II. DIFFERENTIABLE MANIFOLDS

Washington Mio

Anuj Srivastava and Xiuwen Liu

(Illustrations by D. Badlyans)

CENTER FOR APPLIED VISION AND IMAGING SCIENCES

Florida State University

WHY MANIFOLDS?

Non-linearity underlies many mathematical structures and representations used in computer vision. For example, many **image and shape descriptors** fall in this category, as discussed in the introduction to this tutorial.

On the other hand, **differential calculus** is one of the most useful tools employed in the study of phenomena modeled on **Euclidean n-space** $\mathbb{R}^n$. The study of **optimization and inference problems, dynamics, and geometry** on linear spaces can all be approached with the tools of calculus.

A natural question arises:

What spaces can the techniques of differential calculus be extended to?

Since ***derivative*** is a local notion, it is reasonable to expect that differential calculus can be developed in spaces that locally “look like” Euclidean spaces. These spaces are known as ***differentiable manifolds***.

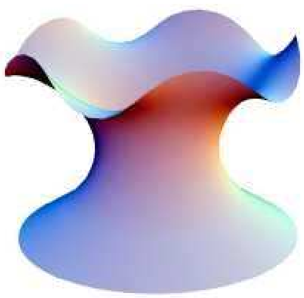

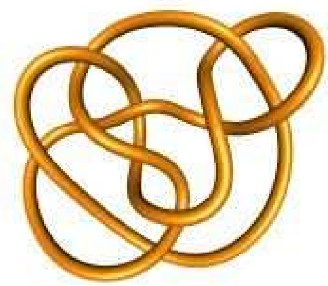

Manifolds Underlying Data Points

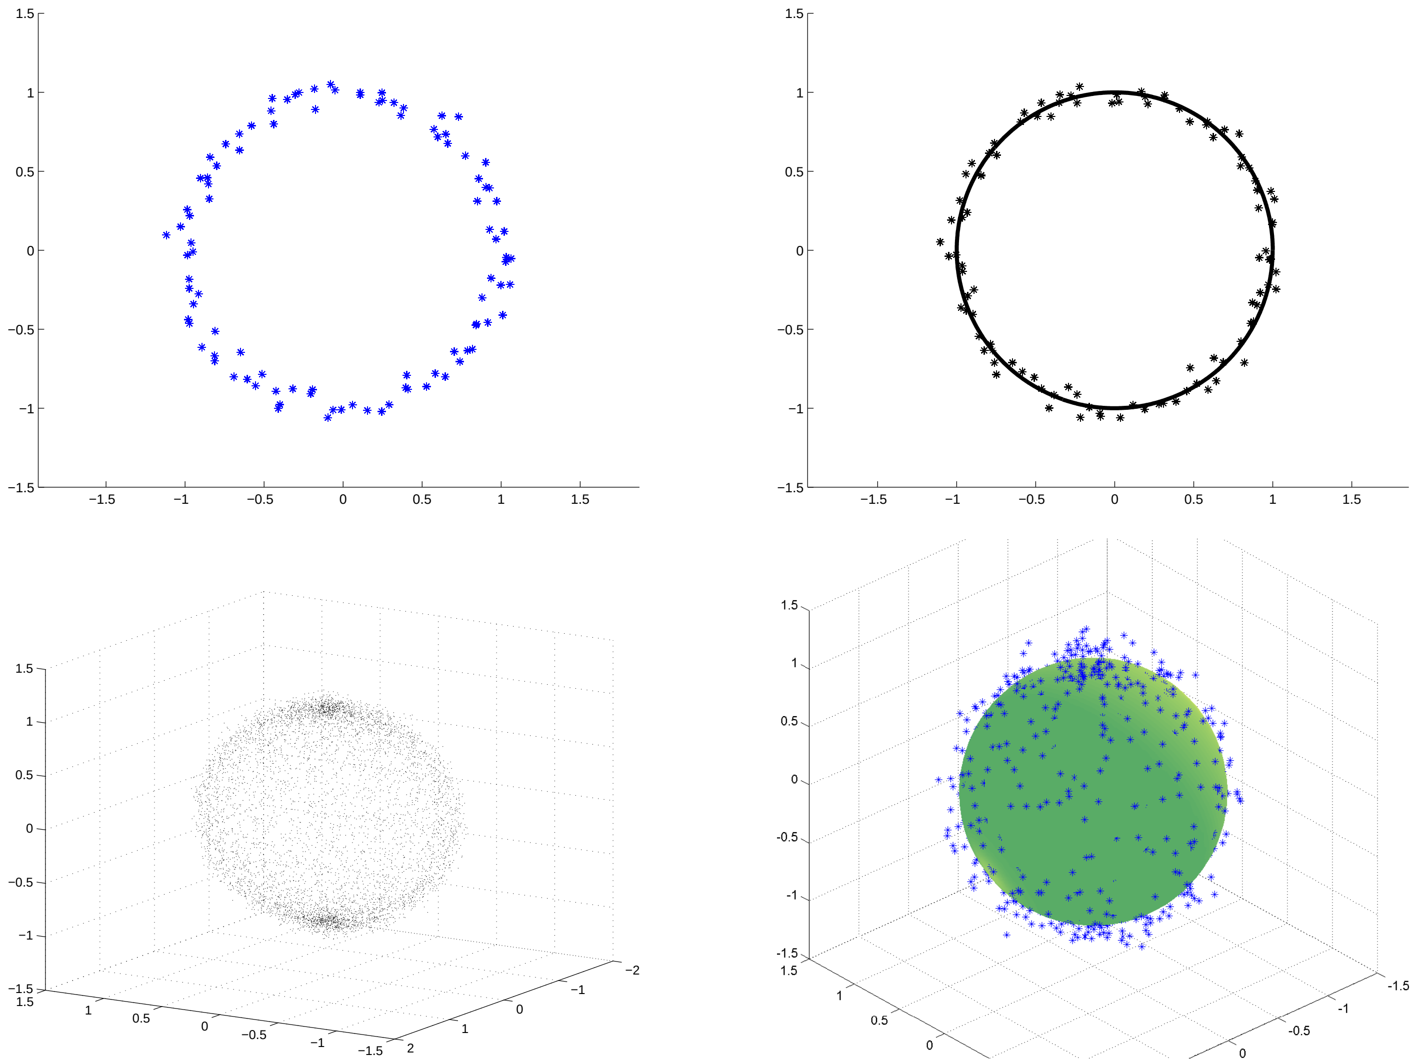

For simplicity, we only consider ***manifolds in Euclidean spaces***, but a more general approach may be taken. Also, most concepts to be discussed extend to ***infinite-dimensional manifolds***, whose relevance to shape and vision problems already has been indicated in the introduction.

Our discussion will be somewhat informal. A few **references** to more complete and general treatments:

- **W. Boothby**, *An Introduction to Differentiable Manifolds and Riemannian Geometry*, Academic Press, 2002.
- **F. Warner**, *Foundations of Differential Geometry and Lie Groups*, Graduate Texts in Mathematics, Springer-Verlag, 1983.
- **J. Lee**, *Introduction to Smooth Manifolds*, Springer-Verlag, 2002.

How to make sense of “locally similar” to an Euclidean space?

A map $\varphi\colon U \to \mathbb{R}^m$ defined on an open region $U \subseteq \mathbb{R}^n$, $n \leq m$, is said to be a ***parameterization*** if:

(i) φ is a ***smooth*** (i.e., infinitely differentiable), ***one-to-one*** mapping.

