is not possible.



## Chapter 3

# Clustered shape primitives

In this chapter, clustered shape primitives are formally introduced – this represents the major theoretical contribution of the thesis. In Section 3.1, a general framework is introduced. The various ways in which the local shape may be captured at a point are discussed in Sections 3.1.1 and 3.1.2. Section 3.2 comprises a discussion of the significance and contribution of CSPs towards shape analysis, including a discussion of the context in which the method would be implemented most effectively and efficiently.

### 3.1 Theoretical framework

Let  $X$  denote a binary function on  $\mathbb{R}^n$ , and let  $\Omega$  denote the set

$$\{p \in \mathbb{R}^n : X(p) = 1\}.$$

For example, if  $n = 2$ , then  $X$  is a binary image. The set  $\Omega$  is the object of interest appearing inside this image.

The objective is to characterise a collection of objects based on shape and compare these with other collections of objects. To that end, consider a collection of objects  $\Omega_i$ ,  $i = 1, 2, \dots, I$ , each with a group label chosen from  $m$  groups  $G_g$ ,  $g = 1, 2, \dots, m$ . In other words, consider a collection of pairs  $(\Omega_i, G_i)$ ,

where  $G_i$  is the corresponding group label chosen from  $\{G_g : g = 1, 2, \dots, m\}$ .

The local shape at a point  $p \in \Omega$  is represented by a shape primitive vector  $\mathbf{v}_p$ , which may be computed in various ways (Sections 3.1.1–3.1.2). A schematic example of the computation of a single shape primitive vector  $\mathbf{v}_p$  in the context of the yeast colony data (Section 4.1.1) is shown in Figure 4.6.

The collection of all shape primitives in a data group  $G_g$  is

$$F_g = \{\mathbf{v}_p : p \in \Omega, X \in G_g\}, \quad (3.1)$$

and is called the representation space for group  $g$ . Shape primitives close to each other in space represent similar shape patterns. Hence, common local shape patterns occurring in the collection  $G_g$  may be identified by clustering (Section 2.3) the vectors in  $F_g$ . The resulting cluster centres are referred to as clustered shape primitives (CSPs), and may be viewed as quantitative shape descriptors learned automatically from the data. The CSPs for data group  $G_g$  are vectors  $C_k^g$ ,  $k = 1, 2, \dots, K$ , of the same dimension as the shape primitives  $\mathbf{v}_p$ . With  $m$  groups and  $K$  CSPs per group, there are  $M = Km$  CSPs in total. For convenience, these are labelled sequentially as  $C_j$ , where  $C_{(g-1)K+k} = C_k^g$ ,  $g = 1, 2, \dots, m$  and  $k = 1, 2, \dots, K$  (Algorithm 3).

The CSP map  $Y$  associated with  $X$  is obtained by replacing every point  $p \in \Omega$  by the label of the CSP closest to  $\mathbf{v}_p$ . That is,

$$Y(p) = \arg \min_{j \in \{1, 2, \dots, M\}} \{\|C_j - \mathbf{v}_p\|\},$$

where  $\|\cdot\|$  may be any norm, but for the purpose of this thesis is chosen to be the Euclidean norm. The normalised histogram of CSP labels associated with  $X$  is

$$\mathbf{h}_X = (h_1, h_2, \dots, h_M), \quad (3.2)$$

where  $h_j$  is the proportion of points in  $\Omega$  with label  $j$  (Algorithm 4). The histogram  $\mathbf{h}_X$  represents the shape content of  $X$  and numbers  $h_j$  are viewed as

---

**Algorithm 3** Computation of CSPs

---

**Input:** collection of pairs  $(\Omega_i, G_i)$ ,  $i = 1, 2, \dots, I$ ,  
 $G_i$  chosen from group labels  $\{G_g : g = 1, 2, \dots, m\}$ ,  
number  $K$  of CSPs per class

**Output:**  $M = Km$  cluster centres (CSPs)

```

initialise array  $C$ 
for  $g = 1$  to  $m$  do
  initialise array  $F_g$ 
  for all  $\Omega_i$  in group  $G_g$  do
    for all  $p$  in  $\Omega_i$  do
      compute  $\mathbf{v}_p$  (Sections 3.1.1–3.1.2)
      append  $\mathbf{v}_p$  to  $F_g$ 
    end for
  end for
   $C_k^g, k = 1, 2, \dots, K \leftarrow$  clustering on  $F_g$  with  $K$  clusters
end for
append  $C_k^g$  to  $C$ 
return  $C$ 

```

---

features either for classifying  $X$  (Section 2.4) or as a predictor for a particular characteristic of  $X$  (for example, in a regression model). A schematic diagram of the full method is shown in Figure 3.1.

The general method described thus far is made specific by defining the methods for setting the shape primitive vectors  $\mathbf{v}_p$  at every point  $p$ . The entire process will later be demonstrated and evaluated on examples comprising real biomedical data (Chapter 4).

The shape primitive vectors  $\mathbf{v}_p$  are an extension of the texture primitive vectors  $\mathbf{v}_p$  discussed in Section 2.2. Here, local shape (as opposed to texture) information is captured.

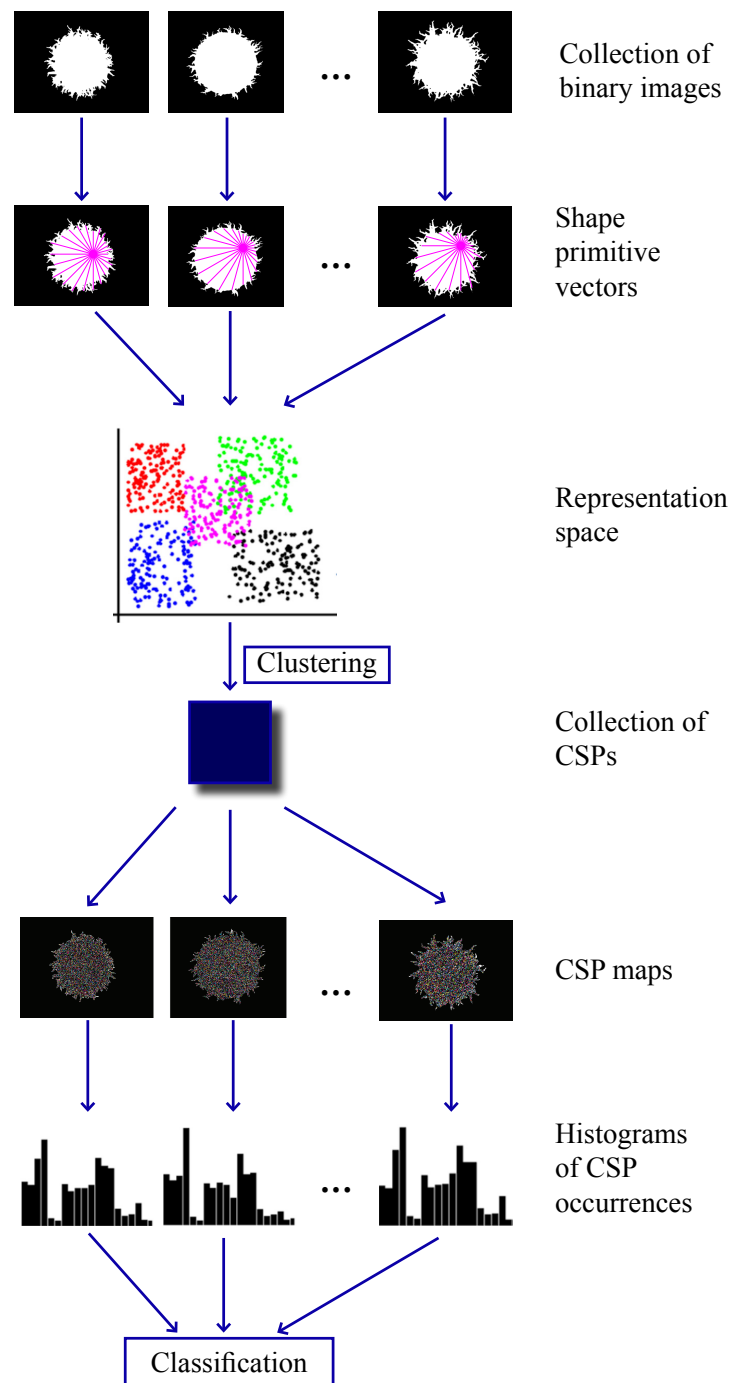


Figure 3.1: Workflow for the theoretical framework for generating CSPs, illustrating computation of shape primitives, representation space, CSPs, CSP maps, and histograms of CSP occurrences. The histograms represent the shape content of each image and are viewed as features for classifying the images or as a predictor for some characteristic of each image.



---

**Algorithm 4** Computation of histograms of CSP labels

---

**Input:** collection of objects  $\Omega_i$ ,  $i = 1, 2, \dots, I$ ,

collection of  $M = Km$  CSPs (Algorithm 3)

**Output:** collection of histograms  $\mathbf{h}_i$ ,  $i = 1, 2, \dots, I$

```

for  $i = 1$  to  $I$  do
  initialise  $\hat{\mathbf{h}}_i$ 
  for  $p$  in  $\Omega_i$  do
    compute  $\mathbf{v}_p$  (Sections 3.1.1–3.1.2)
    find index  $j$  of CSP closest to  $\mathbf{v}_p$ 
    replace  $p$  with index  $j$ 
  end for
  for  $j = 1$  to  $M$  do
    count number of times index  $j$  occurred in  $\Omega_i$ 
    append number of occurrences to  $\hat{\mathbf{h}}_i$ 
  end for
   $\mathbf{h}_i \leftarrow$  set sum of  $\hat{\mathbf{h}}_i$  to equal one
end for
return all  $\mathbf{h}_i$ ,  $i = 1, 2, \dots, I$ 

```

---

### 3.1.1 Oriented thickness measures

The first method for computing shape primitives is called oriented thickness measures. For a point  $p \in \Omega$ , the shape primitive is given by

$$\mathbf{v}_p = (v_{p,1}, v_{p,2}, \dots, v_{p,D}), \quad (3.3)$$

where  $v_{p,d}$ ,  $d = 1, \dots, D$ , is the length of the longest line segment through  $p$  in direction  $\theta_d$  lying entirely inside  $\Omega$ . The intention is to measure the local thickness of the object  $\Omega$  in  $D$  directions. The question of choosing the angles  $\theta_d$ ,  $d = 1, \dots, D$ , at which to make thickness measurements is not trivial. Ideally, the angles would be equally spaced, which reduces to evenly distributing  $D$  points on the surface of an  $n$ -dimensional sphere. However, this problem is also not straightforward [56, 63, 100, 121]. For this reason, the ways in which the angles are chosen are discussed in the context of each data set separately (Section 4.3).




Figure 3.2: Process of fitting an ellipse in two dimensions. A slice taken from the cancellous bone data set (Section 4.1.2) is used as an example. The object  $\Omega$  is shown in white. (a) First, a disk  $B_r^p$  (blue) is drawn at  $p$  (blue cross). Here, the region  $\Omega \cap B_r^p$  (shaded) contains more than one connected component. (b) Second, the connected component  $\Omega^p$  containing the largest number of occupied points (now shaded) is found. This represents the dominant local shape pattern at  $p$ . (c) Third, an ellipse  $E^p$  (now blue) is fit to  $\Omega^p$  (shaded). The ellipse  $E^p$  and the connected component  $\Omega^p$  have the same volume.

### 3.1.2 Features derived from best-fitting ellipsoids

For a point  $p \in \Omega$ , let  $B_r^p$  denote the open ball of radius  $r$  centred at  $p$ . Let  $\Omega^p$  denote the largest connected component of  $\Omega \cap B_r^p$  and let  $E^p$  denote the ellipsoid that best fits  $\Omega^p$  in the sense that the  $n$ -dimensional volume of  $E^p$  and  $\Omega^p$  are equal and that  $\text{vol}(E^p \cap \Omega^p)$  is maximal subject to the condition of equal volume (Figure 3.2). The best fitting ellipsoid  $E^p$  is easily found by applying principal components on the coordinates of the points in  $\Omega^p$ . The directions of the  $n$  axes of the ellipsoid are given by the unit eigenvectors  $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n$ , each of which has  $n$  components,  $\mathbf{u}_i = (u_{i,1}, u_{i,2}, \dots, u_{i,n})$ . The lengths of each axis  $\ell_i$ ,  $i = 1, \dots, n$ , are given by the corresponding eigenvalues, normalised to ensure the volume of  $E^p$  is equal to the volume of  $\Omega^p$ . Let  $L_i$ ,  $i = 1, \dots, n$ , be the length of the longest line segment in the direction given by  $\mathbf{u}_i$ , passing through the centroid of  $E^p$ , and lying entirely inside the object  $\Omega$  (Figure 3.3).

Four definitions of shape primitive vectors  $\mathbf{v}_p$  based on the constructions above are considered in this study.




Figure 3.3: The difference between the length measurements  $\ell_i$ ,  $i = 1, 2$ , and  $L_i$ ,  $i = 1, 2$ . The ellipse fit in Figure 3.2 is used as an example. (a) The lengths  $\ell_1$  and  $\ell_2$  correspond to the lengths of the major and minor axes of the ellipse, respectively. (b) The lengths  $L_1$  and  $L_2$  are the lengths of the longest line segments lying entirely inside  $\Omega$  and passing through the centroid of the ellipse, measured in the directions of the major and minor axes, respectively. The thickness of the object, rather than the ellipse, is measured here.

**Definition 1** (Oriented ellipsoid parameters).

$$\mathbf{v}_p = (\ell_1 u_{1,1}, \ell_1 u_{1,2}, \dots, \ell_1 u_{1,n}, \ell_2 u_{2,1}, \ell_2 u_{2,2}, \dots, \ell_2 u_{2,n-1}, \dots, \ell_n u_{n,1}).$$

Oriented ellipsoid parameter shape primitives retain all orientation and length information embodied in  $E^p$  to represent the local shape of  $\Omega^p$ . Since the unit vectors are orthogonal, the ellipsoid may be specified by a total of  $n(n+1)/2$  parameters instead of the full  $n^2$  components needed to specify  $n$  vectors of length  $n$ .

**Definition 2** (Oriented object parameters).

$$\mathbf{v}_p = (L_1 u_{1,1}, L_1 u_{1,2}, \dots, L_1 u_{1,n}, L_2 u_{2,1}, L_2 u_{2,2}, \dots, L_2 u_{2,n-1}, \dots, L_n u_{n,1}).$$

Oriented object parameter shape primitives capture information about the local orientation and thickness of the segment of  $\Omega$  more directly. This method has the drawback that the centroid of the ellipsoid is not guaranteed to fall

inside the object  $\Omega$ , in which case  $L_i = 0$  for all  $i = 1, \dots, n$ .

**Definition 3** (Ellipsoid shape parameters).

$$\mathbf{v}_p = (\ell_1, \ell_2, \dots, \ell_n),$$

Ellipsoid shape parameter shape primitives retain only information about the lengths of the axes of the ellipsoids.

**Definition 4** (Object shape parameters).

$$\mathbf{v}_p = (L_1, L_2, \dots, L_n).$$

Object shape parameter shape primitives retain only information about the thickness of the object locally at  $p$ .

The purpose of recording shape primitives  $\mathbf{v}_p$  based on a collection of best-fitting ellipsoids is to develop a method for sensibly choosing subsampled points and to ascertain which variation of recording local shape information (Definitions 1–4) leads to the best discrimination between groups of objects. Oriented thickness measures have the shortcoming that the placement of subsampled points  $p$  is not controlled in any way, and that thickness measurements are generally made with respect to laboratory coordinates (Figure 3.4). By finding a collection of best-fitting ellipsoids, the locations and directions at which thickness measurements are made are learned directly from the data rather than being chosen at random.

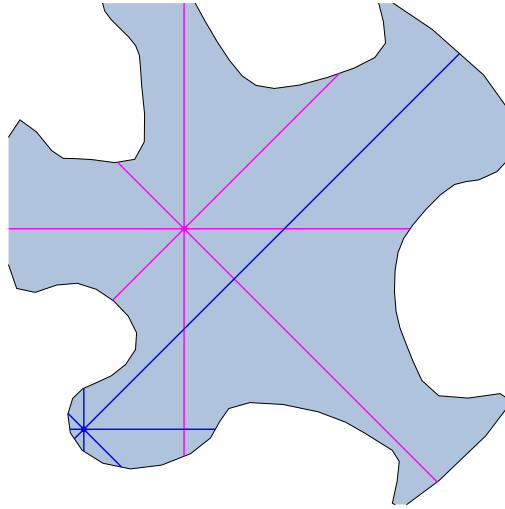


Figure 3.4: Computation of two shape primitives  $\mathbf{v}_{p,1}$ ,  $\mathbf{v}_{p,2}$  at points  $p_1$  (red) and  $p_2$  (blue), respectively, as in the two-dimensional example given in Figure 2.5. The information recorded about the local shape is different depending on the placement of the subsampled point  $p$ .

## 3.2 Significance and contribution

The idea of CSPs is similar to the notion of textons, but there are various notable differences. Importantly, CSPs may be used to characterise the shape of objects whenever only binary data is available. The notion of textons does not extend naturally to binary data because taking linear filter responses (Section 2.2.1) does not make sense. Similarly, recording pixel or voxel values in an  $N \times N$  neighbourhood about the central pixel (Section 2.2.2) is not a useful way of capturing textural or shape information because the only possible responses for pixels in that patch are 0 or 1. Clustering is not an effective method to find common patterns in such a representation space. On the other hand, CSPs provide a useful method for shape analysis even in cases where image intensity data is not available. However, this also means that evaluating CSP-based methods through comparison with baseline approaches such as texton-based methods (Section 2.2) and AAMs (Section 2.1) does not make sense because the latter methods do not readily apply to binary data.

The most significant contribution of CSPs is that complex shape patterns

are learned automatically. This is especially important in cases where the shape patterns are highly irregular (Section 2.6). Methods involving SSMs are not applicable in such cases because the manual placement of landmarks at equivalent points between samples is effectively impossible, whereas in CSP-based methods there is no need for such a step. This means that using CSPs to analyse the shape of highly irregular objects could be an improvement on previous methods, but that comparison with such methods is difficult.

Another advantage of the CSP-based methods proposed here is that the number of features can be kept small, since the most important features are learned and selected automatically. This avoids the need to define a very large number of complex geometric features and ensures that important features are not overlooked – these are examples of problems that may arise when a large list of features is defined *a priori* [119]. The idea of CSPs may be compared to textons in the sense that the purpose of textons is to learn important textural features automatically from greyscale data. Analogously, in CSP-based methods important shape features are learned automatically from binary data.

## Chapter 4

# Experimental details

In order to evaluate whether CSPs provide a valuable means of extracting shape features for classification and prediction, the theoretical framework presented previously (Chapter 3) was applied to the three example data sets. This chapter serves to present the experimental details. In Section 4.1, the technical details of the data sets, including image acquisition and resolution, are discussed. Section 4.2 comprises a review of the spatial indices used to quantitatively describe the spatial patterns in yeast colonies by Binder *et al.* [6]. This is presented here because the spatial indices were reviewed and re-computed in detail as part of the current study.

Section 4.3 comprises details of the implementation of CSPs based on oriented thickness measures. For the yeast colonies, oriented thickness measures and a rotationally invariant variation of oriented thickness measures were used to classify by growth stage groups, and oriented thickness measures were used to classify by strain and nutrient conditions. For the cancellous bone, a variation of the method was implemented that separates shape and density information to determine whether or not shape and density information is complementary. For the marbling in beef, shape features were extracted using oriented thickness measures and then used in a regression to predict the proportion of marbling in the striploin. Section 4.4 describes the implementation of features derived from

best-fitting ellipsoids on yeast colony and cancellous bone data.

All computational steps were carried out using Matlab R2015a–R2017b. Samples of the code used may be found on the author’s Github page: [https://github.com/ameliagontar/PhD\\_project\\_2018.git](https://github.com/ameliagontar/PhD_project_2018.git).

