Geometric Learning

Motivations and Background

Credits:
Bronstein et al, “Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges”, 2021

Rodolà, “Geometric Deep Learning”, 2020

1

WHY GEOMETRIC LEARNING?

Geometry (non-Euclidean)

Pixels (Euclidean)

2

WHY GEOMETRIC LEARNING?

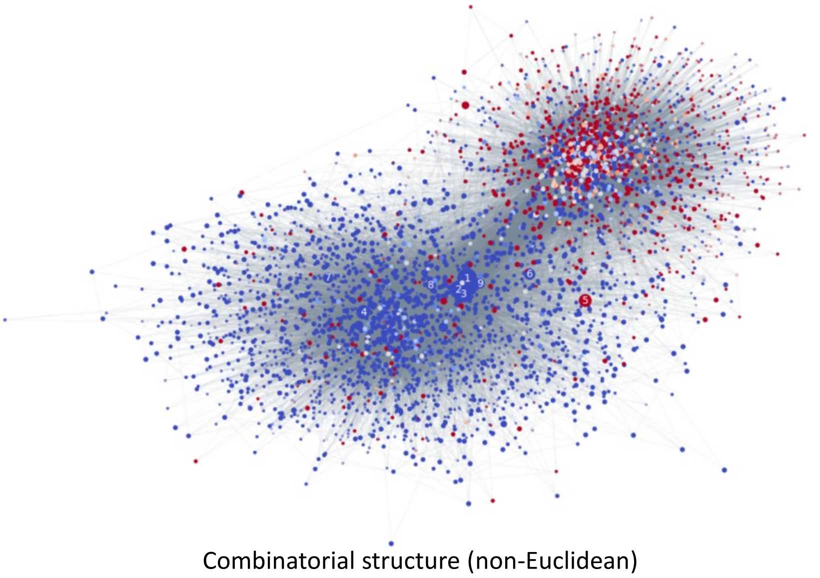

Combinatorial structure (non-Euclidean)

PRELIMINARIES: REPRESENTATION

What is a shape?

"There can be no such thing as a mathematical theory of shape. The very notion of shape belongs to the natural sciences."[1]

For us, shapes are mathematical objects (specifically, manifolds) that can be modeled in the continuous setting and brought to the digital world by translating to the discrete setting.

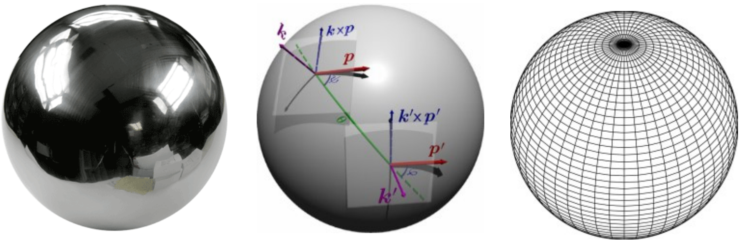

[1]J. Koenderink, "Solid Shape". MIT Press 1990

5

5

Other non-Euclidean objects

Geometric Learning is not just about shapes.

- Graphs
- Point clouds
- Splines
- Tet-meshes

However, with shapes we touch upon several key notions that are useful for understanding all other non-Euclidean structures.

- Manifold hypothesis
- Laplacian and other operators
- Functional spaces
- Parametrization vs. embedding
- …