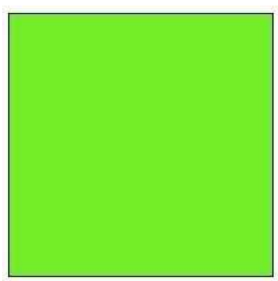

$$\xrightarrow{\varphi}$$

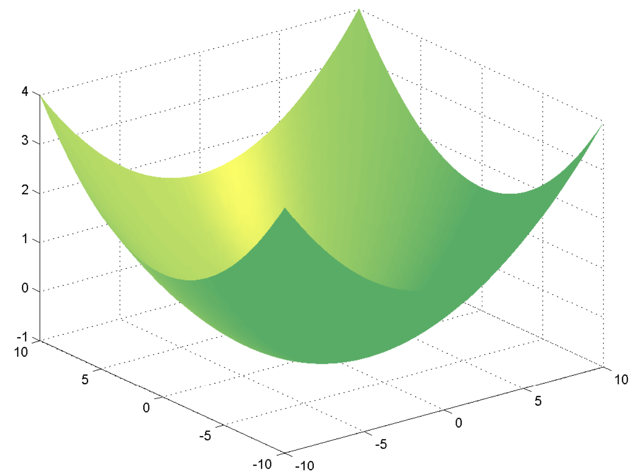

$U \subset \mathbb{R}^2$

This simply says that $V = \varphi(U)$ is produced by ***bending and stretching*** the region U in a gentle, elastic manner, ***disallowing self-intersections***.

(ii) The $m \times n$ ***Jacobian*** matrix $J(x) = \left[\frac{\partial \varphi_i}{\partial x_j}(x) \right]$ has ***rank*** $\boldsymbol{n}$, for every $x \in U$.

Here, $x = (x_1, \ldots, x_n)$ and $\varphi = (\varphi_1, \ldots, \varphi_m)$.

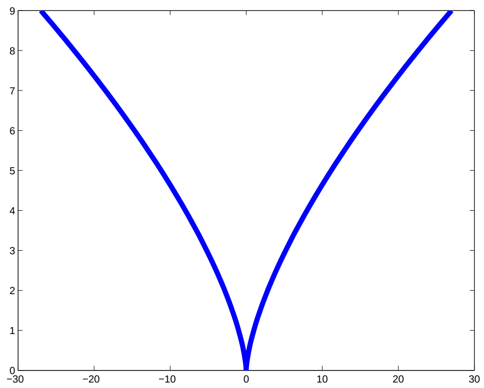

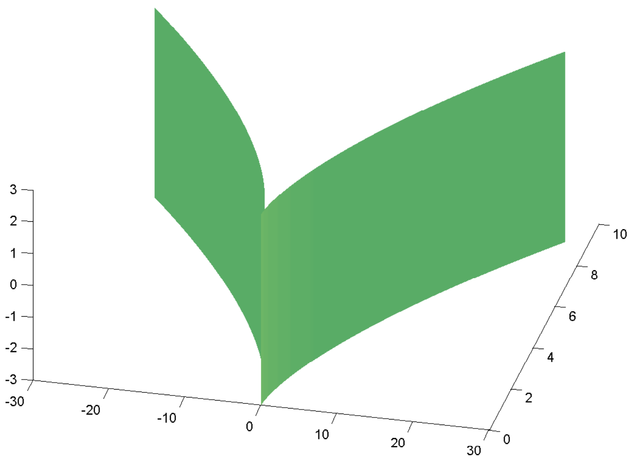

(a) $\boldsymbol{\varphi(t) = (t^3, t^2)}$

(b) $\boldsymbol{\varphi(t, s) = (t^3, t^2, s)}$

This condition further ensures that V has ***no sharp bends, corners, peaks***, or other singularities.

(ii) The ***inverse*** map $\boldsymbol{x} = \varphi^{-1}: V \to U$ is ***continuous***.

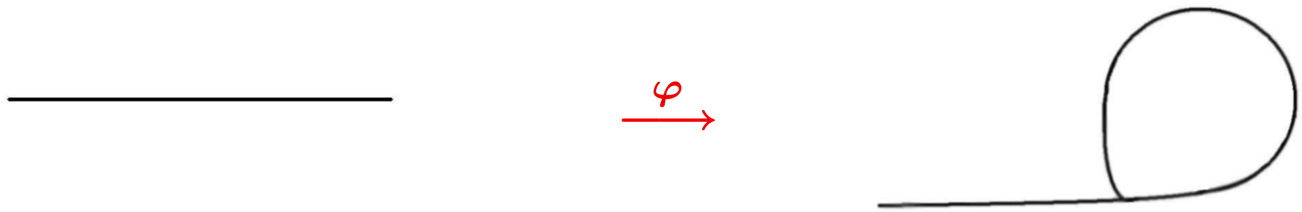

This condition has to do with the fact that we should be able to recover $\boldsymbol{U}$ from $\boldsymbol{V}$ with a ***continuous deformation***. In the illustration above, φ^{-1} is not continuous.

One often refers to φ as a ***parameterization*** of the set $\boldsymbol{V = \varphi(U)}$, and to the inverse mapping $\boldsymbol{x = \varphi^{-1}}$ as a ***local chart***, or a ***local coordinate system*** on V.

Definition of Differentiable Manifolds

A subspace $M \subseteq \mathbb{R}^m$ is an ***n-dimensional differentiable manifold*** if every point has a neighborhood that admits a parameterization defined on a region in $\mathbb{R}^n$. More precisely, for every point $x \in M$, there are a neighborhood $V_x \subseteq M$ of x, and a parameterization $\varphi\colon U \subseteq \mathbb{R}^n \to \mathbb{R}^m$, with $\varphi(U) = V_x$.

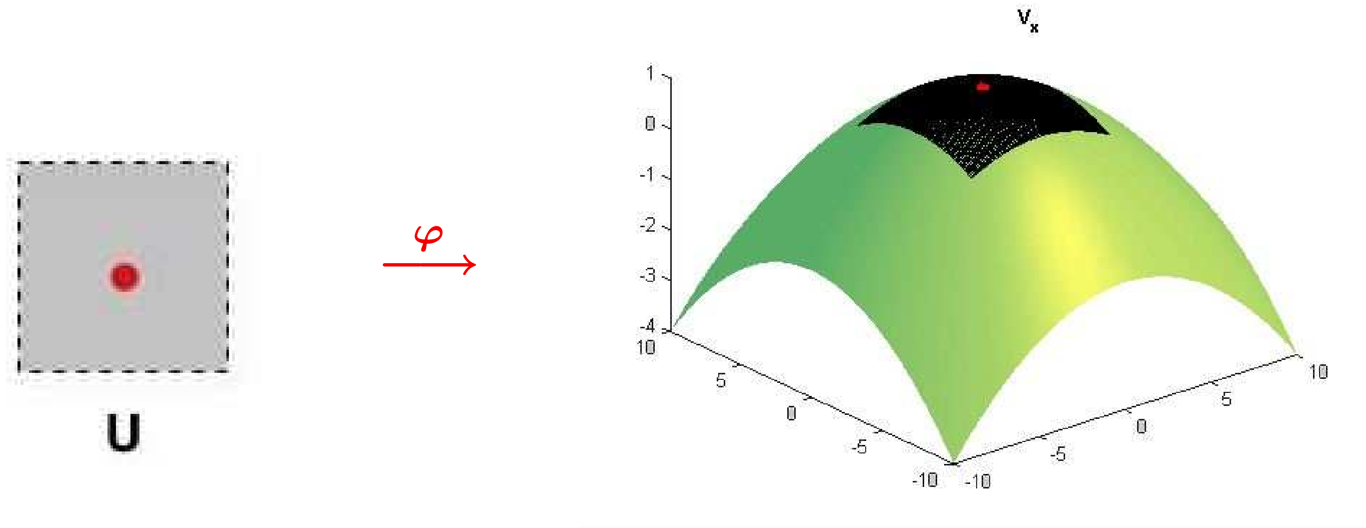

Examples of Manifolds

- **The Space of Directions in $\mathbb{R}^n$**

A direction is determined by a unit vector. Thus, this space can be identified with the unit sphere

$$\mathbb{S}^{n-1} = \{x \in \mathbb{R}^n \ : \ \|x\| = 1\}$$

Notice that $\mathbb{S}^{n-1}$ can be described as the ***level set*** $F^{-1}(1)$ of the function $F(x_1, \ldots, x_n) = x_1^2 + \ldots + x_n^2$.

