Medical Image Analysis 17 (2013) 959–973

Contents lists available at SciVerse ScienceDirect

Medical Image Analysis

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

Posterior shape models

Thomas Albrecht *, Marcel Lüthi, Thomas Gerig, Thomas Vetter

Department of Mathematics and Computer Science, Universität Basel, Bernoullistrasse 16, 4056 Basel, Switzerland

ARTICLE INFO

Article history:
Received 31 August 2012
Received in revised form 16 May 2013
Accepted 27 May 2013
Available online 15 June 2013

Keywords:
Statistical shape model
Conditional shape model
Posterior shape model
Image segmentation
Trochlear dysplasia

ABSTRACT

We present a method to compute the conditional distribution of a statistical shape model given partial data. The result is a "posterior shape model", which is again a statistical shape model of the same form as the original model. This allows its direct use in the variety of algorithms that include prior knowledge about the variability of a class of shapes with a statistical shape model. Posterior shape models then provide a statistically sound yet easy method to integrate partial data into these algorithms. Usually, shape models represent a complete organ, for instance in our experiments the femur bone, modeled by a multivariate normal distribution. But because in many application certain parts of the shape are known a priori, it is of great interest to model the posterior distribution of the whole shape given the known parts. These could be isolated landmark points or larger portions of the shape, like the healthy part of a pathological or damaged organ. However, because for most shape models the dimensionality of the data is much higher than the number of examples, the normal distribution is singular, and the conditional distribution not readily available. In this paper, we present two main contributions: First, we show how the posterior model can be efficiently computed as a statistical shape model in standard form and used in any shape model algorithm. We complement this paper with a freely available implementation of our algorithms. Second, we show that most common approaches put forth in the literature to overcome this are equivalent to probabilistic principal component analysis (PPCA), and Gaussian Process regression. To illustrate the use of posterior shape models, we apply them on two problems from medical image analysis: model-based image segmentation incorporating prior knowledge from landmarks, and the prediction of anatomically correct knee shapes for trochlear dysplasia patients, which constitutes a novel medical application. Our experiments confirm that the use of conditional shape models for image segmentation improves the overall segmentation accuracy and robustness.

1. Introduction

Statistical shape models have become an indispensable tool in medical image analysis. In essence, statistical shape models can be seen as a probability distribution (usually a normal distribution), which assigns the anatomically normal shapes of an anatomical structure a high probability, while pathological and other shapes that do not correspond to the modeled anatomical structure are assigned a low probability. Their power and versatility can be explained by the fact that they provide a quantitative answer to two fundamental questions in medicine: *(1) How does a normal instance of a given anatomical structure look like? (2) Is a specific anatomical structure normal or pathological?* Statistical shape models thus allow us to develop algorithms whose solution space is restricted to anatomically normal shapes. Such a strong prior on the solution makes the algorithm more robust, leads to easier optimization problems, and even allows us to infer a solution when only partial data is given. Consequently, applications such as implant design, surgery planning, or even medical image segmentation, for which it is clear that the result has to be a normal shape, have been shown to greatly benefit from the use of shape models (Heimann and Meinzer, 2009). In this paper we show how we can build a statistical shape model that even better restricts the solution space for the case when a part of the solution shape is already known. The new model answers the question: *Given a part of an anatomical structure, how does a normal instance of the full shape look like?* Knowledge of a part of the structure is often immediately available in practice. In surgery planning for example, a part of a shape may be missing due to a trauma or tumor, but the remaining part of the shape is known to be intact. It is thus a priori known that the solution needs to correspond to the shape of the part that is still intact. Another typical scenario is that a number of landmark points are available, which need to be matched by an algorithm.

In the following we sketch the main idea behind our method: A PCA-based statistical shape models is a generative model of the form:

$$\mathbf{s} = \mathbf{s}(\boldsymbol{\alpha}) = \boldsymbol{\mu} + \mathbf{Q}\boldsymbol{\alpha}, \tag{1}$$

* Corresponding author. Tel.: +41 61 267 0442.
E-mail address: thomas.albrecht@unibas.ch (T. Albrecht).

1361-8415/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.media.2013.05.010

where $\boldsymbol{\mu} \in \mathbb{R}^p$ is a vector that represents the mean shape and $\mathbf{Q} = (\mathbf{q}_1, \ldots, \mathbf{q}_n) \in \mathbb{R}^{p \times n}$ is a matrix of principal components $\mathbf{q}_i$, derived from training examples. By assuming that the coefficients $\boldsymbol{\alpha} \in \mathbb{R}^n$ in (1) follow a standard normal distribution $p(\boldsymbol{\alpha}) \sim \mathcal{N}(0, \mathbf{I}_n)$, a probability distribution $p(\mathbf{s}) \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{Q}\mathbf{Q}^T)$ is induced on the shape space. For a known, given part of a shape $\mathbf{s}_g \in \mathbb{R}^q$, we wish to compute a new (normal) distribution $p(\boldsymbol{\alpha}|\mathbf{s}_g) \sim \mathcal{N}(\boldsymbol{\eta}, \boldsymbol{\Lambda})$. Using this distribution as a model for the coefficients $\boldsymbol{\alpha}$ in (1) yields a new shape model, which represents shapes whose fixed part corresponds to $\mathbf{s}_g$. The new model, whose mathematical form is identical to that of the original model, can thus be used to strengthen the prior assumptions for any method that uses shape models.