6

6

FUNCTIONS OVER $\mathbb{R}$

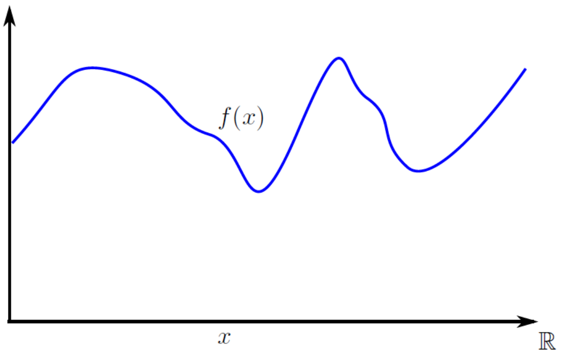

$f: \mathbb{R} \mapsto \mathbb{R}$

7

7

FUNCTIONS OVER $\mathbb{R}^2$

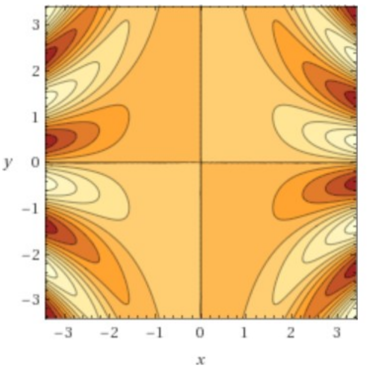

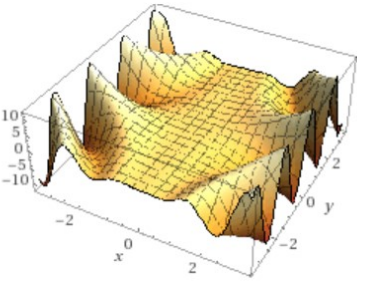

$f: \mathbb{R}^2 \mapsto \mathbb{R}$
$f: (x, y) \mapsto x^2 \sin(\mathrm{xy})$

8

8

Functions over a surface

$$f: \mathcal{M} \mapsto \mathbb{R}^3$$

9

9

Functions over a graph

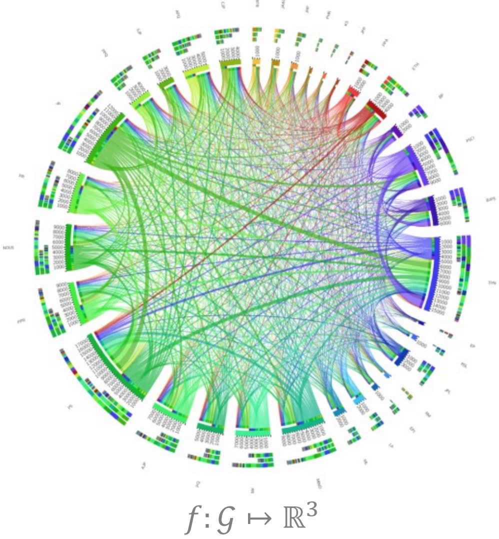

$$f: \mathcal{G} \mapsto \mathbb{R}^3$$

10

10

SHAPES VS IMAGES: DOMAIN

Euclidean (flat)

non-Euclidean (curved)

11

11

SHAPES VS IMAGES: REPRESENTATION

Array of pixels
(uniform grid)

Splines

Graph

Point cloud

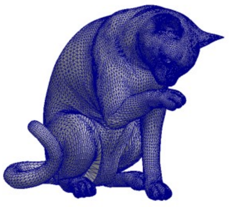

Triangle mesh

12

12

SHAPES VS IMAGES: PARAMETRIZATION

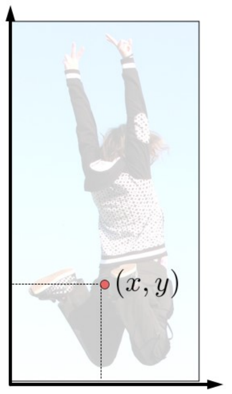

Global

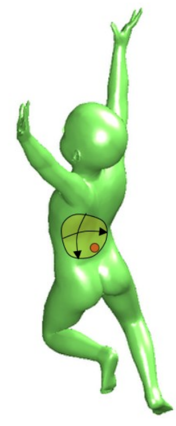

Local

13

13

SHAPES VS IMAGES: TRANSFORMATIONS

Perspective

Affine

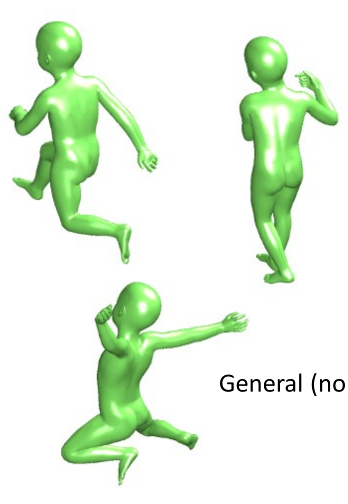

General (non rigid)

14

14

SHAPES VS IMAGES: CALCULUS

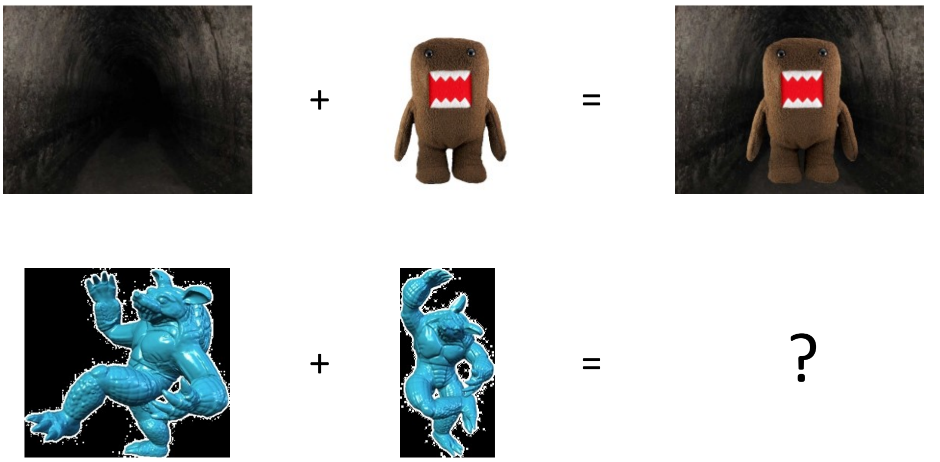

15

15

SHAPE REPRESENTATION

Popular representations:

- Triangle meshes
- Point clouds
- Implicit surfaces

Others exist:

- Polygonal meshes (polygons are not restricted to triangles)
- Parametric surfaces
- Voxel grids (Euclidean)
- ---

The choice of a representation is crucial and can make the success or failure of solving a task.