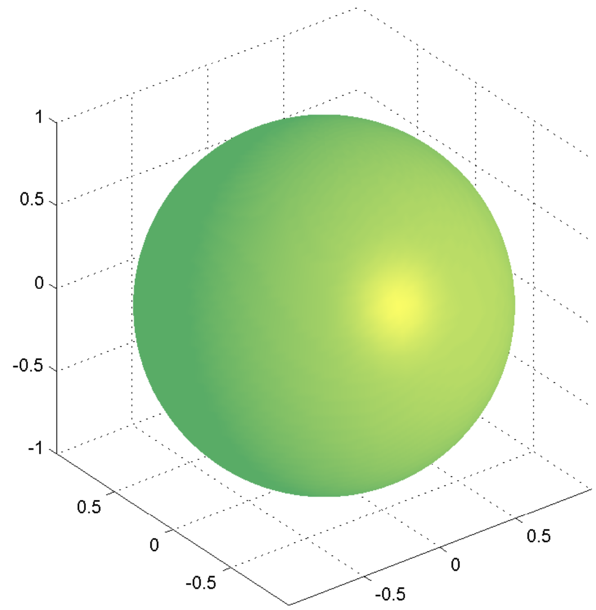

This implicit description of the sphere is a special case of the following general construction.

Let $F\colon \mathbb{R}^m \to \mathbb{R}^k$, $k < m$, be a smooth mapping, and let $a \in \mathbb{R}^k$. If the $k \times m$ **Jacobian matrix** $J(x)$ of F has **rank k**, for every $x \in F^{-1}(a)$, then

$$\boldsymbol{M = F^{-1}(a) \subset \mathbb{R}^m}$$

is an $(m-k)$-dimensional manifold. When this condition on J is satisfied, $\boldsymbol{a \in \mathbb{R}^k}$ is said to be a ***regular value*** of F.

- **The Graph of a Smooth Function**

Let $g\colon \mathbb{R}^n \to \mathbb{R}^\ell$ be smooth. Then,

$$\boldsymbol{M = \{(x,y) \in \mathbb{R}^{n+k} \,:\, y = g(x)\}}$$

is an n-dimensional manifold. Here, $F(x,y) = y - g(x)$ and $M = F^{-1}(0)$.

• The Group $O(n)$ of Orthogonal Matrices

Let $\mathcal{M}(n) \cong \mathbb{R}^{n^2}$ be the space of all $n \times n$ real matrices, and $\mathcal{S}(n) \cong \mathbb{R}^{n(n+1)/2}$ the subcollection of all symmetric matrices. Consider the map $\boldsymbol{F\colon \mathcal{M}(n) \to \mathcal{S}(n)}$ given by $F(A) = AA^t$, where the superscript t indicates transposition.

Orthogonal matrices are those satisfying $AA^t = I$. Hence,

$$\boldsymbol{O(n) = F^{-1}(I).}$$

It can be shown that I is a regular value of F. Thus, the ***group of $n \times n$ orthogonal matrices*** is a manifold of dimension

$$\boldsymbol{n^2 - \frac{n(n+1)}{2} = \frac{n(n-1)}{2}.}$$

Remark. $O(n)$ has two components formed by the orthogonal matrices with positive and negative determinants, resp. The ***positive component*** is denoted $\boldsymbol{SO(n)}$.

• Lines in $\mathbb{R}^n$ Through the Origin

A line in $\mathbb{R}^n$ through the origin is completely determined by its intersection with the unit sphere $\mathbb{S}^{n-1}$. This intersection always consists a pair of antipodal points $\{x, -x\} \subset \mathbb{S}^{n-1}$. Thus, the ***space of lines through the origin*** can be modeled on the sphere $\mathbb{S}^{n-1}$ with antipodal points identified. This is an $(n-1)$-dimensional manifold known as the ***real projective space*** $\mathbb{RP}^{n-1}$.

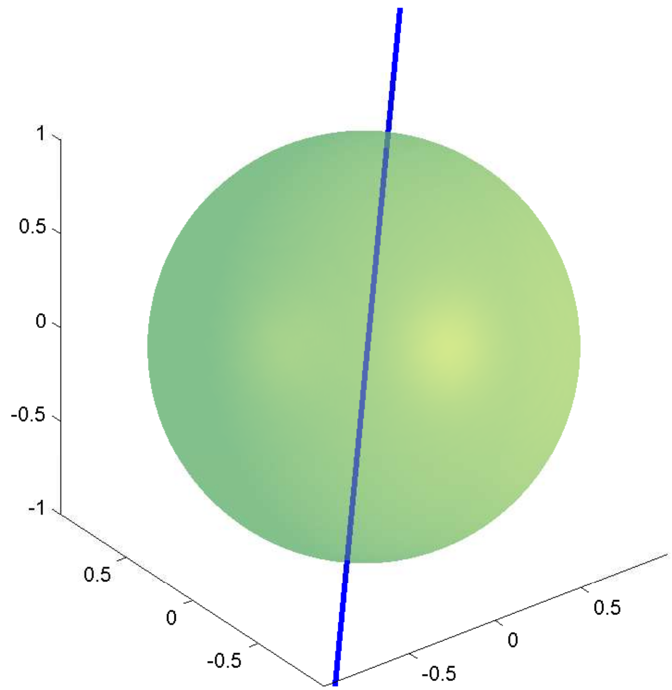

• Grassmann Manifolds

The spaces $\boldsymbol{G(n,k)}$ of all k-planes through the origin in $\mathbb{R}^n$, $k \leq n$, generalize real projective spaces. $G(n,k)$ is a manifold of dimension $k(n-k)$ and is known as a ***Grassmann manifold***.

A k-plane in $\mathbb{R}^n$ determines an $(n-k)$-plane in $\mathbb{R}^n$ (namely, its orthogonal complement), and vice-versa. Thus, $\boldsymbol{G(n,k) \cong G(n,n-k)}$.

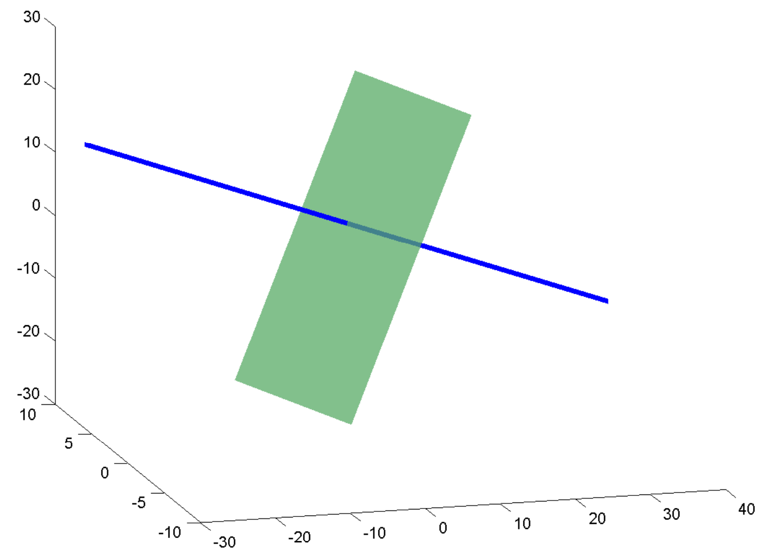

What is a Differentiable Function?