In most shape models, the number of examples n is less than the dimensionality of the shape space p. This makes the covariance matrix $\mathbf{Q}\mathbf{Q}^T$ and therefore the normal distribution $\mathcal{N}(\boldsymbol{\mu}, \mathbf{Q}\mathbf{Q}^T)$ singular and the conditional distribution is more difficult to compute than it seems at first glance. We follow the most common approach from prior work, which is to regularize the covariance matrix, or the part of it corresponding to the given data, by adding a small multiple of the identity matrix $\sigma^2\,\mathbf{I}$. This can be interpreted as modeling the noise or deviation from the model in the partial data. We show the connection of this approach to probabilistic principal component analysis (PPCA) and Gaussian process regression.

We demonstrate two prototypical application of this model in medical applications. The first application targets an atlas-based segmentation of the femur bone from CT images using statistical model fitting. Here, a sparse set of landmark points is used to constrain the shape space, and thus simplify the actual fitting task. The second application targets operation planning for trochlear dysplasia patients. Trochlear dysplasia is a deformity of the knee that is treated surgically by remodeling the joint surface. Our application demonstrates how the shape model can be used to infer the normal shape of the pathological region from the intact part. This constitutes a novel application of (posterior) shape models for surgery planning.

In summary, we have the following main contributions: (1) We show how to efficiently compute the conditional distribution $p(\boldsymbol{\alpha}|\mathbf{s}_g)$ and the resulting posterior shape model, which is again a statistical shape model of the form (1). (2) We show the connection of this method to Probabilistic PCA (Tipping and Bishop, 1999) and Gaussian Process regression Rasmussen and Williams (2006). (3) We provide novel applications of our method to two problems in medical image analysis. (4) We provide a C++ implementation, as an integrated part of the freely available *statismo* library (Lüthi et al., 2012).[1]

1.1. Related work

Since their invention, statistical shape models have been used to infer the full shape from partial or "sparse" data. Often, only the maximum a posterior solution (MAP), i.e. the single most probable shape given the partial data is sought. Of the many papers computing the MAP, we only mention Blanz and Vetter (2002), as it is closest to this work. It uses a regularization term of the form $\sigma^2\mathbf{I}$ to compute a conditional distribution, but only computes the MAP and not the full posterior.

We are interested in computing this posterior model. In previous work (Albrecht et al., 2008), we derived a statistical model matching the given data using a heuristic method. In Lüthi et al. (2009), a similar model was more rigorously derived as the conditional probability of a PPCA formulation (Tipping and Bishop, 1999) of the statistical model given the partial data. The derivation of the conditional models we present here is similar, but it simplifies the formulation by separating the modeling of the partial data and the concept of PPCA models.

Other research groups have also investigated partially determined shape models. In Liu et al. (2004), canonical correlation analysis (CCA) is used to predict an unknown or diseased part of a shape from the healthy part. In Blanc et al. (2009) the given data is not a part of the shape, but given in the form of "surrogate variables" such as weight, sex, or age of a patient. In Blanc et al. (2012), this model is extended to also include partial shape data. In Blanc and Szekely (2012), the confidence of the reconstruction is evaluated, with a focus on including also the uncertainty involved in estimating correspondence between the given data and the model. These last two papers mention conditional shape models in the form we consider here in passing, but do not discuss the technical details or compute the actual shape model of the posterior distribution.

De Bruijne et al. (2007) compute a conditional shape model of a human vertebra given its neighboring vertebrae. They compute the conditional distribution with a regularization term of the form $\sigma^2\mathbf{I}$ and use the posterior shape model to classify fractures of the vertebrae. This posterior model seems very similar to our approach, but no details of its computation, especially for datasets larger than 2D vertebra shapes are given. In Baka et al. (2010) and Tomoshige et al. (2012) the simple regularization term $\sigma^2\,\mathbf{I}$ is replaced with a more general matrix reflecting the uncertainty for each given value individually. No explicit form of the posterior shape model is given in these papers. Their idea of replacing the regularization term can be employed in our approach, if individual uncertainty estimates for the given values are available. For our experiments, however, we use the standard regularization term.

Metz et al. (2010) use a combined model of shape and motion to infer cardiac motion from given shapes. They do not use a regularization term but compute the conditional distribution "after applying PCA", which amounts to simply projecting the given data onto the span of the example data and ignoring how far it actually is from this span. No posterior model is computed. Petersen et al. (2011) aim at computing the conditional distribution of a combined model of shape and rigid alignment, given partial data like landmark points. By including the rigid alignment, their conditional model becomes a non-linear manifold. This is then again linearized using a Laplace approximation (see Bishop (2006) for instance), in order to draw samples from the distribution. While this method has the advantage of incorporating the alignment into the model, no analytic expression of the model and no explicit posterior shape models are given.

To sum up, while all of these papers introduce some form "conditional model", the detailed derivation, explicit and efficient computation of the posterior shape model in the form of a standard shape model, are novel.

The viewpoint of interpreting a shape model as a Gaussian processes has been put forward by Joshi et al. (1997). A very comprehensive overview of their group's approach to shape modeling can be found in Grenander and Miller (1998). The use of Gaussian Process Regression for incorporating additional prior information, or the computation of conditional shape modes has to the best of our knowledge not been discussed, neither the connection to PPCA.

Regarding the applications we present in this paper, the surgical treatment of trochlear dysplasia is presented in Verdonk et al. (2005). In Pfirrmann et al. (2000) a statistical study of trochlear dysplasia is performed based on manual measurements of a few selected geometric criteria. The use of statistical shape model in this area is novel.

Statistical shape models have been used in the context of image segmentation since their invention, see Heimann and Meinzer (2009) for a recent and extensive review. In the terminology of

[1] available at: http://www.statismo.org.