Biomedical Signal Processing and Control 13 (2014) 345–356

Contents lists available at ScienceDirect

Biomedical Signal Processing and Control

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

An implementation of independent component analysis for 3D statistical shape analysis

Jia Wu [a], Katharine G. Brigham [b], Marc A. Simon [c,d], John C. Brigham [d,e,*]

[a] Department of Radiology, University of Pennsylvania, Philadelphia, PA 19104, USA
[b] Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA
[c] Cardiovascular Institute, School of Medicine, University of Pittsburgh, Pittsburgh, PA 15261, USA
[d] Department of Bioengineering, University of Pittsburgh, Pittsburgh, PA 15219, USA
[e] Department of Civil and Environmental Engineering, University of Pittsburgh, Pittsburgh, PA 15261, USA

A R T I C L E I N F O

Article history:
Received 25 March 2014
Received in revised form 27 May 2014
Accepted 2 June 2014
Available online 9 July 2014

Keywords:
Right ventricle
Statistical shape analysis
Function assessment
Computer-aided diagnosis

A B S T R A C T

An implementation of the independent component analysis (ICA) technique for three-dimensional (3D) statistical shape analysis is presented. The capabilities of the ICA approach to uncover inherent shape features are first demonstrated through analysis of sets of artificially generated surfaces, and the nature of these features is compared to a more traditional proper orthogonal decomposition (POD) technique. For the surfaces generated, the ICA approach is shown to consistently extract surface features that closely resembled the original basis surfaces used to generate the artificial dataset, while the POD approach produces features that clearly mix the original basis. The details of an implementation of the ICA approach within a statistical shape analysis framework are then presented. Results are shown for the ICA decomposition of a collection of clinically obtained human right ventricle endocardial surfaces (RVES) segmented from cardiac computed tomography imaging, and these results are again compared with an analogous statistical shape analysis framework utilizing POD in lieu of ICA. The ICA approach is shown to produce shape features for the RVES that capture more localized variations in the shape across the set compared to the POD approach, and overall, the ICA approach produces features that represent the RVES variation throughout the set in a considerably different manner than the more traditional POD approach, providing a potentially useful alternate to statistically analyze such a set of shapes.

1. Introduction

Statistical shape analysis has been shown to provide a powerful means to efficiently represent a large variety of shapes for various applications, and especially for applications in medical image analysis. In particular, statistical shape analysis has been shown to be capable of providing shape-based features that can be used to build effective metrics for classification and diagnostic purposes in medicine [1–6]. In general, the key components of statistical shape analysis frameworks for classification purposes involve building a correspondence between the given set of shapes (as could typically be obtained by medical image segmentation from a population), statistically decomposing the shape set into fundamental shape components, and then building features from the shape components that are suitable to cluster the shape set into various groupings and/or build a classifier associated with the application (e.g., pathological state) of interest. There have been a wide variety of techniques developed and employed within statistical shape analysis frameworks depending on the specific features, restrictions, and/or objectives of the particular applications. Furthermore, the statistical shape analysis work to-date has been largely focused on developing the mathematical representation of the shapes in a given set along with the preprocessing methods necessary to build a correspondence between these shapes, including aspects of topological mapping, shape alignment/registration, and parameterization, while much less consideration has been given to the method of statistically decomposing the shape set once this correspondence is set. The vast majority of statistical shape analysis work thus far has used some form of principal component analysis (PCA) (interchangeably referred to as proper orthogonal decomposition (POD) or by other names depending upon the specific formulation and/or application) to decompose the shape sets into fundamental components (i.e., modes or basis functions). PCA

* Corresponding author. Tel.: +1 4126249047; fax: +1 4126240135.
E-mail addresses: jia.wu@uphs.upenn.edu (J. Wu), brigham@pitt.edu (J.C. Brigham).

http://dx.doi.org/10.1016/j.bspc.2014.06.003
1746-8094/

can be viewed as providing the orthogonal basis of the specific order that is optimal in an average sense for representing the given dataset. PCA has been shown to be useful in several of the examples referenced above for extracting significant shape features, and yet, PCA provides only one perspective to view the components of the shape set of interest and is subject to the constraints of its formulation (e.g., orthogonality and average L_2 optimality), which may or may not be optimally suitable for the given application.

An alternative statistical decomposition technique known as independent component analysis (ICA) was established in [7] that utilizes a considerably different approach compared with PCA. Generally, ICA seeks to uncover the inherent patterns in a given signal dataset by identifying the fundamental components that can represent the dataset in a linear combination and are maximally statistically independent from one another [8–11]. The derivation of ICA has been approached with several different concepts and tools, such as information maximization [12], maximum likelihood estimation [13], and utilization of artificial neural networks [14]. Furthermore, ICA has been applied to extract independent features in a diverse range of research fields, including image processing [15,16], electroencephalogram (EEG) signal analysis [17,18], and audio signal processing [19], among others. A large portion of the prior work has been focused on processing one-dimensional signals, yet, formulations have also been presented to consider multi-dimensional signals, namely the multilinear ICA, which were formulated to assess multiple "modes" of discrete signals within a dataset [20–23]. ICA has shown promising capabilities in terms of extracting substantially distinct features and utilizing these features for further classification purposes compared with PCA, thus leading to ICA being considered an alternative to PCA that may be preferable in some instances depending upon the application. ICA and PCA have been directly compared in several works, including applications in face recognition [24,25] and EEG signal processing [17,18], with ICA being preferred in someinstances and PCA in others.