## 4.1 Data sets

First, a detailed description of the three data sets is given, including the (retrospective) experimental methods, image acquisition, resolution of the data, and some illustrative example images from each of the data sets.

### 4.1.1 Yeast colonies

The yeast colony study is retrospective, based on data collected in Binder *et al.* [6]. Briefly, single cells of *S. cerevisiae* were used to initiate the growth of several individual yeast colonies. The colonies were grown in ammonium sulphate at a fixed concentration, and each of the colonies was imaged successively over time. The images were converted into two-dimensional binary images using customised software [1, 4], where pixels were designated as either “occupied” by yeast cells or “unoccupied”. The data comprises three separate experiments, corresponding to different concentrations of ammonium sulphate nutrient and yeast strains (Table 4.1). The data was also imaged at different resolutions, depending on the strain and observation time (Table 4.2).

Table 4.1: Details of the three yeast groups used in this study. Shown are the strains, concentrations of ammonium sulphate (in  $\mu\text{M}$ ), number of colonies, and observation times (in hours) after initiation of growth. The name column refers to the abbreviated name of the group used hereafter.

Name	Strain	Conc.	Trials	Observation times
A7-50	AWRI 796	50	10	{23, 48, 73, 87, 115, 162, 211, 233}
A7-500	AWRI 796	500	9	{25, 49, 73, 97, 121, 145, 169, 193, 212, 240}
AR-50	AWRI R2	50	10	{25, 49, 73, 100, 121, 145, 168, 192, 212, 237}



Table 4.2: Details of the resolution of the yeast colony images. The resolution (per pixel) is given according to strain and observation time (in hours) after initiation of growth.

Name	Observation times	Resolution
A7-50	{23, 48}	$0.61 \mu\text{m} \times 0.61 \mu\text{m}$
	{73, 87, 115, 162, 211, 233}	$1.52 \mu\text{m} \times 1.52 \mu\text{m}$
A7-500	{25}	$0.62 \mu\text{m} \times 0.62 \mu\text{m}$
	{49, 73, 97, 121, 145, 169, 193, 212, 240}	$1.55 \mu\text{m} \times 1.55 \mu\text{m}$
AR-50	{25, 49}	$0.61 \mu\text{m} \times 0.61 \mu\text{m}$
	{73, 100, 121, 145, 168, 192, 212, 237}	$1.53 \mu\text{m} \times 1.53 \mu\text{m}$




Figure 4.1: Example images from the three groups of yeast colonies. The images are taken at 233, 240, and 237 hours after initiation of growth, respectively. The AWRI 796, 50  $\mu\text{M}$ , example image was published by Binder *et al.* [6] and is licenced under CC BY 4.0. The AWRI R2 and AWRI 796, 500  $\mu\text{M}$ , images are reproduced courtesy of the Wine Innovation Cluster, based at the Waite Research Precinct at the University of Adelaide.

An example image corresponding to each of the three strains at the last observation time and before processing is shown in Figure 4.1. The temporal evolution of one colony of the AWRI 796 strain grown in 50  $\mu\text{M}$  ammonium sulphate (A7-50) and after image processing is shown in Figure 4.2. More details of the experiments may be found in Binder *et al.* [6].

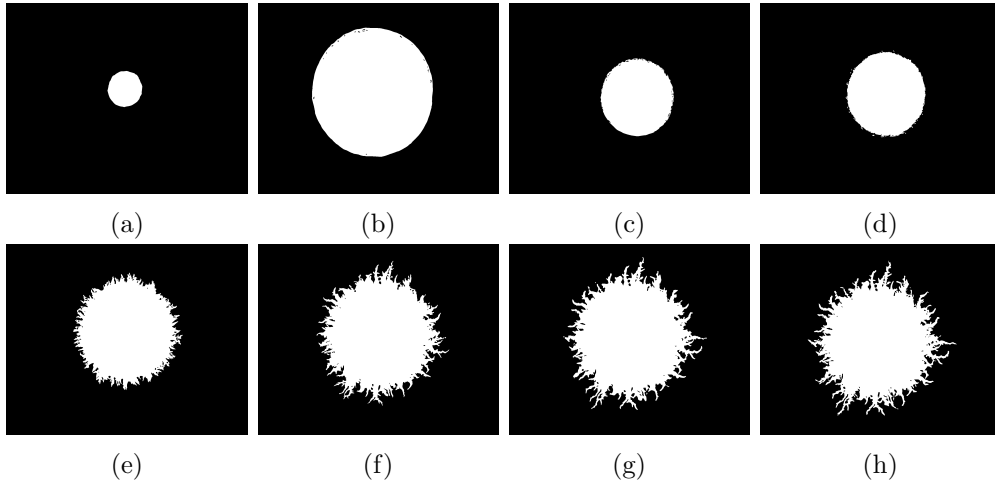


Figure 4.2: AWRI 796 colony number five grown in  $50 \mu\text{M}$  ammonium sulphate. (a)–(h) The colony was imaged at  $t = \{23, 48, 73, 87, 115, 162, 211, 233\}$  hours, respectively. For  $t = 23, 48$  hours, the resolution is approximately  $0.61 \mu\text{m} \times 0.61 \mu\text{m}$  per pixel. For  $t = 73, \dots, 233$  hours, the resolution is approximately  $1.52 \mu\text{m} \times 1.52 \mu\text{m}$  per pixel.

#### 4.1.2 Cancellous bone

Material in this section appears in Gontar *et al.* [45] © 2016 IEEE.

The cancellous bone study is a retrospective study based on data collected by Fazzalari *et al.* [37]. The data set contains 30  $\mu$ -CT scans of cancellous bone obtained from rat tibiae. The rats were randomly assigned to three experimental groups, 10 in each group. At the start of the study (week 0), the sham group underwent a sham surgery, where an incision was made but no organs were removed. The ovx group underwent a full ovariectomy, and the ovx+zol group underwent a full ovariectomy and subsequently started treatment with zoledronic acid at week 2. The rats in the other two groups received saline injections at the same time as the rats in the ovx+zol group received the zoledronic acid injections. The ovariectomy induced oestrogen deprivation and subsequent bone loss.

Micro-CT scans of the right tibia of each rat were obtained at 0, 2, 4, 8, and 12 weeks after surgery. The resolution was  $8.702 \mu\text{m} \times 8.702 \mu\text{m} \times 8.702 \mu\text{m}$  per voxel. Since two rats died between weeks 8 and 12, only the scans for week 8 were used




Figure 4.3: Examples of cancellous bone from rat tibiae taken eight weeks after the beginning of the study. (a) The sham group corresponds to the control group. (b) For the ovx group, bone loss was induced through ovariectomy. (c) For the ovx+zol group, bone loss was induced through ovariectomy and then treated with zoledronic acid. The bottom of each block represents the slice closest to the growth plate (slice 1). Reprinted from Gontar *et al.* [45] © 2016 IEEE.

in this study. For each of the 30 rats, a rectangular block of size  $121 \times 121 \times 400$  voxels, corresponding to a size of approximately  $1 \text{ mm} \times 1 \text{ mm} \times 3.5 \text{ mm}$ , was manually segmented from the  $\mu$ -CT reconstruction. The longest edge of each block was roughly parallel to the long axis of the tibia, and the first cross-sectional slice was approximately 1.2 mm from the growth plate of the bone. Each voxel was assigned as “bone present” or “bone absent” using a threshold, resulting in three-dimensional binary rectangular blocks (Figure 4.3). Details are discussed in Fazzalari *et al.* [37].

Each block was further divided into three sub-blocks. The first sub-block comprised the first 100 slices counting from the growth plate end, the second sub-block comprised the next 100 slices, and the third sub-block comprised slices 200–300. The last 100 slices of each block were not used in the study because the bone was too sparse for structure characterisation in some cases



Figure 4.4: Splitting a block from the ovx group into four sub-blocks. (a) The original block (week 8), with the slice closest to the growth plate at the bottom of the image. (b) From bottom to top: sub-blocks 1–4. Block 1 is closest to the growth plate, corresponding to the “newest” bone, and block 4 is furthest away. Sub-block 4 was not included in the experiment since, in some cases, the bone was too sparse for meaningful characterisation of the structure. Reprinted from Gontar *et al.* [45] © 2016 IEEE.

(Figure 4.4). Thus, a total of nine groups were used in the study, corresponding to three experimental groups and three distances from the growth plate. The rats in the study were growing and new bone was being formed during endochondral ossification at the growth plate. Hence, different structure patterns could be expected to occur at different distances from the growth plate, in addition to differences in the structure between the three experimental groups.

For each of the three experimental groups, five of the 10 rats were randomly assigned to the training set and the remaining five rats were assigned to the testing set for classification.

### 4.1.3 Marbling in beef

The marbling data used in this study has been described previously [69, 127]. Briefly, 20 Angus steers with a high propensity for marbling were placed in feedlot facilities at 12 months of age. At the beginning of the study, the steers were randomly separated into two groups of 10 animals each: vitamin A supplemented (A+) and non-supplemented (A-). All animals were fed a standard feedlot diet and kept in the same pen for the duration of the study. Once per week, the steers in the A+ group received a dosage of vitamin A through a supplement added to their feed for six hours. Throughout the study, the animals were monitored for vitamin A deficiency.

After slaughter, striploins quartered at the 12–13<sup>th</sup> rib were collected, trimmed, vacuum packed and frozen at  $-20^{\circ}\text{C}$ . Two drill holes were made through each of the striploins. Twenty five consecutive slices, each 4 mm thick, were cut from each of the striploins, cleaned, and stored. These slices were photographed. For each slice, the rib eye region was extracted manually from the image and the area of the rib eye region was calculated. Then, each image was converted to a binary image, with each pixel converted to “marbling present” or “marbling absent” using an empirically determined threshold (Figure 4.5).

The 25 individual slices were then aligned using the drill holes, to form a three-dimensional binary representation of the marbling in the striploin. The marbling proportion

$$\text{MP} = \frac{\text{total volume of marbling}}{\text{total volume of striploin}} \quad (4.1)$$

was recorded for each striploin.

The resolution of each slice is approximately  $0.143 \text{ mm} \times 0.143 \text{ mm}$  per pixel. If the data is considered in three dimensions by aligning the slices, then each voxel has resolution approximately  $0.143 \text{ mm} \times 0.143 \text{ mm} \times 4 \text{ mm}$ .



Figure 4.5: One slice taken from one striploin in the marbling data set. White represents marbling. Shown is slice number 13 from the striploin collected from steer number one.

## 4.2 Spatial indices

Binder *et al.* defined three indices based on spatial metrics to describe filamentous growth in colonies of *S. cerevisiae*: the radial index, angular index of filamentation, and pair-correlation index of aggregation [6]. In the original study, spatial indices were introduced to characterise shape. In Section 4.3.1 the method proposed in this thesis will be benchmarked against spatial indices and so this method is presented here. (The notation here is different to that in the original study to match the notation used throughout this thesis. The symbols  $i, j, k$  are re-used as counters.)

Let the set of position vectors for the occupied sites within the colony be

$$P = \{\mathbf{w}(p) : p \in \Omega\}.$$

The maximal radial distance of the colony is

$$R = \arg \max_{p \in \partial\Omega} \{ \|c - p\| \},$$

where  $c$  is the centroid of the colony, and  $\partial\Omega$  is the set of boundary pixels of  $\Omega$ .

The mean-field density is

$$\rho = \frac{|\Omega|}{\pi R^2}.$$

In order to compute the radial metric, the magnitudes of the position vectors are defined in terms of subsets

$$S_r(i) = \{\mathbf{w} : \Delta_r(i-1) \leq |\mathbf{w}| < \Delta_r i, \mathbf{w} \in P\},$$

where  $\Delta_r = \frac{R}{N_r}$ ,  $i = 1, \dots, N_r$ , and  $N_r$  is the number of equally spaced partitions defined on the domain  $[0, R)$ . The radial metric is

$$F_r(i) = \frac{c_r(i)}{\pi \rho \Delta_r^2 (2i-1)},$$

where  $c_r(i) = |S_r|$  is the count of the radial distances of the position vectors in each bin. The denominator is a normalisation factor that ensures the radial metric is scaled with respect to circular domains populated uniformly at random.

When  $F_r(i) > 1$ , the spatial domain is aggregated and the probability of finding an occupied site is greater than that for a state of complete spatial randomness (CSR). When  $F_r(i) < 1$ , the spatial domain is dispersed and the probability of finding an occupied site is less likely than for the CSR state. For  $F_r(i) = 1$ , the probability of finding an occupied site is equal to that at the CSR state, and the radial distance at which this occurs is called  $R_{\text{CSR}}$ . The radial index is

$$I_r = 1 - \frac{R_{\text{CSR}}}{R}. \quad (4.2)$$

and is a measure of the colony's filamentation in the radial direction. The radial index may be used to compare filamentous growth in the radial direction between two or more yeast colonies, and may be used to distinguish between strains exhibiting different filamentous growth patterns.

In order to define the angular index of filamentation, the principle arguments

of the position vectors are defined in terms of subsets

$$S_\theta(j) = \{\mathbf{w} : \pi + \Delta_\theta(j-1) \leq \arg(\mathbf{w}) < -\pi + \Delta_\theta j, \mathbf{w} \in P\},$$

where  $\Delta_\theta = \frac{2\pi}{N_\theta}$ ,  $j = 1, \dots, N_\theta$ , and  $N_\theta$  is the number of bins defined on the domain  $[-\pi, \pi)$ . The counts of the arguments of the position vectors are given by  $c_\theta(j) = |S_\theta(j)|$  and the angular metric is

$$F_\theta(j) = \frac{c_\theta(j)}{\frac{1}{2}\rho\Delta_\theta R^2},$$

where the denominator is again a normalisation factor. The angular metric is analogous to the radial metric, except that the spatial domain is separated into segments rather than annuli. This metric describes the angles at which the spatial domain is aggregated and can detect the orientation of filaments. The discrete Fourier transform of the angular metric is given by

$$\hat{f}_k = \frac{1}{N_\theta c_k} \sum_{j=0}^{N_\theta-1} F_\theta(j) e^{-ikx_j}, \quad k = 0, \pm 1, \dots, \pm \frac{N_\theta}{2},$$

where  $x_j = -\pi + \frac{2\pi j}{N_\theta}$ ,  $j = 0, \dots, N_\theta$ ,  $c_k = 2$  for  $k = \pm \frac{N_\theta}{2}$ , and  $c_k = 1$  for  $k \neq \pm \frac{N_\theta}{2}$ . The angular index describes the evolution of the spectrum  $|\hat{f}_k|$  of the angular metric, and is given by

$$I_\theta = \sum_{k=1}^{N_\theta/2} |\hat{f}_k|^2. \quad (4.3)$$

The angular index provides the means to compare angular filamentous growth between different colonies. This index may be used to distinguish between colonies growing in different environments where the growth is directionally biased towards a particular nutrient source.

Finally, to define the pair-correlation index of aggregation, the angles be-



tween pairs of position vectors are defined in terms of the subsets

$$S_{\Theta}(k) = \left\{ (\mathbf{w}_1, \mathbf{w}_2) : \Delta_{\Theta}(k-1) \leq \cos^{-1} \left( \frac{\mathbf{w}_1 \cdot \mathbf{w}_2}{|\mathbf{w}_1||\mathbf{w}_2|} \right) < \Delta_{\Theta}k, \mathbf{w}_1, \mathbf{w}_2 \in P \right\},$$

where  $\Delta_{\Theta} = \frac{\pi}{N_{\Theta}}$ ,  $k = 1, \dots, N_{\Theta}$ , and  $N_{\Theta}$  is the number of bins defined on the domain  $[0, \pi)$ . The counts of the angles between position vectors are given by  $c_{\Theta}(k) = |S_{\Theta}(k)|$ , and the (angular) pair-correlation metric [1, 4, 5] is given by

$$F_{\Theta}(k) = \frac{c_{\Theta}(k)}{\frac{1}{2}\rho\Delta_{\Theta}R^2(\pi\rho R^2 - 1)},$$

where the denominator is a normalisation factor. The (angular) pair-correlation metric measures the frequency of the occurrence of angles between position vectors. The value of  $F_{\Theta}(1)$  is the (normalised) number of pairs separated by an angle in the domain  $[0, \frac{\pi}{N_{\Theta}})$ , and hence corresponds to localised aggregation. The pair-correlation index of aggregation is defined as

$$I_{\Theta} = F_{\Theta}(1) - 1 \tag{4.4}$$

and is a quantitative measure of localised aggregation in the yeast colony. Binder *et al.* provided extensive working examples to illustrate the implementation of the spatial indices [6].

Binder *et al.* showed that these indices are an effective means of quantifying the spatial patterning in yeast colonies [6]. In general, as time evolved, the average radial index, average angular index, and average pair-correlation index across all samples increased, for all three strains. This is consistent with the observation of increasing filamentous growth with time for each of the strains.

Here, the spatial indices were computed using Equations 4.2–4.4 with  $N_r = 178$ ,  $N_{\theta} = 200$ , and  $N_{\Theta} = 200$ , to match the values used in the original study [6].

### 4.3 Implementation of oriented thickness measures

This section describes how CSPs based on oriented thickness measures were implemented on the three data sets using variations of the method as appropriate to the data set. These variations are discussed in detail in this section. In each case, the analysis of a single point  $p$  is scaled up to the entire collection of images as shown in the workflow diagram in Figure 3.1.

#### 4.3.1 Yeast colonies

For the yeast colony data, since the shape patterns appearing in the filaments are of interest, only pixels on the boundary  $\partial\Omega$  of  $\Omega$  were analysed. The directions were defined by

$$\theta_d = (d - 1)\frac{\pi}{D}, \quad d = 1, \dots, D, \quad (4.5)$$

where  $0 \leq \theta_d < \pi$ , and  $D$  is the number of directions. For each pixel  $p \in \partial\Omega$ , let  $\hat{v}_{p,d}$  be the length of the longest line segment through  $p$  in direction  $\theta_d$  lying entirely inside  $\Omega$  (Figure 4.6). In order to sensibly compare shape characteristics of colonies at different growth stages and to consider shape independent of the orientation of the colonies, the distance  $\hat{v}_{p,d}$  was converted to a new distance  $v_{p,d}$  measured in terms of pixels, but accounting for the difference in resolution between the images in the A7-50 group (Table 4.2). For characterising shape by strain and nutrient conditions, the distance  $\hat{v}_{p,d}$  in pixels was converted to a physical distance  $v_{p,d}$ , measured in  $\mu\text{m}$ , so that images at different resolutions could be compared. Then, the shape primitives  $\mathbf{v}_p$  were set as in Equation 3.3. Since the number of pixels in  $\partial\Omega$  varied between approximately 12,000 and 24,000, shape primitives were computed at all  $p \in \partial\Omega$ .

Various classification problems were considered for the yeast colony data in order to test the effectiveness of oriented thickness measures in distinguishing between shape patterns. First, only the A7-50 group was considered and classified according to observation time (growth stage). Second, the yeast colonies




Figure 4.6: The computation of a shape primitive for  $D = 4$ . For a region  $\Omega$  (shaded squares), shape primitives are computed for each pixel  $p$  on the boundary  $\partial\Omega$  (dark grey). The length of the longest line segment  $\hat{v}_{p,d}$  lying entirely inside  $\Omega$  and passing through  $p$  is measured, once each at angles  $\phi_d = (d-1)\pi/D$ , where  $d = 1, \dots, D$ , to the horizontal axis. The associated shape primitive is the collection of these lengths converted to  $\mu\text{m}$ ,  $\mathbf{v}_p = (v_{p,1}, v_{p,2}, \dots, v_{p,D})$ .

were classified at the last observation time by strain and nutrient conditions. These experiments are described in the following subsections.

### Classification by growth stage

Initially, only the A7-50 strain was considered, with the data set comprising images taken at all observation times after initiation of growth. Thus, the data set comprised 80 images from the A7-50 set and each observation time was considered to be one experimental group. The purpose was to use shape features to distinguish between observation times, and to test whether or not oriented thickness measures may successfully be used to predict the growth stage of the yeast colonies.

