Journal of Computational and Applied Mathematics 348 (2019) 120–129

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

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

Learning shape metrics with Monte Carlo optimization

Serdar Cellat, Yu Fan, Washington Mio, Giray Ökten *

Department of Mathematics, Florida State University, Tallahassee, FL 32306, USA

ARTICLE INFO

Article history:
Received 8 March 2018
Received in revised form 16 July 2018

Keywords:
Landmark models
Simulated annealing
Shape metrics

ABSTRACT

Quantifying and modeling shape variation within a population, identifying morphological contrasts across groups, and categorizing individuals or objects according to morphological similarity are central problems in numerous domains of science and applications. In this paper, we present an approach to optimal shape categorization through a new family of metrics for shapes presented as a finite collection of labeled landmark points. We develop a technique to learn metrics that optimally differentiate and categorize shapes using Monte Carlo optimization methods. We discuss the theory and the practice of the methods and apply them to the analysis of synthetic data and the classification of multiple species of fruit flies based on the shape of their wings.

1. Introduction

The problems of modeling, classifying and recognizing shapes permeate many domains of science and applications. To exemplify, the investigation of organismal phenotypic variation and its genetic underpinnings often leads to complex problems in shape analysis (cf. [1–3]). The intricate morphology of pollen grains is of great interest to paleontologists as they form the most abundant and extensive record of plant diversity (cf. [4]). In ecology, there is evidence that variation in shape of spatial vegetation patterns might signal critical changes in ecosystems (cf. [5,6]). These are just a few examples in the broad landscape of problems that involve shape quantification and analysis.

Loosely speaking, the shape of an object is its geometry modulo position, orientation and scale, although depending on the context scale might not be filtered out. The first formal mathematical treatment of shape, due to Kendall, adopts a shape representation based on labeled landmark points [7]. This is illustrated in Fig. 1 that shows a fruit fly wing represented by twelve landmark points placed at vein crossings [8]. Kendall [9] constructed a shape space equipped with a metric that quantifies morphological similarity and contrast, providing a setting for systematic statistical analysis of shape variation (cf. [10]).

In the original model, all landmark points were treated as equally important in the process of constructing a shape space and defining a shape metric. Oftentimes, however, the morphological differences that most sharply contrast two or more shape populations, such as different species of fruit flies, are concentrated in particular regions. This suggests the investigation of inhomogeneous variants that highlight regions of interest. Statistical approaches to "weighted" landmark systems have been investigated in [11,12]. In this paper, we develop a shape space formulation that yields a family of shape spaces and metrics parameterized by $n \times n$ positive definite, symmetric matrices, where n is the number of landmark points. The matrices encode weights assigned to landmarks or linear combinations thereof, thus having the effect of attributing different importance to different parts of a shape.

This family of shape spaces and metrics at hand, in practice, one now has the problem of selecting a model that is "optimally" suited to an application. In this context, an important goal is to have computationally feasible ways of choosing

* Corresponding author.
E-mail address: okten@math.fsu.edu (G. Ökten).

https://doi.org/10.1016/j.cam.2018.08.043

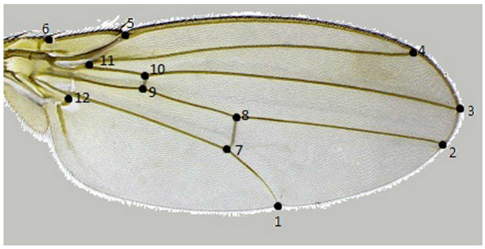

Fig. 1. Fruit fly wing with twelve landmark points.

a shape model that is particularly effective in solving a shape categorization problem such as classifying species of fruit flies. As the search space (of symmetric matrices) is typically high dimensional, this paper also investigates Monte Carlo methods to learn shape models for effective shape categorization. We formulate the question as an optimization problem in the vector space of trace-zero, symmetric matrices that we approach with simulated annealing. We illustrate the approach with synthetic data and validate the method through the classification of multiple species of fruit flies based on rather subtle differences in the shape of their wings.

We note that since the introduction of the landmark model there has been an explosion in developments in the field of shape analysis, particularly over the last two decades. Since the literature is vast, we just provide a few references that may help the interested reader navigate some of the shape literature. In many approaches, various types of morphological signatures or features are associated with objects and shape analysis is performed in feature space. These signatures typically are of a geometric, statistical or topological nature (cf. [13–17]). In other approaches, shape spaces whose elements are curves, surfaces or point clouds representing the shape of objects are employed and a variety of shape metrics have been investigated using techniques from differential and metric geometry (cf. [18–23]). Spectral geometry methods also have been used in multiple studies of shape (cf. [24,25]). These approaches have been applied to problems in a broad variety of areas. Nonetheless, development of techniques for learning a shape model that is well suited to a particular application is still incipient. This paper addresses the problem in the context of supervised learning of landmark models.

The remainder of the paper is organized as follows. In Section 2, we first review a formulation of the standard landmark model and then construct a family of shape spaces with weighted landmarks. In Section 4, we use Monte Carlo methods to solve optimization problems associated with learning shape models that optimize shape classification. A set of experiments with synthetic data are discussed in Section 5 to illustrate the method and the gains using the simulated annealing approach to the learning problem. Applications are presented in Section 6, where we apply our method to taxonomic classification of fruit flies based on the shape of their wings. We close with some additional discussion and remarks.

2. Landmark models of shape

In a landmark model, a shape in $\mathbb{R}^k$ is represented by a labeled collection of n points $p_1, \ldots, p_n \in \mathbb{R}^k$, which we encode as a $k \times n$ matrix $P = \begin{bmatrix} p_1 & \ldots & p_n \end{bmatrix}$. (Throughout the paper, we write vectors in $\mathbb{R}^k$ as column vectors and transposition of matrices is indicated by a superscript T.) The only restriction imposed is that not all landmark points be the same.

2.1. The classical model

Before introducing the general weighted model, we briefly review a formulation of the classical model that is based on the Euclidean metric on the space $\mathbb{R}^{k\times n}$ of all $k \times n$ matrices, induced by the usual inner product

$$\langle P, Q\rangle = \sum_{j=1}^{n} \langle p_j, q_j\rangle = \sum_{i=1}^{k}\sum_{j=1}^{n} p_{ij}q_{ij}. \tag{1}$$

(We abuse notation by also writing $\langle\,,\rangle$ for the usual dot product on $\mathbb{R}^k$.) The corresponding norm is the Frobenius norm $\|P\| = \langle P, P\rangle^{1/2}$.

The first step in defining shape is to obtain a representation that is insensitive to translations. This is done by restricting P to the $(kn-k)$-dimensional subspace $M(k, n)$ of centered matrices; that is, the subspace of matrices that satisfy $p_1+\cdots+p_n = 0$. We refer to the orthogonal projection $\pi : \mathbb{R}^{k\times n} \to M(k, n)$ as the centering operation. If $c = (p_1 + \cdots p_n)/n$ is the centroid of P, then centering is given by $\pi(P) = \begin{bmatrix} p_1 - c & \ldots & p_n - c \end{bmatrix}$.

For scale invariance, we simply restrict centered matrices to the unit sphere $\mathcal{P}(k, n)$ in $M(k, n)$. The corresponding scaling operation is given by $P \mapsto P/\|P\|$. Note that $P \neq 0$ because we are assuming that not all landmarks coincide. Centered matrices of unit Frobenius norm are known as *pre-shapes*. The pre-shape space $\mathcal{P}(k, n)$ is thus a (unit) sphere of dimension $kn - k - 1$.