16

16

EXAMPLE

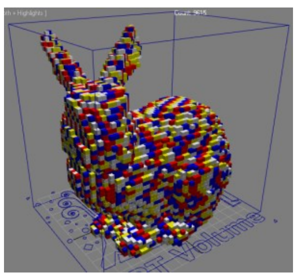

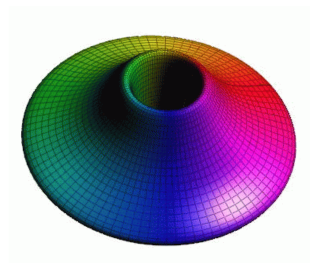

17

SHAPE REPRESENTATION: TRIANGLE MESH

A triangle mesh is a collection of connected triangles.

The incidence relations of triangles defines the mesh connectivity (also referred to as mesh topology).

In this example, the same underlying surface is discretized with meshes having different connectivity:

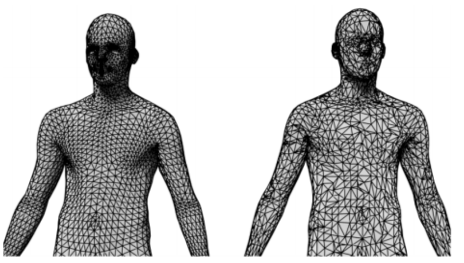

18

SHAPE REPRESENTATION: TRIANGLE MESH

A triangle mesh is a collection of connected triangles.

We will only consider oriented manifold meshes.

- Each triangle has an orientation
- All triangles should be consistently oriented (e.g., counter-clockwise)
- Each triangle has a normal (i.e., orthogonal) vector consistent with the triangle orientation

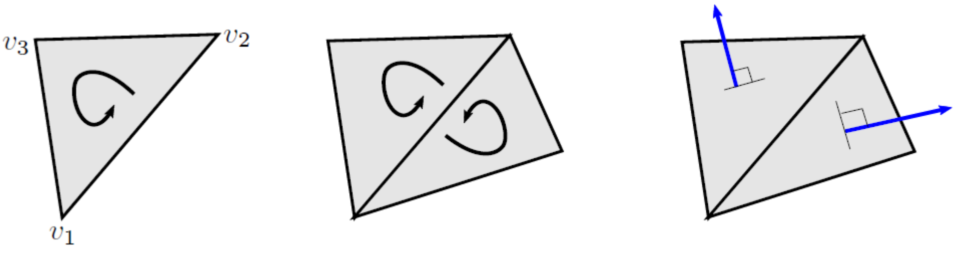

19

19

SHAPE REPRESENTATION: TRIANGLE MESH

We will only consider oriented manifold meshes.

- All edges have at most two incident triangles
- Edges with only one incident triangle form the mesh boundary
- The faces incident to a vertex form a closed or an open fan

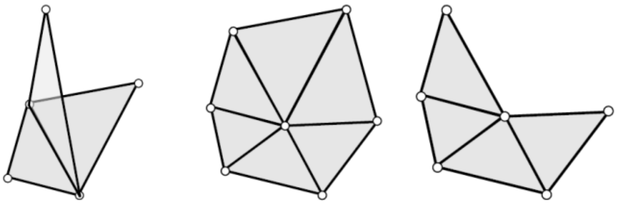

3 triangles incident to 1 edge
⇒ non-manifold

closed fan

open fan

20

20

EXAMPLE: NON-MANIFOLD MESHES

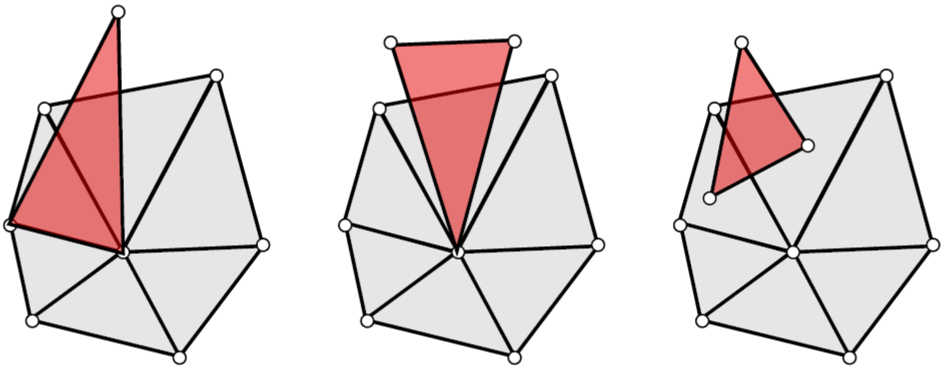

NORMALS

Normals are unit vectors that are orthogonal to each face

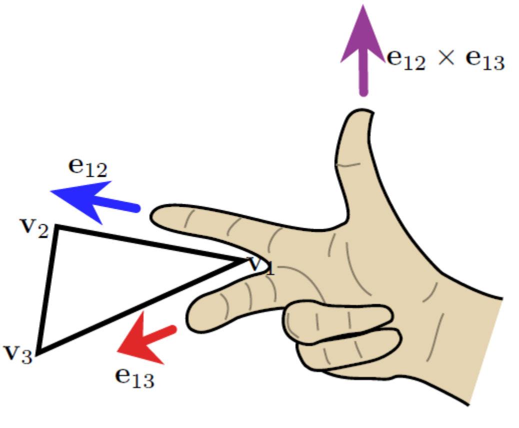

Normals

Italian right-hand rule

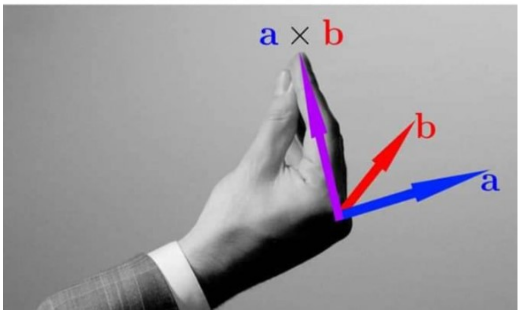

$$\hat{n} = \frac{e_{12} \times e_{13}}{\|e_{12} \times e_{13}\|}$$

23

23

Shape representation: Point cloud

A point cloud is a collection of points in 3D space.

