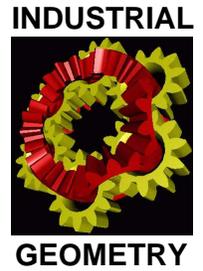


Forschungsschwerpunkt S92

Industrial Geometry

<http://www.ig.jku.at>



FSP Report No. 51

Regularized Reconstruction of Shapes with Statistical A Priori Knowledge

M. Fuchs and O. Scherzer

first revision, October 2007

(original version: May 2007)

FWF

Der Wissenschaftsfonds.



REGULARIZED RECONSTRUCTION OF SHAPES WITH STATISTICAL A PRIORI KNOWLEDGE

MATTHIAS FUCHS AND OTMAR SCHERZER

ABSTRACT. The reconstruction of geometry or, in particular, the shape of objects is a common issue in image analysis. Starting from a variational formulation of such a problem on a shape manifold we introduce a regularization technique incorporating statistical shape knowledge. The key idea is to consider a Riemannian metric on the shape manifold which reflects the statistics of a given training set. We investigate the properties of the regularization functional and illustrate our technique by applying it to region-based and edge-based segmentation of image data. In contrast to previous works our framework can be considered on arbitrary (finite-dimensional) shape manifolds and allows the use of Riemannian metrics for regularization of a wide class of variational problems in image processing.

1. INTRODUCTION

This work is concerned with the problem of detecting geometries in imaging. In a very general setting this means the identification of N -dimensional hypersurfaces in an $(N + 1)$ -dimensional space. We expect this surface to be the minimizing argument of an energy functional, which depends on the specific application we have in mind. In case of image segmentation it might be the Mumford-Shah functional [28, 29] or the “Snakes” energy [22]. In case of more general inverse problems we can think of regularization functionals like Tychonoff-regularization [11]. All these energies incorporate some kind of regularization to ensure the well-posedness of the corresponding variational problems. Common choices are the inclusion of the area of the hypersurface or its distance to some fixed argument in the energy functional. This forces these quantities to stay bounded and thus imposes regularity on them.

As an example consider the case of segmenting planar images. There the notion of the area of the hypersurface corresponds to the length of the curve defining the segmentation. If we additionally assume that the segmentation is defined by a single continuous curve of finite length, then this curve can be modeled as a differentiable map from a bounded interval into the image domain. The norms of this map and its derivative are one possibility to measure the distance between the argument of the segmentation functional and a fixed curve known beforehand.

A second basic problem is the right representation of the hypersurface in implementations. There exists no canonical approach to model such objects, but the right choice depends on the expected topology and regularity of the solution.

We propose the use of *intelligent* shape models to represent the geometry we want to detect and to regularize the underlying detection problem. These shape models share the following two characteristics:

- A shape model is associated with a finite dimensional parameter manifold. An element of this manifold corresponds to an instance of the shape model. Furthermore, we assume a function which computes the hypersurface corresponding to specific parameters. We frequently call this hypersurface the *visual appearance* or *representation* of the shape.
- A shape model can be associated with statistical data which describes how frequently individual instances of the shape model occur. In contrast to the parametrization which determines which shapes are technically feasible, the statistical part of the shape model tells which shapes are likely to actually appear in real-world applications.

We employ the two properties above to define a regularization technique which takes into account a priori information on the expected solution. As illustrated by two examples this enables us to detect geometries even if the original data is perturbed (e.g. some parts of it are missing), or allows us to significantly reduce the complexity of reconstruction methods.

Both, the statistical analysis of shapes and the use of a priori knowledge for applications in segmentation have been investigated by numerous researchers. Chen et al. [6, 5] and Gastaud et al. [18] evolve an active contour and penalize its distance from a reference curve, which is obtained as the mean shape of some training data. A similar approach in terms of level set functions instead of explicit curves is used by Cremers et al. [9]. Because of the level set approach the level set function of the mean shape and the evolving level set function have to be aligned. Further properties of the training statistics, namely the principal directions of its distribution and the variance within them, are not considered in these papers.

Leventon et al. [24] combine a maximum a posteriori (MAP) approach with geodesic active contours. They perform a Principal Component Analysis (PCA) to obtain a low dimensional representation of the distance functions of the training data and assume a Gaussian distribution of the shapes in each principal component. After each iteration of the evolution of the geodesic active contour they compute the MAP estimate of the current level set function and evolve the level set function into this direction. Thus, this approach takes into account the information about the variance of the training data. Fang and Chan [12] recently have adopted a similar approach but adaptively adjust the influence of the shape prior.

A different idea to incorporate statistical a priori information was proposed by Rousson and Paragios [31]. They assume the level set representation of the shapes to be normally distributed in each pixel. Considering aligned distance functions of the training data, the authors compute maximum likelihood estimators of the mean values and the variances in each pixel (under the condition of some regularity constraints). Then they combine the geodesic active contour evolution with a force which pulls the evolving contour towards its MAP estimator. Fritscher and Schubert [16] represent the training set and the evolving shape as 3D deformations of a reference shape.

These deformations consist of a rigid global transformation and non-rigid local deformations.

In [10, 8] Cremers et al. perform a PCA of the control points of quadratic B-spline representations of the training shapes. They give a variational formulation of the segmentation problem (based on the Mumford-Shah functional) and add the squared Mahalanobis distance obtained from the PCA of the training data to the segmentation energy. This penalizes deviations from the mean shape depending on the direction of the deviation. As for the MAP approach the distance term forces the minimizer to be “probable” in terms of the distribution of the training data. Compared to level sets, the representation of the shapes as vectors of spline control points is relatively sparse.

Tsai et al. [32] perform a PCA of the distance functions of the training shapes and consider the segmentation problem on the significant principal components only. This stabilizes the segmentation process and automatically incorporates the a priori data but does not reflect the variance of the training data within the principal components.

In all the stated level set approaches the pointwise difference between aligned level set functions forms the basis of the respective statistical approach. The mathematical relation between the linear vector space structure of level set functions and geometric shape spaces is not known. If spline curves are used, their control points always have to be labeled in the right order to compute distances between two shape instances. Moreover, the geometric interpretation of linear combinations of spline control point vectors is again not clear in a purely geometric setting.

Fletcher et al. [15, 14] propose the use of parametric representations of the medial axis transform of shapes (M-Reps). These representations are elements in a parameter manifold rather than in a vector space and the geodesic difference between two such objects corresponds much better to the visual perception of the “difference” between shapes. In these works, the authors generalize the concept of the PCA to the principal geodesic analysis (PGA, cf. Section 5) on manifolds.

The general idea of statistics of shapes on tangent spaces of Riemannian shape manifolds goes back Le and Kendall [23]. In the infinite setting the idea of replacing shapes on a manifold by the corresponding tangent vectors was investigated by Vaillant et al. and Miller et al. [33, 27].

In our work we combine the idea of adding a statistically motivated regularization term to a variational problem with the use of advanced shape models on manifolds. From this point of view it generalizes the approaches [12, 10, 8, 16] to arbitrary (finite-dimensional) shape manifolds and to a general class of variational formulations of shape reconstruction problems. In particular, it provides a framework to use M-Reps [15, 14] as a starting point for a Mahalanobis regularization of variational segmentation problems. As noted above the formulation of the regularization problem on a manifolds seems to be more natural for the reconstruction of shapes than the vector space settings. To our knowledge this is the first paper which combines the idea of M-Reps with statistical segmentation functionals. An extension of M-Reps to a shape space with a purely *geometric* metric as e.g. described

in [26, 25] is discussed in the last section of this work and also fits into our framework.

In this paper all the shapes we look at can be modeled as elements of a finite-dimensional Riemannian manifold. All the results in Sections 3–7 hold in the more general setting of such Riemannian manifolds. I.e. we do not rely on the fact that these manifolds are related to actual shapes in some sense. This fact is emphasized by the following introductory definition:

DEFINITION 1.1. By M we denote a finite dimensional and complete Riemannian C^∞ manifold.

In our opinion the main contributions of this paper are the following:

- We extend the notion of the Mahalanobis distance to M .
- This distance is used to define a class of segmentation functionals on M . We investigate the regularization properties of these functionals.
- For the concluding applications we perform a tangent space PCA of training data and incorporate the resulting statistics in the segmentation functional. We show that in case of product manifolds, the weighting of the factor manifolds does not influence the behavior of the regularization functional.

In the following we give a short overview of the contents of the individual sections. In Section 2 we review some basic properties of Riemannian manifolds and introduce some notation. The next section is devoted to regularization functionals on Riemannian manifolds. We define the notion of functionals and regularization functionals on finite-dimensional Riemannian manifolds. In Section 4 we generalize the idea of the Mahalanobis distance on vector spaces to manifolds. The Mahalanobis distance on manifolds is induced by the Mahalanobis metric which depends on given statistical data. These data are represented by a mean element and an orthogonal frame in the tangent space in the mean element. In the following section we show how we can compute these data from a training set by performing a Principal Geodesic Analysis (PGA) of the training data. We conclude this part by focusing on the Mahalanobis distance and the PGA on products of manifolds in Section 6.

In Section 7 we combine the results of the preceding part to formulate a regularization functional on shape spaces which meets the above requirements of an intelligent shape model and additionally defines a well-posed regularization problem. The two final sections are devoted to concrete examples. We first introduce a very simple version of the parametric medial axis shape model and show how it fits into our framework. Subsequently we look at two examples based on this shape model. The first one is concerned with region-based segmentation of planar images. The second one treats the edge-based detection of multiple objects in a given 2-dimensional image.

2. PRELIMINARIES AND NOTATION

In the following we will state some properties of M , i.e. of a complete C^∞ Riemannian manifold of finite dimension. For further details we refer to [19, 1]. The dimension of M is always denoted as $N \geq 1$. For $p \in M$

let T_pM be the tangent space at p . The tangent bundle of M is denoted as $TM := \dot{\bigcup}_{p \in M} T_pM$.

The manifold M is a metric space (in the topological sense), in which the distance $d_M(p, q)$, $p, q \in M$, is the infimum of the length of all piecewise differentiable curves in M which connect p and q . The topology on M induced by this metric is the same as its manifold topology. Further, if M is complete, then each bounded and closed subset of M is compact.

Let $p \in M$ and $V \in T_pM$ a tangent vector in the tangent space at p . Then there exists a unique geodesic $\gamma : [0, 1] \rightarrow M$, such that $\gamma(0) = p$ and $\dot{\gamma}(0) = V$. The *Exponential map* $\text{Exp}_p : T_pM \rightarrow M$ in p is defined by

$$\text{Exp}_p(V) := \gamma(1).$$

If for $p, q \in M$ there exists a geodesic $\gamma(0) = p$ and $\gamma(1) = q$ such that the length of γ equals $d_M(p, q)$, then we call γ a *geodesic segment connecting p and q* . On a complete connected manifold any two points can be connected by a geodesic segment. We further know that there exists an open neighborhood \mathcal{V}_p of $0 \in T_pM$ such that

- Exp_p is a diffeomorphism on \mathcal{V}_p , and
- there is a *unique* geodesic segment γ connecting any two points $q, q' \in \text{Exp}_p(\mathcal{V}_p) =: \mathcal{U}_p$ and the image of γ is contained in \mathcal{U}_p [19, Chapter I, Theorem 6.2], [1, Chapter VII, Theorem (7.2)].