CSPs based on oriented thickness measures were computed by  $K$ -means clustering (Section 2.3.1) on the representation space  $F_g$  derived from the training images in each group  $g$ . The number of clusters was  $K = 10$ , with 100 iterations and one replicate, resulting in a CSP dictionary of size  $M = 8K = 80$ . A normalised histogram  $\mathbf{h}_i$  of CSP labels (Equation 3.2) was computed for each im-

age  $X_i$ . Linear discriminant analysis (Section 2.4.1) was used to classify images into each of the eight groups, with feature vectors as described below. Exhaustive search was the method for feature selection. That is, for three features, all combinations of three features were tested and those corresponding to the best classification accuracy were chosen. The classification process was repeated 30 times, with the stochastic step being  $K$ -means clustering. The entire process was repeated five times with  $D = 2, 4, 8, 12$ , and 18.

The classification process proceeded as described here, with three alternatives for the feature vectors input into the classifier. First, the normalised histograms  $\mathbf{h}_i$  of CSP labels (Equation 3.2) were input into the classifier. Second, to make use of both the oriented thickness measures and spatial indices (Section 4.2), the augmented feature vector for colony  $X_i$  was defined to be

$$\mathbf{h}_i^* = \left( h_{i,1}, h_{i,2}, \dots, h_{i,M}, \frac{I_r}{M}, \frac{I_\theta}{M}, \frac{I_\Theta}{M} \right), \quad (4.6)$$

where the  $h_{i,j}$  are as in Equation 3.2 and  $M = 80$  is the number of CSPs. The factor  $1/M$  ensures that all features are of the same magnitude. Third, only the three indices  $I_r, I_\theta$  and  $I_\Theta$  (Section 4.2) were used for classification. More specifically, the feature vector representing colony  $X_i$  was

$$\mathbf{h}_i^* = \left( \frac{I_r}{M}, \frac{I_\theta}{M}, \frac{I_\Theta}{M} \right), \quad (4.7)$$

with all three features used for classification. Since these feature vectors did not require CSPs to be computed and there was no feature selection step, this classification was only conducted once.

The purpose of this part of the study was to quantify how well the indices  $I_r, I_\theta$  and  $I_\Theta$  describe and discriminate between each of the growth stages, and to ascertain whether including the indices increases the classification score.

In addition, for each value of  $D$  (Equation 4.5) the entire classification process was repeated with randomised data. Randomising the data involved per-

muting all of the image labels (including sample number and time) 30 times. For each permutation, the classification process was applied with the feature vectors chosen to be as in Equation 3.2. Ten clusters per class were used to compute the CSPs, and three features were used for classification using linear discriminant analysis (Section 2.4.1). For each permutation, only one trial of the classification process was conducted. Conducting more than one trial was unnecessary because the 30 permutations were random.

### Rotation invariance

The yeast colonies considered in this study do not appear to grow with bias in any particular direction (Section 4.1.1). Thus, ideally, a classification method should be equally effective regardless of the orientation of the colony inside the image frame. In particular, a classification process is desired that recognises the same pattern, but occurring at different orientations, as a single pattern. Recording shape primitives using oriented thickness measures as described in Section 4.3.1 does not achieve rotation invariance (Figure 4.7). The purpose of this part of the study was to introduce a rotationally invariant modification of oriented thickness measures for the classification of yeast colonies, and to evaluate whether this improves the classification accuracy.

Consider a point  $p \in \partial\Omega$ , and let  $\phi$  be the orientation of line joining the centroid of the colony to  $p$ . Let the shape primitive  $\mathbf{v}_p$  at  $p$  be as in Equations 3.3 and 4.5, where the angle  $\theta_d$  is instead given by

$$\theta_d = \phi + (d - 1)\frac{\pi}{D}, \quad d = 1, \dots, D. \quad (4.8)$$

Shape primitives  $\mathbf{v}_p$  computed using the orientations in Equation 4.8 are rotation invariant (Figure 4.8). Classification of the colonies with these shape primitives proceeded as described in the previous subsection.




Figure 4.7: Recording oriented thickness measures in an identical image, rotated. The process of measuring oriented thicknesses does not achieve rotation invariance, in the sense that the same pattern appearing at two different orientations is not recognised as the same pattern. A section cropped from A7-50 colony number five at  $t = 233$  hours is used as an example. White pixels represent “occupied” pixels. The shape primitive  $\mathbf{v}_p$  is computed at  $p$  (black point). The lengths of the coloured lines represent the thickness measurements, with blue, orange, yellow, and purple representing  $v_{p,1}$ ,  $v_{p,2}$ ,  $v_{p,3}$ , and  $v_{p,4}$ , respectively. (a) The original shape pattern and the associated thickness measures. Here,  $\mathbf{v}_p = (23.6, 18.2, 30.0, 86.8)$ , with lengths recorded in  $\mu\text{m}$ . (b) The entire image has been rotated by  $\pi/2$  radians, with the point  $p$  being kept consistent in relation to the object. The shape primitive here is  $\mathbf{v}_p = (30.0, 86.8, 23.6, 19.3)$ , which is recognised as a different shape pattern to the one in (a). Note that these images have been generated using the actual program used to compute the shape primitives, and the discrepancies in measurements between images (a) and (b) are due to sampling. In practice, the shape primitives were computed at points on the boundary  $\partial\Omega$ , but here a point was chosen inside  $\Omega$  for clarity.




Figure 4.8: Rotationally invariant method for computing shape primitives using oriented thickness measures. The example image from Figure 4.7 is used. Here, the first thickness measurement is always in the direction of the line joining the centroid of the colony (not pictured) and the point  $p$ . (a) The original image. (b) The image has been rotated by  $\pi/2$  radians. The shape primitives are given by  $\mathbf{v}_p = (111.4, 31.1, 18.2, 24.6)$  for both examples, demonstrating that this method for recording shape primitives results in the same pattern occurring at a different orientation being recognised as the same pattern.

### Classification by strain and nutrient conditions

The author would like to acknowledge and thank B. J. Binder and H. Tronolone for their helpful discussion regarding the work presented in this section.

For each of the three strains, only images taken at the last observation time were considered, resulting in a data set of 10 A7-50 images, nine A7-500 images, and 10 AR-50 images. The purpose of conducting the classification was to test how well shape features could distinguish between two different strains with the same nutrient concentration (using A7-50 and AR-50), between the same strain grown at different nutrient concentrations (using A7-50 and A7-500), and between any of the three groups of yeast given no prior information about the strain or nutrient concentration (a three-group classification problem comprising A7-50, A7-500, and AR-50). In this part of the study, both oriented thickness measures (Section 4.3.1) and spatial indices (Section 4.2) were used to classify the yeast colonies.

To make use of both the oriented thickness measures and spatial indices, the augmented feature vector  $\mathbf{h}_i^*$  for colony  $X_i$  was as in Equation 4.6, where  $M$  is the number of CSPs. Using the augmented feature vectors  $\mathbf{h}_i^*$  as input, linear discriminant analysis (Section 2.4.1) was used for classification. One feature at a time was used for classification in order to avoid overfitting, since the number of images in the testing sets was small.

Within each classification problem, three methods for testing and analysis were considered. First, the images from each group were split as evenly as possible between training and testing sets. The classifier was trained and tested once for each individual feature, and the best feature and the corresponding highest classification score were recorded. The  $K$ -means clustering step (Section 2.3.1) was repeated  $n = 30$  times. The mean and standard deviation of the classification accuracies, and the number of times each of the spatial indices  $I_r$ ,  $I_\theta$ , and  $I_\Theta$  (Section 4.2) were chosen as a best feature were recorded. Since the CSP labels and values were different for each run, it was not possible to record the number of times these were chosen as a best feature. The purpose was to ascertain whether CSPs or spatial indices tend to be more important in distinguishing between yeast colonies, and to check if a combination of both can result in high classification accuracies.

Second, leave-one-out cross-validation (LOOCV) was conducted [51, 68]. For each image that was left out, the full process of computing CSPs was performed on the remaining images. The single best feature for classifying the left-out image was chosen by exhaustive search. This resulted in  $S$  binary classification accuracies (correct or incorrect), where  $S$  is the total number of images. The total accuracy out of  $S$  samples and the number of times each spatial index was chosen as a best feature were recorded. The leave-one-out method implemented here was designed to test the performance of the full process, as opposed to recommending an algorithm for classifying yeast colonies. For the latter, one would have to apply LOOCV based on a fixed set of CSPs.



The above two methods for testing the classification process give classification scores, but do not necessarily shed any light on which shape patterns differentiate well between colonies. For the third method for evaluating the classification process (hereafter called the feature analysis method), the training and testing method was repeated only once ( $n = 1$ ), and the features that gave a high classification score were recorded. If a particular CSP is the best feature for classification, this means that the CSP occurs more frequently in one experimental group than another. A given CSP may be visualised by drawing a histogram of the lengths  $v_{p,d}$  corresponding to the associated shape primitive  $v_p$ . Each bin represents one angle  $\theta_d$ , where  $d = 1, \dots, D$ , and the height of the bin represents the corresponding length  $v_{p,d}$ . Although this method for evaluating the classification process is not as robust as training and testing, or LOOCV, it allows for characterisation and visualisation of the important shape patterns that discriminate well between yeast colonies.

For each of the three classification problems, the training and testing, and LOOCV methods were repeated with the number of angles used to compute the CSPs set at  $D = 4$  and 12, and with  $K = 5$  and 10 clusters per class. The feature analysis method was then applied to the best combination of  $D$  and  $K$  from the training and testing method to produce images representing the best shape features. A rotationally invariant variation of the method was not considered in this case. This choice, along with the choice to use  $D = 4$  and 12, was made because of the results on the A7-50 group individually – the best classification accuracies were achieved using  $D = 4, 12$ , and without rotation invariance incorporated into the method (Section 5.1.1).

Three classification tasks were conducted: (1) distinguish between strains A7-50 and AR-50; (2) distinguish by nutrient concentration for A7-50 and A7-500; and (3) simultaneous three way classification of A7-50, A7-500, and AR-50. Within each problem, the images from each data set were split evenly between the training and testing groups at random, except for A7-500, for which five

images were used for training with four left for the testing set. A classification task with  $m$  groups and  $K$  clusters resulted in  $M = Km$  CSPs, while the number of images used for LOOCV was equal to the total number of images in the classification problem.

### 4.3.2 Cancellous bone

Since the cancellous bone data is in the form of discrete arrays, thicknesses were measured in 13 directions, defined by the 26 voxels connected to the central voxel  $p$ . Formally, 13 unit vectors were defined as

$$u_d = \frac{e_d}{\|e_d\|}, \quad d = 1, \dots, 13,$$

where the vectors  $e_d$  were given by

$$\begin{aligned} e_1 &= (1, 0, 0), & e_2 &= (0, 1, 0), & e_3 &= (1, 1, 0), & e_4 &= (1, -1, 0), \\ e_5 &= (1, 0, 1), & e_6 &= (-1, 0, 1), & e_7 &= (0, 1, 1), & e_8 &= (0, -1, 1), \\ e_9 &= (1, 1, 1), & e_{10} &= (-1, 1, 1), & e_{11} &= (1, -1, 1), & e_{12} &= (-1, -1, 1), \\ e_{13} &= (0, 0, 1). \end{aligned}$$

Then, each angle  $\theta_d$  was defined by the direction of the vectors  $u_d$ ,  $d = 1, \dots, 13$ . The direction defined by  $e_{13} = (0, 0, 1)$  was approximately the direction of the major axis of the tibia.

The shape primitives  $\mathbf{v}_p$  were computed as in Equation 3.3 using the angles  $\theta_d$ ,  $d = 1, \dots, 13$ , at approximately 10 000 points randomly subsampled from the object  $\Omega$ .  $K$ -means clustering with  $K = 10$  (Section 2.3.1) was applied to each representation space  $F_g$ ,  $g = 1, \dots, 9$  (Section 3.1), so the number of CSPs was  $M = 9K = 90$ . CSP maps were obtained for each bone block, and the frequencies of CSP occurrences (Equation 3.2) were viewed as features for classifying the bone blocks into nine classes. Linear discriminant analysis (Section 2.4.1) was used for classification, and the three best features were chosen by exhaustive search – for each of the  $\binom{90}{3}$  combinations of the three features, a classification accuracy was found, and the best accuracy was chosen as the clas-

sification accuracy. The  $K$ -means clustering step (Section 2.3.1) was repeated 30 times. These steps were used in all variations of the shape analysis described in the following subsections.

Four separate versions of shape analysis were considered, in which different versions of the shape primitives were computed. Due to the nature of the data set, measuring the thickness of the bone in a number of directions inherently captures both shape and density information, where density refers to bone volume per unit tissue volume (BV/TV). Rather than considering rotation invariance here, methods for separating shape and density information were explored.

### Shape and density information, unseparated

CSPs were computed in the directions  $\theta_d$ ,  $d = 1, \dots, 13$ , according to Equation 3.3. Classification was performed as described at the beginning of Section 4.3.2. This shape analysis appeared in a preliminary study by Martin and Bottema [91]. In that paper, shape information and density information were naturally included in the shape primitives in the sense that a given shape pattern at a point  $p$  consisted of the lengths of bone segments in the 13 directions encoded in the shape primitive  $\mathbf{v}_p$ . If  $\mathbf{v}_p$  is a pattern in sparse bone and  $\alpha\mathbf{v}_p$  is a pattern in dense bone, where  $\alpha > 1$ , then the shape patterns may be viewed as the same but scaled by the amount of bone. Thus, two identical shape patterns at different scales are treated as two distinct patterns. The methods of computing shape primitives presented in the following subsections are new. The purpose is to separate shape information and density information in the sense that, wherever shape information is considered, two identical patterns at different scales are treated as an identical pattern. For some versions of the classification method, density information was included separately, or only density information was considered.

### Shape information only

The shape primitive  $\mathbf{v}_p$  (Equation 3.3) was replaced by the normalised shape primitive

$$\mathbf{w}_p = \begin{cases} \frac{\mathbf{v}_p - \text{mean}(\mathbf{v}_p)}{\text{std}(\mathbf{v}_p)} & \text{if } \text{std}(\mathbf{v}_p) \geq T, \\ \mathbf{v}_p - \text{mean}(\mathbf{v}_p) & \text{if } \text{std}(\mathbf{v}_p) < T, \end{cases}$$

where  $T = 10^{-13}$  (the “tolerance”) was set to avoid round-off errors. Classification was performed as described at the beginning of Section 4.3.2.

### Shape and density information, separated

In order to consider both shape patterns and size, but as distinct attributes of each sub-block, the normalised histogram  $\mathbf{h}_i$  of CSP labels (Equation 3.2) for sub-block  $X_i$  was augmented with a component representing a density parameter. For sub-block  $X_i$ , define  $S_i$  by

$$S_i = \frac{\text{number of on voxels in sub-block } X_i}{\text{total volume of sub-block } X_i}. \quad (4.9)$$

The augmented feature vector representing  $X_i$  was

$$\mathbf{h}_i^* = \left( h_{i,1}, h_{i,2}, \dots, h_{i,M}, \frac{S_i}{M} \right)$$

where the  $h_{i,j}$  are as in Equation 3.2 and  $M = 90$  is the total number of CSPs.

The classification proceeded as described at the beginning of Section 4.3.2.

### Density information only

The single feature  $S_i/M$  (Equation 4.9) was used to classify the sub-blocks. Since the classifier did not involve any CSPs, there was no clustering step and no feature selection step.

### 4.3.3 Marbling in beef

In a preliminary study on using CSP-based features to classify the steers (Section 4.1.3) according to treatment group (A+ versus A-), the results showed no clear difference between the two groups (M. Bottema, personal communication, 3 May 2017). For this reason, the focus here was to determine whether the shape patterns in the marbling vary as the amount of the marbling increases. Thus, in this part of the study, all of the striploins were treated as one experimental group, and the purpose was to use the occurrences of the CSPs as features for linear regression to predict marbling proportion. Let  $MP_i$  be the marbling proportion corresponding to animal  $X_i$  (Equation 4.1).

The shape primitive computed at  $p \in \Omega$  was

$$\mathbf{v}_p = (v_{p,1}, v_{p,2}, v_{p,3}, v_{p,4}, v_{p,5}),$$

where  $v_{p,d}$  is the length of the longest line segment passing through  $p$  in the direction  $\theta_d = (d - 1)\frac{\pi}{4}$ , for  $d = 1, \dots, 4$ . The angles  $\theta_d$  were measured within each slice cut from each striploin (Section 4.1.3). The measurement  $v_{p,5}$  was the length of the longest line segment passing through  $p$  in the direction perpendicular to the slices. That is, thicknesses of the marbling were measured in four directions within each slice, and in one direction between the slices. This was necessary since the distance between slices was much larger than the resolution within the slices (Section 4.1.3).

In practice, all length measurements were converted to absolute lengths in mm. The shape primitives were computed at approximately 10 000 randomly subsampled from the marbling  $\Omega$  in each striploin, resulting in approximately 200 000 shape primitives in total. A  $K$ -means clustering algorithm (Section 2.3.1) was applied to the space of all shape primitives, with  $K = 10$ , resulting in  $M = 10$  CSPs in total. For each striploin  $X_i$ , the normalised histogram

of CSP occurrences

$$\mathbf{h}_i = (h_{i,1}, h_{i,2}, \dots, h_{i,10}) \quad (4.10)$$

was computed (Section 3.1), where the  $h_{i,j}$  are as in Equation 3.2. The frequencies of the occurrence of each CSP,  $h_{i,j}, j = 1, \dots, 10$ , were then viewed as features for linear regression to predict marbling proportion MP.

A linear regression model was fit using one, two, and three features. The best feature (or combination of features) was chosen by exhaustive search. That is, for each combination of features, a linear regression model was fit. The model with the highest correlation value  $R$  was chosen as the best model. For each number of features, the best  $R^2$  value, best feature (or combination of features), and the vector corresponding to the best feature (or combination of features) were recorded. (Here,  $R$  refers to the Pearson correlation coefficient.)

#### 4.4 Implementation of features derived from best-fitting ellipsoids

In order to implement the method of best-fitting ellipsoids (Section 3.1.2), two data sets were considered: the A7-50 yeast colony data with colonies taken from all eight observation times (each observation time was again considered to be one experimental group, and the data set comprised 80 images in total), and the cancellous bone data. Features derived from best-fitting ellipsoids were not implemented on the marbling data set, predominantly because of the difference in resolution between slices and within each slice.

In each case, the analysis of a single point  $p$  is scaled up to the entire collection of images as shown in the workflow diagram in Figure 3.1.

In the case of the yeast colony data, let  $\partial\Omega$  be the subset of  $\Omega$  comprising boundary pixels, that is, pixels in  $\Omega$  with at least one unoccupied neighbour.




Figure 4.9: A7-50 colony number five imaged at  $t = 233$  hours. (a) An illustration of the largest disk  $B_{R'}^c$  (red) centred at the centroid  $c$  of the colony (red cross) that fits entirely inside the colony. The radius of the disk is  $R'$ . (b) The subset  $\Omega_{\text{fil}}$  corresponding to the filamentous part of the colony, inside which the ellipse fitting process was initialised.

Let  $c$  be the centroid of the colony and let

$$R' = \arg \min_{p \in \partial\Omega} \{ \|c - p\| \}.$$

That is,  $R'$  is the radius of the largest disk  $B_{R'}^c$  centred at  $c$  that fits entirely inside the colony. Let  $\Omega_{\text{fil}} = \Omega \setminus B_{R'}^c$ , so named because this is the “filamentous” part of the colony (Figure 4.9). Ellipses were initialised at 10 000 sampled points  $p \in \Omega_{\text{fil}}$ , since the shape of the filamentous part of the colony was of interest. In the case of the cancellous bone data, ellipsoids were initialised at a randomly sampled set of points from  $\Omega$ .