- An oriented point cloud also has a normal vector for each point
- Point clouds are interpreted as point-wise samplings of an underlying unknown continuous surface...
- ...in practice, they come from depth sensors and can be very noisy!

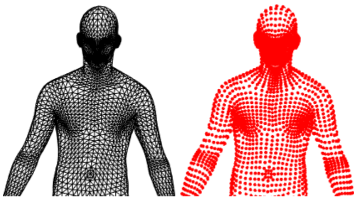

Synthetic point cloud (obtained by removing mesh connectivity)

Real-world Kinect scan

24

24

IMAGE REPRESENTATION

We represent images as uniform grids of pixels.
Each pixel is associated with a number or a color vector.

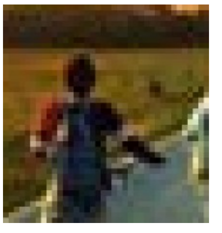

25

25

THE CHALLENGE

26

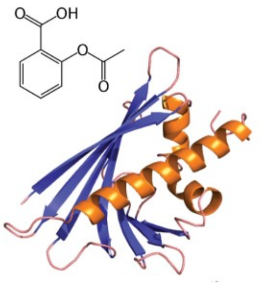

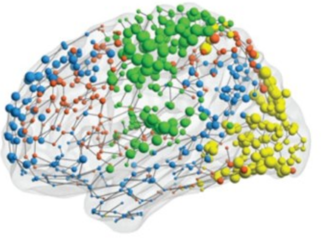

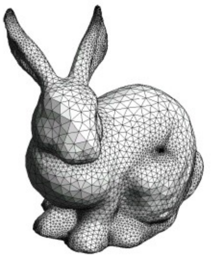

27

Applications of geometric learning

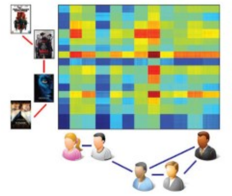

Recommender system

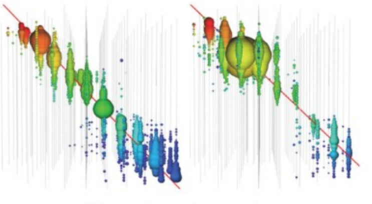

Neutrino detection

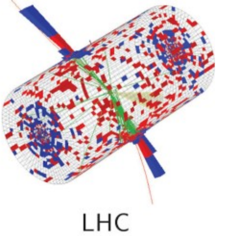

LHC

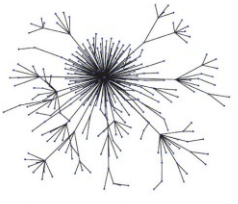

Fake news detection

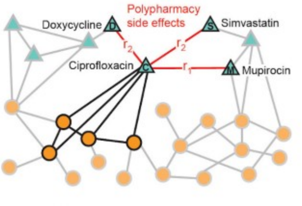

Drug repurposing

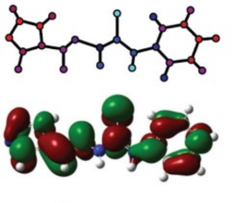

Chemistry

28

APPLICATION: PROTEIN-PROTEIN INTERACTION

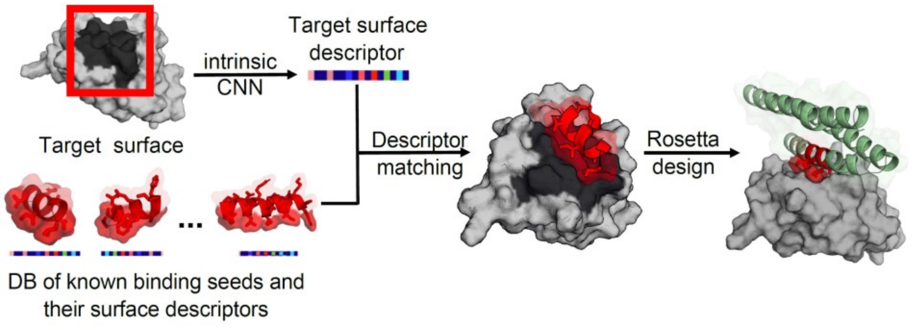

Designing protein binder for the PD-L1 protein target

Gainza et al, “Deciphering interaction fingerprints from protein molecular surfaces using geometric deep learning", Nature Methods 2020