We call \mathcal{U}_p a *normal neighborhood* of p and define the *Logarithmic map* as $\text{Log}_p := \text{Exp}_p^{-1} : \mathcal{U}_p \rightarrow \mathcal{V}_p$. For $q \in \mathcal{U}_p$ the logarithmic map Log_q is defined on \mathcal{U}_p and for $q' \in \mathcal{U}$ the distance from q to q' is given by $d(q, q') = |\text{Log}_q(q')|$.

To improve readability we tried to use as little indices as possible, but due to the nature of this work some still remained. We always use lower-case letters for running indices and uppercase letters as their upper bounds. Moreover, we use Einstein sum convention for the indices i and j exclusively. Whenever they (either one of them or both) appear, they denote the implicit sum over $1 \leq i, j \leq N$. Implicit summation only appears in combinations of basis vectors such as $V = v^i E_i$ and transformations of coefficients (i.e. $w_n = a_i^n v^i$). In case of coefficients v^1, \dots, v^n , the vector $v \in \mathbb{R}^N$ refers to $v = (v^1, \dots, v^n)$. E.g. the inner product of $v^i E_i$ and $w^i E_i$ can be expressed by $v^t w$, if $(E_n)_{1 \leq n \leq N}$ is an orthonormal basis.

We denote the Riemannian metric of M as $\langle \cdot, \cdot \rangle$ and as $\langle \cdot, \cdot \rangle_p$ for a specific $p \in M$. Accordingly $\| \cdot \|_p$ denotes the norm on T_pM defined by $\|V\|_p^2 = \langle V, V \rangle_p$, $V \in T_pM$. By $S(N)$ we mean the symmetric $(N \times N)$ -matrices.

3. REGULARIZATION FUNCTIONALS ON RIEMANNIAN MANIFOLDS

We start with the problem of minimizing a proper, continuous functional defined on M . I.e. we assume

$$(3.1) \quad F : M \rightarrow [0, \infty], \quad F \neq +\infty, \quad F \text{ continuous.}$$

Then the infimum of F on M is finite, but a minimizer of F may not exist. The idea is to define a *regularization functional* I_α by adding a *regularization term* R_α to F . The resulting regularization functional should satisfy the following two properties (cf. [11]):

- (1) I_α should attain a minimum on M .
- (2) Provided that a minimizer \bar{p} of the original functional F exists, the minimizers of I_α should converge to \bar{p} in a suitable sense as α tends to 0.

First we introduce an exact notion of the regularization term. In the following we denote $\text{dom}(J) := \{p \in M : J(p) \neq \infty\}$ for a map $J : M \rightarrow [0, \infty]$.

DEFINITION 3.1. Let $R_\alpha : M \rightarrow [0, \infty]$, $\alpha > 0$, be a family of continuous maps. Define

$$\mathcal{D} = \bigcap_{\alpha > 0} \text{dom}(R_\alpha).$$

We call $(R_\alpha)_{\alpha > 0}$ a *regularization on M* , if it satisfies the following conditions:

- R_α vanishes for small parameters α , i.e. for every $p \in \mathcal{D}$

$$(3.2) \quad \lim_{\alpha \rightarrow 0} R_\alpha(p) = 0.$$

- There exists $p^0 \in M$ such that for every sequence $(p_k)_{k \in \mathbb{N}}$, $p_k \in M$, satisfying

$$\lim_{k \rightarrow \infty} d_M(p^0, p_k) = \infty,$$

the regularization is unbounded, i.e. for every $\alpha > 0$

$$(3.3) \quad \lim_{k \rightarrow \infty} R_\alpha(p_k) = \infty.$$

- For every $\bar{p} \in \mathcal{D}$ there exists $C > 0$, such that for all $p \in M$ and all $\alpha > 0$

$$(3.4) \quad R_\alpha(p) \leq R_\alpha(\bar{p}) \quad \Rightarrow \quad d_M(p^0, p) \leq C.$$

We call \mathcal{D} the *domain* of the regularization $(R_\alpha)_{\alpha > 0}$.

Now assume F as above and define

$$I_\alpha = F + R_\alpha : M \rightarrow \bar{\mathbb{R}}.$$

Then we can state the following theorem:

THEOREM 3.1. *The functional I_α attains a minimum in M for every α . I.e. there exists $p' \in M$ such that*

$$I_\alpha(p') = \inf_{p \in G} I_\alpha(p).$$

Proof. This is a simple consequence of the fact that R_α forces any minimizing sequence of I_α to be bounded and that I_α is continuous. Let $(p_k)_{k \in \mathbb{N}}$ be such that $\lim_{k \rightarrow \infty} I_\alpha(p_k) = \inf_{p \in M} I_\alpha(p) =: c$. Because F and R_α are bounded from below, c is finite. Because of property (3.3) the sequence $(p_k)_{k \in \mathbb{N}}$ is bounded in M .

By extracting a converging subsequence (M is complete according to Definition 1.1) and defining $p' = \lim_{k \rightarrow \infty} p_k$ we obtain

$$c = \lim_{k \rightarrow \infty} I(p_k) = I(p').$$

□

Next we prove that I_α is indeed a regularization functional.

THEOREM 3.2. *Let F , R_α and I_α for $\alpha > 0$ be as in Theorem 3.1 and \mathcal{D} the domain of $(R_\alpha)_{\alpha>0}$. Assume that a minimizer of F exists in $M \cap \mathcal{D}$. For $\alpha > 0$ denote the set of all minimizers of I_α as $\operatorname{argmin}_{p \in M} I_\alpha(p)$. Then for every sequence $(\alpha_k)_{k \in \mathbb{N}}$ satisfying $\lim_{\alpha \rightarrow \infty} \alpha_k = 0$ and every sequence*

$$(p_k)_{k \in \mathbb{N}}, \quad p_k \in \operatorname{argmin}_{p \in M} I_{\alpha_k}(p),$$

there exists $p' \in M$ satisfying

$$F(p') = \inf_{p \in G} F(p),$$

and a subsequence $(p_\ell)_{\ell \in \mathbb{N}}$, $p_\ell \in M$ such that

$$\lim_{\ell \rightarrow \infty} p_\ell = p'.$$

Proof. Let $\alpha > 0$, $p^\alpha \in \operatorname{argmin}_{p \in M} I_\alpha(p)$ and $\bar{p} \in (\operatorname{argmin}_{p \in M} F) \cap \mathcal{D}$. From

$$F(\bar{p}) + R_\alpha(\bar{p}) = I_\alpha(\bar{p}) \geq I_\alpha(p^\alpha) = F(p^\alpha) + R_\alpha(p^\alpha) \geq F(\bar{p}) + R_\alpha(p^\alpha)$$

it follows that

$$R_\alpha(p^\alpha) \leq R_\alpha(\bar{p}).$$

In consequence, by property (3.4), the regularized solutions are bounded. Let $(\alpha_k)_{k \in \mathbb{N}}$ be a real sequence converging to 0. Define $(p_k)_{k \in \mathbb{N}}$ by choosing $p_k \in \operatorname{argmin}_{p \in M} I_{\alpha_k}(p)$ for every $k \in \mathbb{N}$. Since the sequence $(p_k)_{k \in \mathbb{N}}$ is bounded we can extract a convergent subsequence $(p_\ell)_{\ell \in \mathbb{N}}$ and define its limit as

$$p' = \lim_{\ell \rightarrow \infty} p_\ell.$$

From the definition of a regularized solution and the fact that $R_\alpha \geq 0$ it follows that

$$\begin{aligned} (3.5) \quad F(p') &= F\left(\lim_{\ell \rightarrow \infty} p_\ell\right) = \lim_{\ell \rightarrow \infty} F(p_\ell) \\ &= \lim_{\ell \rightarrow \infty} (F(p_\ell) + R_{\alpha_\ell}(p_\ell) - R_{\alpha_\ell}(p_\ell)) \\ &= \lim_{\ell \rightarrow \infty} (I_{\alpha_\ell}(p_\ell) - R_{\alpha_\ell}(p_\ell)) \\ &\leq \lim_{\ell \rightarrow \infty} \inf_{p \in M} I_{\alpha_\ell}(p). \end{aligned}$$

Moreover,

$$\inf_{p \in M} I_{\alpha_\ell}(p) = \inf_{p \in M} ((F(p) + R_{\alpha_\ell}(p)) \leq F(\bar{p}) + R_{\alpha_\ell}(\bar{p}).$$

This, together with (3.5) and the fact that $F(\bar{p}) = \inf_{p \in M} F(p)$, shows that

$$F(p') \leq \inf_{p \in M} F(p) + \lim_{\ell \rightarrow \infty} R_{\alpha_\ell}(\bar{p}).$$

Moreover, from (3.2) it follows that

$$\lim_{\ell \rightarrow \infty} R_{\alpha_\ell}(\bar{p}) = 0,$$

which proves the claim. \square

4. MAHALANOBIS DISTANCE ON M

In the following we generalize the notion of the *Mahalanobis distance* to Riemannian manifolds. We start with a review of the Mahalanobis distance on \mathbb{R}^N . Assume that we are given a vector $\mu \in \mathbb{R}^N$ and a symmetric and positive definite matrix $\Sigma \in S(N)$. Usually μ and Σ correspond to the mean and the covariance of a given probability distribution on \mathbb{R}^N . Then the Mahalanobis distance $d_\Sigma(x, y)$ of $x, y \in \mathbb{R}^N$ is defined by

$$(4.1) \quad d_\Sigma^2(x, y) = (x - y)^t \Sigma^{-1} (x - y).$$

Note that in some literature the term ‘‘Mahalanobis distance’’ refers to the square of d_Σ . We prefer (4.1) because it also defines a metric on \mathbb{R}^N . If we define the inner product $\langle \cdot, \cdot \rangle_\Sigma$ on \mathbb{R}^N by

$$\langle v, w \rangle_\Sigma = v^t \Sigma^{-1} w,$$

then we can rewrite (4.1) as

$$d_\Sigma^2(x, y) = \langle x - y, x - y \rangle_\Sigma.$$

Our goal is the extension of the above definition to Riemannian manifolds. Since the notion of a ‘‘symmetric, positive definite’’ matrix does not make sense on manifolds we replace Σ by a symmetric, positive definite bilinear form on the tangent bundle of M . In Section 5 we see that it makes sense to assume a fixed element $\mu \in M$ and an inner product on $T_\mu M$ as a starting point for the definition of a statistically motivated distance on M . For a given μ , we shall agree that E_1, \dots, E_N is a fixed orthonormal basis of $T_\mu M$ with respect to the metric $\langle \cdot, \cdot \rangle_\mu$. We further assume a symmetric, positive definite matrix $\Sigma \in S(N)$. This matrix defines an inner product $\langle \cdot, \cdot \rangle_\Sigma$ on $T_\mu M$ by

$$(4.2) \quad \langle v^i E_i, w^i E_i \rangle_\Sigma = v^t \Sigma^{-1} w$$

for two tangent vectors $v^i E_i$ and $w^i E_i$ in $T_\mu M$.