Let $f\colon M^n \subset \mathbb{R}^m \to \mathbb{R}$ be a function. Given $u_0 \in M$, let $\varphi\colon U \to \mathbb{R}^m$ be a parameterization such that $\varphi(x_0) = u_0$. The map f is said to be ***differentiable at the point u_0*** if the composite map

$$f \circ \varphi\colon U \subset \mathbb{R}^n \to \mathbb{R},$$

given by $f \circ \varphi(x) = f(\varphi(x))$, is differentiable at x_0.

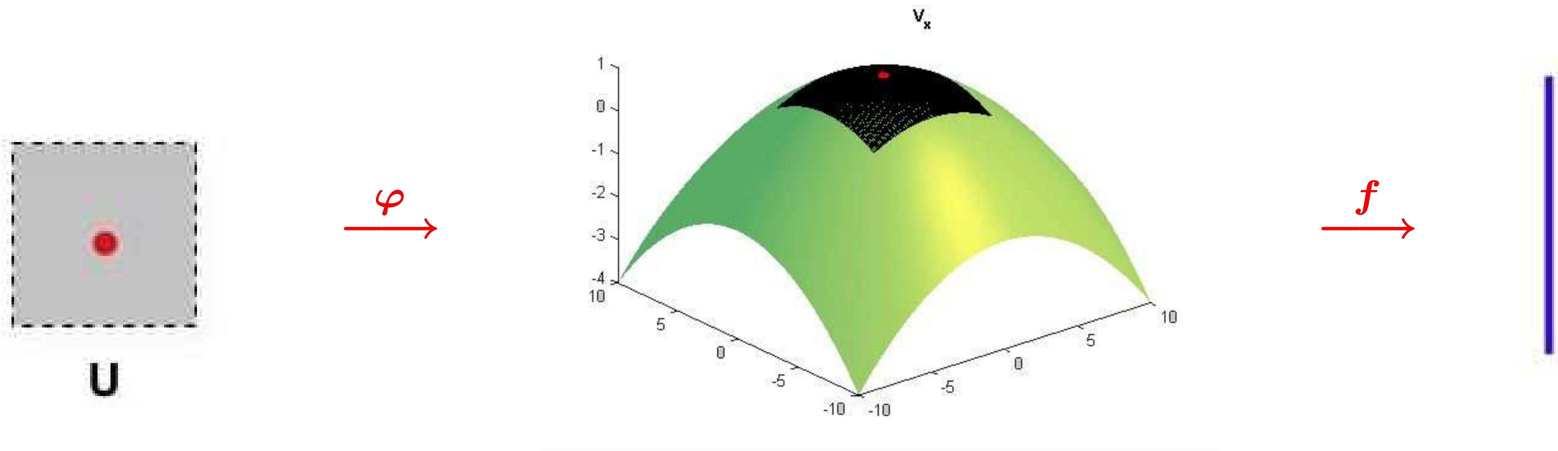

Tangent Vectors

A ***tangent vector*** to a manifold $M^n \subset \mathbb{R}^m$ at a point $p \in M$ is a vector in $\mathbb{R}^m$ that can be realized as the velocity at p of a curve in M.

More precisely, a vector $\boldsymbol{v} \in \mathbb{R}^{\boldsymbol{m}}$ is tangent to M at p if there is a smooth curve $\alpha\colon (-\varepsilon, \varepsilon) \to M \subseteq \mathbb{R}^m$ such that $\boldsymbol{\alpha(0) = p}$ and $\boldsymbol{\alpha'(0) = v}$.

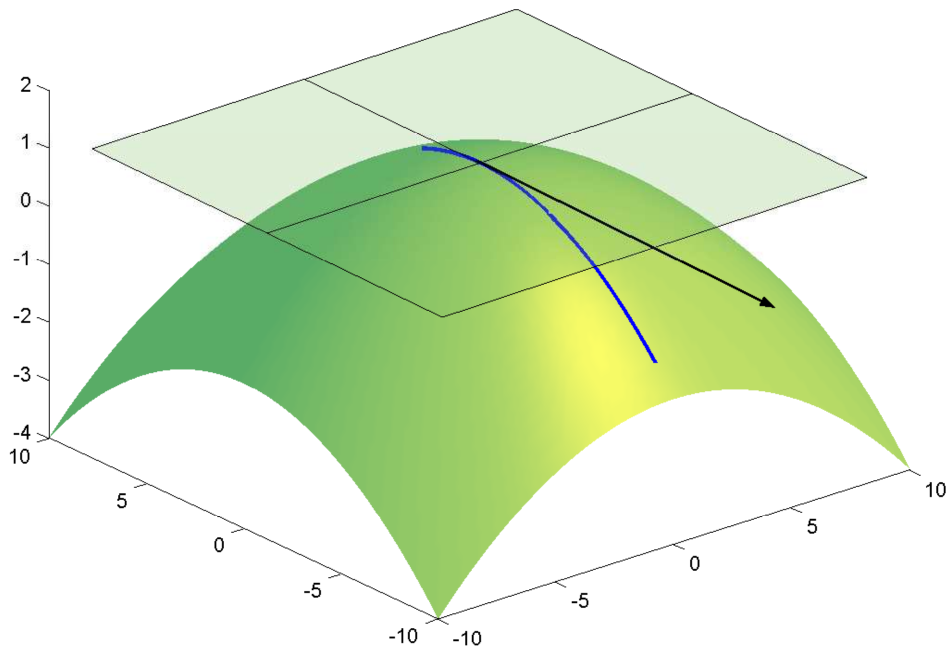

Tangent Spaces

The collection of all tangent vectors to an n-manifold $M^n \subset \mathbb{R}^m$ at a point p forms an n-dimensional real vector space denoted $\boldsymbol{T_p M}$.

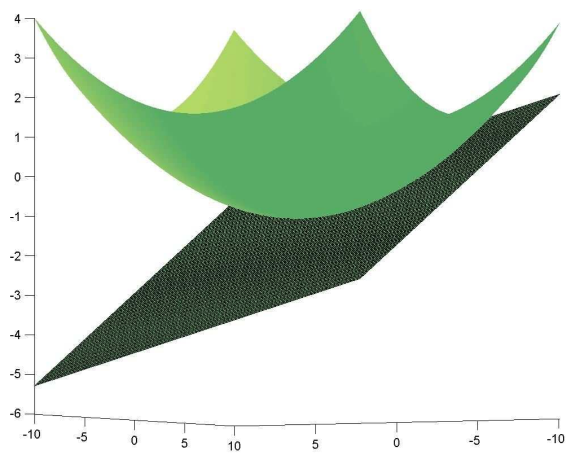

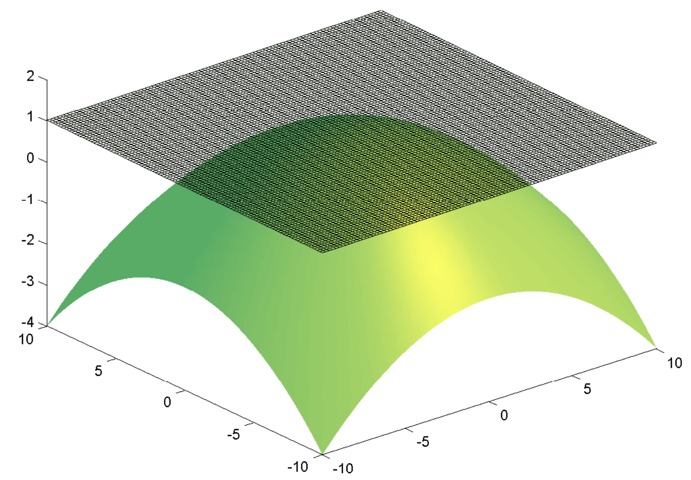

If $\varphi\colon U \subset \mathbb{R}^n \to M$ is a parameterization with $\varphi(a) = p$, then the set

$$\left\{ \frac{\partial \varphi}{\partial x_1}(a), \ldots, \frac{\partial \varphi}{\partial x_n}(a) \right\}$$

is a basis of $T_p M$.