29

29

MOLECULE PROPERTY PREDICTION

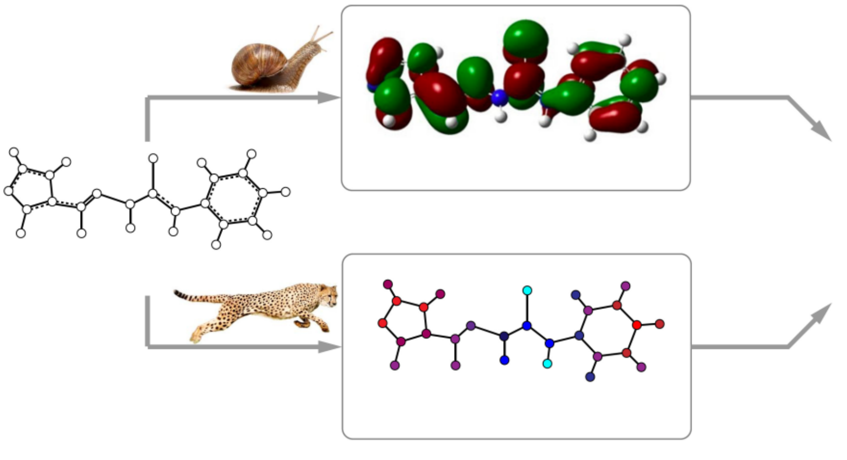

Duvenaud et al, “Convolutional Networks on Graphs for Learning Molecular Fingerprints", NIPS 2015; Gomez-Bombarelli et al, “Automatic chemical design using a data-driven continuous representation of molecules", ACS Cent. Sci. 2018

30

30

GENERATIVE MODELS

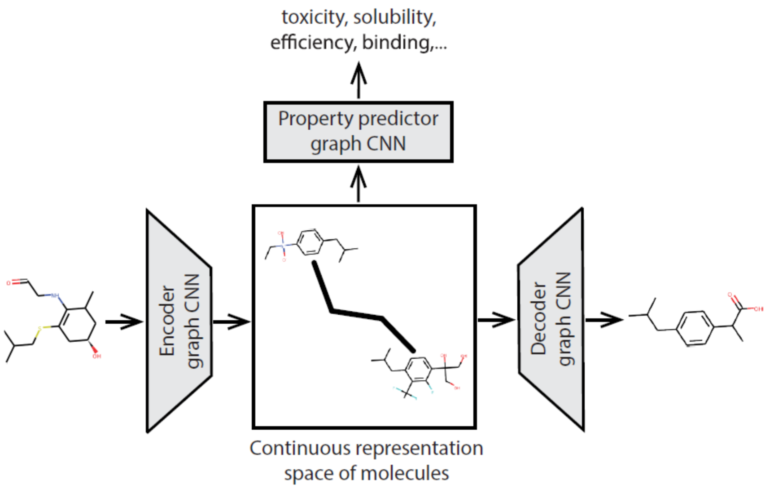

31

31

MOLECULE GENERATION

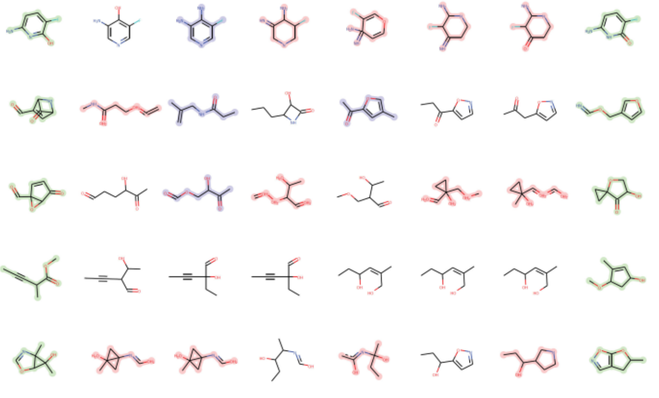

Molecules generated with a graph VAE

Simonovsky and Komodakis, “Graphvae: Towards generation of small graphs using variational autoencoders", 2017
De Cao and Kipf, “MolGAN: An implicit generative model for small molecular graphs", 2018

32

32

Face from DNA

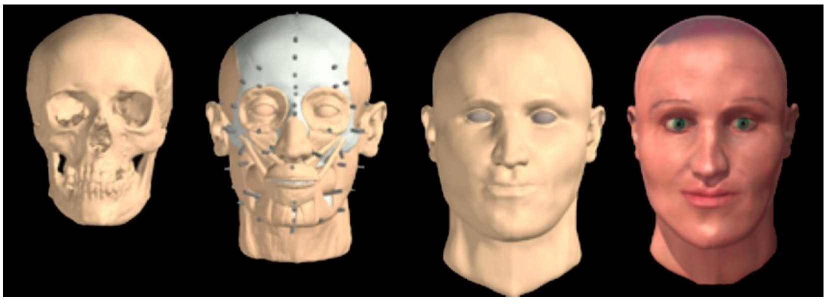

Claes et al, "Facial recognition from DNA using face-to-DNA classifiers", Nature Communications 2019