Our next step is to transport this inner product to larger parts of the manifold M . This is done by utilizing the *parallel transport* along geodesics on Riemannian manifolds as described in [1, Chapter VII]. Let

$$(4.3) \quad \mathcal{U} := \{p \in M : \text{there exists a unique geodesic connecting } \mu \text{ and } p\}.$$

Then we can state the following result:

THEOREM 4.1. *Let b be a bilinear form on $T_\mu M$. Then there exists a unique family $(b_p)_{p \in \mathcal{U}}$ of bilinear forms on the tangent bundle of \mathcal{U} which satisfies the following properties:*

- (1) *The family $(b_p)_{p \in \mathcal{U}}$ extends b , i.e. $b = b_\mu$.*
- (2) *Let $V_\mu, W_\mu \in T_\mu M$ be tangent vectors, $p \in \mathcal{U}$ and V_p, W_p the unique parallel transport of V_μ, W_μ along the (unique) geodesic connecting μ and p . Then*

$$(4.4) \quad b(V_\mu, W_\mu) = b_p(V_p, W_p).$$

Proof. Let E_1, \dots, E_N be an orthonormal basis of $T_\mu M$, $D \in \mathbb{R}^{N \times N}$ the matrix defined by

$$D_{ij} = b(E_i, E_j),$$

and $p \in M$. Using the parallel transport along the unique geodesic connecting μ and p we move the frame $(E_n)_{1 \leq n \leq N}$ into $T_p M$. Denote the transported frame as $(E_{n,p})_{1 \leq n \leq N}$ and define the bilinear form b_p by

$$(4.5) \quad b_p(v^i E_{i,p}, w^i E_{i,p}) = v^t D w.$$

Because the parallel transport of an arbitrary $v^i E_i \in T_\mu M$ to $T_p M$ can be expressed by $v^i E_{i,p}$, Equation (4.4) is satisfied.

Assume two families $(b_p)_{p \in \mathcal{U}}$ and $(b'_p)_{p \in \mathcal{U}}$ of bilinear forms on the tangent bundle of \mathcal{U} satisfying the properties in the theorem. Transporting a basis from μ to each $p \in \mathcal{U}$ immediately shows that the matrices of b_p and b'_p agree and thus the two bilinear forms are equal. \square

The above considerations motivate the following definition:

DEFINITION 4.1. Let $\mu \in M$, \mathcal{U} as in (4.3) and $\langle \cdot, \cdot \rangle_\Sigma$ as in (4.2). We call the transport of this inner product as in Theorem 4.1 the *Mahalanobis metric* $\langle \cdot, \cdot \rangle_{p,\mu,\Sigma}$ with respect to μ and Σ on $T_p M$.

The *Mahalanobis distance* $d_{M,\mu,\Sigma}$ from p to q on \mathcal{U} with respect to μ and Σ is given by

$$(4.6) \quad d_{M,\mu,\Sigma}(p, q) = \inf_{\substack{\gamma: [0,1] \rightarrow \mathcal{U} \\ \gamma(0)=p \\ \gamma(1)=q}} \int_0^1 \langle \dot{\gamma}, \dot{\gamma} \rangle_{\gamma,\mu,\Sigma}^{\frac{1}{2}} dt.$$

As in the definition of d_M we assume γ to be a piecewise differentiable curve.

Note that \mathcal{U} is not empty, because it includes at least a normal neighborhood of μ . Moreover, \mathcal{U} contains not only the points uniquely connected to μ but also the corresponding geodesics themselves. Also keep in mind that the Mahalanobis distance on manifolds depends not only on Σ but also on the mean μ . This is not the case in the vector space setting.

REMARK 4.1. If \mathcal{U} is open (this can always be achieved by replacing \mathcal{U} by its interior), then Definition 4.1 means that \mathcal{U} with the metric $\langle \cdot, \cdot \rangle_{p,\mu,\Sigma}$ is again a Riemannian manifold.

Below we rewrite $\langle \cdot, \cdot \rangle_{p,\mu,\Sigma}$ in terms of the eigenvector decomposition of Σ . This will be used to derive some properties of the Mahalanobis distance in Theorem 4.2.

REMARK 4.2. Let μ , \mathcal{U} and $\langle \cdot, \cdot \rangle_{p,\mu,\Sigma}$ be as in Definition 4.1. As in the introduction we assume an orthogonal frame $(E_n)_{1 \leq n \leq N}$ in $T_\mu M$ and let Σ^{-1} be the matrix of $\langle \cdot, \cdot \rangle_{p,\mu,\Sigma}$ with respect to this frame. Let v_1, \dots, v_N be an orthonormal basis of eigenvectors of Σ and $\sigma_1^2, \dots, \sigma_N^2$ the corresponding eigenvalues. As in the proof of Theorem 4.1 $E_{n,p}$ denotes of the parallel transport of E_n to $T_p M$, $1 \leq n \leq N$. Then for $V = v^i E_{i,p}$ and $W = w^i E_{i,p} \in T_p M$

$$\langle V, W \rangle_{p,\mu,\Sigma} = \sum_{n=1}^N \frac{(v^t v_n)(w^t v_n)}{\sigma_n^2}.$$

In particular

$$(4.7) \quad d_{M,\mu,\Sigma}(p, q) = \inf_{\substack{\gamma: [0,1] \rightarrow \mathcal{U} \\ \gamma(0)=p \\ \gamma(1)=q}} \int_0^1 \left(\sum_{n=1}^N \frac{(v(t)^t v_n)^2}{\sigma_n^2} \right)^{\frac{1}{2}} dt,$$

with

$$\dot{\gamma}(t) = v^i(t) E_{i,\gamma(t)}.$$

Before stating the next theorem we remind ourselves of the fact that the parallel transport of the orthonormal frame $(E_n)_{1 \leq n \leq N}$ is an orthonormal system in every point $p \in \mathcal{U}$ with respect to the Riemannian metric on M . In particular, the inner product of $v^i E_{i,p}$ and $w^i E_{i,p}$ in $T_p M$ is

$$(4.8) \quad \langle v^i E_{i,p}, w^i E_{i,p} \rangle_p = v^t w.$$

THEOREM 4.2 (Properties). *Assume μ , Σ and \mathcal{U} as in Definition 4.1.*

- (1) *For $p \in \mathcal{U}$*

$$d_{M,\mu,\Sigma}(p, p) = 0.$$

- (2) *$d_{M,\mu,\Sigma}$ is symmetric, i.e. for $p, q \in \mathcal{U}$*

$$d_{M,\mu,\Sigma}(p, q) = d_{M,\mu,\Sigma}(q, p).$$

- (3) *The triangle inequality holds, i.e. for $p, q, r \in \mathcal{U}$*

$$d_{M,\mu,\Sigma}(p, r) \leq d_{M,\mu,\Sigma}(p, q) + d_{M,\mu,\Sigma}(q, r).$$

- (4) *Denote $\sigma_- := \min(\sigma_1, \dots, \sigma_N)$ and $\sigma_+ := \max(\sigma_1, \dots, \sigma_N)$ and assume $p, q \in \mathcal{U}$. Then*

$$(4.9) \quad \sigma_+^{-1} d_{\mathcal{U}}(p, q) \leq d_{M,\mu,\Sigma}(p, q) \leq \sigma_-^{-1} d_{\mathcal{U}}(p, q).$$

Here $d_{\mathcal{U}}$ denotes the metric on \mathcal{U} which is induced by the Riemannian metric on M .

- (5) *Moreover,*

$$(4.10) \quad \sigma_+^{-1} d_M(\mu, p) \leq d_{M,\mu,\Sigma}(\mu, p) \leq \sigma_-^{-1} d_M(\mu, p).$$

Proof.

- (1) – (3) Follows from Definition 4.1.

- (4) Let $\sigma_- > 0$ be as in the theorem. As in Remark 4.2 we write $\dot{\gamma}(t) = v^i(t) E_i(\gamma(t))$ for $t \in [0, 1]$. Let further be $(v_n)_{1 \leq n \leq N}$ be the eigenvectors of Σ . Then $V_n(p) = v_n^i E_i(p)$, $1 \leq n \leq N$, is an orthonormal basis of $T_p M$ for $p \in \mathcal{U}$. From (4.7) it follows that for

$p, q \in \mathcal{U}$

$$\begin{aligned}
d_{M,\mu,\Sigma}(p, q) &= \inf_{\substack{\gamma:[0,1] \rightarrow \mathcal{U} \\ \gamma(0)=p \\ \gamma(1)=q}} \int_0^1 \left(\sum_{n=1}^N \frac{(v(t)^t v_n)^2}{\sigma_n^2} \right)^{\frac{1}{2}} dt \\
&\geq \sigma_+^{-1} \inf_{\substack{\gamma:[0,1] \rightarrow \mathcal{U} \\ \gamma(0)=p \\ \gamma(1)=q}} \int_0^1 \left(\sum_{n=1}^N \langle \dot{\gamma}, V_n(\gamma) \rangle_{\gamma}^2 \right)^{\frac{1}{2}} dt \\
&\geq \sigma_+^{-1} \inf_{\substack{\gamma:[0,1] \rightarrow \mathcal{U} \\ \gamma(0)=p \\ \gamma(1)=q}} \int_0^1 |\dot{\gamma}| dt = \sigma_+^{-1} d_{\mathcal{U}}(p, q)
\end{aligned}$$

In the second line we made use of the equality (4.8). Analogous for the second part.

(6) For $p \in \mathcal{U}$

$$\begin{aligned}
d_{M,\mu,\Sigma}(\mu, p) &= \inf_{\substack{\gamma:[0,1] \rightarrow \mathcal{U} \\ \gamma(0)=\mu \\ \gamma(1)=p}} \int_0^1 \left(\sum_{n=1}^N \frac{\langle \dot{\gamma}, v_n(\gamma) \rangle_{\gamma}^2}{\sigma_n^2} \right)^{\frac{1}{2}} dt \\
&\geq \sigma_+^{-1} \inf_{\substack{\gamma:[0,1] \rightarrow \mathcal{U} \\ \gamma(0)=\mu \\ \gamma(1)=p}} \int_0^1 \left(\sum_{n=1}^N \langle \dot{\gamma}, v_n(\gamma) \rangle_{\gamma}^2 \right)^{\frac{1}{2}} dt \\
&= \sigma_+^{-1} d_M(\mu, p)
\end{aligned}$$

because $d_M(\mu, p)$ is attained by the length of a geodesic which is entirely in \mathcal{U} (by the definition of \mathcal{U}). Again the second part is analogous. \square

Note that estimate (4.9) means that the Mahalanobis distance is bounded from above and below only by $d_{\mathcal{U}}$ but not by the original Riemannian distance d_M . In contrast, the Mahalanobis distance of an arbitrary point to μ is compared to d_M in (4.10). It is easy to see that (4.10) does not hold for every pair p, q on M (think e.g. of the cylinder $S^1 \times \mathbb{R} \subseteq \mathbb{R}^3$ with the induced metric).

In the following we investigate, if $d_{M,\mu,\Sigma}$ as defined in Definition 4.1 is consistent with existing definitions, i.e.

- if it corresponds to the classical Mahalanobis distance in the Euclidean setting, and
- if it corresponds to the standard Riemannian metric on M in the case $\sigma_1^2 = \dots = \sigma_N^2 = 1$.