Example

Let $F\colon \mathbb{R}^m \to \mathbb{R}$ be a differentiable function, and let $a \in \mathbb{R}$ be a regular value of F. The manifold $\boldsymbol{M = F^{-1}(a)}$ is $(n-1)$-dimensional and the tangent space to M at p consists of all vectors ***perpendicular to the gradient vector*** $\nabla F(x)$; that is,

$$\boldsymbol{T_x M = \{v \in \mathbb{R}^m \ : \ \langle \nabla F(x), v\rangle = 0\}\,.}$$

Here, $\langle\ ,\ \rangle$ denotes the usual inner (or dot) product on $\mathbb{R}^n$.

- **The Sphere**

If $F(x_1, \ldots, x_n) = x_1^2 + \ldots + x_n^2$, the level set $\mathbb{S}^{n-1} = F^{-1}(1)$ is the $(n-1)$-dimensional ***unit sphere***. In this case, $\nabla F(x) = 2(x_1, \ldots, x_n) = 2x$.

Thus, the tangent space to the $\mathbb{S}^{n-1}$ at the point x consists of vectors orthogonal to x; i.e.,

$$\boldsymbol{T_x \mathbb{S}^{n-1} = \{v \in \mathbb{R}^n \ : \ \langle x, v\rangle = 0\}\,.}$$

Derivatives

Let $f\colon M \longrightarrow \mathbb{R}$ be a differentiable function and $\boldsymbol{v \in T_p M}$ a tangent vector to M at the point p.

Realize v as the velocity vector at p of a parametric curve $\alpha\colon(-\varepsilon,\varepsilon) \longrightarrow M$. The ***derivative of f*** at p ***along v*** is the derivative of $f(\alpha(t))$ at $t=0$; that is,

$$\boldsymbol{df_p(v) = \frac{d}{dt}\left(f \circ \alpha\right)(0).}$$

The derivatives of f at $p \in M$ along vectors $v \in T_p M$ can be assembled into a single linear map

$$\boldsymbol{df_p\colon T_p M \longrightarrow \mathbb{R} \qquad v \mapsto df_p(v),}$$

referred to as the ***derivative*** of f at p.

Differential Geometry on M

Given a manifold $M^n \subset \mathbb{R}^m$, one can ***define geometric quantities*** such as the length of a curve in M, the area of a 2D region in M, and curvature.

The length of a parametric curve $\alpha\colon [0,1] \to M$ is defined in the usual way, as follows. The velocity vector $\alpha'(t)$ is a tangent vector to M at the point $\alpha(t)$. The ***speed*** of α at t is given by

$$\|\alpha'(t)\| = [\langle \alpha'(t), \alpha'(t) \rangle]^{1/2},$$

and the ***length*** of α by

$$L = \int_0^1 \|\alpha'(t)\| \, dt.$$

Geodesics

The study of geodesics is motivated, for example, by the question of finding the ***curve of minimal length*** connecting two points in a manifold M. This is important in applications because minimal-length curves can be used to define an intrinsic ***geodesic metric*** on M, which is useful in a variety of ways. In addition, geodesics provide interpolation and extrapolation techniques, and tools of statistical analysis on non-linear manifolds.

In order to avoid further technicalities, think of geodesics as length-minimizing curves in M (locally, this is always the case). A more accurate way of viewing geodesics is as ***parametric curves with zero intrinsic acceleration***.

A curve $\alpha\colon I \to M \subset \mathbb{R}^m$ is a geodesic if and only if the ***acceleration vector*** $\boldsymbol{\alpha''(t)} \in \mathbb{R}^m$ ***is orthogonal*** to $\boldsymbol{T_{\alpha(t)}M}$, for every $t \in I$.

Numerical Calculation of Geodesics

1. Geodesics with prescribed initial position and velocity

For many manifolds that arise in applications to computer vision, one can write the differential equation that governs geodesics explicitly, which can then be integrated numerically using classical methods.

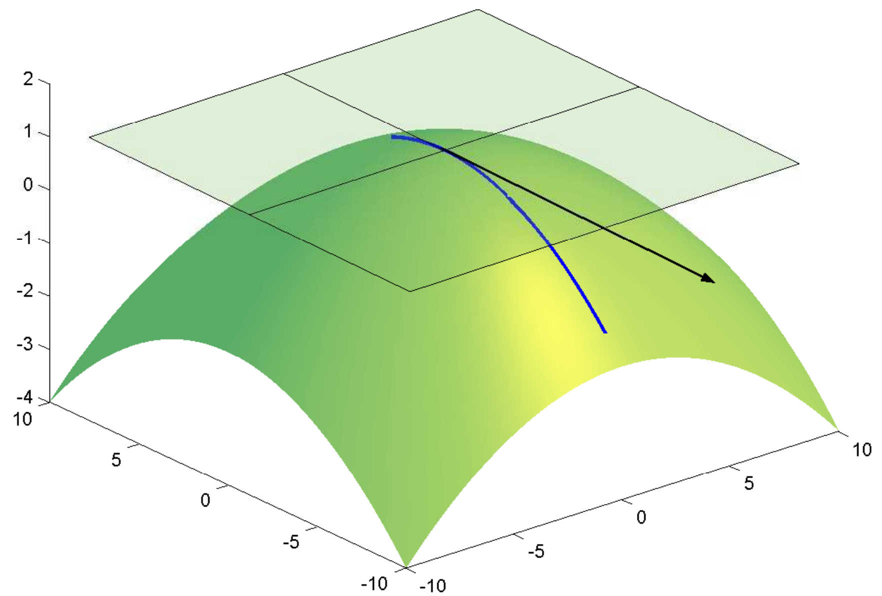

2. Geodesics between points $a, b \in M$

One approach is to use a ***shooting strategy***. The idea is to shoot a geodesic from a in a given direction $v \in T_aM$ and follow it for unit time. The terminal point $\exp_a(v) \in M$ is known as the ***exponential of*** v. If we **hit the target** (that is, $\exp(v) = b$), then the problem is solved.

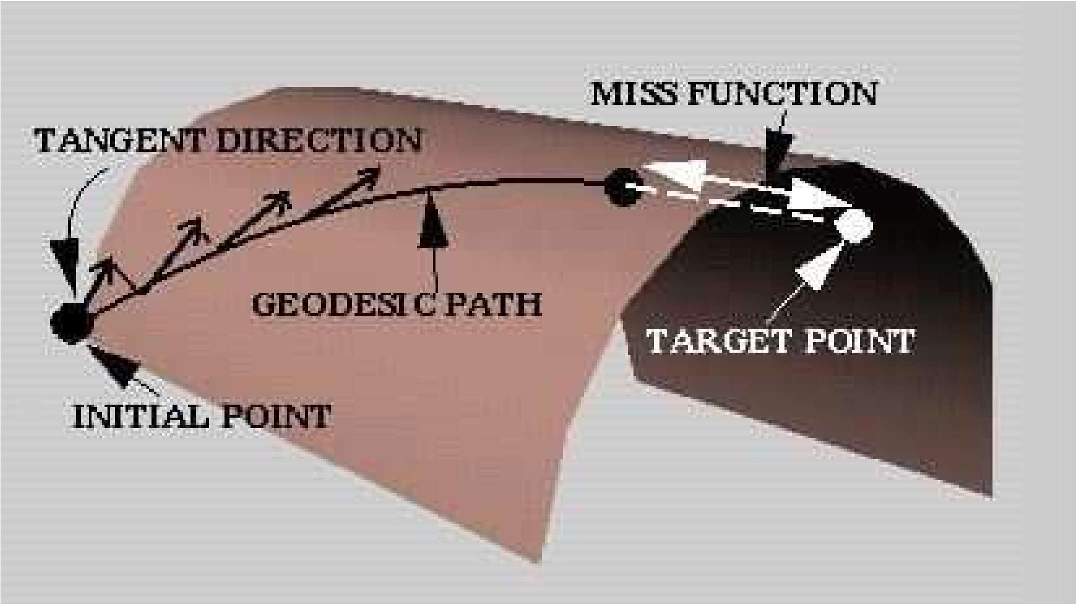

Otherwise, we consider the ***miss function*** $\boldsymbol{h(v) = \|\exp(v) - b\|^2}$, defined for $v \in T_aM$.