In both cases, the best fitting ellipsoid  $E^p$  to  $\Omega^p$  was found as described in Section 3.1.2. Shape primitives  $\mathbf{v}_p$  were constructed according to the general construction given in Definitions 1–4 of Section 3.1.2. For the oriented object parameter (Definition 2) and object shape parameter (Definition 4) shape primitives, if  $L_i = 0$  for all  $i = 1, \dots, n$ , then those shape primitives were omitted from the representation space. The proportion of centroids of all ellipsoids fit to all images or arrays falling outside the object  $\Omega$  was recorded in order to ascertain whether an unreasonable proportion of shape primitives was being

Table 4.3: Numerical details of the classification process for each data set. # pts is the number of sampled points at which the ellipsoids were initialised,  $n$  is the dimension of the binary data array,  $K$  is the number of clusters obtained using  $K$ -means clustering,  $M$  is the number of CSPs, # fts is the number of best features selected during classification, and  $r$  is the radius of the balls (in  $\mu\text{m}$ ) used to initialise the ellipsoids.

Data set	# pts	$n$	$K$	$M$	# fts	$r$
Yeast colonies	10 000	2	10	80	3	3.03, 15.2, 75.6, 303, 909
Cancellous bone	10 000	3	10	90	3	17.4, 43.5, 131, 261, 435

omitted.

CSPs were computed by clustering each representation space  $F_g$  (Equation 3.1) constructed from the training images for each group  $g$ . The  $K$ -means algorithm (Section 2.3.1) with 100 iterations and one replicate was used for clustering, with  $K = 10$  clusters per group. This resulted in  $M = 80$  or  $M = 90$  CSPs in total for yeast colonies and cancellous bone, respectively (Table 4.3). Normalised histograms of CSP labels were computed for all images or arrays in the data set. Linear discriminant analysis (Section 2.4.1) was used during the training and testing steps, and the best combination of three features was chosen by exhaustive search. Each combination of features was tested individually and the combination that gave the highest classification accuracy was selected for the testing step. Thirty trials were conducted, with the stochastic step being the  $K$ -means clustering algorithm.

For each definition of the shape primitives, the entire process was repeated five times with the parameter  $r$  corresponding to the radius of the ball  $B_r$  (Section 3.1.2) set at varying lengths (Table 4.3). These lengths were chosen to be whole numbers of pixels or voxels and converted to  $\mu\text{m}$  during the implementation of the process.



## Chapter 5

# Results and analysis

In this chapter, the results of the experiments described in Chapter 4 are presented. These results are interpreted and analysed throughout. Section 5.1 gives detailed results for the implementation of CSPs using oriented thickness measures. The results for CSPs computed using best-fitting ellipsoids are given in Section 5.2.

All  $p$ -values given in this chapter correspond to a two-tailed  $t$ -test. The values of  $g$  correspond to Hedges'  $g$  used to measure the level of effect [46].

### 5.1 Oriented thickness measures

First, the results on the yeast colonies are discussed, including classifying only the A7-50 group by growth stage, classifying the A7-50 group by growth stage using a rotationally invariant adaptation of oriented thickness measures, and classifying the colonies by strain and nutrient conditions (Section 5.1.1). Second, the results of classifying the cancellous bone data into the nine groups defined by the three experimental groups and three distances from the growth plate are presented (Section 5.1.2). Third, the results on the marbling data are given, where the objective was to use CSP-based features in a regression model to predict marbling proportion. The important predictive features are

also visualised and discussed (Section 5.1.3).

### 5.1.1 Yeast colonies

#### Classification by growth stage

When the A7-50 yeast strain was classified into experimental groups defined by the eight observation times after initiation of growth, using features based on CSPs only, the highest classification accuracy was achieved with  $D = 4$  angles (Table 5.1). When the spatial indices were included in the augmented feature vector (Equation 4.6), the classification accuracy increased overall, with the best accuracy being achieved with  $D = 12$  (Table 5.1). The best classification accuracy when spatial indices were included was statistically significantly different from the best classification accuracy when only CSPs were used ( $p = 0.011$ ). Only the radial index was chosen as a best feature during classification (Table 5.2). Out of the 30 trials, the radial index was chosen 27 times (Table 5.2). The number of times each feature was chosen was independent of the order in which the features appeared in the feature vector (the experiment was repeated in full for each of the six possible permutations of three features, and the angular and pair-correlation indices were never chosen, with the radial index chosen approximately 20 times on average – detailed results are not shown).

These results suggest that the information extracted by the radial index significantly contributes towards the classification. This is not surprising, since the indices were designed to capture temporal changes in the yeast colonies [6]. The radial index detects the outward growth of the colonies, whereas the angular index captures angular growth of the colonies and the pair-correlation index may be thought of as a measure of aggregation. The angular and pair-correlation indices are designed to capture differences in directional growth, for example if a yeast colony is growing towards a nutrient source. For a set of colonies that are growing radially outwards, evenly in all directions, the result that the radial index best captures the differences between colonies makes sense (Figure 4.2).

Table 5.1: Classification accuracy as a function of  $D$ , for the A7-50 strain with each of the eight observation times considered to be one experimental group. CSPs refers to the classification accuracies obtained when only features based on CSPs were used (Equation 3.2). Spatial indices refers to the classification accuracies obtained when spatial indices were included (Equation 4.6), and Randomised shows classification accuracies for the corresponding classification on randomised data. For each classifier,  $M = 80$  CSPs were computed and  $n = 30$  trials were conducted. The highest classification accuracy for each (non-randomised) classification process is shown in bold.

$D$	CSPs	Spatial indices	Randomised
2	$\mu = 0.793, \sigma = 0.025$	$\mu = 0.838, \sigma = 0.013$	$\mu = 0.330, \sigma = 0.034$
4	$\mu = \mathbf{0.864}, \sigma = \mathbf{0.033}$	$\mu = 0.879, \sigma = 0.025$	$\mu = 0.323, \sigma = 0.029$
8	$\mu = 0.854, \sigma = 0.028$	$\mu = 0.872, \sigma = 0.014$	$\mu = 0.325, \sigma = 0.027$
12	$\mu = 0.860, \sigma = 0.026$	$\mu = \mathbf{0.883}, \sigma = \mathbf{0.021}$	$\mu = 0.327, \sigma = 0.033$
18	$\mu = 0.841, \sigma = 0.026$	$\mu = 0.877, \sigma = 0.024$	$\mu = 0.325, \sigma = 0.028$

Table 5.2: The number of times (out of 30) that each of the spatial indices were chosen as one of the three best features for classification, as a function of  $D$ , for the classification of the A7-50 strain with each of the eight observation times considered to be one experimental group.

$D$	$\frac{I_r}{M}$	$\frac{I_\theta}{M}$	$\frac{I_\Theta}{M}$
2	27	0	0
4	14	0	0
8	15	0	0
12	25	0	0
18	20	0	0

When the feature vectors described in Equation 4.7 (that is, spatial indices only) were input into the classifier, the classification accuracy was 0.525. This suggests that, while spatial indices capture some information about the temporal evolution of the colony, the information captured by the CSPs and the spatial indices complements each other. The combination of CSPs and spatial indices results in better prediction of the growth stage of the yeast.

In addition, for all values of  $D$ , the classification accuracy achieved by the CSPs or by the combination of CSPs and spatial indices were significantly higher than for the corresponding experiment with randomised class labels (Table 5.1),

suggesting that these features did not attain such high classification accuracies by chance alone.

### Rotation invariance

When the A7-50 strain was classified according to growth stage, with the shape primitives computed using Equations 3.3 and 4.8, the best classification accuracy was achieved with  $D = 12$  angles (Table 5.3). For all values of  $D$ , the CSPs computed without rotation invariance and the CSPs without rotation invariance combined with the spatial indices outperformed the CSPs computed with rotation invariance (Tables 5.1 and 5.3). The best classification accuracy obtained using the CSPs without rotation invariance ( $D = 4$ ) was significantly different from the best classification accuracy obtained using the CSPs computed with rotation invariance ( $D = 12$ ) ( $p < 0.001$ ). Similarly, the best classification accuracy obtained using the CSPs computed without rotation invariance plus the spatial indices ( $D = 12$ ) was significantly different from the best classification accuracy obtained using the CSPs computed with rotation invariance ( $D = 12$ ) ( $p < 0.001$ ).

Table 5.3: Classification accuracy as a function of  $D$ , for the A7-50 strain with each of the eight observation times considered to be one experimental group and oriented thickness measures modified to be rotationally invariant (Equation 4.8). For each value of  $D$ ,  $M = 80$  CSPs were computed and  $n = 30$  trials were conducted. The best classification accuracy is shown in bold.

$D$	Classification accuracy
2	$\mu = 0.729, \sigma = 0.027$
4	$\mu = 0.778, \sigma = 0.031$
8	$\mu = 0.819, \sigma = 0.030$
<b>12</b>	$\mu = \mathbf{0.835}, \sigma = \mathbf{0.027}$
18	$\mu = 0.816, \sigma = 0.030$

These results are somewhat surprising. The angles at which the thickness measurements were made were modified so that two patterns occurring in the pseudohyphae at two different angles (with respect to the centroid of the colony)

would be recognised as the same pattern. The yeast colonies do not appear to have grown with bias towards any particular direction (for instance, towards a nutrient source), instead appearing to grow radially outwards at an even rate in all directions (Figure 4.1). For this reason, in theory, CSPs computed with rotation invariance should improve the classification score, since the same shape pattern should be recognised as the same, regardless of the angle at which it appears. However, the results do not support such a theory, instead suggesting that directional information does contribute towards the classification of the A7-50 colonies into groups defined by growth stage.

For a majority of points  $p$ , the rotationally invariant variation of the method would inherently result in the largest length measurement being recorded as the first element of the shape primitive. In the rotationally invariant method, the classification results could possibly be due to filament length, rather than the complex shape patterns present in the pseudohyphae. This may be an explanation for the lower classification accuracies obtained using this variation of the method.

### **Classification by strain and nutrient conditions**

The author would like to acknowledge and thank B. J. Binder and H. Tronolone for their helpful discussion regarding the work presented in this section.

#### *(1) Classification by strain:*

When the data sets A7-50 and AR-50 were classified according to strain, the best accuracy score was achieved using  $D = 4$  and  $K = 10$  (Table 5.4). However, this classification accuracy was not significantly different to those obtained using  $D = 4, K = 5$  or  $D = 12, K = 10$ . The results suggest that measuring the lengths of the yeast colonies at only four angles provides sufficient information towards classification using features based on CSPs. Choosing a higher number of clusters per class appears to be more important than using a higher number

Table 5.4: The mean and standard deviation ( $n = 30$ ) as a function of the number of clusters  $K$  per class and number of angles  $D$  chosen to compute the CSPs when classifying A7-50 and AR-50 by strain. The best classification accuracy is shown in bold.

$K$	4 angles	12 angles
5	$\mu = 0.987, \sigma = 0.035$	$\mu = 0.907, \sigma = 0.037$
10	$\mu = \mathbf{0.997}, \sigma = \mathbf{0.018}$	$\mu = 0.987, \sigma = 0.035$

Table 5.5: The accuracy score as a function of the number of clusters  $K$  per class and number of angles  $D$  used to compute the CSPs when performing LOOCV for the classification of A7-50 and AR-50 by strain.

$K$	4 angles	12 angles
5	1.00	1.00
10	1.00	1.00

of angles at which to make length measurements.

During training and testing, the spatial indices were never chosen as one of the best features for classification. At approximately 240 hours after initiation of growth, the values of the indices did not discriminate well between A7-50 and AR-50, which is consistent with the results in Binder *et al.* [6]. However, Binder *et al.* showed that the spatial indices took on different values at approximately 170 hours after the start of growth [6]. If the classification process presented here was repeated using yeast colonies imaged at 170 hours instead of 240 hours, then the spatial indices may be more useful towards the classification. The tradeoff is that the CSPs may not capture as much information because the pseudohyphae in either colony have not had as much time to develop.

During LOOCV, all 20 images were classified correctly regardless of the number of angles  $D$  at which the CSPs were computed and the number of clusters  $K$  per class (Table 5.5). This means that, for all 20 images, there was at least one feature (either based on a CSP or spatial index) that could classify that image correctly.

Table 5.6: The number of times (out of 20 runs) each of the spatial indices was chosen as one of the best features when classifying A7-50 and AR-50 by strain at a single time point and performing LOOCV.

Index	Times chosen
$I_r$	0
$I_\theta$	6
$I_\Theta$	8

The angular index and angular pair-correlation index were picked as one of the best features six and eight times, respectively (Table 5.6). This does not necessarily mean that these spatial indices were the only features that correctly classified these images, since there may have been more than one feature that gave a correct classification. The number of times each index achieves a classification accuracy of 1.00 is independent of the choices for  $D$  and  $K$ , since these affect the computed CSPs only. The results demonstrate that a combination of shape features based on CSPs and spatial indices can classify yeast colonies by strain to a very high degree of accuracy.

Although not a robust test of the classification process, the feature analysis method (Section 4.3.1) allows for visualisation of the shape patterns that result in good discrimination between the groups. In order to analyse the features that gave the best discrimination between strains, the training and testing method was repeated once with  $D = 4$  and  $K = 10$ , since these gave the best classification accuracy during training and testing with  $n = 30$  (Table 5.4). Out of 23 features tested, two gave an equal best classification score of 1.00. These were both CSPs rather than spatial indices, and are given by

$$C_6 = (59.7, 52.1, 59.3, 977.6)$$

and

$$C_8 = (54.1, 968.9, 55.0, 44.3),$$




Figure 5.1: The shape patterns that gave a perfect classification of A7-50 and AR-50 by strain. Each bin represents an angle, and the height of the bin represents the length measurement made at that angle (after clustering).

where each entry in the CSP vectors represents a length in  $\mu\text{m}$ , measured at angles of  $0, \pi/4, \pi/2$ , and  $3\pi/4$ , respectively. These may be represented by histograms (Figure 5.1). The histograms provide a means of characterising the shape patterns that are learned by the classification method as the features that discriminate best between strains. The numerical labels for the CSPs do not provide any information other than providing a convenient means of keeping track of the CSPs for each individual run of the classification process.

For each of the features that gave a perfect accuracy for one run of the classification process, the largest length measurement occurred at  $3\pi/4$  and  $\pi/4$ , respectively, with the two features being at a comparable scale. The largest length measurement was approximately  $970 \mu\text{m}$  for both shape features, which suggests that these features were picking out longer filaments present in the yeast colonies at an angle of  $3\pi/4$  and  $\pi/4$ . The filaments of the A7-50 colonies appear to be typically longer than those of AR-50 (Figure 4.1), which likely enables the CSPs to distinguish between strains.

(2) *Classification by nutrient concentration:*

When the colonies from A7-50 and A7-500 were classified by nutrient concentration by splitting the data into training and testing sets and conducting  $n = 30$



Table 5.7: The mean and standard deviation ( $n = 30$ ) as a function of the number of clusters  $K$  per class and number of angles  $D$  used to compute the CSPs when classifying A7-50 and A7-500 by nutrient concentration.

$K$	4 angles	12 angles
5	$\mu = 1.00, \sigma = 0.00$	$\mu = 1.00, \sigma = 0.00$
10	$\mu = 1.00, \sigma = 0.00$	$\mu = 1.00, \sigma = 0.00$

Table 5.8: The accuracy score obtained using LOOCV when classifying A7-50 and A7-500 by nutrient concentration, as a function of the number of clusters  $K$  per class and number of angles  $D$  used to compute the CSPs.

$K$	4 angles	12 angles
5	1.00	1.00
10	1.00	1.00

trials, all combinations of values for  $D$  and  $K$  gave an equal best classification accuracy of 1.00, while the spatial indices were never chosen as a best feature (Table 5.7). Classifying by nutrient concentration appears to be an easier classification problem than classifying by strain. The growth patterns in the A7-500 group are closer to the regular growth pattern than the A7-50 group, exhibiting pseudohyphae that are less developed (Figure 4.1). The classification process picks out these differences effectively.

Similarly, during LOOCV, all choices for  $D$  and  $K$  gave a perfect classification accuracy (Table 5.8). For each of the 19 images, there was at least one individual feature that classified that image correctly.

The radial index was selected over two times more often than either of the other indices (Table 5.9). This is again not surprising as the growth patterns in the A7-500 group appear to be more circular in shape with shorter pseudohyphae than those in the A7-50 group, so the radial distance  $R_{\text{CSR}}$  is expected to be higher for the colonies in the A7-500 group. In general, radial information appears to be more important than angular information when using the spatial indices to discriminate between yeast colonies by nutrient concentration.

Table 5.9: The number of times (out of 19 runs) each of the spatial indices was chosen as one of the best features for classification when performing LOOCV for the classification of A7-50 and A7-500 by nutrient concentration.

Index	Times chosen
$I_r$	15
$I_\theta$	6
$I_\Theta$	5

Since all combinations of values of  $D$  and  $K$  performed equally well during training and testing,  $D = 4$  and  $K = 10$  were chosen for feature analysis for consistency with the results for classification by strain above. Of the 23 features tested, 14 of the CSPs gave a tied best classification accuracy of 1.00 (Figure 5.2). None of the spatial indices achieved such a high classification accuracy. The largest length measurements of several of the best CSP-based features occurred at  $\pi/4$  and  $3\pi/4$ , which is similar to the two best features chosen when classifying by strain (Figure 5.1). In addition, the features were at a similar scale. However, many additional features gave a perfect classification here, which suggests that classification by nutrient concentration is an easier classification problem than classifying by strain.

### *(3) Classification by strain and nutrient concentration*

When the yeast colonies were classified by both strain and nutrient concentration using all three data sets (A7-50, A7-500, and AR-50) by splitting the data into training and testing sets and conducting  $n = 30$  trials, the best classification score was achieved with  $D = 12$  and  $K = 10$  (Table 5.10). However, this score was not significantly different from those obtained using  $D = 4, K = 5$  or  $D = 4, K = 10$ . Increasing the information content of the CSPs by increasing the number of angles at which the CSPs are computed and the number of clusters per class slightly increases the effectiveness of the classification process. The spatial indices were never chosen as a best feature, which suggests that



Figure 5.2: The shape patterns that gave a perfect classification of A7-50 and A7-500 by nutrient concentration. In each histogram, the bins represent the angles  $0, \pi/4, \pi/2, 3\pi/4$  (left to right), and the height of each bin represents the length measurement made at that angle (after clustering). In each panel, the range of the vertical axis is 0 to 2000  $\mu\text{m}$ .

Table 5.10: The mean and standard deviation ( $n = 30$ ) of the classification accuracy as a function of the number of clusters  $K$  per class and number of angles  $D$  used to compute the CSPs for the three-group classification problem using A7-50, A7-500, AR-50. The best classification accuracy is shown in bold.

$K$	4 angles	12 angles
5	$\mu = 0.957, \sigma = 0.052$	$\mu = 0.898, \sigma = 0.064$
10	$\mu = 0.955, \sigma = 0.064$	$\mu = \mathbf{0.969}, \sigma = \mathbf{0.041}$

Table 5.11: The accuracy score obtained using LOOCV for the three-group classification problem as a function of the number of clusters  $K$  per class and number of angles  $D$  used to compute the CSPs using A7-50, A7-500, and AR-50.

$K$	4 angles	12 angles
5	1.00	1.00
10	1.00	1.00

shape-based features are more important when classifying yeast colonies when no information about the strain or nutrient concentration is known beforehand.

During LOOCV, all 29 images were classified correctly regardless of the choices made for  $D$  and  $K$  (Table 5.11). This means there was at least one feature that could classify each image correctly, suggesting that a classification process that uses a combination of features based on CSPs and spatial indices extracts useful information from the binary yeast colonies.

The number of times that each spatial index was chosen as a best feature is shown in Table 5.12. These results are not necessarily achieved exclusively by that feature, as other (CSP-based) features may have resulted in an equal best classification. The number of times each spatial index gave an accuracy score of 1.00 was independent of the choices for  $D$  and  $K$ .

In order to analyse the features,  $D = 12$  and  $K = 10$  were chosen for one run of the classification process, since this combination achieved the best accuracy during training and testing. Of the 33 features tested, two CSPs tied for a classification accuracy of 1.00 (Figure 5.3). The spatial features did not produce a perfect accuracy score.