THEOREM 4.3. *Assume that μ and Σ are as in Definition 4.1. For $\mu \in \mathbb{R}^N$ and $\Sigma \in S(N)$ the distance $d_{\mathbb{R}^N,\mu,\Sigma}(p, q)$ is defined for every pair $(p, q) \in \mathbb{R}^N \times \mathbb{R}^N$ (provided that Σ is non-singular) and*

$$d_{\mathbb{R}^N,\mu,\Sigma}(p, q)^2 = (p - q)^t \Sigma^{-1} (p - q).$$

Moreover, assume M , μ and \mathcal{U} as in Definition 4.1 and let $\Sigma = \text{Id}_N$. Then for every $p \in M$

$$d_{M,\mu,\text{Id}_N}(\mu, p) = d_M(\mu, p).$$

Proof. Obviously \mathcal{U} in Definition 4.1 equals \mathbb{R}^N . We further observe that $\dot{\gamma}(t) = \dot{\gamma}^i(t)E_i$ (i.e. the parallel transport of the standard basis gives the standard basis in every point). By definition

$$\begin{aligned} d_{\mathbb{R}^N,\mu,\Sigma}(p, q) &= \inf_{\substack{\gamma:[0,1] \rightarrow \mathcal{U} \\ \gamma(0)=p \\ \gamma(1)=q}} \int_0^1 \langle \dot{\gamma}, \dot{\gamma} \rangle_{\gamma,\mu,\Sigma}^{\frac{1}{2}} dt \\ &= \inf_{\substack{\gamma:[0,1] \rightarrow \mathcal{U} \\ \gamma(0)=p \\ \gamma(1)=q}} \int_0^1 (\dot{\gamma}^t \Sigma^{-1} \dot{\gamma})^{\frac{1}{2}} dt. \end{aligned}$$

Computing the optimality condition of this variational problem gives $\ddot{\gamma} = 0$. This again shows that the optimal γ is the straight line segment connecting p and q . The second claim is obvious by the definition of the Mahalanobis metric. \square

In the proof of Theorem 4.3 we state that the geodesics with respect to the standard metric on \mathbb{R}^N and with respect to the Mahalanobis distance coincide (in both cases they are line segments). This is not true in general. In other words, if a minimizer exists for the variational problem (4.6), it can be different from the geodesic with respect to the Riemannian metric. A simple example is the 2-dimensional sphere S^2 , $\sigma_1 > \sigma_2 > 0$ and $\Sigma = \text{diag}(\sigma_1^2, \sigma_2^2)$. The Mahalanobis metric with respect to Σ and with respect to any $\mu \in S^2$ and orthonormal frame $E_1, E_2 \in T_\mu M$ induces geodesics which are not great circles.

Note that Theorem 4.3 holds only for the distance from μ to points in \mathcal{U} but not for arbitrary pairs $p, q \in \mathcal{U}$. This is because in general \mathcal{U} does not completely cover the manifold M . If, however, p and q are such that at a sequence of curves, whose lengths converge to $d_M(p, q)$, lies in \mathcal{U} , then the original Riemannian distance and the Mahalanobis distance coincide. This is e.g. the case on the 2-dimensional sphere S^2 but not on the cylinder $S^1 \times \mathbb{R}$.

If Σ is the covariance matrix of a given probability distribution then the Mahalanobis distance $d_{M,\mu,\Sigma}$ reflects the shape of this distribution. Roughly spoken, it assigns large distances from a reference point to elements which are located in “improbable directions”, whereas the distance between two “probable elements” is small. If the Mahalanobis distance is used as a regularization, this has the beneficial effect that we penalize deviations from a reference element in an intelligent way, i.e. depending on how probable they are.

5. PRINCIPAL GEODESIC ANALYSIS

In this section we want to outline the idea of the *Principal Geodesic Analysis* (PGA) on manifolds as proposed in [14, 13]. A more sophisticated analysis of this topic can be found in [20].

The PGA is based on the idea of the Principal Component Analysis (PCA) on vector spaces. The PCA of a given set of data points in a vector space V is an orthonormal basis of V such that the first basis vector points into the direction of the largest variance of the data, the second one into the direction of the second-largest variance and so on. Thus, the PCA can also be interpreted as sequence of orthogonal linear subspaces $V_n \subseteq V$ given by the linear hull of the n -th basis vector. On manifolds the concept of linear subspaces is replaced by *geodesic submanifolds*. These are submanifolds of the manifold M such that geodesics in the submanifolds are also geodesics in M .

Let $p_1, \dots, p_S \in M$ be data points on M . We start by giving a formal definition of a PGA of the data p_1, \dots, p_S . This definition motivates the introduction of the *approximated Principal Geodesic Analysis* later in this section. The *mean* μ' of p_1, \dots, p_S is defined by

$$(5.1) \quad \mu' := \operatorname{argmin}_{p \in M} \sum_{s=1}^S d_M(p_s, p)^2.$$

Note that we do not state any results concerning existence and uniqueness of μ' , but provide them later for the approximated mean instead. Assuming that μ' is well-defined, we proceed with the definition of the principal geodesics. First, we formally define the projection on a closed subset $H \subseteq M$ as

$$\pi_H : M \rightarrow H, \quad p \mapsto \operatorname{argmin}_{h \in H} d_M(h, p)^2.$$

Now let

$$(5.2) \quad V'_1 := \operatorname{argmax}_{\substack{V \in T_\mu M \\ |V|=1}} \sum_{s=1}^S d_M(\mu, \pi_{H_1(V)}(p_s))^2.$$

Here for $V \in T_\mu M$

$$H_1(V) = \operatorname{Exp}_\mu(\langle V \rangle),$$

i.e. the exponential of the linear hull of V . In other words, $H_1(V)$ is the image of the geodesic γ defined by $\gamma(0) = \mu$ and $\dot{\gamma}(0) = V$. We proceed by recursively defining

$$(5.3) \quad V'_n := \operatorname{argmax}_{\substack{V \in T_\mu M \\ |V|=1 \\ V \perp_{\mathbb{R}} \langle V'_1, \dots, V'_{n-1} \rangle}} \sum_{s=1}^S d_M(\mu, \pi_{H_n(V)}(p_s))^2, \quad 2 \leq n \leq N,$$

where for $V \in T_\mu M$

$$H_n(V) = \operatorname{Exp}_\mu(\langle V'_1, \dots, V'_{n-1}, V \rangle),$$

i.e. the exponential of the linear hull of the vectors V'_1, \dots, V'_{n-1} and V . Again the existence and uniqueness of V'_1, \dots, V'_N is not clear in this setting but is shown for the corresponding approximations below. We call μ, V'_1, \dots, V'_N the Principal Geodesic Analysis of the data p_1, \dots, p_S .

From (5.1), (5.2) and (5.3) we now derive an approximation of the Principal Geodesic Analysis. In Section 2 we stated that $d(q, q') = |\operatorname{Log}_q(q')|$ provided that q and q' are elements of a *common* normal neighborhood as

we introduced it in Section 2. Assume that also p lies in this neighborhood. Then we can approximate $d_M(q, q') = |\text{Log}_q(q')|$ by

$$(5.4) \quad |\text{Log}_p(q) - \text{Log}_p(q')|.$$

This observation leads us to the following introduction of the approximated PGA which eventually can be reduced to a standard PCA in tangent spaces.

We now assume that p_1, \dots, p_S are contained in a common normal neighborhood. Let further $p \in M$ be an element in this normal neighborhood. Applying (5.4) to (5.1) leads to the definition of the *approximated mean* μ of p_1, \dots, p_S as

$$(5.5) \quad \mu := \underset{p' \in M}{\operatorname{argmin}} \sum_{s=1}^S (\text{Log}_p(p_s) - \text{Log}_p(p'))^2.$$

This immediately yields

$$\mu = \text{Exp}_p \left(\frac{1}{S} \sum_{s=1}^S \text{Log}_p(p_s) \right).$$

Provided that the data points p_1, \dots, p_S are close enough to each other and that also p is chosen close enough to them μ is well-defined and unique.

Having fixed μ we apply (5.4) to (5.2) and (5.3). This yields the following equations for the approximations V_1, \dots, V_N of V'_1, \dots, V'_N :

$$(5.6) \quad V_1 := \underset{\substack{V \in T_\mu M \\ |v|=1}}{\operatorname{argmax}} \sum_{s=1}^S \langle V, \text{Log}_\mu(p_s) \rangle_\mu^2,$$

and

$$(5.7) \quad V_n := \underset{\substack{V \in T_\mu M \\ |V|=1 \\ V \perp_{\mathbb{R}} \langle V_1, \dots, V_{n-1} \rangle}}{\operatorname{argmax}} \sum_{s=1}^S \left(\sum_{k=1}^{n-1} \langle V_k, \text{Log}_\mu(p_s) \rangle_\mu^2 + \langle V, \text{Log}_\mu(p_s) \rangle_\mu^2 \right),$$

for $2 \leq n \leq N$. Looking these equations we observe that they define a PCA of $\text{Log}_\mu(p_1), \dots, \text{Log}_\mu(p_S) \in T_\mu M$. This means that we get V_1, \dots, V_N by computing the eigenvalues of the covariance matrix of $\text{Log}_\mu(p_1), \dots, \text{Log}_\mu(p_S)$. Thus, we choose an orthonormal basis E_1, \dots, E_N of $T_\mu M$ and write $w_s^i E_i = \text{Log}_\mu(p_s)$, $1 \leq s \leq S$. The covariance matrix of $\text{Log}_\mu(p_1), \dots, \text{Log}_\mu(p_S)$ with respect to E_1, \dots, E_N is given by

$$(5.8) \quad \Sigma = \frac{1}{S} \sum_{s=1}^S w_s w_s^t \in S(N).$$

Let v_1, \dots, v_N be eigenvectors of Σ and $\sigma_1^2, \dots, \sigma_N^2$ its eigenvalues. Then $V_i = v_n^i E_i$, $1 \leq n \leq N$.

The above considerations are summarized in the following definition:

DEFINITION 5.1. Assume M and data points $p_1, \dots, p_S \in M$ as above. Let μ be as in (5.5) and E_1, \dots, E_N an orthonormal basis of $T_\mu M$. Let further be Σ as in (5.8). We call (μ, Σ) the *approximated Principal Geodesic Analysis (approximated PGA)* of the data p_1, \dots, p_S with respect to the basis E_1, \dots, E_N .

Obviously, the approximated PGA reduces to a tangent space PCA centered around the approximated mean μ which is again obtained in the tangent space setting. This means, that in principle this approach is limited to training data which can be mapped to a common tangent space, i.e. which are close enough to each other. To our experience this does not pose a problem in many applications. In particular the parameter manifolds we consider in Section 8 can be mapped to a single tangent space with the exception of a set of lower dimension only.

Note that the approximated PGA with respect to E_1, \dots, E_N provides not only the mean and the principal directions of the data p_1, \dots, p_S but also its variation into the principal directions. In other words, the eigenvalues $(\sigma_1, \dots, \sigma_N)$ of Σ can be interpreted as the standard deviation within the corresponding submanifold. I.e.

$$\text{Exp}_\mu(-\sigma_n V_n), \quad \mu, \quad \text{Exp}_\mu(\sigma_n V_n)$$

correspond to the mean and to the elements of the n -th geodesic submanifold whose distance to μ is σ . We call the above triple the n -th mode of the data p_1, \dots, p_S .

6. MAHALANOBIS DISTANCE ON PRODUCTS OF MANIFOLDS

A number of parametric shape spaces can be represented as products of manifolds. E.g. circles can be parametrized by the position of their center and the radius, which means that $M = \mathbb{R}^2 \times \mathbb{R}^+$. Ellipses might be considered in $M = \mathbb{R}^2 \times S^1 \times \mathbb{R}^+ \times \mathbb{R}^+$ and planar spline curves in $\mathbb{R}^2 \times \dots \times \mathbb{R}^2$. The parametric medial axis representation we introduce in Section 8 is a further prominent example of a product manifold. There are various ways of defining Riemannian metrics on such products. Perhaps the most simple approach is to use the product metric defined by the metrics on the single factor manifolds. This metric is unique except for a constant scaling of each component. Note that the choice of these weighting factors can not be avoided but is inherently done whenever the product metric is utilized. The obvious idea of setting each weight to 1 just pushes the decision about the scaling from the product metric to the metric of each of the factor. The choice of these scaling coefficients influences the distance on the product manifold but not the shape of geodesics between them (Theorem 6.1). But still in general, the scaling of the individual metrics influences the performance of minimization algorithms on the product manifolds and the question of optimal weighting factors in this respect is a difficult one.