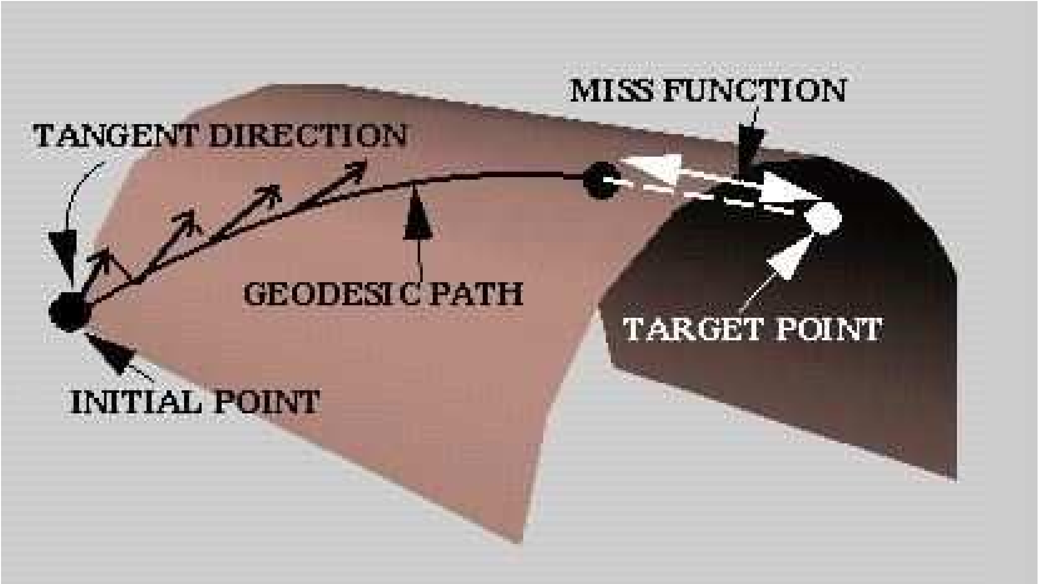

The goal is to minimize (or equivalently, annihilate) the function h. This ***optimization problem*** on the linear space T_aM can be approached with standard gradient methods.

Optimization Problems on Manifolds

How can one find ***minima*** of functions defined on nonlinear manifolds ***algorithmically***? How to carry out a ***gradient search***?

Example. Let $x_1, \dots, x_k$ be sample points in a manifold M. How to make sense of the ***mean*** of these points?

The arithmetic mean $\mu = (x_1 + \dots + x_k)/k$ used for data in linear spaces does not generalize to manifolds in an obvious manner. However, a simple calculation shows that μ can be viewed as the ***minimum*** of the ***total variance function***

$$V(x) = \frac{1}{2} \sum_{i=1}^{k} \|x - x_i\|^2.$$

This minimization problem can be posed in more general manifolds using the geodesic metric.

Given $x_1, \dots, x_k \in M$, consider the total variance function

$$V(x) = \frac{1}{2}\sum_{i=1}^{k} d^2(x, x_i),$$

where d denotes geodesic distance. A ***Karcher mean*** of the sample is a local minimum $\mu \in M$ of V.

Back to Optimization

Gradient search strategies used for functions on $\mathbb{R}^n$ can be adapted to ***find minima*** of functions $f\colon M \to \mathbb{R}$ defined on nonlinear manifolds. For example, the calculation of Karcher means can be approached with this technique. We discuss gradients next.

What is the Gradient of a Function on M?

The ***gradient*** of a function $F\colon M \to \mathbb{R}$ at a point $p \in M$ is the unique vector $\nabla_M F(p) \in T_pM$ such that

$$dF_p(v) = \langle \nabla_M F(p), v \rangle \, ,$$

for every $v \in T_pM$. ***How can it be computed in practice?*** Oftentimes, F is defined not only on M, but on the larger space $\mathbb{R}^m$ in which M is embedded. In this case, one first calculates

$$\nabla F(p) = \left(\frac{\partial F}{\partial x_1}, \ldots, \frac{\partial F}{\partial x_1} \right) \in \mathbb{R}^m.$$

The orthogonal projection of this vector onto T_pM gives $\nabla_M F(p)$.

Gradient Search

At each $p \in M$, calculate $\nabla_M f(p) \in T_pM$. The goal is to search for minima of f by integrating the negative gradient vector field $-\boldsymbol{\nabla_M f}$ on M.

- **Initialize** the search at a point $p \in M$.
- **Update** p by infinitesimally **following the unique geodesic** starting at p with initial velocity $-\nabla_M f(p)$.
- **Iterate** the procedure.

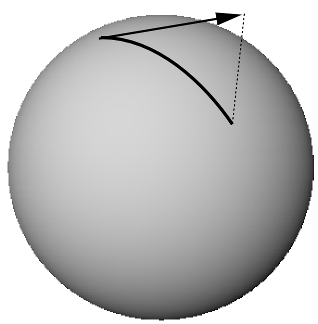

Remark. The usual convergence issues associated with gradient methods on $\mathbb{R}^n$ arise and can be dealt with in a similar manner.

Calculation of Karcher Means

If $\{x_1, \ldots, x_n\}$ are sample points in a manifold M, the negative gradient of the total variance function V is

$$\nabla_M V(x) = v_1(x) + \ldots + v_n(x),$$

where $v_i(x)$ is the initial velocity of the geodesic that connects x to x_i in unit time.

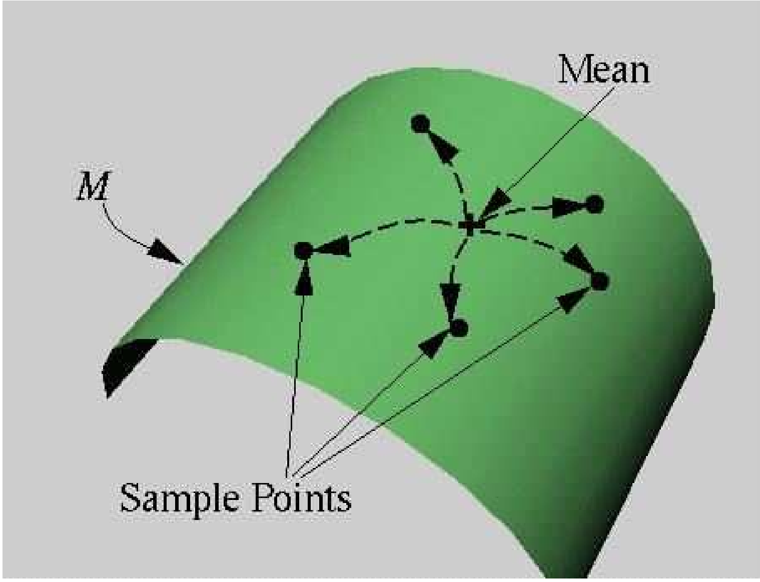

Applications to the Analysis of Planar Shapes

To illustrate the techniques discussed, we apply them to the study of shapes of planar contours. We consider two different approaches, namely:

- The ***Procrustean Shape Analysis*** of Bookstein and Kendall.
- The ***Parametric-Curves Model*** of Klassen, Srivastava, Mio and Joshi.

