Computers in Biology and Medicine 43 (2013) 1045–1052

Contents lists available at SciVerse ScienceDirect

Computers in Biology and Medicine

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

Predicting knee cartilage loss using adaptive partitioning of cartilage thickness maps

Dan R. Jørgensen [a,b,*], Erik B. Dam [b], Martin Lillholm [b]

[a] Department of Computer Science, University of Copenhagen, Denmark
[b] Biomediq A/S, Copenhagen, Denmark

ARTICLE INFO

Article history:
Received 1 November 2012
Accepted 18 May 2013

Keywords:
Osteoarthritis
Knee MRI
Cartilage thickness
Clinical trials
Machine learning
Classification
Spatial data mining
Dimensionality reduction

ABSTRACT

This study investigates whether measures of knee cartilage thickness can predict future loss of knee cartilage. A slow and a rapid progressor group was determined using longitudinal data, and anatomically aligned cartilage thickness maps were extracted from MRI at baseline. A novel machine learning framework was then trained using these maps. Compared to measures of mean cartilage plate thickness, group separation was increased by focusing on local cartilage differences. This result is central for clinical trials where inclusion of rapid progressors may help reduce the period needed to study effects of new disease-modifying drugs for osteoarthritis.

1. Introduction

Knee osteoarthritis (OA) is a degenerative joint disease that, among others, is characterised by structural changes in the subchondral bone, osteophytes, bone marrow lesions, meniscal tears and loss of articular cartilage. OA affects a large proportion of the elder population and is a prevalent cause of disability [1,2]. Common symptoms include pain, joint stiffness, and decreased motor function, leading to an impaired quality of life. To date, there are no approved disease-modifying drugs for OA and treatment is limited to controlling symptoms and, in severe cases, joint replacement.

Currently, the gold standard for assessing OA progression in clinical trials is to measure the joint space narrowing from radiographs. However, this is only an indirect assessment because soft tissue, such as cartilage, is not directly visible on radiographs. As a consequence, joint space narrowing is insensitive to subtle cartilage changes which are thought to occur in the earlier stages of the disease [3]. Development of more sensitive biomarkers that can accurately monitor disease progression is therefore essential in the discovery and clinical testing of new disease-modifying drug for OA.

Magnetic resonance imaging (MRI) can give an accurate three-dimensional representation of all diarthrodial tissues, such as bone, cartilage and ligaments [4]. For the past decade, quantitative measures of cartilage morphology (thickness, volume, surface area) from knee MRI have received much attention as a way to identify risk factors and monitor structural changes over time [5,6]. However, results from early longitudinal studies evaluating quantitative cartilage measures in entire cartilage plates have varied greatly, showing annual cartilage loss ranging from 0.5% to −8% [7,8]. In a number of recent studies, cartilage have been examined using either detailed thickness maps [9,10] or a small number of predefined subregions [11–14]. These studies have shown that cartilage thickness changes are highly heterogeneous and that regional cartilage thickness may actually increase in the earlier stages of OA, likely due to hypertrophy or swelling [15,16]. This could help to explain some of the variations seen in earlier studies and suggests that regional based biomarkers may be better suited for monitoring disease progression than biomarkers based on measures of global cartilage change.

To date, most studies of cartilage thickness have either focused on cross sectional differences between healthy and OA knees or investigated longitudinal changes in knees at certain disease stages. Some work has also been done to analyse if changes in subregional cartilage thickness are significantly different between healthy knees and OA knees [16,17]. Yet, only a small number of studies have evaluated MRI derived cartilage morphology measures as predictive biomarkers [18–20], and these have relied on

* Corresponding author at: Biomediq A/S, Fruebjergvej 3 2100 Copenhagen, Denmark. Tel.: +45 391 79450.
E-mail addresses: drj@biomediq.com, danjoe@diku.dk (D.R. Jørgensen).

0010-4825/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.compbiomed.2013.05.012

measures of global cartilage thickness change. The ability to separate slow or non-progressors from rapid OA progressors is, however, of particular interest in clinical trials where the inclusion of fast OA progressors may reduce the observation period and/or population size needed to show the effect of new disease-modifying drugs [21].

In this study, we investigate whether measures of knee cartilage thickness can predict future loss of knee cartilage. Using longitudinal MRI data, a slow and a rapid progressor group is first determined based on 1-year loss of cartilage volume. The analysis is then performed using a novel supervised machine learning framework trained on anatomically aligned medial tibial (MT) and medial femoral (MF) cartilage thickness maps extracted from MRI images at baseline. To validate the new method, it is also tested on synthetic data and benchmarked against two well known classifiers. To the best of our knowledge, this is the first study to investigate the connection between focal cartilage thickness differences and longitudinal loss of cartilage.