Table 5.12: The number of times (out of 29 runs) each of the spatial indices was chosen as one of the best features for classification, using LOOCV for the three-group classification problem using A7-50, A7-500 and AR-50.

Index	Times chosen
$I_r$	6
$I_\theta$	1
$I_\Theta$	2

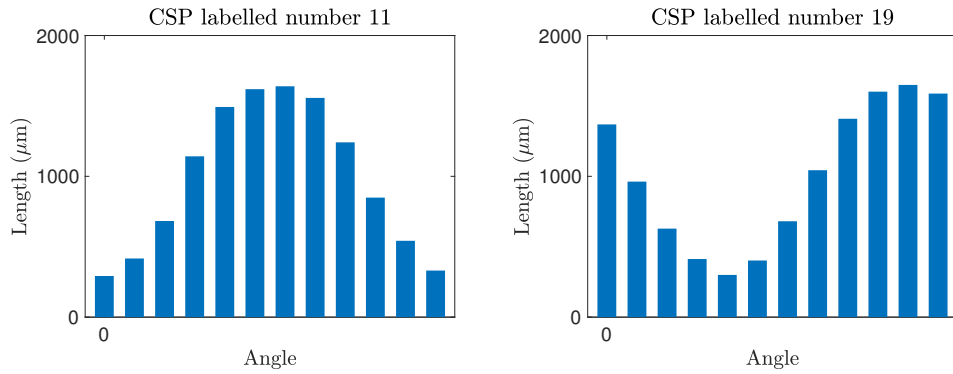


Figure 5.3: The shape patterns that gave a perfect classification when classifying A7-50, A7-500 and AR-50 by both strain and nutrient concentration. Each bin represents an angle, with the leftmost bin corresponding to an angle of 0 and the rightmost bin corresponding to  $11\pi/12$ . The increments between angles are  $\pi/12$ . The height of each bin represents the length measurement made at the corresponding angle (after clustering).

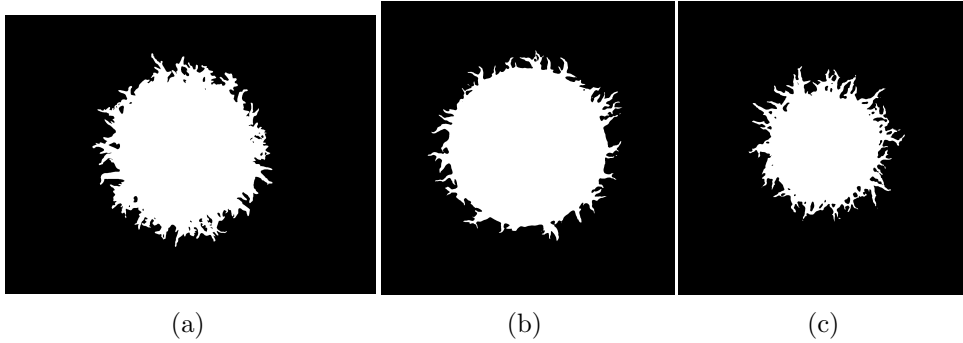


Figure 5.4: One example colony from each of the three groups in the classification by strain and nutrient concentration, with the sizes of the colonies relative to each other displayed correctly so that colonies images at different resolutions can be compared. (a) A7-50. (b) A7-500. (c) AR-50. There are differences in filament lengths and overall colony size between the three groups.

The shapes corresponding to the CSP-based features cannot be compared directly to those from the other classification problems (Figures 5.1 and 5.2) because different values of  $D$  and  $K$  were used to compute the features, although similar patterns can be seen. For instance, the CSP labelled number 19 corresponds to a similar shape to that of CSP labelled number 6 in Figure 5.1, with longer lengths around an angle of  $3\pi/4$  and shorter lengths around  $\pi/4$ . Overall, for the classification of yeast colonies by strain and nutrient concentration, a shape pattern with the largest length measurement oriented at an angle of  $3\pi/4$  appears to give most discrimination between classes, possibly since these measurements are picking out differences in filament lengths between the groups (Figure 5.4).

#### *Comparison with previous work*

Previous studies have considered the classification of *S. cerevisiae* colonies into two morphological groups (smooth versus fluffy) based on 427 image features using regularised logistic regression [119]. Of these features, only six were chosen as important for classification and were derived directly from the images, including texture features derived from the interior of the colony. The classification

method presented here has been developed based on less information, namely binary information only rather than greyscale images. The average classification accuracy for the two-group classification problem where the yeast colonies from A7-50 and AR-50 were classified by strain was slightly better than the two-group morphological classification problem of smooth and fluffy colonies by [119] (0.997 compared with 0.988), indicating that the method presented here is equally powerful.

Previous work on the quantification of yeast colonies required extensive lists of features to be identified as input to the classification in order to identify which were useful. Despite the size of such lists, this approach may result in important features being overlooked and thus not considered by the classification process. The method presented here has the advantage of selecting the best features automatically, avoiding the need to specify features, which risks missing key attributes.

### 5.1.2 Cancellous bone

When CSPs were used to classify cancellous bone based on oriented thickness measures, the best classification accuracy was achieved when shape and density information was input into the classifier separately (Table 5.13). However, this classification accuracy was not significantly higher than the accuracy obtained when both shape and density information was used but input together ( $p = 0.056$ ). Overall, including both shape and density information increased the performance of the classification process.

When shape and density were separated, the density feature was chosen as one of the three best features in all of the 30 trials. The classification process based on these CSPs outperformed the shape only CSPs (Table 5.13,  $p < 0.001, g = 3.41$ ) and the density only CSPs (Table 5.13,  $p < 0.001, g = 6.71$ ). The CSPs based on shape only outperformed the CSPs based on density only (Table 5.13,  $p = 0.092, g = 1.69$ ).

Table 5.13: Performance as a function of the classification process on cancellous bone data. The second column shows the classification accuracies for each method for implementing CSPs with  $K = 10$  clusters per class. For the first three implementations,  $n = 30$  trials were conducted. The best classification accuracy is shown in bold.

Implementation	Classification accuracy
Shape and density, unseparated	$\mu = 0.637, \sigma = 0.021$
Shape only	$\mu = 0.575, \sigma = 0.025$
<b>Shape and density, separated</b>	$\mu = \mathbf{0.647}, \sigma = \mathbf{0.017}$
Density only	0.533

In general, including shape information significantly increased the classification accuracy, regardless of whether the shape and density information was separated. Using CSPs containing shape information only gave a better classification accuracy than using bone density as a single feature, although the level of effect was not as high and the difference was not statistically significant at the 0.05 level. In the context of cancellous bone in rats, shape and density information appear to be complementary. Classification appears to rely more strongly on the shape information than the density information, which suggests that there are indeed differences in the microstructure of the cancellous bone between sham, ovx, and ovx+zol rats.

### 5.1.3 Marbling in beef

The CSPs and histograms of CSP occurrences  $\mathbf{h}_i$  were computed on the marbling data (Section 4.3.3). One feature chosen from  $h_{i,j}$ ,  $j = 1, \dots, 10$  (Equation 4.10) was used to fit the linear regression model to predict marbling proportion MP. The coefficient of determination  $R^2$  was recorded for each feature (Table 5.14).

The occurrence of CSP labelled number 10 was the single best feature for regression (Table 5.14). When combinations of two or three features were used, the occurrences of CSPs labelled 2 and 4 also contributed to predicting marbling proportion (Table 5.15). The best combination of two or three features did not correspond to the best single features, with the exception of the occurrence of



Table 5.14: The coefficient of determination  $R^2$  as a function of CSP label. The best three CSPs are shown in bold.

CSP label	$R^2$
1	0.213
2	0.027
<b>3</b>	<b>0.579</b>
4	0.010
5	0.139
6	0.021
7	0.096
<b>8</b>	<b>0.365</b>
9	0.129
<b>10</b>	<b>0.793</b>

Table 5.15: The best coefficient of determination achieved using linear regression, and the corresponding best combination of features as a function of the number of features used to fit the regression model.

Number of features	Best $R^2$	Combination of features
1	0.793	10
2	0.889	4, 10
3	0.909	2, 4, 10

CSP labelled number 10. The presence of the shape pattern represented by CSP number 10 appears to be strongly correlated with marbling proportion. The values of CSPs 2, 3, 4, 8, and 10 are shown in Table 5.16.

Each of the vectors shown in Table 5.16 corresponds to marbling thicknesses in mm, with the first four thicknesses being measured in four directions within each slice in the striploin, and the fifth thickness being measured in the direction perpendicular to the slices. The thicknesses within each slice were measured with respect to the  $x$  axis, as determined by the orientation at which each slice was photographed. In general, as the occurrence of CSP number 10 increased, the marbling proportion decreased (Figure 5.5).

The physical manifestations of CSPs labelled 2, 3, 4, 8, and 10 are shown

Table 5.16: The values of the CSPs labelled 2, 3, 4, 8, and 10. The CSP labelled number 10 is shown in bold, since this appears to be the most important indicator of marbling proportion. The vectors shown here correspond to the values of the CSPs whose occurrences resulted in the best features for the linear regression.

CSP label	Value
2	(5.60, 7.50, 17.4, 7.36, 11.9)
3	(6.05, 6.04, 4.91, 4.71, 22.1)
4	(3.40, 3.18, 2.68, 2.61, 12.7)
8	(7.14, 18.9, 7.49, 5.74, 13.1)
<b>10</b>	<b>(2.26, 2.26, 2.12, 1.96, 5.66)</b>




Figure 5.5: Marbling proportion plotted as a function of the occurrence of CSP labelled number 10. The correlation is  $R^2 = 0.793$ .

in Figure 5.6, where each CSP is represented by a histogram, and Figure 5.7, where each CSP is represented in the form of a three-dimensional drawing. In the three-dimensional drawing, each thickness measurement is represented by a line segment of length proportional to the thickness in the corresponding direction. Figure 5.6 is an “honest” representation of each CSP in the sense that, in each of the histograms, the height of each bin represents the thickness of the marbling in a given direction. Figure 5.7 is given here to facilitate the visualisation of each CSP, but is not necessarily a true representation of the CSPs because the line segments representing the thickness measurements have been assumed to intersect at their midpoints. Actually, information about the point of intersection of each of the line segments has been lost.

Of the five CSPs that have been identified as most important for predicting marbling proportion, CSP labelled number 10 corresponds to the most “round” shape, with the five thicknesses being relatively uniform (Figures 5.6 and 5.7). The CSPs labelled 2, 3, 4, and 8 each exhibit at least one thickness that is much larger than the others. Thus, some measure of “roundness” appears to be a good predictor of marbling proportion. In addition, CSP labelled number 10 corresponds to a relatively small shape overall compared with the other CSPs shown in either Figure 5.6 or 5.7, so it is not surprising that a large marbling proportion corresponds to a low frequency of the occurrence of CSP 10 (Figures 5.5 and 5.8).

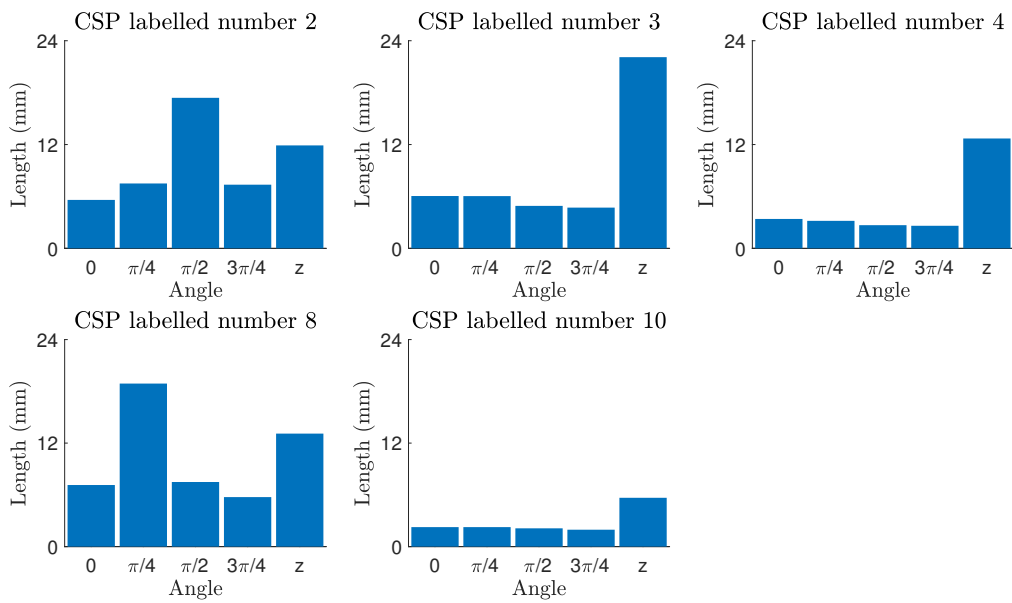


Figure 5.6: The CSPs identified as important for predicting marbling proportion, each represented using a histogram. For each image, the height of each bin represents the thickness of the marbling in a given direction. The left-most four bins represent the angle within the slices of the striploin, and the bin labelled  $z$  represents the angle perpendicular to the slices, or the  $z$  direction.

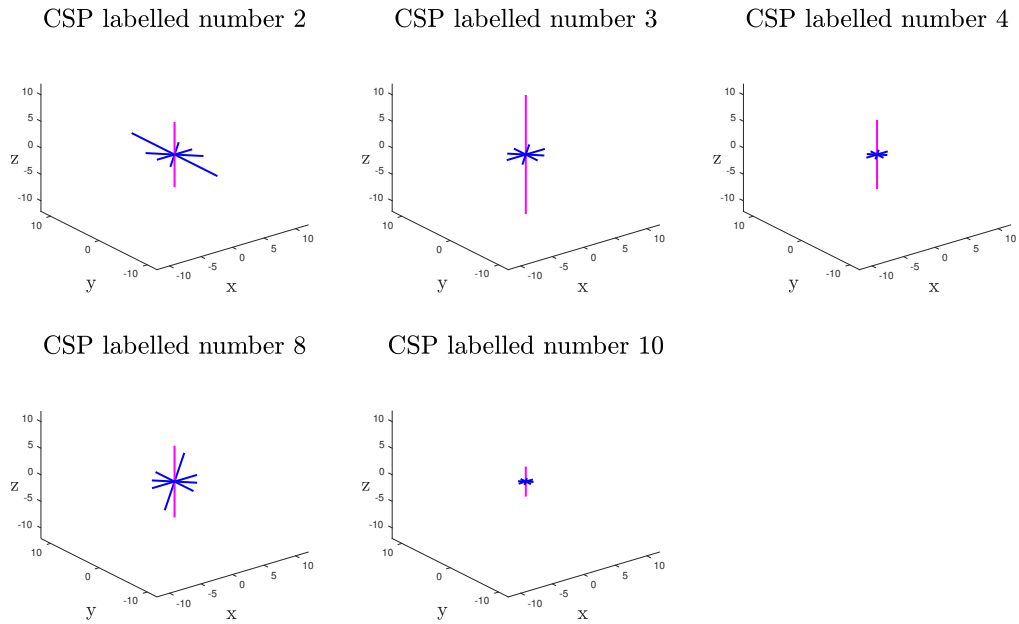



Figure 5.7: The CSPs identified as important for predicting marbling proportion, each represented by a drawing in three dimensions. The length of each line segment is proportional to the thickness measurement in the corresponding direction. The measurements in the  $xy$  plane are shown in blue and the measurement in the  $z$  direction is shown in red, for clarity. These drawings are not necessarily a true representation of each CSP because, for the purpose of visualisation, each of the line segments has been assumed to intersect at the midpoint. Actually, information about the point of intersection of the line segments has been lost. CSP number 10 is smaller and “rounder” than the other CSPs shown.



(a)



(b)

Figure 5.8: Examples of one slice of striploin taken from steers with (a) minimum ( $MP = 0.090$ , represented by slice 13 taken from steer number nine) and (b) maximum ( $MP = 0.286$ , represented by slice 13 taken from steer number 14) marbling proportions. The frequency of the occurrence of CSP labelled number 10 can be used as a feature to predict marbling proportion  $MP$ .

## 5.2 Features derived from best-fitting ellipsoids

Here, the results obtained using CSPs computed with features derived from best-fitting ellipsoids are presented. First, the results of classifying the A7-50 yeast strain by growth stage are given (Section 5.2.1). Second, the results of classifying the cancellous bone blocks into the nine groups as defined by the three experimental groups and three distances from the growth plate are presented (Section 5.2.2). The results are compared with relevant previous work (Section 5.2.3) and briefly summarised (Section 5.2.4).

### 5.2.1 Yeast colonies

When fitting ellipses to the yeast colony data, the largest proportion of ellipse centroids that fell outside the colony was  $3.11 \times 10^{-2}$  for  $r = 75.6 \mu\text{m}$ , with the proportions dropping for smaller and larger values of  $r$  (Table 5.17). In the context of yeast colonies, the possibility of centroids of the ellipses falling outside  $\Omega$  did not pose a serious problem, since this caused relatively few feature vectors to be excluded from the representation space.

When the shape primitives were recorded using oriented ellipsoid parameters (Definition 1), the highest classification accuracy was achieved using  $r = 303 \mu\text{m}$  (Table 5.18). Here, information about the orientation and size of the ellipses was included in the shape primitives. Using disks  $B_r$  with  $r < 303 \mu\text{m}$  resulted in ellipses that captured information at the local level. Although there appears to be some filamentation at the local level, these small filaments seem to be similar in shape regardless of the development of the larger filaments. These small ellipses failed to capture the shape of the larger filaments. For  $r = 909 \mu\text{m}$ ,

Table 5.17: The proportion of ellipses fit to all 80 images whose centroids fell outside the colony as a function of  $r$  (in  $\mu\text{m}$ ) when ellipses were fit to yeast colony data.

$r$	3.03	15.2	75.6	303	909
Proportion	$6.00 \times 10^{-4}$	$8.27 \times 10^{-3}$	$3.11 \times 10^{-2}$	$7.28 \times 10^{-4}$	0




Figure 5.9: An ellipse fit to a yeast colony. The ellipse (blue) was fit using a ball  $B_r$  of radius  $r = 909 \mu\text{m}$  to a sampled point in colony number five at  $t = 233$  hours after initiation of growth. The connected component  $\Omega^p$  is shaded.

the ellipses were very large and captured information about a large area of the colony, leading to information without much discriminative power (Figure 5.9). The ellipses fit using disks  $B_r$ ,  $r = 303 \mu\text{m}$ , resulted in a good tradeoff between picking out information at the local and the global level.

When oriented object parameters were recorded in the shape primitives (Definition 2), the best classification score was achieved when the ellipses were fit using disks of radius  $r = 15.2 \mu\text{m}$ . Here, the orientation of each ellipse was recorded, along with the length of the colony itself in the major and minor directions of the ellipse. When disks  $B_r$  with  $r < 15.2 \mu\text{m}$  were used to fit the ellipses, the ellipses fell too close to the boundary of the colony, resulting in length measurements that did not adequately capture the shape patterns of the colony. For  $r > 15.2 \mu\text{m}$ , the ellipses were large, giving centroids and directions that result in length measurements with little discriminatory power (Figure 5.9). Using disks with  $r = 15.2 \mu\text{m}$  resulted in the centroids of the ellipsoids falling far from the boundaries of the filaments, meaning that oriented object parameters extracted useful information about the length and width of the filaments (Figure 5.10).

When only ellipsoid shape parameters were recorded in the shape primitives






Figure 5.10: An ellipse fit to a portion of a yeast colony. The ellipse (red) was fit using a ball  $B_r$  of radius  $r = 15.2\ \mu\text{m}$  to one filament (white) isolated from colony number five at  $t = 233$  hours. The largest connected component  $\Omega^p$  that intersects with  $B_r^p$  is shaded in grey. The blue and orange lines show the length measurements in the major and minor directions of the ellipse, respectively. Oriented object parameters appear to adequately capture the shape of the filament at this point, whereas the thickness of the ellipse itself does not extract as much information about the shape of the colony.