33

33

Prototypical non-Euclidean objects

Manifolds

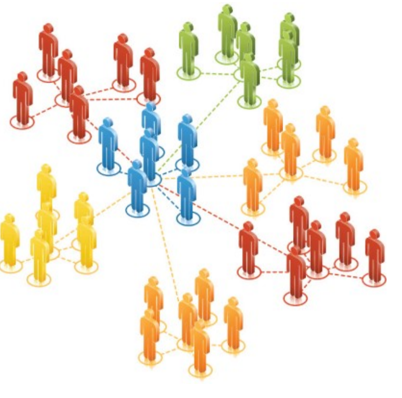

Graphs

34

34

DOMAIN STRUCTURE VS DATA ON A DOMAIN

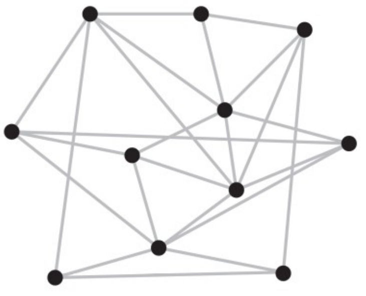

Domain structure

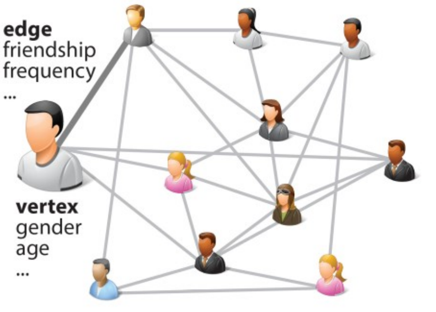

Data on a domain

35

35

DOMAIN STRUCTURE VS DATA ON A DOMAIN

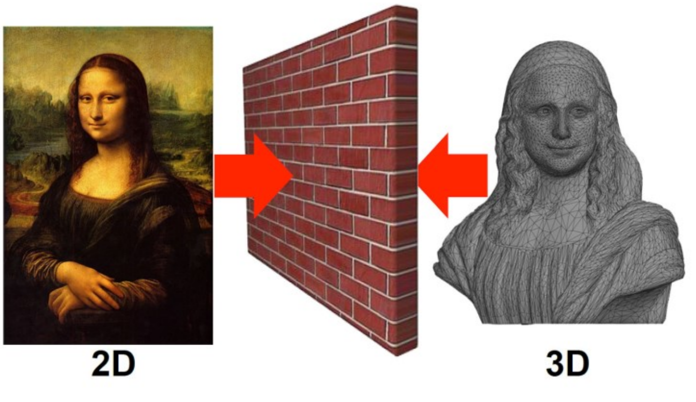

only data

only structure

36

36

FIXED VS DIFFERENT DOMAIN

Social network (fixed graph)

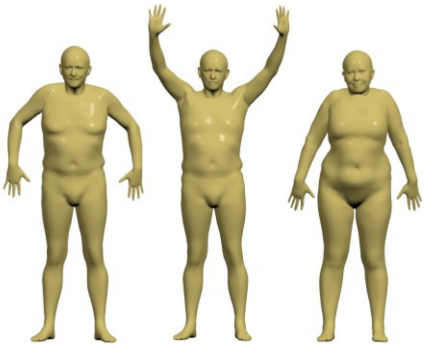

3D shapes (different manifolds)

37

37

GEOMETRIC LEARNING ≠ MANIFOLD LEARNING

In manifold learning, we seek for a (possibly high-dimensional) manifold that justifies a given set of data points:

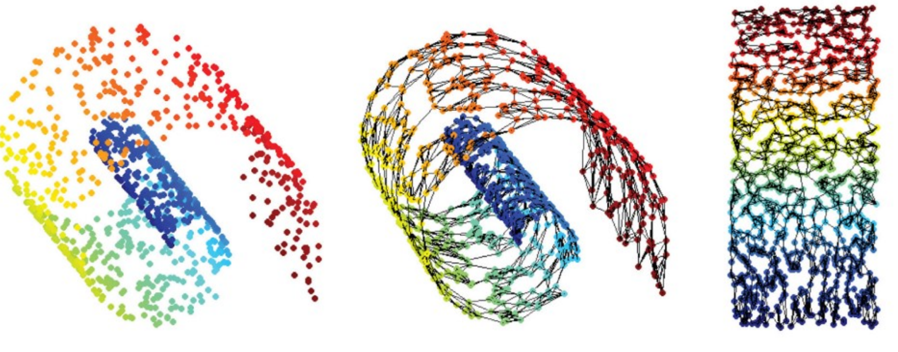

In geometric learning, the data has a known geometric structure.

38

38

...do we actually need geometric learning?