ICA has seen minimal application in the area of statistical shape analysis to-date, with the current contributions including [26,27], which both sought to identify patterns related to the shape variation of cardiac structures (e.g., left and right ventricles) with two-dimensional cross-sectional analysis (i.e., only analyzing line segments from the cardiac structures). There has yet to be work considering three-dimensional (3D) statistical shape analysis with ICA. A likely contributor to the lack of work integrating ICA into 3D statistical shape analysis is the lack of the necessary ICA formulation to generally accommodate multi-dimensional and/or continuously distributed shapes (i.e., signals), with the applications thus far only considering single-dimension uniformly distributed (i.e., sampled) discretized signals.

This work presents a formulation of ICA that is generally applicable for analyzing multi-dimensional continuous signals, and that is appropriate for 3D statistical shape analysis. The formulation presented is focused on one particular and popular implementation of ICA known as FastICA [28,29], but the extensions presented could easily be similarly applied to other existing ICA methods as well. Sets of artificially generated shapes were analyzed to verify the ICA algorithm and the results are presented and compared with a standard PCA (i.e., POD) algorithm. The ICA approach was then implemented within a statistical shape analysis framework (as previously developed by the authors and reported in [30]). The shape analysis framework with ICA was applied to analyze a clinically obtained set of 3D right ventricle endocardial surfaces (RVES), and the results are again presented in comparison to the shape analysis approach with POD as the statistical decomposition strategy to show the substantial differences in the two techniques and their outcomes as related to statistical shape analysis. As such, the core contributions of this work are: (1) a uniquely formulated and universally applicable ICA derivation and algorithm that can be directly applied to 3D statistical shape analysis, and (2) the presentation and examination of an ICA-based statistical shape analysis workflow applied to analyze human RVES in comparison to the previously built more standard PCA/POD-based approach [30]. The work presented in this paper broadens the state of the art of ICA-based signal processing to analyze not only vectorized discrete signals, but also more complex continuous and/or non-uniformly sampled signals, and, in particular, provides a useful alternate path to statistical shape analysis for future disease diagnostic purposes.

Section 2 presents the details of the ICA algorithm developed. Section 3 shows the analysis of artificially generated shapes with both ICA and POD, and outlines the statistical shape analysis framework incorporating ICA along with the results of analyzing the clinically obtained set of right ventricle endocardial surfaces with both ICA and POD, which is followed by the concluding remarks in Section 4.

2. ICA for shape analysis

The following formulation assumes that a given collection of n shapes to be analyzed have already been parameterized and a correspondence has been built (see [4] or similar works for examples of how such a correspondence can be built), such that $\vec{u}_i(\vec{x})$ defines the three-dimensional (3D) cartesian coordinates of the surface of the ith shape in terms of the coordinates $\vec{x}$ in the common domain of the parameterization Ω, such that $\vec{x} \in \Omega$. The fundamental assumption of ICA is that each shape function is a linear combination of m hidden (latent) statistically independent basis functions (i.e., independent components) as

$$\vec{u}_i(\vec{x}) = \sum_{j=1}^{m} a_{ij}\vec{s}_j(\vec{x}), \quad \text{for} \quad i = 1, 2, \ldots, n, \tag{1}$$

where $\left\{\vec{s}_j(\vec{x})\right\}_{j=1}^{m}$ is the set of m independent components and a_{ij} is the modal coefficient corresponding to the ith shape and jth independent component. Assuming each shape is arranged as a column vector, the entire collection of the vector-valued functions can be defined as

$$\begin{bmatrix} \vec{u}_1^T(\vec{x}) \\ \vdots \\ \vec{u}_n^T(\vec{x}) \end{bmatrix} = \begin{bmatrix} a_{11} & \cdots & a_{1m} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nm} \end{bmatrix} \begin{bmatrix} \vec{s}_1^T(\vec{x}) \\ \vdots \\ \vec{s}_m^T(\vec{x}) \end{bmatrix}, \tag{2}$$

or

$$[U] = [A][S]. \tag{3}$$

Therefore, to determine an estimate of the set of independent components, $[S]$, it is only necessary to estimate the left pseudoinverse of the mixing matrix, $[W] \approx [A]^{-1}$, leading to

$$[\widetilde{S}] = [W][U], \tag{4}$$

where $[\widetilde{S}]$ is used to denote the set of estimated independent components.

2.1. ICA algorithm

The first challenge of any ICA algorithm is to find an effective way of quantitatively measuring statistical independence. Non-Gaussianity is commonly considered as an equivalent means to measure the independence under the Central Limit Theorem [11,9]. In brief, the Central Limit Theorem states that the distribution