In this section we will show that the Mahalanobis metric on product manifolds is invariant to the choice of these coefficients if the underlying statistics have been computed with respect to the same metric. I.e. for all problems formulated on M equipped with the Mahalanobis metric the initial choice of the weighting of the metrics of the factor manifolds is negligible. In particular, minimization routines utilizing the Mahalanobis metric perform identically for each sequence of scaling factors.

Assume M to be the product of K Riemannian manifolds, i.e. $M = M_1 \times \dots \times M_K$. Let further $c = (c_1, \dots, c_K)$ be a K -tuple of positive coefficients.

We define the metric $\langle \cdot, \cdot \rangle_p^c$ in $p = (p_1, \dots, p_K) \in M$ by

$$(6.1) \quad \langle V, W \rangle_p^c = c_1 \langle V_1, W_1 \rangle_{p_1} + \dots + c_K \langle V_K, W_K \rangle_{p_K},$$

for two tangent vectors $V = (V_1, \dots, V_K)$ and $W = (W_1, \dots, W_K)$ in $T_p M$. Moreover we denote the distance on M associated with $\langle \cdot, \cdot \rangle_p^c$ as d_M^c .

Assume data points $p_1, \dots, p_S \in M$ and their approximated PGA (μ, Σ) with respect to the orthonormal basis $E_1, \dots, E_N \in T_\mu M$. Let $d_{M, \mu, \Sigma}$ be the Mahalanobis distance on M . Now we replace the Riemannian metric $\langle \cdot, \cdot \rangle$ on M by $\langle \cdot, \cdot \rangle^c$ and do the same computations again. We denote the resulting approximated PGA with respect to the orthonormal basis $E_1^c, \dots, E_N^c \in T_{\mu^c} M$ as (μ^c, Σ^c) and the corresponding Mahalanobis distance as d_{M, μ^c, Σ^c}^c . Note that the orthonormality of E_1^c, \dots, E_N^c is also understood with respect to $\langle \cdot, \cdot \rangle^c$. In the following we will prove that for any choice of coefficients c , the maps $d_{M, \mu, \Sigma}$ and d_{M, μ^c, Σ^c}^c are defined on the same subset of M and that

$$d_{M, \mu, \Sigma} = d_{M, \mu^c, \Sigma^c}^c.$$

We remember the following fact about geodesics on products of Riemannian manifolds:

THEOREM 6.1. *Let $M = M_1 \times \dots \times M_K$ be as above and $c = (c_1, \dots, c_K)$, $c_i > 0$, $1 \leq k \leq K$. Denote the metric defined by the coefficients c as $\langle \cdot, \cdot \rangle^c$. Now assume $p, q \in M$ such that p and q can be connected by a unique geodesic segment γ with respect to $\langle \cdot, \cdot \rangle$, i.e. γ is geodesic, $\gamma(0) = p$, $\gamma(1) = q$ and $L(\gamma) = d_M(p, q)$.*

Then p and q are also connected by a unique geodesic segment with respect to the scaled metric $\langle \cdot, \cdot \rangle^c$. Moreover, these two geodesic segments coincide.

Proof. This is a straightforward consequence of the fact that the geodesic segment γ minimizes the kinetic energy of paths connecting p and q (cf. [30]). Let $\gamma = (\gamma_1, \dots, \gamma_K)$. According to (6.1) the kinetic energy of γ is given by

$$\int_0^1 \langle \dot{\gamma}, \dot{\gamma} \rangle_\gamma dt = \sum_{k=1}^K \int_0^1 \langle \dot{\gamma}_k, \dot{\gamma}_k \rangle_{\gamma_k} dt.$$

The latter sum is minimal if each of the summands are minimal, which proves the claim. \square

As a consequence we can state the following result:

COROLLARY 6.1. *Let $M = M_1 \times \dots \times M_K$ and $c = (c_1, \dots, c_K)$, $c_k > 0$ be as in Theorem 6.1. Again we look at M equipped with the metrics $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle^c$ respectively. For $p \in M$ denote the corresponding exponential and logarithmic maps as Exp_p , Exp_p^c and Log_p , Log_p^c respectively. Then*

$$\text{Exp}_p = \text{Exp}_p^c \quad \text{and} \quad \text{Log}_p = \text{Log}_p^c.$$

This finally leads us to the main result of this section:

THEOREM 6.2. *Let $M = M_1 \times \dots \times M_K$ and $c = (c_1, \dots, c_K)$, $c_k > 0$ be as in Theorem 6.1. Assume data $p_1, \dots, p_S \in M$. As in the beginning of this section let (μ, Σ) be the approximated Principal Geodesic Analysis of these data with respect to the orthonormal basis E_1, \dots, E_N of $T_\mu M$ and*

$$d_{M, \mu, \Sigma} : \mathcal{U} \times \mathcal{U} \rightarrow [0, \infty[$$

the corresponding Mahalanobis distance. With respect to the modified metric $\langle \cdot, \cdot \rangle_c$ we accordingly get the approximated Principal Geodesic Analysis (μ^c, Σ^c) with respect to the orthonormal basis E_1^c, \dots, E_N^c of $T_{\mu^c}M$ and the Mahalanobis distance

$$d_{M, \mu^c, \Sigma^c}^c : \mathcal{U}^c \times \mathcal{U}^c \rightarrow [0, \infty[.$$

Then $\mu = \mu^c$, $\mathcal{U} = \mathcal{U}^c$ and for $p, q \in \mathcal{U}$

$$d_{M, \mu, \Sigma}(p, q) = d_{M, \mu^c, \Sigma^c}(p, q)^c.$$

Proof. Applying Corollary 6.1 to (5.5) shows $\mu = \mu^c$. Because the geodesics with respect to $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle^c$ coincide (cf. Theorem 6.1), it follows that $\mathcal{U} = \mathcal{U}^c$.

Let $a = (a_j^i)_{ij}$ be the matrix transforming the basis $(E_n^c)_{1 \leq n \leq N}$ to $(E_n)_{1 \leq n \leq N}$, i.e. $E_n = a_n^i E_i^c$, $1 \leq n \leq N$. If

$$w_s^i E_i = \text{Log}_\mu(p_s), \quad 1 \leq s \leq S,$$

are the coordinate representations of the data p_s , then

$$\text{Log}_\mu(p_s) = w_s^i a_i^j E_j^c, \quad 1 \leq s \leq S.$$

Hence, from formula (5.8) it follows that

$$\Sigma^c = a \Sigma a^t.$$

Assume $p \in \mathcal{U}$. Let $E_{n,p}$, $1 \leq n \leq N$, be the parallel transport of $(E_n)_{1 \leq n \leq N}$ to $T_p M$. From (4.5) it follows that for $p \in \mathcal{U}$ and for $V = v^i E_{i,p}$ and $W = w^i E_{i,p}$ in $T_p M$

$$(6.2) \quad \langle V, W \rangle_{p, \mu, \Sigma} = v^t \Sigma^{-1} w.$$

Let $E_{n,p}^c$, $1 \leq n \leq N$ the parallel transport of $(E_n^c)_{1 \leq n \leq N}$ to p with respect to $\langle \cdot, \cdot \rangle^c$. In principle this transport depends on the metric on M (as the metric defines the geodesics), but Theorem 6.1 shows that in our case the geodesics on M with respect to $\langle \cdot, \cdot \rangle$ and $\langle \cdot, \cdot \rangle^c$ coincide. Hence, the orthonormal bases

$$(E_{n,p})_{1 \leq n \leq N} \quad \text{and} \quad (E_{n,p}^c)_{1 \leq n \leq N}$$

are transported along the same geodesics. This means that the transformation between them is constant, i.e.

$$E_{n,p}^c = a_n^i E_{n,p} \quad \text{for every } p \in U.$$

This finally leads to

$$\begin{aligned} \langle V, W \rangle_{p, \mu^c, \Sigma^c}^c &= \langle v^i a_i^j E_{j,p}^c, w^i a_i^j E_{j,p}^c \rangle_{p, \mu^c, \Sigma^c} \\ &= (av)^t (a \Sigma a^t)^{-1} (av) = v^t \Sigma^{-1} v. \end{aligned}$$

Together with Definition 4.1 and (6.2) this proves the claim. \square

7. THE REGULARIZATION FUNCTIONAL

Let M be N -dimensional, $p_1 \dots, p_S \in M$ a set of training data and (μ, Σ) its approximated Principal Geodesic Analysis. We assume that Σ is non-singular, i.e. all its singular values are positive. In the following theorem we show that the Mahalanobis distance qualifies as a regularization function with the properties introduced in Section 3.

Let μ, Σ and \mathcal{U} be as in Definition (4.1), i.e. the Mahalanobis distance $d_{M,\mu,\Sigma}$ is defined on \mathcal{U} . Define for $p \in M$

$$(7.1) \quad R_\alpha(p) = \begin{cases} \alpha d_{M,\mu,\Sigma}(\mu, p)^2 (d_M(\partial\mathcal{U}, p)^{-1} + 1) & \text{for } p \in \mathcal{U}, \\ \infty & \text{for } p \notin \mathcal{U}, \end{cases}$$

where $\partial\mathcal{U}$ is the topological boundary of \mathcal{U} . Then $R_\alpha \geq 0$ and clearly (3.2) is satisfied for every $p \in \mathcal{U}$. Because of (4.10) in Theorem 4.2 also (3.3) and (3.4) hold. Thus, R_α qualifies as a regularization for any functional F with $\text{dom}(F) = \mathcal{U}$. The regularization term forces minimizers of I_α to stay in \mathcal{U} and consequently ensures that the Mahalanobis distance is well-defined.

Note that in many applications \mathcal{U} is chosen as the manifold M . This holds in particular in all the cases we present in the next section. In this case (7.1) reduces to

$$(7.2) \quad R_\alpha(p) = \alpha d_{M,\mu,\Sigma}(\mu, p)^2.$$

8. THE PARAMETRIC MEDIAL AXIS REPRESENTATION

In this section we introduce the parametric shape space we used in the examples in the following section. It is based on the medial axis representation and was developed in [14, 15, 13, 21]. The authors' idea was to parametrize the medial axis transform of shapes instead of the actual boundary curve. We chose a simple model which is defined as follows:

DEFINITION 8.1. Assume a tree G (i.e. a connected graph without cycles) with m vertices and $m - 1$ edges. Define

$$M = \mathbb{R}^2 \times \mathbb{R}_+^m \times (S^1 \times \mathbb{R}_+)^{m-1}.$$

Obviously M is the direct product of Riemannian manifolds. On \mathbb{R}^2 we use the Euclidean metric, on S^1 the metric induced by the embedding of S^1 into \mathbb{R}^2 . There exists a unique (up to scalar multiplication) metric on \mathbb{R}_+ which is invariant with respect to multiplication [1, Chapter VI, Corollaries (3.5), (3.7)]. We denote the product metric on M as $\langle \cdot, \cdot \rangle$. We emphasize that according to Theorem 6.2 the Mahalanobis regularization (7.2) (with respect to some training data) is invariant with respect to the choice of the weighting coefficients of the metrics on \mathbb{R}^2 , S^1 and \mathbb{R}_+ .