Table 5.18: Best classification accuracy as a function of  $r$  ( $\mu\text{m}$ ) for yeast colony data. Columns two to five correspond to the accuracy obtained using each of the four methods for recording the shape primitives (Definitions 1–4, respectively). For each method, 80 CSPs were computed and 30 trials were conducted. The highest accuracy achieved by each method is shown in bold.

$r$	Oriented ellipsoid parameters	Oriented object parameters	Ellipsoid shape parameters	Object shape parameters
3.03	$\mu = 0.722,$ $\sigma = 0.014$	$\mu = 0.819,$ $\sigma = 0.025$	$\mu = 0.744,$ $\sigma = 0.022$	$\mu = \mathbf{0.785},$ $\sigma = \mathbf{0.018}$
15.2	$\mu = 0.702,$ $\sigma = 0.023$	$\mu = \mathbf{0.843},$ $\sigma = \mathbf{0.021}$	$\mu = 0.776,$ $\sigma = 0.019$	$\mu = 0.773,$ $\sigma = 0.025$
75.6	$\mu = 0.763,$ $\sigma = 0.033$	$\mu = 0.808,$ $\sigma = 0.020$	$\mu = \mathbf{0.840},$ $\sigma = \mathbf{0.028}$	$\mu = 0.633,$ $\sigma = 0.025$
303	$\mu = \mathbf{0.827},$ $\sigma = \mathbf{0.020}$	$\mu = 0.815,$ $\sigma = 0.023$	$\mu = 0.755,$ $\sigma = 0.010$	$\mu = 0.576,$ $\sigma = 0.021$
909	$\mu = 0.478,$ $\sigma = 0.009$	$\mu = 0.485,$ $\sigma = 0.013$	$\mu = 0.413,$ $\sigma = 0.019$	$\mu = 0.471,$ $\sigma = 0.016$

(Definition 3), the highest classification accuracy was obtained with  $r = 75.6 \mu\text{m}$ . This suggests the ellipses fit the data closely, being of a comparable length and width to that of the filaments themselves. When only object shape parameters were recorded (Definition 4), the best classification accuracy occurred at  $r = 3.03 \mu\text{m}$  (Table 5.18). This suggests the ellipses did not fit the data closely, but their centroids corresponded to a set of subsampled points at which informative object shape parameter measurements were made.

The highest classification accuracy overall ( $\mu = 0.843$ ,  $\sigma = 0.020$ ) was achieved by recording oriented object parameters, after fitting ellipses using disks of radius  $r = 15.2 \mu\text{m}$  (Table 5.18). This result suggests that interrogating the data directly, and including directional information in conjunction with this, is the most effective way of extracting information from the yeast colony data.

### 5.2.2 Cancellous bone

When ellipsoids were fit to the cancellous bone data, a large proportion of centroids fell outside of the bone itself. The proportion increased with  $r$ , with the largest proportion being 0.340 when balls of radius  $r = 435 \mu\text{m}$  were used to fit the ellipsoids (Table 5.19).

Table 5.19: The proportion of ellipsoids fit to all 90 sub-blocks whose centroids fall outside the bone as a function of  $r$  (in  $\mu\text{m}$ ), for cancellous bone data.

$r$	17.4	43.5	131	261	435
Proportion	$7.28 \times 10^{-4}$	$6.33 \times 10^{-3}$	0.083	0.236	0.340

When oriented ellipsoid parameters were used for the shape primitives (Definition 1), the highest classification accuracy was achieved at  $r = 131 \mu\text{m}$  (Table 5.20). If balls of smaller radius were used, then the subset  $\Omega^p$  was small and did not contain much meaningful information about the local structure. In fact, the balls  $B_r^p$  were likely to fall entirely inside the bone, for most points  $p$ , resulting in ellipsoids that were also balls. If balls of larger radius were used, then each  $\Omega^p$  was likely to be large and complicated, and the ellipsoid properties failed to capture orientation and thickness information at the local level. Choosing  $r = 435 \mu\text{m}$  meant that the diameter of each ball was almost as large as the length of the two shortest sides of the bone block, so it is not surprising that this choice of  $r$  corresponded to the lowest classification accuracy. The choice  $r = 151 \mu\text{m}$  gave the optimal tradeoff between extracting local versus global information about the trabecular structure.

When oriented object parameters were recorded (Definition 2), the equal highest classification score occurred when the ellipsoids were fit using balls of radius  $r = 17.4$  and  $43.5 \mu\text{m}$ . The standard deviation for  $r = 17.4 \mu\text{m}$  was slightly smaller than that for  $r = 43.5 \mu\text{m}$  (Table 5.20). As  $r$  increased, the proportion of centroids of the ellipsoids that fell outside the bone also increased, and the classification accuracy decreased. This suggests that, although the

ellipsoids fit using balls with  $r = 17.4$  and  $43.5\ \mu\text{m}$  may be small and capture local information only, their centroids most likely did fall close to the centre of the bone and hence resulted in meaningful bone thickness measurements, especially when combined with directional information.

In the case of the cancellous bone, the classification accuracy tended to improve when directional information was removed from the shape primitives. The highest classification accuracy overall ( $\mu = 0.745$   $\sigma = 0.024$ ) occurred at  $r = 131\ \mu\text{m}$  when only object shape parameters (Definition 4) were recorded (Table 5.20). In the context of cancellous bone in rats, including directional information in the features did not increase the discriminatory power of the classification process. However, the lengths of the bone were measured in the directions of the major, middle, and minor axes of the ellipsoid. Hence, the length measurements alone contained some directional information in the sense that the important directions were learned directly from the data, as opposed to making assumptions about the direction of the main axis of the tibia.

Table 5.20: Best classification accuracy as a function of  $r$  ( $\mu\text{m}$ ) for cancellous bone data. Columns two to five correspond to the accuracy obtained using each of the four methods for recording the shape primitives (Definitions 1–4, respectively). For each method, 90 CSPs were computed and 30 trials were conducted. The highest accuracy achieved using each method is shown in bold.

$r$	Oriented ellipsoid parameters	Oriented object parameters	Ellipsoid shape parameters	Object shape parameters
17.4	$\mu = 0.593,$ $\sigma = 0.021$	$\mu = \mathbf{0.676},$ $\sigma = \mathbf{0.016}$	$\mu = 0.633,$ $\sigma = 0.025$	$\mu = 0.670,$ $\sigma = 0.021$
43.5	$\mu = 0.640,$ $\sigma = 0.019$	$\mu = 0.676,$ $\sigma = 0.020$	$\mu = 0.670,$ $\sigma = 0.014$	$\mu = 0.666,$ $\sigma = 0.017$
131	$\mu = \mathbf{0.663},$ $\sigma = \mathbf{0.026}$	$\mu = 0.664,$ $\sigma = 0.021$	$\mu = \mathbf{0.695},$ $\sigma = \mathbf{0.016}$	$\mu = \mathbf{0.745},$ $\sigma = \mathbf{0.024}$
261	$\mu = 0.620,$ $\sigma = 0.024$	$\mu = 0.649,$ $\sigma = 0.029$	$\mu = 0.670,$ $\sigma = 0.024$	$\mu = 0.713,$ $\sigma = 0.020$
436	$\mu = 0.613,$ $\sigma = 0.023$	$\mu = 0.625,$ $\sigma = 0.023$	$\mu = 0.657,$ $\sigma = 0.019$	$\mu = 0.624,$ $\sigma = 0.016$

### 5.2.3 Comparison with previous work

Martin and Bottema classified cancellous bone in rats into the same three experimental groups and three distances from the growth plate as described in this study using a classification process based on oriented thickness measures [91]. However, oriented thickness measures may have the shortcoming that the subsampled points  $p$  are not guaranteed to fall close to the centre of the bone. Recording the shape primitives at a point  $p$  near the edge of the bone could lead to the local shape being inadequately captured [45, 91]. In this part of the study, the intention was to replace each of the sampled points by the centroid of the best fitting ellipsoid in order to relocate the sample away from the edge of the bone. However, a preliminary study showed that a large proportion of centroids typically fell outside of the bone [45]. This result was supported here (Table 5.19), suggesting that the expected outcome of fitting ellipsoids to the data was not achieved.

Nonetheless, the preliminary study gave some promising classification scores on the cancellous bone data, despite being limited in scope [45]. For this reason, in this study, the shape primitives were defined in a more robust way, and the potential of using ellipsoid-based features to classify binary objects based on their shape was more thoroughly investigated.

In the current study, the definitions of the shape primitives are new. In the earlier study, each ellipsoid was represented by the shape primitive

$$\mathbf{v}_p = (u_{1,1}, u_{1,2}, u_{1,3}, u_{2,1}, u_{2,2}, \ell_1, \ell_2, \ell_3),$$

where the lengths  $\ell_i$ ,  $i = 1, 2, 3$ , were measured in terms of voxel units [45]. The elements of those shape primitives were not necessarily of equal magnitude (all values of the elements of the unit direction vectors,  $u_{1,1}, u_{1,2}, u_{1,3}, u_{2,1}, u_{2,2}$ , were in the interval  $[-1, 1]$ , whereas  $\ell_i$ ,  $i = 1, 2, 3$ , could vary by up to two orders of magnitude). Here, all elements of the shape primitives were of comparable magnitude. In addition, all length measurements were made in  $\mu\text{m}$  rather than voxel units, meaning that the resolution of the data was accounted for and length measurements were standardised in some sense.

The best classification scheme for cancellous bone in rats proposed in this study ( $\mu = 0.745$ ,  $\sigma = 0.024$ , 30 trials) significantly outperformed the classification schemes using shape features alone ( $\mu = 0.607$ ,  $\sigma = 0.018$ , 10 trials,  $p < 0.001$ ), and standard attributes alone (classification accuracy of 0.467,  $p < 0.001$ , where the  $p$ -value is the value of the cumulative distribution function of the normal distribution with mean 0.745 and standard deviation 0.024 at  $x = 0.467$ ) in Martin and Bottema [91], as well as the classification scheme proposed using ellipsoids ( $\mu = 0.699$ ,  $\sigma = 0.014$ , 30 trials,  $p < 0.001$ ) in the preliminary study [45]. If features derived from best-fitting ellipsoids were combined with standard attributes [91], the classification accuracy may be improved further.

#### 5.2.4 Summary

For both data sets, the highest classification accuracy was achieved using either oriented object parameters or object shape parameters, suggesting that interrogating the data directly is advantageous. For the yeast data, recording oriented object parameters resulted in the highest classification accuracy, suggesting that information about the thickness and orientation of the colony is important towards the classification. For the cancellous bone data, the highest classification score was achieved using object shape parameters (Definition 4). Although directional information was not specifically included in these shape primitives, the process of fitting ellipsoids has the advantage that the directions in which length is measured do not rely on laboratory coordinates or on assumptions being made about the main axis of the tibia. The important directions in which to make length measurements are learned directly from the data, and the results suggest that these choices are sensible.





## Chapter 6

# Discussion

Here, the study is discussed as a whole and conclusions are drawn. The intention is to compare the two methods for extracting features from the data, namely via oriented thickness measures and best-fitting ellipsoids. To that end, the two methods are first compared in the context of the yeast colonies (Section 6.1), cancellous bone (Section 6.2) and marbling in beef (Section 6.3). These comparisons are then drawn upon to discuss the study overall, including the advantages, disadvantages, and relevance of CSPs (Section 6.4). The significant contributions of this thesis are summarised in Section 6.5.

As in Chapter 5, all  $p$ -values quoted here correspond to the two-tailed  $t$ -test.

### 6.1 Discussion of the yeast colonies

First, the images of the *S. cerevisiae* AWRI 796 strain grown in 50  $\mu\text{m}$  ammonium sulphate nutrient (A7-50) were classified into eight groups defined by the observation time after initiation of growth. The best overall classification accuracy obtained using oriented thickness measures ( $\mu = 0.883$ ,  $\sigma = 0.021$ ,  $n = 30$ ,  $D = 12$  with spatial indices included) significantly outperformed the best overall classification accuracy using features computed from best-fitting ellipsoids ( $\mu = 0.843$ ,  $\sigma = 0.021$ ,  $n = 30$ ,  $r = 15.2 \mu\text{m}$  with oriented object parame-

ters). Here,  $p < 0.001$ . However, in order to compare the effectiveness of oriented thickness measures and features derived from best-fitting ellipsoids more fairly, the classification accuracies for the two methods should be compared when spatial indices were not included. Again, the best classification accuracy obtained using oriented thickness measures ( $\mu = 0.864$ ,  $0.033$ ,  $n = 30$ ,  $D = 4$  with no spatial indices included) outperformed the best classification accuracy obtained using features computed from the best-fitting ellipsoids above ( $p = 0.005$ ).

In general, extracting shape features using oriented thickness measures appears to give the most discriminatory information for distinguishing between A7-50 colonies by growth stage. This makes sense because, by visual observation, the most obvious change as time passes is in the length of the filaments. Thus, it is not surprising that length measurements pick out the lengths of the filaments and predict the observation time effectively.

For both methods for recording shape primitives, including some version of directional information increased the classification accuracy. In the case of the oriented thickness measures, this meant making all thickness measurements with respect to the  $x$  axis, rather than with respect to the direction of the line joining the centroid of the colony to the point  $p$ . For the features derived from best-fitting ellipsoids, this involved recording the directions of the major, middle and minor axes in addition to the thickness of the colony at  $p$  (Definition 2; oriented object parameters). This is somewhat surprising, since it appears that the filaments of the yeast colonies grow radially outward at the same rate in all directions. Since growth does not appear to be biased in any particular direction, this result is counterintuitive. However, for both methods, including more information during computation of the features is definitively advantageous.

Second, the *S. cerevisiae* colonies were classified by strain, by nutrient concentration, and by strain and nutrient concentration when no information was known about either using CSPs computed via oriented thickness measures. The results were very promising in the sense that the classification method was

equally as powerful as previous methods [119], but achieved this classification accuracy using only binary data and a much smaller number of features, all of which were learned automatically from the data. Furthermore, this classification accuracy was obtained using shape features rather than textural features, demonstrating that CSPs represent a viable and promising method for biomedical image analysis wherever only binary data is available.

Since CSPs computed using features derived from best-fitting ellipsoids were not implemented on this particular classification problem, comparing the methods used throughout this thesis (oriented thickness measures and features derived from best-fitting ellipsoids) does not make sense. For a more complete and fair comparison of the two methods, the yeast colonies should also be classified by strain, nutrient concentration, and both strain and nutrient concentration using features derived from best-fitting ellipsoids. This is left as future work.

## 6.2 Discussion of the cancellous bone

In the context of the cancellous bone data, the best classification process using features derived from best-fitting ellipsoids ( $\mu = 0.745$ ,  $0.024$ ,  $n = 30$ ,  $r = 131 \mu\text{m}$  with object shape parameters) significantly outperformed the classification process using oriented thickness measures ( $\mu = 0.647$ ,  $\sigma = 0.017$ ,  $n = 30$  with shape and density information considered separately) ( $p < 0.001$ ). In fact, for each method for recording features using best-fitting ellipsoids (Definitions 1–4), the best classification accuracy (corresponding to the best choice of the radius parameter  $r$ ) was higher than the best classification accuracy obtained using oriented thickness measures. This suggests that, in the case of the cancellous bone, features extracted using best-fitting ellipsoids result in a stronger classification process than features obtained using oriented thickness measures.

When extracting features using best-fitting ellipsoids, recording object shape

parameters (Definition 4) means recording only the thickness of the bone in three directions as specified by the directions of the major, middle, and minor axes of the corresponding ellipsoid. Recording oriented thickness measures, in this case, means recording the thickness of the bone in 13 directions as specified by the laboratory coordinate frame, which depends on the orientation at which the tibia of the rat was placed in the scanner. Thus, in the context of cancellous bone, choosing the directions in a sensible way appears to be much more important than recording more information (that is, recording shape primitive vectors of a higher dimension) without forethought about the structure of the data. The key is that a better classification accuracy is obtained using shape primitives containing less information, but with that information chosen in a more sensible manner.

Although fitting ellipsoids enables the directions at which the thickness measurements are made to be chosen more carefully, recording oriented shape parameters means directional information is not directly included in the shape primitives. Using shape primitives without directional information appears to be advantageous in this case.

### 6.3 Discussion of marbling in beef

A linear regression model was used to predict the marbling proportion (Equation 4.1) of each of the striploins. The features were the occurrences of CSPs computed by treating all of the striploins as one group. The best coefficient of determination  $R^2 = 0.909$  was obtained when a regression model was fit using three features. The best feature overall was the occurrence of CSP labelled number 10. When CSP 10 was used as a single feature in the regression model, the resulting coefficient of determination was  $R^2 = 0.793$ . Overall, the coefficients of determination obtained were high, suggesting that CSP-based features are good predictors of marbling proportion.

For CSP labelled number 10, the thicknesses in all of the five directions were

close to each other, especially when compared to the other CSPs identified as important towards predicting marbling proportion (Table 5.16 and Figure 5.6). This means that a “round” shape is a good predictor of marbling proportion.

The results suggest that shape features can be good predictors of the marbling proportion of the striploins. However, whether or not the shape characteristics change as the marbling proportion increases (or whether the marbling only becomes thicker) remains an unanswered question. More work is required, for example principal component analysis to determine the amount of variation in marbling proportion that is described by shape, or developing methods for counting the average number of branches in each striploin. This is work in progress (M. Bottema, personal communication, 19 October 2017).

In the context of marbling in beef, the idea of deriving features from best-fitting ellipsoids was not considered. This is because of the difference in resolution between and within slices, which would presumably result in elongated ellipsoids in the  $z$  direction and would effectively reduce to measuring thicknesses in the  $x$ ,  $y$ , and  $z$  directions. In this case, measuring the amount of branching or using oriented thickness measures is more intuitive. Thus, oriented thickness measures and features derived from best-fitting ellipsoids cannot be compared here.

The results for the marbling in beef illustrate the versatility of CSPs. Although the emphasis in this thesis has primarily been on classification based on CSPs, this part of the study demonstrates that CSPs may be applied and evaluated in the context of a linear regression model. The results show that CSPs can be good predictors of a specific characteristic of the data set, such as marbling proportion.

## 6.4 Discussion of the overall study

In the context of the yeast colonies, in general, CSPs computed using oriented thickness measures outperformed features derived from best-fitting ellipsoids. In

addition, the classification process gave better results when directional information was included in the shape primitives. On the other hand, in the context of the cancellous bone, features derived from best-fitting ellipsoids outperformed oriented thickness measures. Also, the classification process was stronger when no directional information was recorded in the shape primitives.

Overall, there appears to be no one best method for shape analysis based on CSPs, and instead the choices made should depend on the structure of the objects being analysed. These choices should be made sensibly and with forethought, as demonstrated by the superior classification results obtained in the cancellous bone using thickness measurements made in only three directions chosen sensibly, as opposed to 13 directions chosen by the laboratory coordinate frame (Section 6.2). In addition, this means that testing different methods on a particular data set could act as a form of interrogation in itself, and may reveal information about the shape properties of that data.

This thesis also demonstrates that CSP-based features can be used for a wide variety of different purposes, including classification, and regression for predicting a certain characteristic of the data – as demonstrated by the marbling in beef, where CSP-based features were used to predict the marbling proportion in the striploin (Section 6.3).

Importantly, throughout this study, an automated shape analysis method has been developed and its versatility has been demonstrated. Many existing shape analysis methods, such as ASMs and AAMs, require some user input despite incorporating a training step to learn shape patterns from the data. For example, in ASMs and AAMs, landmark points must be placed on each sample before shape patterns can be learned. However, when the shape patterns appearing in the object under consideration are highly irregular, as is the case with the examples considered in this thesis, the placement of such landmark points is difficult or impossible. For example, for the yeast colonies, landmark points would have to be placed at two equivalent points on two different pseudohyphae.

For the cancellous bone, landmark points would have to be placed at equivalent points on the trabeculae of two different samples. For an object exhibiting highly irregular shape patterns, the question of how to define equivalent points on two different samples is essentially impossible to answer.