MULTI-VIEW CNNS

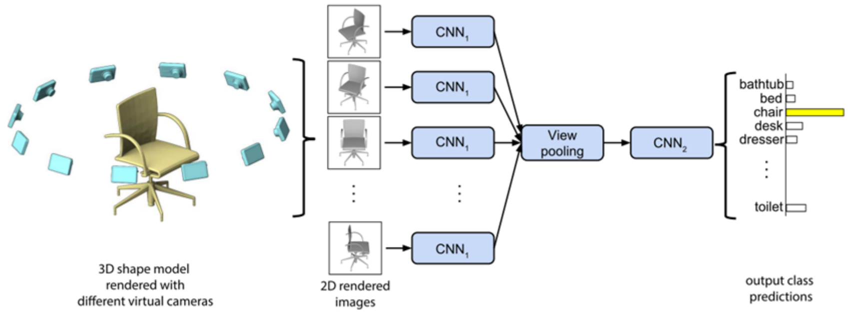

- Represent a 3D object as a collection of range images
- CNN_1: Extract image features (parameters are shared across views)
- Element-wise max pooling across all views
- CNN_2: Produce shape descriptors + final prediction

Su et al, "Multi-view Convolutional Neural Networks for 3D Shape Recognition", 2015

APPLICATIONS OF MULTI-VIEW CNNS

3D shape classification and retrieval

- Pre-trained on ImageNet
- Fine-tuned on 2D views

Sketch classification

- Mimic views by jittering

Sketch-based shape retrieval

- Render views with hand-drawn style (edge maps)

Su et al, "Multi-view Convolutional Neural Networks for 3D Shape Recognition", 2015

41

41

3D SHAPENETS

Volumetric representation (shape = binary voxels on 3D grid)

3D convolutional network

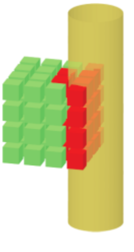

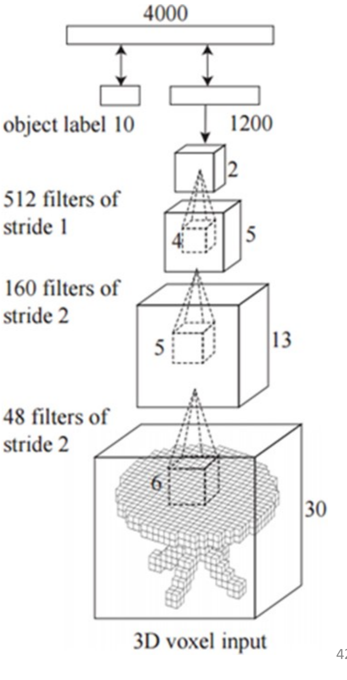

Wu et al, "3D ShapeNets: A Deep Representation for Volumetric Shapes", 2015

42

42

LEARNED FEATURES: 3D PRIMITIVES

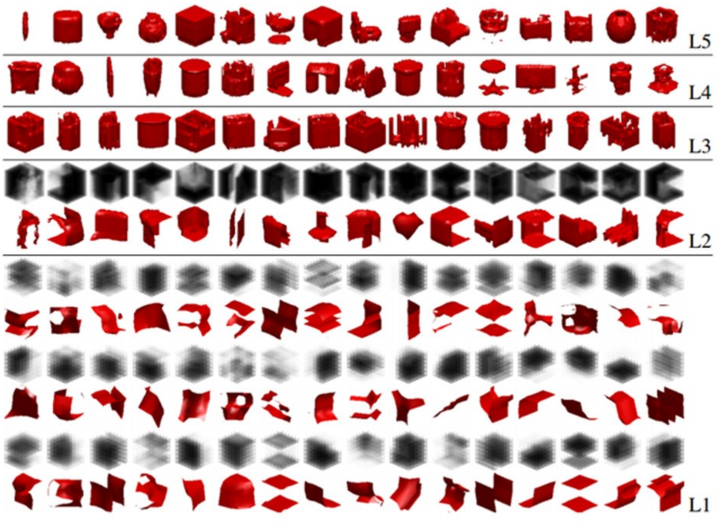

Wu et al, "3D ShapeNets: A Deep Representation for Volumetric Shapes", 2015

43

43

CHALLENGES OF GEOMETRIC LEARNING

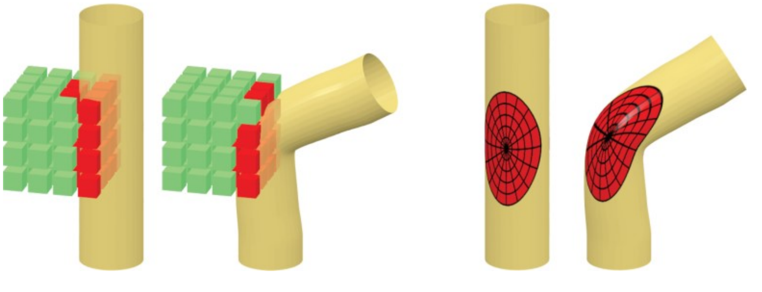

Extrinsic Intrinsic

- How to define convolution?
- How to do pooling?
- How to work fast?

44

44

Extrinsic vs Intrinsic

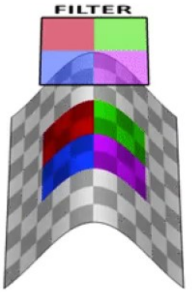

Extrinsic

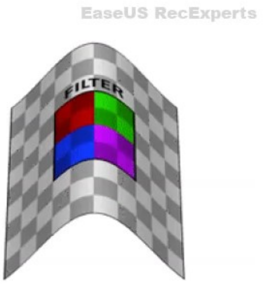

Intrinsic

45

45

Non-Euclidean convolution?