1.1. Machine learning in spatial data

A common approach in current studies aimed at identifying regions affected by OA is to use region-wise statistical analysis. In this setting, differences between patient groups are detected by analysing the changes in corresponding regions across images. This is usually done using a parametric univariate statistical test, such as the t-test or the F-test, and the output is a map of p-values assembled from these tests. To compensate for the multiple testing problem, heuristic techniques, such as False Discovery Rate [22] and the Holm-Bonferroni correction [23], are often used to determine which regions are statistically significant. However, because regions are treated independently, this form of data analysis is sensitive to noise and outliers. This means that region-wise statistical analysis is mainly suited for data exploration and not for the task of identifying regions that optimise the separation between classes, such as slow and fast OA progressors.

Supervised machine learning techniques do not treat regional measurements independently, but rather perform a joint optimisation over all input measurements. In general, this makes them better suited for classification problems than methods based on region-wise statistical tests.

However, machine learning techniques do have some shortcomings when applied to spatial data. One of the main issues is that most supervised classifiers ignore the strong correlation between data in neighbouring regions. When this type of information is incorporated into a classifier, its performance on spatial data is generally improved. This has, for instance, been shown by Stoeckel and Fung [24] and Lee et al. [25] who obtained good results in the field of neuroimaging by including spatial information in the standard SVM formulation. More recently Qazi et al. [26] proposed a spatially regularised version of Fisher LDA which was shown to outperform Tikhonov regularised Fisher LDA on both synthetic data and clinical data from an OA trial.

A second challenge in classifying spatial data is that the dimensionality of the problems is often much larger than the number of available data samples. For instance, while a typical data set is composed of anywhere from tens to a few hundred images, even a small 3D-scan with 128^3 voxels constitutes a 2097152-dimensional problem. This situation, known as the "large p (dimension), small n (sample size)" paradigm [27], makes efficient pattern recognition and accurate classification difficult. To alleviate this problem, feature extraction or dimensionality reduction techniques are commonly used to map the high dimensional input data into to a lower dimensional feature space. However, the dimensionality of the transformed data may still be too high compared to the number of available data points, particularly if the spatial component of the data is to be preserved. To enhance the generalisation capability, many classifiers therefore incorporate various regularisation and sparsity techniques that constrain the space of possible solutions. Although these techniques have been successfully applied to a wide range of problems, they are not always sufficient to obtain robust models.

Here, we propose a novel framework for dealing with the large p, small n problem in spatial data. By combining an adaptive coarse-to-fine data analysis scheme with standard machine learning techniques, the framework seeks to reduce the dimensionality of the classification problem. In addition, the coarse-to-fine analysis serves as an implicit form of spatial regularisation when classifying the data.

2. The dynamic partitioning framework (DPF)

In the following, we first outline the general workings of the framework and then provide specific details for an implementation suited for two- and three dimensional spatial problems. Since the framework can be applied to a number of different problem domains, the description in this section is not specifically tied to the problem of classifying cartilage thickness maps.

A basic assumption throughout this paper is that the input data has been preprocessed in such a way that the biological objects under consideration are spatially aligned and represented in a common frame of reference. Furthermore, each object has been divided into m "atomic" regions defined on a regular grid in n-dimensional space and each region is represented by a k-dimensional feature vector calculated in that region. An atomic region may consist of a single voxel in an MRI scan but can also be composed of a larger set of connected voxels. Similarly, the vector that describes the region may be a scalar, e.g. cartilage thickness, or a more advanced regional descriptor, such as the histogram of oriented gradients [28]. Finally, each object is assumed to have a binary label indicating its class membership. As an example, fast OA progressors in this study are assigned class label 1 and slow/non-progressors are assigned class label 0.

Given such a collection of $m \times k$ dimensional objects, the main idea of the proposed framework is to treat the initial spatial domain as one region which is adaptively bi-partitioned into smaller subregions in such a way that the separation between the two classes is maximised. To alleviate the large p, small n problem, all (sub)regions are represented by a single k-dimensional feature vector. When a (sub)region spans more than one of the original m atomic regions, this vector is calculated by taking some measure of central tendency, i.e. the mean, median, or mode, of the feature vectors representing the involved atomic regions. Prior to the first partitioning of the data, the dimensionality of the classification problem is therefore reduced from $m \times k$ to $1 \times k$. After the first partition of the data, the classification problem becomes $2 \times k$ dimensional and so on. If the data contains spatially coherent regions that are different between the two classes, the number of partitions needed to "box in" these regions is likely to be smaller that the number of atomic regions. As illustrated in Fig. 1, the framework therefore effectively reduces the dimensionality of the original $m \times k$ dimensional classification problem. In addition, the adaptive partitioning strategy also serves as an implicit form of spatial regularisation because the discriminative and the non-discriminative atomic regions are grouped in separate subregions.

In order to locate the discriminative areas in the data, the framework works by combining a search algorithm with a supervised classifier in an iterative scheme. In each iteration of the process, a hyperplane is used to partition a single region in such a way that class separation is maximised. The search algorithm is