The method developed in this thesis avoids the need for manual placement of landmark points, and instead allows for highly complex and irregular shape patterns to be learned automatically from the data. In addition, the user is allowed to understand how the shape patterns are captured and can visualise the resulting CSPs, as demonstrated by the results on the yeast colonies (Section 5.1.1) and marbling in beef (Section 5.1.3). Furthermore, in comparable classification studies, a very large list of pre-defined features is often considered, with the authors presumably defining such a large number of features to avoid important features being overlooked [119]. In this study, the number of features input into the classifier is relatively small, and since these features are learned automatically from the data, the possibility of overlooking important features is alleviated.

The method proposed here is effective and efficient for shape analysis of objects exhibiting highly irregular shape patterns, but may be inefficient for regular shapes. Depending on the position of the sampled point  $p$ , two identical shapes may result in two different shape primitives, and two different shapes may have identical shape primitives (Figure 6.1). At first glance, this appears to be a drawback of the method. However, if the full method is applied, including the computation of shape primitives at a large set of sampled points, clustering, and computation of histograms of CSP occurrences, then the final normalised histograms will be the same (no significant difference) for identical shapes and different for different shapes. Accordingly, the method may be used for regular shapes, but is not efficient. The method is efficient and effective specifically in the context of irregular shapes, such as those often encountered in biomedical data.

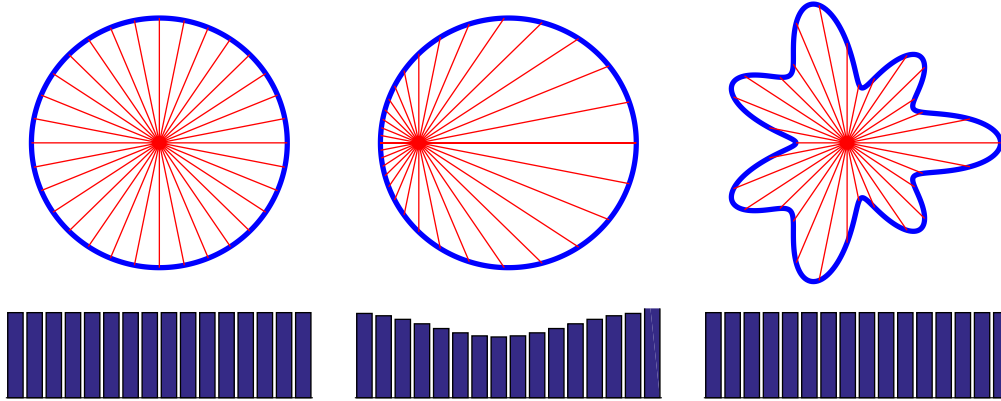


Figure 6.1: Three regular shapes and corresponding shape primitives computed at different sampled points  $p$ . The boundaries of the objects  $\Omega$  are shown in blue. A shape primitive is recorded at  $D = 16$  equally spaced angles (red). The shape primitive is represented by a histogram (dark blue). Left and middle: two identical shapes are represented by two different shape primitives. Left and right: two different objects are represented by the same shape primitive. The shape on the right is regular in the sense that it was generated as a radial function using sine and cosine function with five chosen parameters: one initial radius, two frequencies and two amplitudes.

Finally, the results demonstrate that the CSP-based method may be implemented on a range of data sets, each exhibiting different characteristics, and for different purposes (for example, for classification or for building a regression model). The idea of CSPs is promising, especially for shape analysis of objects exhibiting highly irregular shape patterns. However, since SSMs or methods where geometric parameters are defined *a priori* are not readily applicable for objects with highly irregular shape patterns, evaluating the strength of CSP-based methods by benchmarking with similar shape analysis methods is impossible in practice. Despite this, the thesis demonstrates that strong classification and regression results can be achieved using CSPs and illustrates the versatility of the method.

The classifiers chosen throughout the study were very simple. In particular, a linear classifier was chosen for classification, with exhaustive search chosen as the feature selection method. These were chosen as opposed to support vector machines, artificial neural networks, deep learning, *etc.*, because of the small number of samples in each of the data sets, and because the purpose



of the study was to evaluate the discriminatory power of the features arising from the shape primitives and to compare the different methods of computing features. A simple linear classifier gives more certainty that positive results were achieved due to the features rather than overclassification. In addition, the number of features chosen for the classifier will always be small because the most important features are learned automatically from the data and there is no need to define a large number of features. Hence, despite its computational complexity, exhaustive search remains a reasonable choice for feature selection in the context of this study.

## 6.5 Summary of the contributions of the thesis

The significant contributions of this thesis are as follows.

1. A general method has been developed for shape analysis, in particular for objects exhibiting highly irregular shape patterns.
2. Shape patterns are learned automatically from the data. This means that the most important descriptive features are learned, as opposed to using a large list of features defined *a priori*, many of which may be redundant. This also means that important features are not inadvertently overlooked.
3. There is no need to place landmark points at equivalent locations on each sample. This means the method may be used for the analysis of highly irregular shape patterns, and for this reason is advantageous compared to methods such as SSMs for characterising biomedical data. However, the method cannot feasibly be evaluated by comparison with baseline approaches such as SSMs.



## Chapter 7

# Conclusions and future work

In this chapter, the conclusions that may be drawn from the results of this study are summarised (Section 7.1). Some shortcomings of the work are briefly discussed and ideas for future work are presented (Section 7.2).

### 7.1 Conclusions

The conclusions drawn from this study may be summarised as follows.

1. Features may be computed by taking oriented thickness measures, or by finding best-fitting ellipsoids and extracting features based on these ellipsoids. The method of computing features must be chosen sensibly, and certain choices achieve better results depending on the nature of the data.
2. The proposed method has been developed specifically in the context of highly irregular binary data. The method may not be efficient for shape analysis of regular or semi-regular objects, but has been shown to achieve competitive classification results wherever only binary data was available and the shape patterns exhibited were complex (Sections 5.1.1 and 5.2.3). This is a promising method for the analysis of binary biomedical data, and pursuing the ideas developed in this thesis further is worthwhile.

## 7.2 Future work

Although this study proposes a method for shape analysis, considers different ways in which the method can be implemented, and validates the method by providing classification and regression results in a variety of contexts, there is more work that could be done. Some current shortcomings and ideas for future work are listed below.

1. In a few instances throughout the thesis, the  $K$ -means algorithm was run only once, and thus only one set of CSPs was computed. For example, this occurred when the yeast colonies were classified by strain and nutrient conditions and the features were visualised (Section 4.3.1), and during regression to predict marbling proportion (Section 4.3.3). Although visualising the CSPs that are important for classification or regression allows the user to understand the important discriminatory or predictive shape patterns, this is not a rigorous way of evaluating the method. In these cases, the CSPs could be evaluated further by repeating the process  $n$  times, recording the highly discriminatory CSPs each time, and then clustering the collection of these CSPs. This would allow for analysis of the clusters, including the within-cluster distance, and thus an understanding of how consistently similar shape patterns are being learned.
2. When shape primitives are recorded, both using oriented thickness measures and best-fitting ellipsoids, the points  $p$  at which to record the shape primitives tend to be chosen arbitrarily. Choosing the points  $p$  in a more sensible way was originally the motivation for fitting the ellipsoids to the object  $\Omega$ , with  $p$  being taken to be the centroids of the ellipsoids. However, a reasonably large proportion of the centroids were shown to fall outside  $\Omega$  (Sections 5.2.1 and 5.2.2). In addition, there is no obvious way, in general, to evaluate whether the remaining points fall far enough from the boundary  $\partial\Omega$ . For this reason, a method for choosing a more satisfactory point  $p$

is desired. Finding such points may be possible using an idea based on convex sets, and the resulting points  $p$  could be taken to compute oriented thickness measures or to fit ellipsoids.

3. The yeast colonies were not classified by strain and nutrient conditions by deriving features from best-fitting ellipsoids. This was because, when the A7-50 group was classified by growth stage, the oriented thickness measures significantly outperformed the features derived from best-fitting ellipsoids. However, exploring the idea of best-fitting ellipsoids in this context could be worthwhile. In the context of classifying the colonies by time, the oriented thickness measures could have worked exceptionally well because the lengths of the filaments were increasing with time. In this case, ellipsoids may work more effectively by capturing the more subtle differences in shape patterns. In addition, being able to identify the strain of a yeast colony at the termination of growth is of more practical importance than being able to predict its growth stage [6].
4. In the case of the marbling in beef, the idea of using CSP-based features to classify animals into the A+ and A- groups was not pursued. In previous runs, such a classification process achieved approximately 14 correct assignments out of 20 (M. Bottema, personal communication, 3 May 2017). Although this is a positive result, in practice it is not sufficiently reliable to be considered useful. In addition, the data set is not sufficiently large for the results to be validated rigorously. Classifying animals into experimental groups defined by their diet could be worth pursuing, but ideally would require larger data sets, including animals of different breeds, genotypes, ages, dietary conditions, and so on.
5. The work presented in the thesis could extend to continuous image data. Mathematically, there is no difference between the work presented here and the framework for continuous image data. In fact, the general setting as

described in Chapter 3 is presented in the continuous setting. There may be practical difficulties in measuring the thickness at various orientations depending on how the continuously defined object is presented.

# References

- [1] D. J. G. Agnew, J. E. F. Green, T. M. Brown, M. J. Simpson, and B. J. Binder. Distinguishing between mechanisms of cell aggregation using pair-correlation functions. *J. Theor. Biol.*, 352:16–23, 2014.
- [2] A. Andreopoulos and J. K. Tsotsos. Efficient and generalizable statistical models of shape and appearance for analysis of cardiac MRI. *Med. Image Anal.*, 12:335–357, 2008.
- [3] R. A. Barton. Consumer preferences and the classification and grading of beef carcasses. In A. G. Campbell, editor, *New Zealand beef production, processing and marketing*, pages 423–443. New Zealand Institute of Agricultural Science (Inc.), 1970.
- [4] B. J. Binder and M. J. Simpson. Quantifying spatial structure in experimental observations and agent-based simulations using pair-correlation functions. *Phys. Rev. E*, 88:022705, 2013.
- [5] B. J. Binder and M. J. Simpson. Spectral analysis of pair-correlation bandwidth: application to cell biology images. *R. Soc. Open Sci.*, 2:140494, 2015.
- [6] B. J. Binder, J. F. Sundstrom, J. M. Gardner, V. Jiranek, and S. G. Oliver. Quantifying two-dimensional filamentous and invasive growth spatial patterns in yeast colonies. *PLOS Comput. Biol.*, 11(2):e1004070, 2015.

- 
- [7] F. Bookstein, K. Schäfer, H. Prossinger, H. Seidler, M. Fieder, C. Stringer, G. W. Weber, J.-L. Arsuaga, D. E. Slice, F. J. Rohlf, W. Recheis, A. J. Mariam, and L. F. Marcus. Comparing frontal cranial profiles in archaic and modern *Homo* by morphometric analysis. *Anat. Rec. B New Anat.*, 257:217–224, 1999.
- [8] M. J. Bottema, Z. A. Kruk, W. S. Pitchford, and C. D. K. Bottema. Vitamin A, marbling and connected sets. In *Proceedings of Farm Animal Imaging, FAIM III*, pages 94–97, Copenhagen, Denmark, 2014.
- [9] J. E. M. Brouwers, R. M., B. v. Rietbergen, and M. L. Bouxsein. Determination of rat vertebral bone compressive fatigue properties in untreated intact rats and zoledronic-acid-treated, ovariectomized rats. *Osteoporos. Int.*, 20:1377–1384, 2009.
- [10] J. E. M. Brouwers, B. van Rietbergen, and M. L. Bouxsein. Influence of early and late zoledronic acid administration on vertebral structure and strength in ovariectomized rats. *Calcif. Tissue. Int.*, 83:186–191, 2008.
- [11] J. Camargo Neto, G. E. Meyer, D. D. Jones, and A. K. Samal. Plant species identification using Elliptic Fourier leaf shape analysis. *Comput. Electron. Agr.*, 50:121–134, 2006.
- [12] G. M. Campbell, H. R. Buie, and S. K. Boyd. Signs of irreversible architectural changes occur early in the development of experimental osteoporosis as assessed by in vivo micro-CT. *Osteoporos. Int.*, 19:1409–1419, 2008.
- [13] L. Chen, J. Noorbakhsh, R. M. Adams, J. Samaniego-Evans, G. Agollah, D. Nevozhay, J. Kuzdzal-Fick, P. Mehta, and G. Balázs. Two-dimensionality of yeast colony expansion accompanied by pattern formation. *PLOS Comput. Biol.*, 10(12):e1003979, 2014.
- [14] S. Y. Y. Chen, P. E. Lestrel, W. J. S. Kerr, and J. H. McColl. Describing



- shape changes in the human mandible using elliptical Fourier functions. *Eur. J. Orthodont.*, 22:201–216, 2000.
- [15] J. V. Coles. Household buyers choose beef. *Calif. Agr.*, 10(5):3–10, 1956.
- [16] C. Cooper. Review of UK data on the rheumatic diseases–6. Osteoporosis. *Brit. J. Rheumatol.*, 30(2):135–137, 1991.
- [17] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active Appearance Models. *IEEE T. Pattern Anal.*, 23(6):681–685, 2001.
- [18] T. F. Cootes, A. Hill, C. J. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. In B. H.H. and G. A.F., editors, *Information Processing in Medical Imaging*, Lecture Notes in Computer Science 687, pages 28–37. Springer, Berlin, Heidelberg, IPMI 1993.
- [19] T. F. Cootes and C. J. Taylor. Active Shape Models - ‘Smart Snakes’. In *Proceedings of the British Machine Vision Conference*, pages 266–275. BMVA Press, 1992.
- [20] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active Shape Models – Their training and application. *Comput. Vis. Image Und.*, 61(1):38–59, 1995.
- [21] T. F. Cootes, C. J. Taylor, A. Lanitis, D. H. Cooper, and J. Graham. Building and using flexible models incorporating grey-level information. In *Proceedings of the Fourth International Conference on Computer Vision*, pages 242–246, 1993.
- [22] O. G. Cula and K. J. Dana. 3D texture recognition using bidirectional feature histograms. *Int. J. Comput. Vision*, 59(1):33–60, 2004.
- [23] P. J. Cullen. Investigating filamentous growth and biofilm/mat formation in budding yeast. *Cold Spring Harb. Protoc.*, 3:235–238, 2015.

- 
- [24] P. J. Cullen and G. F. Sprague Jr. The regulation of filamentous growth in yeast. *Genetics*, 190:23–49, 2012.
- [25] S. R. Cummings and L. J. Melton III. Epidemiology and outcomes of osteoporotic fractures. *Lancet*, 359:1761–1767, 2002.
- [26] C. J. Darroch, R. M. R. Barnes, and J. Dawson. Circulating antibodies to *Saccharomyces cerevisiae* (bakers’/brewers’ yeast) in gastrointestinal disease. *J. Clin. Pathol.*, 52:47–53, 1999.
- [27] J. G. Daugman. Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *J. Opt. Soc. Am.*, 2(7):1160–1169, 1985.
- [28] V. R. Deepthi Loka and S. Putheti. Classification of normal, benign and malignant tissues using fuzzy texton and support vector machine in mammographic images. *IJCA*, 82(15):36–39, 2013.
- [29] A. R. Dill, M. D. Levine, and P. B. Noble. Multiple resolution skeletons. *TPAMI*, 9(4):495–504, 1987.
- [30] V. E. Dincel, M. Şengelen, V. Sepici, T. Çavuşoğlu, and B. Sepici. The association of proximal femur geometry with hip fracture risk. *Clin. Anat.*, 21:575–580, 2008.
- [31] H. G. Dolezal, G. C. Smith, J. W. Savell, and Z. L. Carpenter. Comparison of subcutaneous fat thickness, marbling and quality grade for predicting palatability of beef. *J. Food Sci.*, 47:397–401, 1982.
- [32] M. Du, Y. Huang, A. K. Das, Q. Yang, M. S. Duarte, M. V. Dodson, and M.-J. Zhu. Meat Science and Muscle Biology Symposium: Manipulating mesenchymal progenitor cell differentiation to optimize performance and carcass value of beef cattle. *J. Anim. Sci.*, 91:1419–1427, 2013.

- [33] R. Duda and P. Hart. *Pattern Classification and Scene Analysis*. John Wiley & Sons, 1973. Data mining, inference, and prediction.
- [34] A. F. Egan, D. M. Ferguson, and J. M. Thompson. Consumer sensory requirements for beef and their implications for the Australian beef industry. *Aust. J. Exp. Agr.*, 41:855–859, 2001.
- [35] A. Enache-Angoulvant and C. Hennequin. Invasive *Saccharomyces* infection: A comprehensive review. *Clin. Infect. Dis.*, 41:1559–1568, 2005.
- [36] R. H. K. Eng, R. Drehmel, S. M. Smith, and E. J. C. Goldstein. *Saccharomyces cerevisiae* infections in man. *Sabouraudia*, 22:403–407, 1984.
- [37] N. L. Fazzalari, B. L. Martin, K. J. Reynolds, T. M. Cleek, A. Badiei, and M. J. Bottema. A model for the change of cancellous bone volume and structure over time. *Math. Biosci.*, 240:132–140, 2012.
- [38] M. M. Fernández-Carrobles, G. Bueno, O. Déniz, J. Salido, M. García-Rojo, and L. González-López. Frequential versus spatial colour textons for breast TMA classification. *Comput. Med. Imag. Grap.*, 42:25–37, 2015.
- [39] R. Fisher. The use of multiple measurements in taxonomic problems. *Ann. Eugen.*, 7(2):179–188, 1936.
- [40] H. Fujikawa and M. Matsushita. Fractal growth of *Bacillus subtilis* on agar plates. *J. Phys. Soc. Jpn*, 58(11):3875–3878, 1989.
- [41] J. A. Gasser, P. Ingold, A. Venturiere, V. Shen, and J. A. Green. Long-term protective effects of zoledronic acid on cancellous and cortical bone in the ovariectomized rat. *J. Bone Miner. Res.*, 23(4):544–551, 2008.
- [42] M. H. Giaffer, A. Clark, and C. D. Holdsworth. Antibodies to *Saccharomyces cerevisiae* in patients with Crohn’s disease and their possible pathogenic importance. *Gut*, 33:1071–1075, 1992.