To use this parameter manifold in actual applications, we have to provide ways of computing the actual appearance of a shape $p \in M$. This leads us to the following definition:

DEFINITION 8.2. Let $\mathcal{V} \subseteq M$ an open subset of M and

$$\psi : \mathcal{V} \rightarrow \mathcal{C}^1(S^1, \mathbb{R}^2),$$

such that for every $p \in \mathcal{V}$ the curve $\psi(p)$ is a Jordan curve, i.e. $\psi(p)$ is closed and injective. Then we call (M, ψ) a *shape model*.

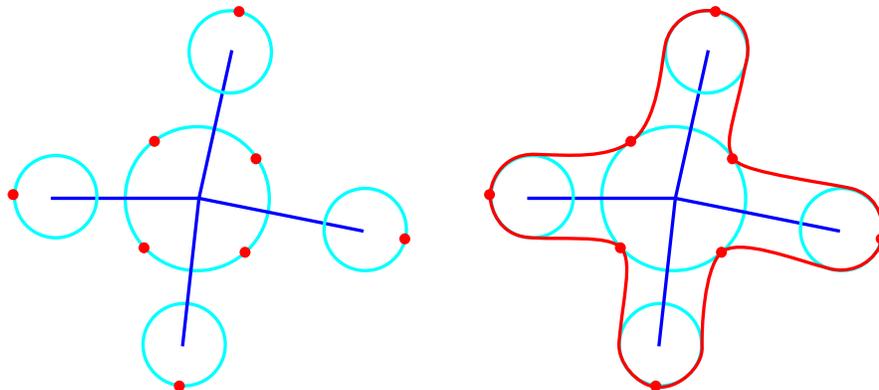


FIGURE 1. An instance of a shape model. The shape parameters $p \in M = \mathbb{R}^2 \times \mathbb{R}_+^5 \times (S^1 \times \mathbb{R}_+)^4$ are mapped to the red outline on the right by ψ .

In Figure 1 we can see an example of a shape model. Here, $M = \mathbb{R}^2 \times \mathbb{R}_+^5 \times (S^1 \times \mathbb{R}_+)^4$. On the left hand side is a visualization of one $p \in M$. Starting from the reference position (in \mathbb{R}^2) we added four edges (in $S^1 \times \mathbb{R}_+$). This results in the dark blue skeleton. The topology of the skeleton is dictated by the graph G . Thus, the skeleton here is in fact the first step of the application of ψ to p , which transforms the abstract data p into its visual representation. Secondly, we add the light blue atoms of different radii (in \mathbb{R}_+) with centers on the vertices of the skeleton. This representation incorporates all the data provided by p .

To obtain the shape appearance of p we first selected the red points on the atom circles. On the right image we interpolated these points with cubic spline segments such that the resulting curve is differentiable and their tangents in the interpolation points agree with the circle tangents. This spline curve is $\psi(p)$. We can see that the dark blue skeleton approximates the medial axis of the red curve, which was the original motivation for this kind of shape model.

In terms of the results in this paper, G is not a part of our shape space but specifies how the shape appearance is computed from parameters $p \in M$. As such it can be thought as a parameter of the map ψ .

9. APPLICATIONS

In this section we present two examples of regularization functionals based on the shape model introduced in the previous section. Both examples are concerned with geometry reconstruction in image data.

Our approach consists of two steps. Starting from a fixed shape model (M, ψ) we manually segment some training data. In other words, we estimate a set of shapes $p_1, \dots, p_S \in M$ which we consider to be a probable collection of results of a segmentation. Then we perform an approximated PGA of this data (Definition 5.1) and use the result to define the regularization R_α as in (7.2).

In the next step, we select a segmentation energy F , which is well suited for the specific kind of image data and which maps a shape p to an energy $F(p)$. Then, we define the regularization functional I_α as in Definition 3.1. If F is bounded from below and continuous, Theorem 3.1 tells us that a minimizer of I_α exists. Furthermore, if F attains a minimum, then, by Theorem 3.2, a subsequence of the minimizers of I_α converges to a minimizer of F for $\alpha \rightarrow 0$. The actual minimization of I_α is done by using a gradient descent technique in the first example and a Monte Carlo method in the second one.

In addition to the properties of the regularization functional I_α mentioned above, the knowledge of the statistics of the training data p_1, \dots, p_S has several other advantages:

- Choosing a large regularization parameter α forces the minimizers of I_α to be close to the mean shape μ . Furthermore, deviations from the mean shape are only allowed into directions where the variation of the data $(p_s)_{1 \leq s \leq S}$ is large. This property can be used to reconstruct perturbed image data. Such perturbations can have the effect that a minimizer of F does not represent a satisfactory result. The regularization term penalizes these outcomes. This is well illustrated in the example in Section 9.1.
- During the iterative minimization of I_α we can make use of the underlying statistics to select search directions which are most probable in terms of the training data. A gradient descent technique e.g. might compute the gradients of I_α with respect to $\langle \cdot, \cdot \rangle_{p, \mu, \Sigma}$ instead of the standard metric on M . In the case of Monte Carlo methods we can confine ourselves to randomly generated shapes p which satisfy $R_\alpha(p) < C$ for some constant C . This is equivalent to limiting the search for optimal shapes to a ball of radius C around μ with respect to $d_{M, \mu, \Sigma}$.

9.1. Region-Based Segmentation. In this example we use the Mahalanobis regularization together with a simplification of the Mumford-Shah functional. We generate a training set by manually segmenting artificial image data and apply the resulting regularization functional to a partially occluded test image.

Assume a given image $f : \Omega \rightarrow \mathbb{R}$, where Ω is a 2-dimensional domain. The Mumford-Shah functional [28, 29] reads as

$$(9.1) \quad I_{\beta_1, \beta_2}^{\text{MS}}(\Gamma, u) = \int_{\Omega} (u - f)^2 dx + \frac{\beta_1}{2} \int_{\Omega \setminus \Gamma} |\nabla u|^2 dx + \beta_2 \mathcal{H}^1(\Gamma),$$

where $u : \Omega \rightarrow \mathbb{R}$ is smooth on $\Omega \setminus \Gamma$ and $\beta_1, \beta_2 > 0$. The goal is to find u' and Γ' such that $I_{\beta_1, \beta_2}^{\text{MS}}(\Gamma', u')$ is minimal. Then u' is smooth everywhere with exception of possible discontinuities along the 1-dimensional set Γ' and approximates the original image data f . If f consists of several regions of low variation in contrast, then we expect u' to approximate f on these regions and Γ' to separate them from each other. Thus, Γ' is a segmentation of f .

Often a simplified version of (9.1) is used for segmentation. The idea is to approximate f by piecewise constant functions instead of piecewise smooth functions. The corresponding functional can be interpreted as the

limit of $I_{\beta_1, \beta_2}^{\text{MS}}$ for $\beta_1 \rightarrow \infty$. If we restrict ourselves to closed Jordan curves $\gamma : [0, 1] \rightarrow \Omega$ instead of general sets Γ we obtain

$$I_{\beta}^{\text{SMS}}(\gamma) = \int_{\mathcal{J}(\gamma)} (u_1(\gamma) - f)^2 dx + \int_{\mathcal{O}(\gamma)} (u_2(\gamma) - f)^2 dx + \beta \text{Length}(\gamma),$$

where

$$u_1(C) = \frac{1}{|\mathcal{J}(\gamma)|} \int_{\mathcal{J}(\gamma)} f dx \quad \text{and} \quad u_2(C) = \frac{1}{|\mathcal{O}(\gamma)|} \int_{\mathcal{O}(\gamma)} f dx$$

are the mean values of f inside and outside of γ . This functional was also proposed by Chan & Vese in [4]. Note that the extension to a collection of curves $\gamma_1, \dots, \gamma_R$, which do not intersect or lie within each other, is straightforward [17].

Now consider a shape model ψ as illustrated in Figure 1. I.e. $M = \mathbb{R}^2 \times \mathbb{R}_+^5 \times (S^1 \times \mathbb{R}_+)^4$ (cf. Definition 8.1). Then we define $F : M \rightarrow [0, \infty]$ by

$$F(p) = \int_{\mathcal{J}(\psi(p))} (u_1(\psi(p)) - f)^2 dx + \int_{\mathcal{O}(\psi(p))} (u_2(\psi(p)) - f)^2 dx.$$

I.e. F maps a shape p to the simplified Mumford-Shah energy associated with its boundary $\psi(p)$ without length regularization (meaning that $\beta = 0$).

We manually segmented hand-drawn images of crosses and obtained 13 training shapes p_1, \dots, p_{13} (Figure 2). The results of the approximated PGA are illustrated in Figure 3. We used the PGA to define R_α and finally $I_\alpha = F(p) + R_\alpha$. Since we consider the actual position and rotation of the shape models not to be important for the statistics we additionally set the corresponding parts of the Mahalanobis distance to zero. In other words, for $p, q \in M$

$$R_\alpha(p) = R_\alpha(q),$$

if p and q differ only in position and rotation.

In case f is continuous and ψ differentiable $M \rightarrow \mathcal{C}^1([0, 1], \mathbb{R}^2)$, then F is differentiable as a map $M \rightarrow [0, \infty]$. If we write $F = G \circ \psi$, then its derivative is

$$DF(\gamma) = DG(\psi(p)) \circ D\psi(p),$$

where

$$DG(\gamma) = (u_1(\gamma) - f \circ \gamma)^2 \mathbf{n}_\gamma - (u_2(\gamma) - f \circ \gamma)^2 \mathbf{n}_\gamma,$$

for a curve γ . Here \mathbf{n}_γ denotes the outer unit normal of γ .

The Mahalanobis distance on M is differentiable on

$$M \setminus \mathbb{R}^2 \times \mathbb{R}_+^5 \times (\{-\mu\} \times \mathbb{R}_+)^4,$$

and hence also R_α and the regularization functional I_α . Thus, we can minimize I_α using steepest descent. As initial value we chose the mean μ of the approximated PGA.

We applied this technique to a cross (not in the training set), where the central part is missing. The result is depicted in the middle of the top row in Figure 4. To illustrate the influence of the regularization parameter we compared this result to the two minimizers we obtained by choosing the regularization parameter α one order of magnitude larger and smaller (left

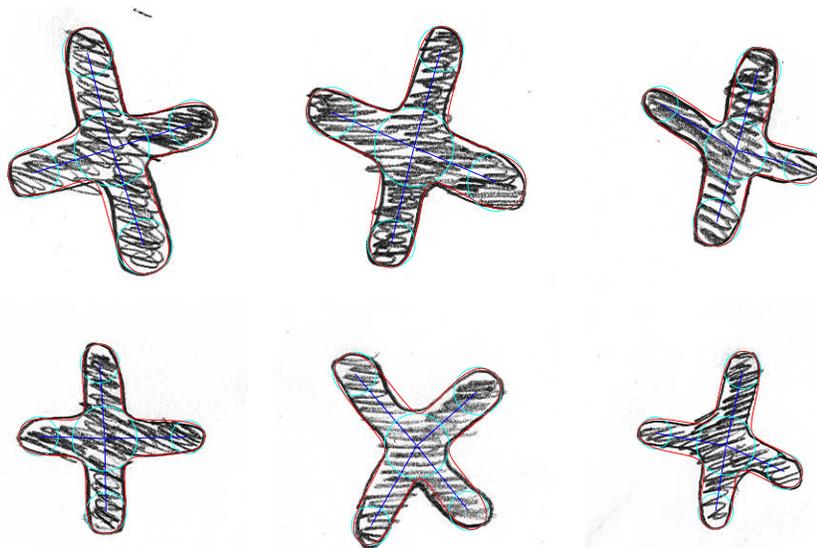


FIGURE 2. 6 training shapes (out of 13).

and right) than in the middle image. The initial shape in the steepest descent algorithm is displayed in the lower row at the left.

