Medical Image Analysis 45 (2018) 55–67

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

How does the femoral cortex depend on bone shape? A methodology for the joint analysis of surface texture and shape

A.H. Gee[a,*], G.M. Treece[a], K.E.S. Poole[b]

[a] Department of Engineering, University of Cambridge, Trumpington Street, Cambridge CB2 1PZ, UK
[b] Department of Medicine, University of Cambridge, Level 5, Addenbrooke's Hospital, Box 157, Hills Road, Cambridge CB2 2QQ, UK

ARTICLE INFO

Article history:
Received 16 June 2017
Revised 24 November 2017
Accepted 12 January 2018
Available online 3 February 2018

Keywords:
Spatial normalization
Textured surfaces
Shape
Statistical parametric mapping

ABSTRACT

In humans, there is clear evidence of an association between hip fracture risk and femoral neck bone mineral density, and some evidence of an association between fracture risk and the shape of the proximal femur. Here, we investigate whether the femoral cortex plays a role in these associations: do particular morphologies predispose to weaker cortices? To answer this question, we used cortical bone mapping to measure the distribution of cortical mass surface density (CMSD, mg/cm^2) in a cohort of 125 females. Principal component analysis of the femoral surfaces identified three modes of shape variation accounting for 65% of the population variance. We then used statistical parametric mapping (SPM) to locate regions of the cortex where CMSD depends on shape, allowing for age. Our principal findings were increased CMSD with increased gracility over much of the proximal femur; and decreased CMSD at the superior femoral neck, coupled with increased CMSD at the calcar femorale, with increasing neck-shaft angle.

In obtaining these results, we studied the role of spatial normalization in SPM, identifying systematic misregistration as a major impediment to the joint analysis of CMSD and shape. Through a series of experiments on synthetic data, we evaluated a number of registration methods for spatial normalization, concluding that only those predicated on an explicit set of homologous landmarks are suitable for this kind of analysis. The emergent methodology amounts to an extension of Geometric Morphometric Image Analysis to the domain of textured surfaces, alongside a protocol for labelling homologous landmarks in clinical CT scans of the human proximal femur.

1. Introduction

Hip fractures are the most common cause of acute orthopaedic hospital admission in older people (Parker and Johansen, 2006), with their annual incidence projected to rise worldwide from 1.7 million in 1990 to 6.3 million in 2050 (Sambrook and Cooper, 2006). Bone mineral density is currently the imaging biomarker of choice for assessing an individual's fracture risk, but although it is specific (Johnell et al., 2005; Kanis et al., 2008) it lacks sensitivity (Kanis et al., 2008; Kaptoge et al., 2008; Sanders et al., 2006), missing the majority who go on to fracture. There is now growing evidence that focal, structural weaknesses may predispose a hip to fracture (Mayhew et al., 2005; Poole et al., 2010; de Bakker et al., 2009), with both trabecular and cortical bone playing a role (Holzer et al., 2009; Verhulp et al., 2008; Poole et al., 2012; Kopperdahl et al., 2014).

Cortical bone mapping (Treece et al., 2010, 2012; Treece and Gee, 2015) is an emerging technique for the quantitative analysis of the cortex using clinical CT data. It measures key properties of the cortex, for instance its thickness and mineral density, with high accuracy at several thousand locations across the proximal femur. Each femur is therefore represented as a *textured* surface, with the scalar texture representing the cortical property of interest. Statistical parametric mapping (SPM) (Friston et al., 1994) can then be used to analyse large cohorts of the textured surfaces (Tucholka et al., 2012; Worsley et al., 2009), in order to deduce, for example, how the cortical property depends on age, sex or group. Analyses of this nature have shed light on focal defects that appear to play a role in fracture risk (Treece et al., 2015; Poole et al., 2017; 2012), and the efficacy of exercise (Allison et al., 2015) and pharmaceuticals (Whitmarsh et al., 2016; Poole et al., 2015; Whitmarsh et al., 2015; Poole et al., 2011) in targeting these defects.

* Corresponding author.
E-mail addresses: ahg13@cam.ac.uk (A.H. Gee), gmt11@cam.ac.uk (G.M. Treece), kp254@medschl.cam.ac.uk (K.E.S. Poole).

https://doi.org/10.1016/j.media.2018.01.001

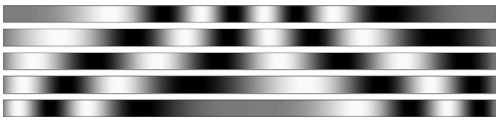

Fig. 1. In this one-dimensional example of a textured ribbon, the figure shows five individuals from a population of 201. The population variance can be explained in its entirety by a single linear shape mode (a squash or expansion around the centre, with the ends fixed) and no variance in the texture. An alternative, though less parsimonious, explanation is that there is no variance in the shape but a complex variance in the texture, requiring three linear texture modes to explain 99% of the variance.