- [43] C. J. Gimeno, P. O. Ljungdahl, C. A. Styles, and G. R. Fink. Unipolar cell divisions in the yeast *S. cerevisiae* lead to filamentous growth: Regulation by starvation and RAS. *Cell*, 68:1077–1090, 1992.
- [44] Y. C. Gong, M. Brady, and S. Petroudi. Texture based mammogram classification and segmentation. In S. M. Astley, M. Brady, C. Rose, and R. Zwiggelaar, editors, *Breast Imaging: 8th International Workshop, IWDM*, Lecture Notes in Computer Science 4046, pages 616–625. Springer, Manchester, UK, 2006.
- [45] A. Gontar, S. Williams, and M. J. Bottema. Characterising 3D structure of cancellous bone. In *Proceedings of the International Conference on Digital Image Computing: Techniques and Applications*, pages 365–371, Gold Coast, Australia, 2016.
- [46] R. J. Grissom and J. J. Kim. *Effect Sizes for Research*. Lawrence Erlbaum Associates, USA, 2005. A Broad Practical Approach.
- [47] K. G. Grunert, L. Bredahl, and K. Brunsø. Consumer perception of meat quality and implications for product development in the meat sector—a review. *Meat Sci.*, 66:259–272, 2004.
- [48] M. Haghghat, S. Zonouz, and M. Abdel-Mottaleb. CloudID: Trustworthy cloud-based and cross-enterprise biometric identification. *Expert Syst. Appl.*, 42:7905–7916, 2015.
- [49] R. M. Haralick, K. Shanmugam, and I. Dinstein. Textural features for image classification. *IEEE Trans. Sys. Man Cyb.*, 3(6):610–621, 1973.
- [50] T. Hastie, R. Tibshirani, and J. Friedman. Datasets for The Elements of Statistical Learning. <https://web.stanford.edu/~hastie/ElemStatLearn/>, February 2009.
- [51] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical*

- learning*. Springer Series in Statistics. Springer, New York, second edition, 2009. Data mining, inference, and prediction.
- [52] G. J. Hausman, U. Basu, M. Du, M. Fernyhough-Culver, and M. V. Dodson. Intermuscular and intramuscular adipose tissues: Bad vs. good adipose tissues. *Adipocyte*, 3(4):242–255, 2014.
- [53] K. C. Hazen. New and emerging yeast pathogens. *Clin. Microbiol. Rev.*, 8(4):462–478, 1995.
- [54] M. Heath, K. Bowyer, D. Kopans, R. Moore, and W. P. Kegelmeyer. The Digital Database for Screening Mammography. In M. Yaffe, editor, *Proceedings of the Fifth International Workshop on Digital Mammography*, pages 212–218. Medical Physics Publishing, 2001.
- [55] T. Heimann and H.-P. Meinzer. Statistical shape models for 3D medical image segmentation: A review. *Med. Image Anal.*, 13:543–563, 2009.
- [56] J. S. Hicks and R. F. Wheeling. An efficient method for generating uniformly distributed points on the surface of an  $n$ -dimensional sphere. *Commun. ACM*, 2(4):17–19, 1959.
- [57] A. Hill, A. Thornham, and C. J. Taylor. Model-based interpretation of 3D medical images. In *Proceedings of the British Machine Vision Conference*, pages 339–348. BMVA Press, 1993.
- [58] D. J. Innes and J. A. Bates. Morphological variation of *Mytilus edulis* and *Mytilus trossulus* in eastern Newfoundland. *Mar. Biol.*, 133:691–699, 1999.
- [59] L. E. Jeremiah. The influence of subcutaneous fat thickness and marbling on beef: Palatability and consumer acceptability. *Food Res. Int.*, 29(5–6):513–520, 1996.

- 
- [60] L. E. Jeremiah, A. K. W. Tong, S. D. M. Jones, and C. McDonell. A survey of Canadian consumer perceptions of beef in relation to general perceptions regarding foods. *Journal of Consumer Studies and Home Economics*, 17:13–37, 1993.
- [61] J. A. Kanis. Diagnosis of osteoporosis and assessment of fracture risk. *Lancet*, 359:1929–1936, 2002.
- [62] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. Comput. Vision*, 1:321–331, 1988.
- [63] A. Katanforoush and M. Shahshahani. Distributing points on the sphere, I. *Exp. Math.*, 12(2):199–209, 2003.
- [64] J. P. Keating, D. Brophy, R. A. Officer, and E. Mullins. Otolith shape analysis of blue whiting suggests a complex stock structure at their spawning grounds in the Northeast Atlantic. *Fish. Res.*, 157:1–6, 2014.
- [65] A. Khademi and S. Krishnan. Medical image texture analysis: A case study with small bowel, retinal and mammogram images. In *Proceedings of the Canadian Conference on Electrical and Computer Engineering (CCECE)*, pages 1949–1954, 2008.
- [66] T. E. Kiehn, F. F. Edwards, and D. Armstrong. The prevalence of yeasts in clinical specimens from cancer patients. *Am. J. Clin. Pathol.*, 73(4):518–521, 1980.
- [67] M. Kleerekoper, R. Villanueva, J. Stanciu, D. Sudhaker Rao, and A. M. Parfitt. The role of three-dimensional trabecular microstructure in the pathogenesis of vertebral compression fractures. *Calcif. Tissue Int.*, 37:594–597, 1985.
- [68] R. Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1995.

- [69] Z. A. Kruk, C. D. K. Bottema, J. J. Davis, B. D. Siebert, G. S. Harper, J. Di, and W. S. Pitchford. Effects of vitamin A on growth performance and carcass quality in steers. *Livest. Sci.*, 119:12–21, 2008.
- [70] A. Lanitis, C. J. Taylor, and T. F. Cootes. A generic system for classifying variable objects using flexible template matching. In *Proceedings of the British Machine Vision Conference*, pages 329–338. BMVA Press, 1993.
- [71] G. N. Y. Lee. *Classification of Masses in Digital Mammograms and Significance of Classification Performance*. PhD thesis, School of Informatics and Engineering, Faculty of Science and Engineering, Flinders University, March 2004.
- [72] T. Leung and J. Malik. Representing and recognizing the visual appearance of materials using three-dimensional textons. *Int. J. Comput. Vision*, 43(1):29–44, 2001.
- [73] X.-Z. Li, S. Williams, and M. J. Bottema. Intensity independent texture analysis in screening mammograms. In A. Maidment, P. Bakic, and S. Gavenonis, editors, *Breast Imaging: 11th International Workshop, IWDM*, Lecture Notes in Computer Science 7361, pages 474–481. Springer, Philadelphia, USA, 2012.
- [74] X.-Z. Li, S. Williams, and M. J. Bottema. Background intensity independent texture features for assessing breast cancer risk in screening mammograms. *Pattern Recogn. Lett.*, 34:1053–1062, 2013.
- [75] X.-Z. Li, S. Williams, and M. J. Bottema. Constructing and applying higher order textons: Estimating breast cancer risk. *Pattern Recogn.*, 47:1375–1382, 2014.
- [76] X.-Z. Li, S. Williams, and M. J. Bottema. Texture and region dependent breast cancer risk assessment from screening mammograms. *Pattern Recogn. Lett.*, 36:117–124, 2014.

- 
- [77] X.-Z. Li, S. Williams, P. Downey, and M. J. Bottema. Temporal breast cancer risk assessment based on higher-order textons. In H. Fujita, T. Hara, and C. Muramatsu, editors, *Breast Imaging: 12th International Workshop, IWDM 2014*, Lecture Notes in Computer Science 8539, pages 565–572. Springer, Gifu City, Japan, 2014.
- [78] X.-Z. Li, S. Williams, G. Lee, and M. Deng. Computer-aided mammography classification of malignant mass regions and normal regions based on novel texton features. In *Proceedings of the 12th International Conference on Control, Automation, Robotics and Vision*, pages 1431–1436, 2012.
- [79] Y. Li, H. Chen, G. Kunde Rohde, C. Yao, and L. Cheng. Texton analysis for mass classification in mammograms. *Pattern Recogn. Lett.*, 52:87–93, 2015.
- [80] H. Liu, C. A. Styles, and G. R. Fink. *Saccharomyces cerevisiae* S288C has a mutation in FL08, a gene required for filamentous growth. *Genetics*, 144:967–978, 1996.
- [81] X. S. Liu, G. Bevill, T. M. Keaveny, P. Sajda, and X. E. Guo. Micromechanical analyses of vertebral trabecular bone based on individual trabeculae segmentation of plates and rods. *J. Biomech.*, 42:249–256, 2009.
- [82] X. S. Liu, P. Sajda, P. K. Saha, F. W. Wehrli, and X. E. Guo. Quantification of the roles of trabecular microarchitecture and trabecular type in determining the elastic modulus of human trabecular bone. *J. Bone Miner. Res.*, 21(10):1608–1616, 2006.
- [83] Y. Liu, D. Jin, C. Li, K. F. Janz, T. L. Burns, J. C. Torner, S. M. Levy, and P. K. Saha. A robust algorithm for thickness computation at low resolution and its application to *in vivo* trabecular bone CT imaging. *IEEE T. Bio-Med. Eng.*, 61(7):2057–2069, 2014.



- [84] X. Llado, A. Oliver, J. Freixenet, R. Marti, and J. Marti. A textural approach for mass false positive reduction in mammography. *Comput. Med. Imag. Grap.*, 33:415–422, 2009.
- [85] S. Loncaric. A survey of shape analysis techniques. *Pattern Recogn.*, 31(8):983–1001, 1998.
- [86] M. C. Lorenz, N. S. Cutler, and J. Heitman. Characterization of alcohol-induced filamentous growth in *Saccharomyces cerevisiae*. *Mol. Biol. Cell*, 11(1):183–199, 2000.
- [87] R.-S. Lu, E. Dennison, H. Denison, C. Cooper, M. Taylor, and M. J. Bottema. Texture analysis based on Gabor filters improves the estimate of bone fracture risk from DXA images. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization*, 2017.
- [88] J. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proc. of the Fifth Berkeley Symposium on Math. Stat. and Prob.*, pages 281–297, 1967.
- [89] J. Main, H. McKenzie, G. R. Yeaman, M. A. Kerr, D. Robson, C. R. Pennington, and D. Parratt. Antibody to *saccharomyces cerevisiae* (bakers’ yeast) in Crohn’s disease. *BMJ*, 297:1105–1106, 1988.
- [90] J. Malik, S. Belongie, J. Shi, and T. Leung. Textons, contours and regions: Cue integration in image segmentation. In *Proceedings of the 7th IEEE International Conference on Computer Vision*, pages 918–925, 1999.
- [91] B. L. Martin and M. J. Bottema. Textons for 3D binary data with applications to classifying cancellous bone. In *Proceedings of the International Conference on Digital Image Computing: Techniques and Applications*, pages 724–729, 2015.
- [92] M. Matsushita and H. Fujikawa. Diffusion-limited growth in bacterial colony formation. *Physica A*, 168(1):498–506, 1990.

- 
- [93] S. G. May, H. G. Dolezal, D. R. Gill, F. K. Ray, and D. S. Buchanan. Effects of days fed, carcass grade traits, and subcutaneous fat removal on postmortem muscle characteristics and beef palatability. *J. Anim. Sci.*, 70:444–453, 1992.
- [94] J. H. McCusker, K. V. Clemons, D. A. Stevens, and R. W. Davis. *Saccharomyces cerevisiae* virulence phenotype as determined with CD-1 mice is associated with the ability to grow at 42°C and form pseudohyphae. *Infect. Immun.*, 62(12):5447–5455, 1994.
- [95] T. McLellan and J. A. Endler. The relative success of some methods for measuring and describing the shape of complex objects. *Syst. Biol.*, 47(2):264–281, 1998.
- [96] L. J. Melton III, E. A. Chrischilles, C. Cooper, A. W. Lane, and B. L. Riggs. How many women have osteoporosis? *J. Bone Miner. Res.*, 7(9):1005–1010, 1992.
- [97] H.-U. Mösch and G. R. Fink. Dissection of filamentous growth by transposon mutagenesis in *Saccharomyces cerevisiae*. *Genetics*, 145:671–684, 1997.
- [98] H.-U. Mösch, R. L. Roberts, and G. R. Fink. Ras2 signals via the Cdc42/Ste20/mitogen-activated protein kinase module to induce filamentous growth in *Saccharomyces cerevisiae*. *Proc. Natl. Acad. Sci. USA*, 93:5352–5356, 1996.
- [99] D. C. Moura and M. A. Guevara Lopez. An evaluation of image descriptors combined with clinical data for breast cancer diagnosis. *Int. J. CARS*, 8:561–574, 2013.
- [100] M. E. Muller. A note on a method for generating points uniformly on  $n$ -dimensional spheres. *Commun. ACM*, 2(4):19–20, 1959.

- [101] C. Muramatsu, T. Hara, T. Endo, and H. Fujita. Breast mass classification on mammograms using radial local ternary patterns. *Comput. Biol. Med.*, 72:43–43, 2016.
- [102] A. Oka, Y. Maruo, T. Miki, T. Yamasaki, and T. Saito. Influence of vitamin A on the quality of beef from the Tajima strain of Japanese Black cattle. *Meat Sci.*, 48(1/2):159–167, 1998.
- [103] A. Oliver, X. Lladó, J. Martí, R. Martí, and J. Freixenet. False positive reduction in breast mass detection using two-dimensional PCA. In J. Martí, J. M. Benedi, A. M. Mendonca, and J. Serrat, editors, *Pattern Recognition and Image Analysis*, Proceedings of the Third Iberian Conference, IbPRIA 2007, Part II, pages 154–161. Springer, Girona, Spain, 2007.
- [104] A. Oliver, J. Martí, R. Martí, A. Bosch, and J. Freixenet. A new approach to the classification of mammographic masses and normal breast tissue. In *18th International Conference on Pattern Recognition (ICPR'06)*, volume 4, pages 707–710, Hong Kong, China, 2006.
- [105] Osteoporosis Australia. What you need to know about osteoporosis. Technical report, Australian Government, August 2014. 2nd edition.
- [106] S. P. Palecek, A. S. Parikh, and S. J. Kron. Sensing, signalling and integrating physical processes during *Saccharomyces cerevisiae* invasive and filamentous growth. *Microbiology*, 148:893–907, 2002.
- [107] Z. Palková and L. Váchová. Yeast cell differentiation: Lessons from pathogenic and non-pathogenic yeasts. *Semin. Cell Dev. Biol.*, 57:110–119, 2016.
- [108] A. M. Parfitt, M. K. Drezner, F. H. Glorieux, J. A. Kanis, H. Malluche, P. J. Meunier, S. M. Ott, and R. R. Recker. Bone histomorphometry: Standardization of nomenclature, symbols, and units. *J. Bone Miner. Res.*, 2(6):595–610, 1987.

- 
- [109] T. Pavlidis. Algorithms for shape analysis of contours and waveforms. *TPAMI*, 2(4):301–312, 1980.
- [110] E. Perilli, V. Le, B. Ma, P. Salmon, K. Reynolds, and N. L. Fazzalari. Detecting early bone changes using in vivo micro-CT in ovariectomized, zoledronic acid-treated, and sham-operated rats. *Osteoporos. Int.*, 21:1371–1382, 2010.
- [111] S. Petroudi and M. Brady. Breast density characterisation using texton distributions. In *Proceedings of the 33rd Annual International Conference of the IEEE EMBS*, pages 5004–5007, 2011.
- [112] S. Petroudi, T. Kadir, and M. Brady. Automatic classification of mammographic parenchymal patterns: a statistical approach. In *Proceedings of the 25th Annual International Conference of the IEEE EMBS*, pages 798–801, 2003.
- [113] S. Piccirillo, T. Kapros, and S. M. Honigberg. Phenotypic plasticity within yeast colonies: differential partitioning of cell fates. *Curr. Genet.*, 62:467–473, 2016.
- [114] T. B. Reynolds and G. R. Fink. Bakers’ yeast, a model for fungal biofilm formation. *Science*, 291:878–881, 2001.
- [115] B. J. Riis, M. A. Hansen, J. A. M., K. Overgaard, and C. Christiansen. Low bone mass and fast rate of bone loss at menopause: Equal risk factors for future fracture: A 15-year follow-up study. *Bone*, 19(1):9–12, 1996.
- [116] S. T. H. Rizvi, G. Cabodi, and P. Gusmao. Gabor filter based image representation for object classification. In *Proceedings of International Conference on Control, Decision and Information Technologies*, 2016.
- [117] F. J. Rohlf. Morphometrics. *Annu. Rev. Ecol. Syst.*, 21:299–316, 1990.

- [118] F. J. Rohlf, A. Loy, and M. Corti. Morphometric analysis of Old World Talpidae (Mammalia, Insectivora) using partial-warp scores. *Syst. Biol.*, 45(3):344–362, 1996.
- [119] P. Ruusuvuori, J. Lin, A. C. Scott, Z. Tan, S. Sorsa, A. Kallio, M. Nykter, O. Yli-Harja, I. Shmulevich, and A. M. Dudley. Quantitative analysis of colony morphology in yeast. *BioTechniques*, 56(1):18, 2014.
- [120] O. Ryan, R. S. Shapiro, C. F. Kurat, D. Mayhew, A. Baryshnikova, B. Chin, Z.-Y. Lin, M. J. Cox, F. Vizeacoumar, D. Cheung, et al. Global gene deletion analysis exploring yeast filamentous growth. *Science*, 337(6100):1353–1356, 2012.
- [121] E. B. Saff and A. B. J. Kuijlaars. Distributing many points on a sphere. *Math. Intell.*, 19(1):5–11, 1997.
- [122] P. K. Saha, Y. Xu, H. Duan, A. Heiner, and G. Liang. Volumetric topological analysis: A novel approach for trabecular bone classification on the continuum between plates and rods. *IEEE T. Med. Imaging*, 29(11):1821–1838, 2010.
- [123] P. Sambrook and C. Cooper. Osteoporosis. *Lancet*, 367:2010–2018, 2006.
- [124] J. W. Savell and H. R. Cross. The role of fat in the palatability of beef, pork, and lamb. In *Designing foods: Animal product options in the marketplace*, pages 345–355. National Academy Press, Washington, D.C., 1988.
- [125] C. Schmid. Constructing models for content-based image retrieval. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 39–45, 2001.
- [126] R. E. Seltzer. Consumer preferences for beef. Phoenix Arizona Survey Bulletin 267, College of Agriculture, University of Arizona, 1955.

- 
- [127] B. D. Siebert, Z. A. Kruk, J. Davis, W. S. Pitchford, G. S. Harper, and C. D. K. Bottema. Effect of low vitamin A status on fat deposition and fatty acid desaturation in beef cattle. *Lipids*, 41(4):365–370, 2006.
- [128] S. Singh and R. Dhir. Recognition of handwritten Gurmukhi numeral using Gabor filters. *IJCA*, 47(1):7–11, 2012.
- [129] M. Stauber and R. Müller. Volumetric spatial decomposition of trabecular bone into rods and plates—A new method for local bone morphometry. *Bone*, 38:475–484, 2006.
- [130] P. K. Strobe, D. A. Skelly, S. G. Kozmin, G. Mahadevan, E. A. Stone, P. M. Magwene, F. S. Dietrich, and J. H. McCusker. The 100-genomes strains, an *S. cerevisiae* resource that illuminates its natural phenotypic and genotypic variation and emergence as an opportunistic pathogen. *Genome Res.*, 25:762–774, 2015.
- [131] D. O. Tambasco Bruno, M. Z. do Nascimento, R. P. Ramos, V. R. Batista, L. A. Neves, and A. S. Martins. LBP operators on curvelet coefficients as an algorithm to describe texture in breast cancer tissues. *Expert Syst. Appl.*, 55:329–340, 2016.
- [132] D. O. Tambasco Bruno, M. Z. do Nascimento, R. P. Ramos, V. R. Batista, L. A. Neves, and A. S. Martins. LBP operators on curvelet coefficients as an algorithm to describe texture in breast cancer tissues. *Expert Syst. Appl.*, 55:329–340, 2016.
- [133] M. Varma and A. Zisserman. Texture classification: Are filter banks necessary? In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 691–698, 2003.
- [134] M. Varma and A. Zisserman. A statistical approach to texture classification from single images. *Int. J. Comput. Vision*, 62(1):61–81, 2005.

- 
- [135] B. Wang, Q. Yang, C. L. Harris, M. L. Nelson, J. R. Busboom, M.-J. Zhu, and M. Du. Nutrigenomic regulation of adipose tissue development – role of retinoic acid: A review. *Meat Sci.*, 120:100–106, 2016.
- [136] J. Wang, B. Zhou, X. S. Liu, A. J. Fields, A. Sanyal, X. Shi, M. Adams, T. M. Keaveny, and X. E. Guo. Trabecular plates and rods determine elastic modulus and yield strength of human trabecular bone. *Bone*, 72:71–80, 2015.
- [137] T. L. Wheeler, L. V. Cundiff, and R. M. Koch. Effect of marbling degree on beef palatability in *Bos taurus* and *Bos indicus* cattle. *J. Anim. Sci.*, 72:3145–3151, 1994.
- [138] J. Yi and F. Su. Histogram of Log-Gabor magnitude patterns for face recognition. In *Proceedings of IEEE International Conference on Acoustic, Speech and Signal Processing*, 2014.
- [139] J. Zupan and P. Raspor. Quantitative agar-invasion assay. *J. Microbiol. Meth.*, 73(2):100–104, 2008.
- [140] J. Zupan and P. Raspor. Invasive growth of *Saccharomyces cerevisiae* depends on environmental triggers: a quantitative model. *Yeast*, 27(4):217–228, 2010.