For further comparison, we computed the minimizer of I_β^{SMS} in (9.1) (lower row, middle). During the iterative minimization of I_β^{SMS} the initial curve γ is split into four curves $\gamma_1, \dots, \gamma_4$ by a topology handling routine. Obviously with this general approach we are not able to recover the original shape but detect the four remaining parts of the cross separately. On the right we can see the minimizer of the map

$$p \mapsto I_\beta^{\text{SMS}}(\psi(p)).$$

I.e. we reformulate (9.1) on M or, in other words, we replace R_α in I_α by the simple regularization $\beta \text{Length}(\psi(p))$. Obviously, the cross is in principle detected as a cross (the shape model leaves no other possibility), but the proportions reflected by the training shapes in Figure 2 are not preserved.

9.2. Edge-Based Segmentation. In this example we are concerned with the detection of multiple yeast cells in microscope images. Because the mean contrast of the objects in question does not significantly differ from the background, we can not use region based segmentation techniques for these image data. Hence, we chose an edge based segmentation technique originally introduced by Kass et al. [22] (“Snakes”). There, the authors propose to minimize

$$I_\beta^{\text{Snakes}}(\gamma) = - \int_\gamma |\nabla f(\gamma(\tau))| d\tau + \beta \text{Length}(\gamma)$$

for a curve γ .

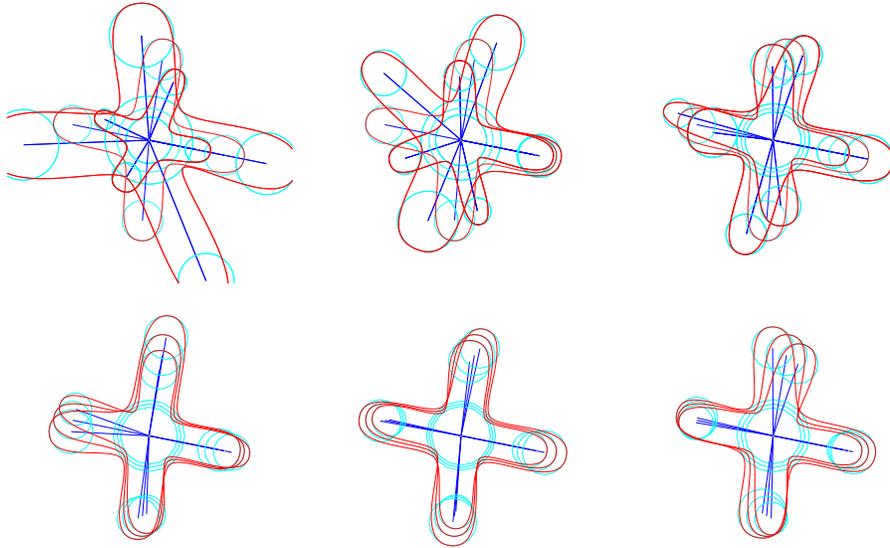


FIGURE 3. Modes 1–6 of the approximated PGA of the training data (including the shapes in Figure 2).

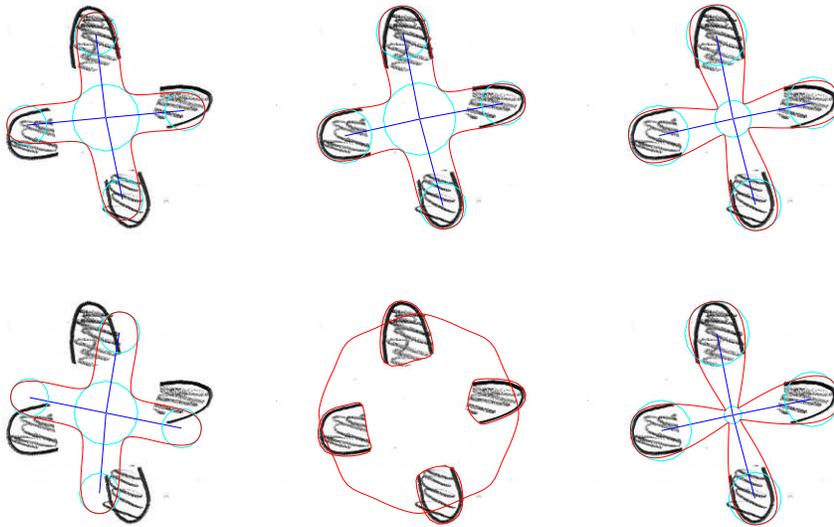


FIGURE 4. Top row: minimizers of the simplified Mumford-Shah energy with (relative) regularization parameters $\alpha = 10$, $\alpha = 1$, $\alpha = 10^{-1}$ (left to right). Lower row: initial shape in the steepest descent algorithm, minimizing spline curve of the simplified Mumford-Shah energy (with initial curve), minimizer of the simplified Mumford-Shah energy with length regularization (left to right).

Again we reformulate this functional for an arbitrary shape model (M, ψ) and define $F : M \rightarrow [0, \infty]$ by

$$F(p) = - \int_{\gamma} |\nabla f(\gamma(\tau))| d\tau.$$

We choose $M = \mathbb{R}^2 \times \mathbb{R}_+^2 \times (S^1 \times R_+)$ (i.e. a 6-dimensional manifold). The associated map ψ is illustrated in Figure 5. Then we perform a manual segmentation of 18 cells in the central part of a sample image (Figure 6, top left) and compute the approximated PGA (μ, Σ) of this data. The results of the PGA are illustrated in Figure 7. Again we do not consider variations in position and rotation in the statistics, which means that only three modes remain. As we can see, the variation of third mode of the cell data is very small. In other words, the shapes of the training data are very close to a 5-dimensional geodesic submanifold of M . This complies with the intuitive perception of the cells as ellipses (which can be considered a 5-dimensional manifold, cf. Section 6).

As in the last example, the PGA leads to the regularization R_α and in consequence to the regularization functional $I_\alpha = F + R_\alpha$. Note that F is not necessarily bounded from below, so only Theorem 3.2 is applicable in this case, but not Theorem 3.1.

To automatically detect cells in image data we choose a fixed number of models (denoted as R) and minimize

$$(p_1, \dots, p_R) \mapsto \sum_{r=1}^R I_\alpha(p_r).$$

For the solution of (9.2) we used the following heuristic algorithm:

- (1) Assume scalars $c_1 = \dots = c_R = C$ for some $C > 0$ and a vector of shapes (p_1, \dots, p_R) .
- (2) Choose a random shape $p' \in B_C := \{p \in M : d_{M, \mu, \Sigma}(\mu, p) \leq C\}$. If $I_\alpha(p') < c_r$ for $1 \leq r \leq R$ and $\psi(p')$ does not overlap with any of the p_1, \dots, p_R , replace p_r by p' and set $c_r = I_\alpha(p')$. If it does overlap with some elements of p_{r_1}, \dots, p_{r_k} but $I_\alpha(p') < \min(c_{r_1}, \dots, c_{r_k})$, then replace p_{r_1} by p' and set $c_{r_1} = I_\alpha(p')$ and $c_{r_1} = \dots = c_{r_k} = C$. Repeat this step.

By “overlapping” we mean that the common area of the bounding boxes of two shapes does not exceed a certain value (relatively to the area of the union of the bounding boxes).

In simple words, we generate random shapes and try to improve the current segmentation by replacing shapes in (p_1, \dots, p_R) by the new shape. This gradually improves the segmentation of the result. This algorithm is far from being optimal and in fact does not necessarily lead to the best result, but it demonstrates two important issues:

- (1) We do not have to evaluate the gradient of I_α in this example. Moreover the danger of being trapped at local minima is much smaller than with gradient descent techniques. This means that we are able to detect different cells even if they are clustered as in our examples.

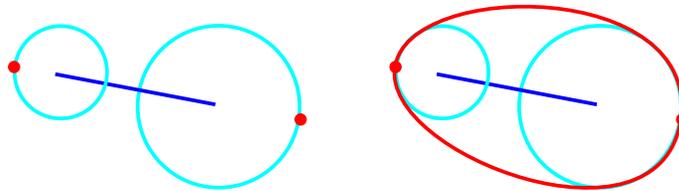


FIGURE 5. An instance of a shape model. The shape parameters $p \in M = \mathbb{R}^2 \times \mathbb{R}_+^2 \times (S^1 \times R_+)$ are mapped to the red outline on the right by ψ .

A simple active contour algorithm [2, 3] can not handle this situation, since features *inside* the clusters have to be considered for a correct segmentation.

- (2) The first property is a characteristic of Monte Carlo optimization techniques in general. However, it is impossible to randomly sample the entire manifold M . One has to agree on some subset of “probable” or “meaningful” shapes. In our case this selection is canonically given by choosing shapes close to the mean shape and varying according to the training data. In other words, by choosing shapes in B_C we generate only shapes, which we expect to actually appear in the image data, as candidates for a correct segmentation.

10. CONCLUSION AND OUTLOOK

In this work we extended the notion of the Mahalanobis distance to Riemannian manifolds and showed that the generalized Mahalanobis distance

- qualifies as a regularization of variational problems on finite dimensional manifolds, and
- enables us to intelligently incorporate statistical knowledge about shapes in such variational problems.

In examples we correctly detected partially occluded objects (Section 9.1) and exemplarily introduced a segmentation method based on Monte Carlo optimization, which can be trained to specific types of image data to optimize its convergence characteristics (Section 9.2). The segmentation functionals used in the examples were a simplified version of the Mumford-Shah energy in the first and the “Snakes” energy in the second example.

In addition we concentrated on products of Riemannian manifolds. We proved that the Mahalanobis distance based on the product metric is *unique*, i.e. invariant with respect to the specific choice of the weighting coefficients in (6.1).

Note that this is not true in the general case. The Mahalanobis distance is invariant with respect to scaling of the factor manifolds, it still depends on the underlying metrics themselves. One might e.g. also equip a shape model (M, ψ) with the local metric

$$(10.1) \quad \langle v, w \rangle_p = \int_{\psi(p)} \langle \psi_*(v), \mathbf{n}_{\psi(p)} \rangle \langle \psi_*(w), \mathbf{n}_{\psi(p)} \rangle ds$$

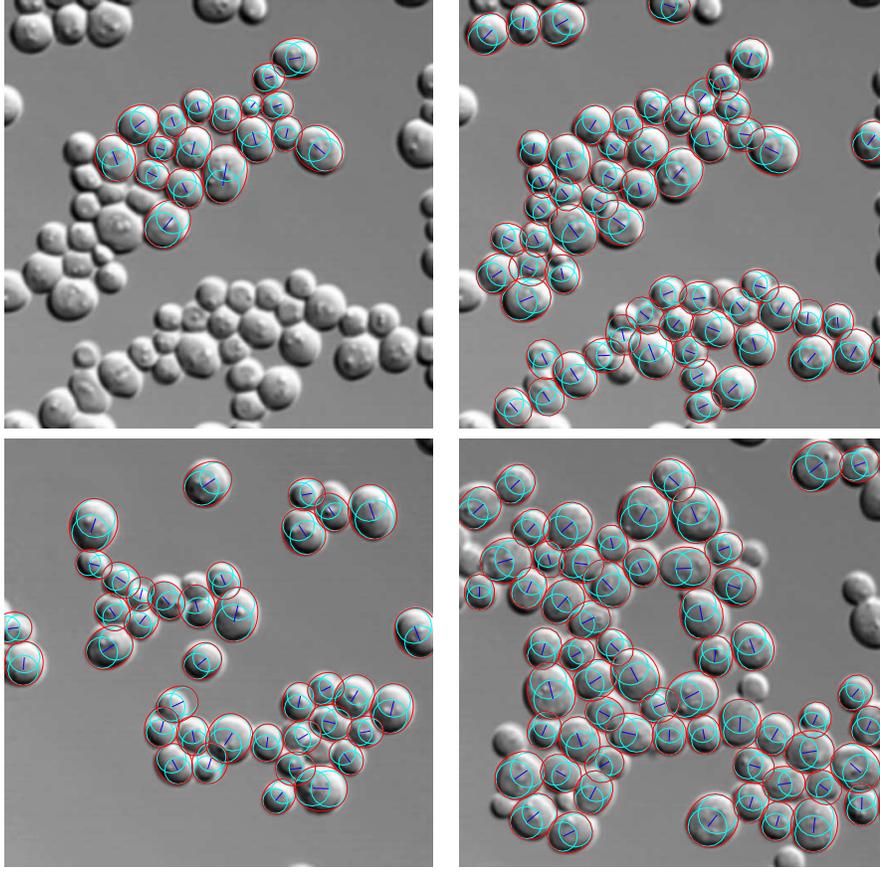


FIGURE 6. Upper left image: Expert segmentation of 18 cells. Upper right and lower row: Automatic detection of cells.