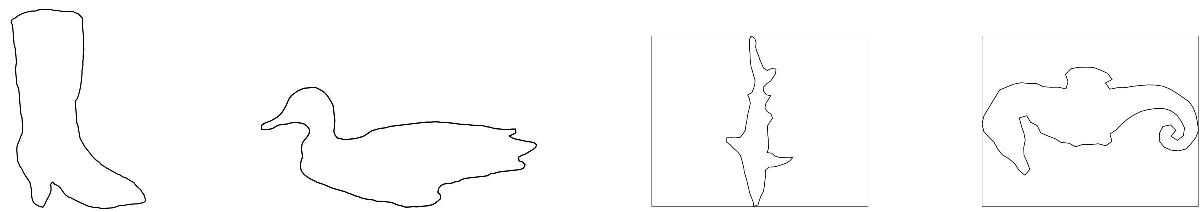

In both cases, a ***shape*** is viewed as an ***element of a shape space***. ***Geodesics*** are used to quantify shape similarities and dissimilarities, interpolate and extrapolate shapes, and develop a statistical theory of shapes.

Procrustean Shape Analysis

A contour is described by a finite, ordered sequence of ***landmark points*** in the plane, say, $p_1, \ldots, p_n$. Contours that differ by ***translations, rotations and uniform scalings*** of the plane are to be viewed as having the same shape. Hence, shape representations should be insensitive to these transformations.

If μ be the centroid of the given points, the vector

$$x = (p_1 - \mu, \ldots, p_n - \mu) \in \mathbb{R}^{2n}.$$

is invariant to translations of the contour. This representation places the centroid at the origin. To account for uniform scaling, we normalize x to have unit length. Thus, we adopt the preliminary representation

$$y = x / \|x\| \in \mathbb{S}^{2n-1}.$$

In this y-representation, rotational effects reduce to ***rotations about the origin***. If $y = (y_1, \ldots, y_n)$, in complex notation, a $\boldsymbol{\theta}$-rotation of $y \in \mathbb{S}^{2n-1}$ about the origin is given by

$$\boldsymbol{y \mapsto \lambda y = (\lambda y_1, \ldots, \lambda y_n),}$$

where $\lambda = e^{j\theta}$, where $j = \sqrt{-1}$.

In other words, vectors $\boldsymbol{y \in \mathbb{S}^{2n-1}}$ that differ by a unit complex scalar represent the same shape. Thus, the ***Procrustean shape space*** is the space obtained from $\mathbb{S}^{2n-1}$ by identifying points that differ by multiplication by a unit complex number. It is manifold of dimension $(2n-2)$ known as the ***complex projective space*** $\mathbb{C}\boldsymbol{P(n-1)}$.

Geodesics in $\mathbb{C}P(n-1)$ can be used to study shapes quantitatively and develop a statistical theory of shapes. This will be discussed in more detail in Part III.

Examples of Procrustean Geodesics

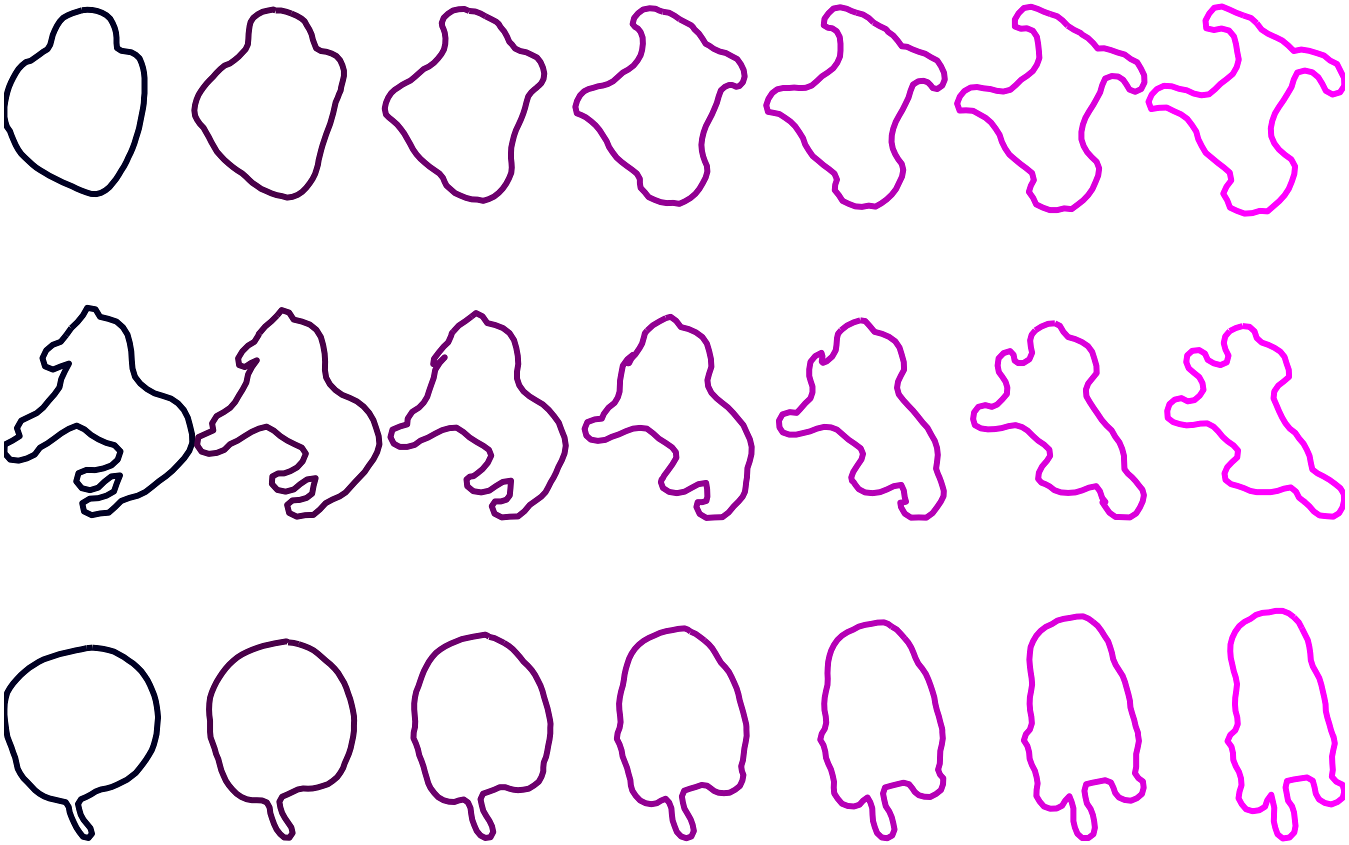

Parametric-Curves Approach

As a preliminary representation, think of a planar shape as a parametric curve $\boldsymbol{\alpha : I \longrightarrow \mathbb{R}^2}$ traversed with constant speed, where $I = [0, 1]$.

To make the representation invariant to uniform scaling, fix the length of α to be 1. This is equivalent to assuming that $\boldsymbol{\|\alpha'(s)\| = 1}$, $\forall s \in I$.

Since α is traversed with unit speed, we can write $\boldsymbol{\alpha'(s) = e^{j\theta(s)}}$, where $j = \sqrt{-1}$.

A function $\theta : I \longrightarrow \mathbb{R}$ with this property is called an ***angle function*** for α. Angle functions are insensitive to translations, and the effect of a rotation is to add a constant to θ.