An important step in the SPM pipeline is to *spatially normalize* the textured surfaces, a process which involves registering each surface to a standardized template. Only once the textures have been expressed on a common mesh, is it possible to fit a general linear model and explain the texture at each vertex in terms of the various regressors. In essence, surface registration involves establishing correspondences between the template's vertices and the vertices of each individual mesh. Inevitably, these correspondences are ambiguous in the barren areas between distinguished features. Different registration algorithms resolve the ambiguity in different ways, in a manner that depends on the surface's shape. Consequently, SPM analysis of the relationship between a surface's texture and its shape is problematic, since shape-dependent misregistration induces shape-dependent texture variation which is seen as statistically significant (Gee and Treece, 2014).

To better understand this phenomenon, consider the contrived example in Fig. 1, which shows some one-dimensional textured surfaces. The surfaces are free to deform in the one dimension, so they are best thought of as elastic ribbons. There is no unique way to explain the evident inter-subject variance. At one extreme, we could say that all the ribbons have precisely the same shape, with no elastic stretching or compression, meaning that all the variance is in the texture. At the other extreme, we could say that all the ribbons have precisely the same texture, meaning that all the variance is in the shape. In between these two extremes are a continuum of explanations which involve some shape variation, and also some texture variation that depends on shape. Given this ambiguity, how could we possibly address questions such as "How does the surface's texture depend on its shape?" And yet such questions are theoretically intriguing and also practically enticing, since femoral shape appears to affect fracture risk (Gregory and Aspden, 2008) and also bone mineral density (Machado et al., 2014). At least in males, the connection between shape and fracture risk is not independent of femoral neck bone mineral density (Ripamonti et al., 2014), hinting at a spatially dependent relationship between gross bone shape and the thickness and density of the cortex.

Returning to the two extreme interpretations of Fig. 1, the shape-only option leads to a compact model that can explain the population variance with a single, linear shape mode: a squash or expansion around the centre, with the ends fixed. This is how the data was generated. In contrast, principal component analysis reveals that the texture-only option requires three texture modes to account for 99% of the population variance. Information parsimony (Davies et al., 2002) is one way to resolve the ambiguity, another being enforced correspondence between distinguished landmarks (Bookstein, 1991). Either way, we need to be clear that any subsequent statistical analysis is entirely predicated on the assumptions used to establish correspondences.

In this paper, we explore these issues in the context of the cortical bone mapping pipeline. Our motivation is to understand how the cortex of the human proximal femur depends on its shape. In Section 2, we review the cortical bone mapping pipeline and describe several different registration algorithms that can be used to spatially normalize the textured surfaces. We design a synthetic data set which sheds light on the systematic misregistration introduced by the various algorithms, and introduce the real human data which we hope to analyse. In Section 3, we perform and discuss a series of experiments on the synthetic data, leading to a novel framework for controlling the correspondence ambiguity. We apply this framework to the real data, producing detailed maps showing the variation of cortical mass with shape across the human proximal femur. After discussing the biomechanical implications of our findings, we draw some conclusions in Section 4.

2. Methods

The context for this work is a pipeline of processes that enables the characterization and statistical analysis of cortical bone from clinical CT images. Although the pipeline can be applied to any bone with cortical and trabecular compartments, in this work we focus exclusively on the human proximal femur. An overview of the pipeline is presented in Fig. 2. Each stage is described in more detail in the following sections.

2.1. Cortical bone mapping

Cortical bone mapping (Treece et al., 2010; 2012; Treece and Gee, 2015) is a technique that estimates the cortical thickness (CTh, cm), cortical bone mineral density (CBMD, mg/cm^3) and cortical mass surface density (CMSD = CTh $\times$ CBMD, mg/cm^2) at thousands of locations distributed over the proximal femoral surface. The most accurate and precise estimates are for CMSD (Treece and Gee, 2015), which is one of the reasons why we focus on this property in the present work. The other reason is that it is likely to play a significant role in local fracture resistance, accounting as it does for both the amount of cortex (CTh) and the mineralization of said cortex (CBMD).

The starting point for cortical bone mapping is an approximate segmentation of the proximal femur, represented by a triangular mesh with $\sim 10^4$ vertices (Fig. 2, step 1). At each vertex, the CT data is sampled along a line passing perpendicularly through the cortex (step 2). A model (step 3, red straight lines), that accounts for the imaging blur, is fitted to the data (step 3, cyan curve) so as to minimize the differences between the blurred model (step 3, red curve) and the data. This is repeated at all vertices. The resulting distributions of CTh, CBMD and CMSD can be visualised as texture maps on the femoral surface (in step 4, red is low CMSD while blue is high CMSD). Software to perform the initial segmentation and cortical bone mapping is available for free download.[1]

2.2. Spatial registration and the parameterization of shape

For a cohort of size n, cortical bone mapping results in n texture distributions like the one in Fig. 2, step 4, each expressed on a different triangular mesh (since each individual femur has a different shape and size). Before we can compare these distributions and test how they depend on various regressors, we must first express each distribution on a common mesh. To this end, a canonical femur with 5580 vertices (step 5, red) is rotated, translated and non-rigidly deformed until it aligns with each individual femur (step 5, green). The choice of the surface registration algorithm, and the implications for the subsequent statistical analysis, are the main focus of this paper. Once aligned, the surface texture is mapped from the individual to the canonical femur and smoothed (step 6). The canonical surface mesh (which was constructed by averaging

[1] www.mi.eng.cam.ac.uk/~rwp/stradwin.