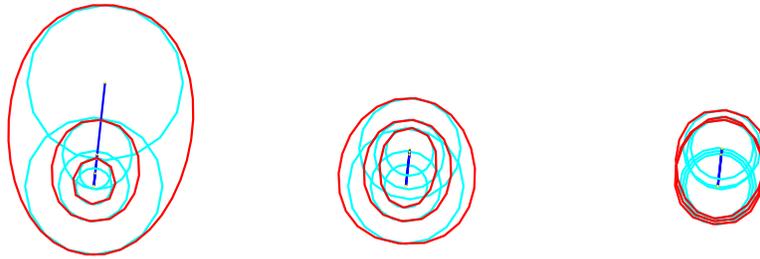


FIGURE 7. Modes 1, 2 and 3 of the approximated PGA of the training data in Figure 6, upper left.)

for $v, w \in T_p M$, $p \in M$. Here the first inner product $\langle \cdot, \cdot \rangle_p$ is the metric on M , whereas the ones inside the integral refer to the standard inner product on \mathbb{R}^2 (note that $\psi_*(v)$ and $\mathbf{n}_{\psi(p)}$ are implicitly evaluated at s in the integral). One could also use an extension of the above concept as suggested by Michor and Mumford [26, 25]. All these metrics are *geometric metrics*, in

the sense that they are invariant with respect to the parametrization of shapes and depend only on the actual shape itself (i.e. on $\psi(p)$ in our terminology).

This means that our framework is flexible in the sense that it applies to a wide range of underlying segmentation energies and parametric shape models. The original versions of both segmentation approaches we presented in Section 9, i.e. the Mumford-Shah and the "Snakes" energy, are regularized by penalizing the boundary area of the detected region. Given the fact that some kind of regularization is actually mandatory for the well-posedness of the corresponding variational problem, it seems natural to use the Mahalanobis regularization on the shape parameter manifold *if* statistical a-priori data is *available*. In contrast to the general formulation of variational segmentation problems using for instance the level set approach or parametric curves, we expect our approach to perform substantially better in case of cluttered or partially occluded data.

Possible further investigations mainly cover two areas. First, in applications the minimization of the segmentation functional requires frequent evaluations of the Mahalanobis distance. These computations are very cheap for the product metric we used in connection with parametric medial axis representation, but can be very expensive if a geometric metric as e.g. (10.1) is used. The development of geometric metric on the M-Rep space, which is yet easy to compute, is subject of current research.

Secondly, our approach is limited to one class of shapes. The variation within this class can be large if this property is reflected by the training data, but the existence of two or more distinct clusters will not be reflected by the Mahalanobis distance. An approach as proposed by Cremers et al. [7] might be used to resolve this shortcoming.

Acknowledgments. This work has been supported by the Austrian Science Foundation (FWF) Projects FSP9202-N12, FSP9203-N12 and FSP9207-N12. We thank Heimo Wolinski, IMB-Graz, Karl-Franzens-Universität, Graz, Austria, and Bettina Heise, FLLL, Linz-Hagenberg, Austria, for the microscope images.

REFERENCES

- [1] William M. Boothby. *An Introduction to Differentiable Manifolds and Riemannian Geometry*, volume 63 of *Pure and Applied Mathematics*. Academic Press, New York, 1975.
- [2] Vicent Caselles, Francine Catté, and Françoise Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66:1–31, 1993.
- [3] Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997.
- [4] Tony F. Chan and Luminita A. Vese. Active contours without edges. *IEEE Transactions on Image Processing*, 10(2):266–277, 2001.
- [5] Yunmei Chen, Hemant D. Tagare, Sheshadri Thiruvenkadam, Feng Huang, David Wilson, Kaundinya S. Gopinath, Richard W. Briggs, and Edward A. Geiser. Using prior shapes in geometric active contours in a variational framework. *International Journal of Computer Vision*, 50(3):315–328, 2002.
- [6] Yunmei Chen, Sheshradi Thiruvenkadam, Hemant D. Tagare, Feng Huang, David Wilson, and Edward A. Geiser. On the incorporation of shape priors into geometric active contours. In *1st IEEE Workshop on Variational and Level Set Methods in Computer Vision*, July 2001.

- [7] Daniel Cremers, Timo Kohlberger, and Christoph Schnörr. Shape statistics in kernel space for variational image segmentation. *Pattern Recognition*, 36(9):1929–1943, 2003.
- [8] Daniel Cremers and Christoph Schnörr. Statistical shape knowledge in variational motion segmentation. *Image and Vision Computing*, 21:77–68, 2003.
- [9] Daniel Cremers, Nir Sochen, and Christoph Schnörr. Towards recognition-based variational segmentation using shape priors and dynamic labeling. In Lewis D. Griffin and Martin Lillholm, editors, *Scale-Space 2003*, volume 2695 of *Lecture Notes in Computer Science*, pages 388–400, Berlin Heidelberg, 2003. Springer-Verlag.
- [10] Daniel Cremers, Florian Tischhäuser, Joachim Weickert, and Christoph Schnörr. Diffusion snakes: Introducing statistical shape knowledge into the mumford-shah functional. *International Journal of Computer Vision*, 50(3):295–313, 2002.
- [11] Heinz W. Engl, Martin Hanke, and Andreas Neubauer. *Regularization of Inverse Problems*, volume 375 of *Mathematics and Its Applications*. Kluwer Academic Publishers, Dordrecht Boston London, 2000.
- [12] Weng Fang and Kap Luk Chan. Incorporating shape prior into geodesic active contours for detecting partially occluded object. *Pattern Recognition*, 40(7):2163–2172, 2007.
- [13] P. Thomas Fletcher and Sarang Joshi. Principal geodesic analysis on symmetric spaces: Statistics of diffusion tensors. In Milan Sonka, Ioannis A. Kakadiaris, and Jan Kybic, editors, *Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis, ECCV 2004 Workshops*, volume 3117 of *Lecture Notes in Computer Science*, pages 87–98, Berlin, 2004. Springer Verlag.
- [14] P. Thomas Fletcher, Conglin Lu, and Sarang Joshi. Statistics of shape via principal geodesic analysis on lie groups. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 1, pages 95–101, 2003.
- [15] P. Thomas Fletcher, Conglin Lu, Sarang Joshi, and Stephen M. Pizer. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE Transactions of Medical Imaging*, 23(8):995–1005, 2004.
- [16] Karl D. Fritscher and Rainer Schubert. 3D image segmentation by using statistical deformation models and level sets. *International Journal of Computer Assisted Radiology and Surgery*, 1(3):123–135, 2006.
- [17] Matthias Fuchs, Bert Jüttler, Otmar Scherzer, and Huaiping Yang. Combined evolution of level sets and B-spline curves for imaging. Technical Report 41, FSP 092: Joint Research Program of Industrial Geometry, January 2007.
- [18] Muriel Gastaud, Michel Barlaud, and Gilles Aubert. Combining shape prior and statistical features for active contour segmentation. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(5):726–734, 2004.
- [19] Sigurdur Helgason. *Differential Geometry, Lie Groups, and Symmetric Spaces*. Number 80 in Pure and Applied Mathematics. Academic Press, 1978.
- [20] Stephan Huckemann and Herbert Ziezold. Principal component analysis for riemannian manifolds, with an application to triangular shape spaces. *Advances in Applied Probability*, 38(2):299–319, 2006.
- [21] Sarang Joshi, Stephen Pizer, P. Thomas Fletcher, Paul Yushkevich, Andrew Thall, and J. S. Marron. Multiscale deformable model segmentation and statistical shape analysis using medial descriptions. *IEEE Transactions on Medical Imaging*, 21(5):538–550, 2002.
- [22] Micheal Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes active contour models. *International Journal of Computer Vision*, 1(4):321–331, 1988.
- [23] Huiling Le and David G. Kendall. The Riemannian structure of Euclidean shape spaces: A novel environment for statistics. *The Annals of Statistics*, 21(3):1225–1271, 1993.
- [24] Micheal E. Leventon, W. Eric L. Grimson, and Olivier Faugeras. Statistical shape influence in geodesic active contours. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 316–323, June 2001.
- [25] Peter W. Michor and David Mumford. An overview of the Riemannian metrics on spaces of curves using the Hamiltonian approach. *Applied and Computational Harmonic Analysis*, 2007.

- [26] Peter W. Michor and David B. Mumford. Riemannian geometries on spaces of plane curves. *Journal of the European Mathematical Society*, 8(1):1–48, 2006.
- [27] Micheal I. Miller, Alain Troué, and Laurent Younes. Geodesic shooting for computational anatomy. *Journal of Mathematical Imaging and Vision*, 24(2):209–228, 2006.
- [28] David Mumford and Jayant Shah. Boundary detection by minimizing functionals. In *Proc. IEEE Conf. Computer Vision Pattern Recognition*, pages 22–26, 1985.
- [29] David Mumford and Jayant Shah. Optimal approximations by piecewise smooth functions and associated variational problems. *Communications on Pure and Applied Mathematics*, 42(4):577–684, 1989.
- [30] Peter Peterson. *Riemannian Geometry*, volume 171 of *Graduate texts in mathematics*. Springer, 1998.
- [31] Mikael Rousson and Nikos Paragios. Shape priors for level set representations. In Anders Heyden, Gunnar Sparr, Mads Nielsen, and Peter Johansen, editors, *Computer Vision - ECCV 2002 : 7th European Conference on Computer Vision, Copenhagen, Denmark, May 28-31, 2002. Proceedings, Part II*, volume 2351 of *Lecture Notes in Computer Science*, pages 78–92. Springer, 2002.
- [32] Andy Tsai, Anthony Yezzi, Clare Tempany, Dewey Tucker, Ayres Fan, W. Eric L. Grimson, and Alan Willsky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging*, 22(2):137–154, 2003.
- [33] Marc Vaillant, Micheal I. Miller, Laurent Younes, and Alain Trouvé. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23:161–169, 2004.

UNIVERSITY OF INNSBRUCK
E-mail address: `matz.fuchs@uibk.ac.at`

UNIVERSITY OF INNSBRUCK
E-mail address: `otmar.scherzer@uibk.ac.at`