To obtain shape representations that are ***insensitive to rigid motions*** of the plane, we fix the average of θ to be, say, π. In other words, we choose representatives that satisfy the constraint

$$\boldsymbol{\int_0^1 \theta(s)\, ds = \pi.}$$

We are only interested in closed curves. The ***closure condition***

$$\alpha(1) - \alpha(0) = \int_0^1 \alpha'(s)\, ds = \int_0^1 e^{j\theta(s)}\, ds = 0$$

is equivalent to the real conditions

$$\boldsymbol{\int_0^1 \cos\theta(s)\, ds = 0} \quad \text{and} \quad \boldsymbol{\int_0^1 \sin\theta(s)\, ds = 0}.$$

These are nonlinear conditions on θ, which will make the geometry of the shape space we consider interesting.

Let $\mathcal{C}$ the collection of all θ satisfying the three constraints above. We refer to an element of $\mathcal{C}$ as a pre-shape.

- Why call $\theta \in \mathcal{C}$ a **pre-shape instead of shape**?

This is because a shape may admit multiple representations in $\mathcal{C}$ due to the choices in the placement of the initial point ($s = 0$) on the curve. For each shape, there is a circle worth of initial points. Thus, our ***shape space*** $\mathcal{S}$ is the quotient of $\mathcal{C}$ by the action of the re-parametrization group $\mathbb{S}^1$; i.e., the space

$$\mathcal{S} = \mathcal{C} / \mathbb{S}^1,$$

obtained by identifying all pre-shapes that differ by a reparameterization.

The manifold $\mathcal{C}$ and the shape space $\mathcal{S}$ are ***infinite dimensional***. Although these spaces can be analyzed with the techniques we have discussed, we consider finite-dimensional approximations for implementation purposes.

We discretize an angle function $\theta\colon I \longrightarrow \mathbb{R}$ using a uniform sampling

$$\boldsymbol{x = (x_1, \ldots, x_m) \in \mathbb{R}^m.}$$

The ***three conditions*** on θ that define $\mathcal{C}$ can be rephrased as:

$$\int_0^1 \theta(s)\, ds = \pi \quad \Longleftrightarrow \quad \frac{1}{m}\sum_{i=1}^m x_i = \pi$$

$$\int_0^1 \cos\theta(s)\, ds = 0 \quad \Longleftrightarrow \quad \sum_{i=1}^m \cos(x_i) = 0$$

$$\int_0^1 \sin\theta(s)\, ds = 0 \quad \Longleftrightarrow \quad \sum_{i=1}^m \sin(x_i) = 0$$

Thus, the ***finite-dimensional analogue*** of $\mathcal{C}$ is a manifold $\mathcal{C}_m \subset \mathbb{R}^m$ of dimension $(m-3)$. Modulo 2π, a reparameterization of θ corresponds to a ***cyclic permutation*** of $(x_1, \ldots, x_m)$, followed by an adjustment to ensure that the first of the three conditions is satisfied. The quotient space by reparameterizations is the finite-dimensional version of the shape space $\mathcal{S}$.

Examples of Geodesics in $\mathcal{S}$

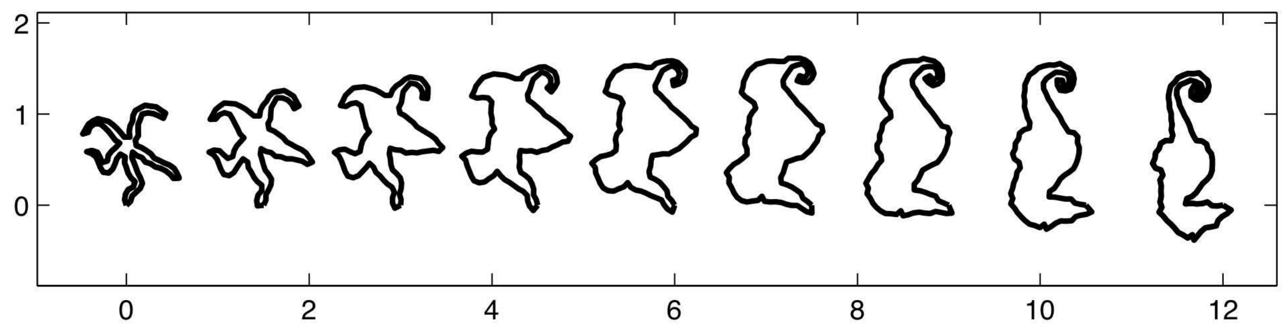

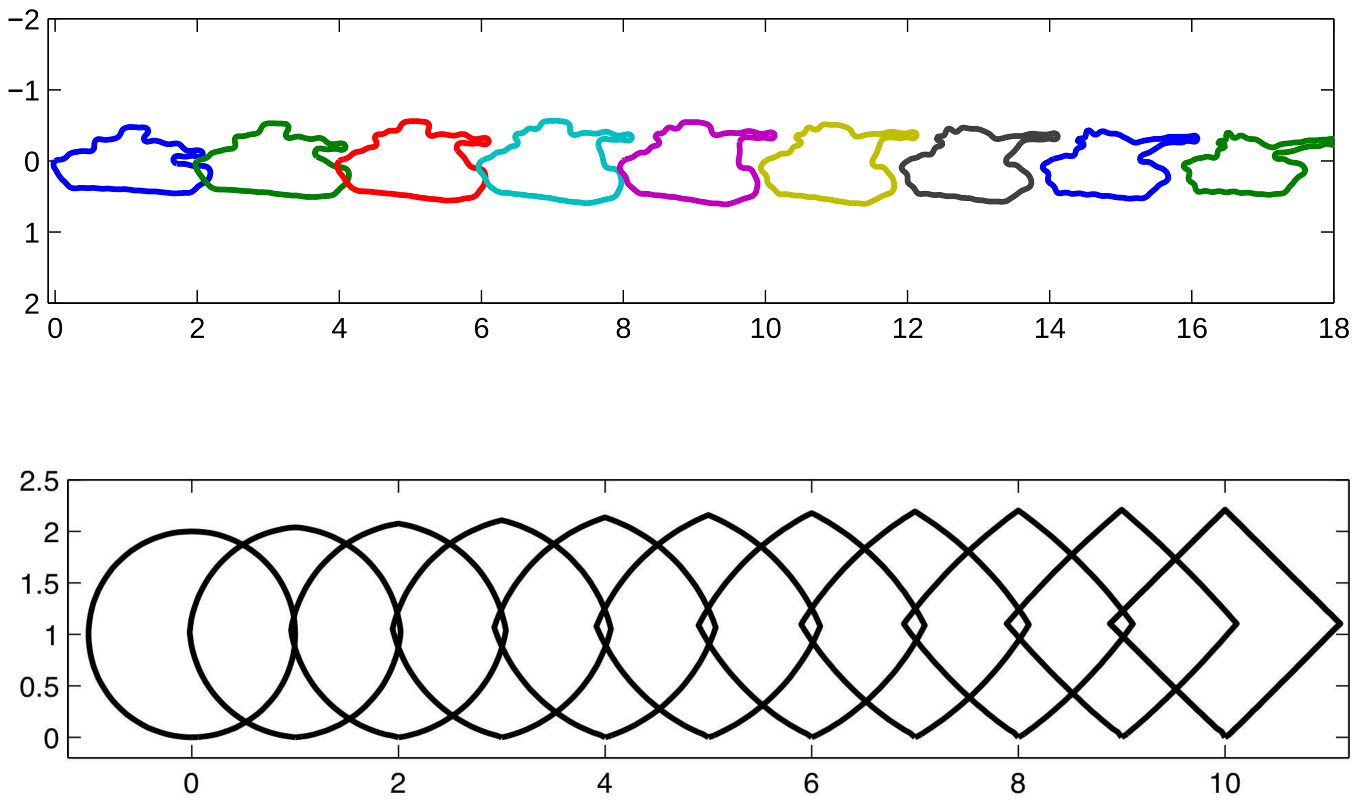
−2
−1
0
1
2
0 2 4 6 8 10 12 14 16 18
2.5
2
1.5
1
0.5
0
0 2 4 6 8 10