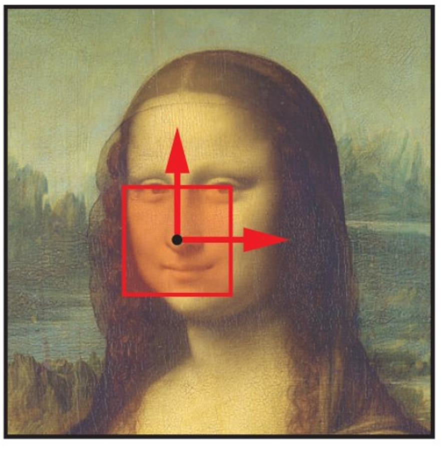

Euclidean

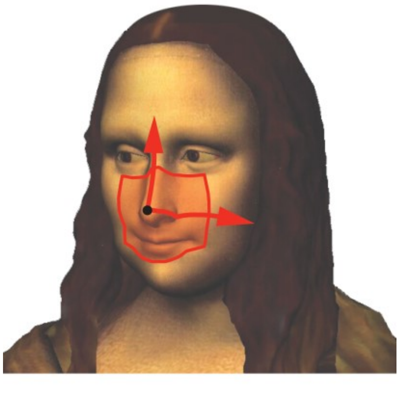

Non-Euclidean

46

46

NON-EUCLIDEAN CONVOLUTION?

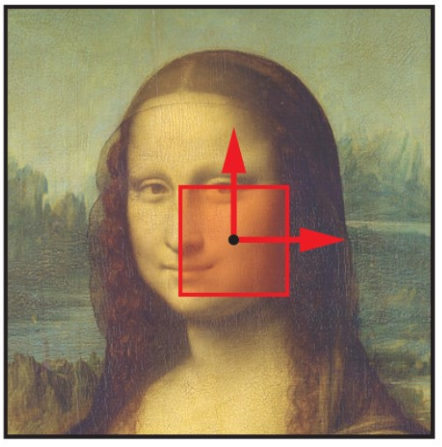

Euclidean

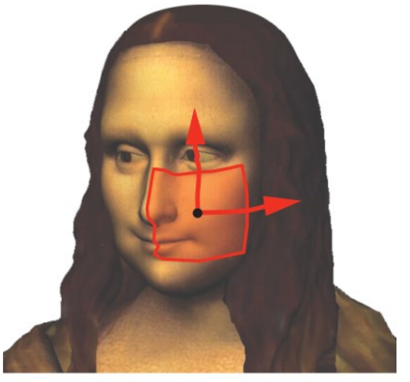

Non-Euclidean

47

47

NON-EUCLIDEAN CONVOLUTION?

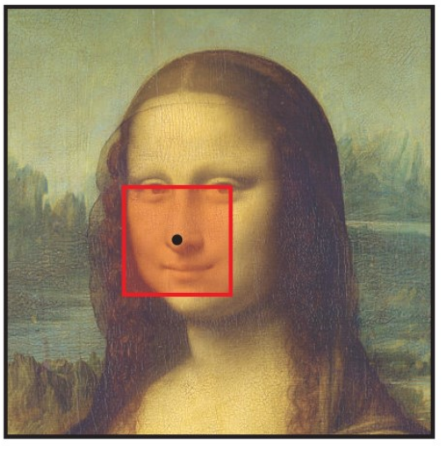

Image

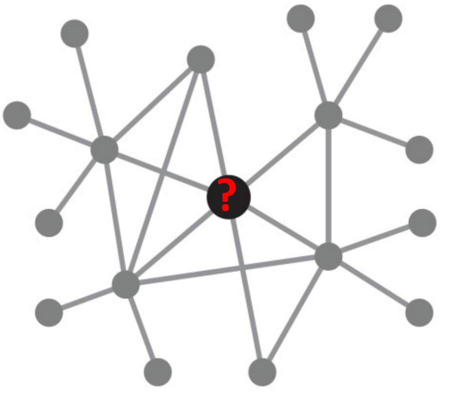

Graph

48

48

Global parametrization

Map the input mesh to some parametric domain (e.g., 2D plane) where operations can be defined more easily.

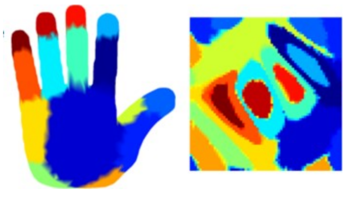

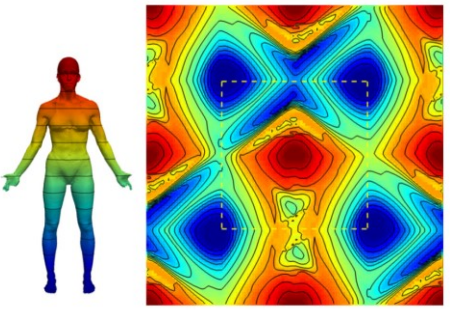

- Can use Euclidean techniques while staying intrinsic
- Provides invariance to certain transformations
- Parametrization may be non-unique
- The map can introduce distortion

Sinha et al, "Deep learning 3d shape surfaces using geometry images", 2016

49

Intrinsic convolution on surfaces

Is shift-invariant convolution on surfaces even possible?

Not in general! Due to singularities in the translation field (Poincare-Hopf or "hairy ball" theorem): "*there is no continuous non-zero vector field tangent to a sphere*".

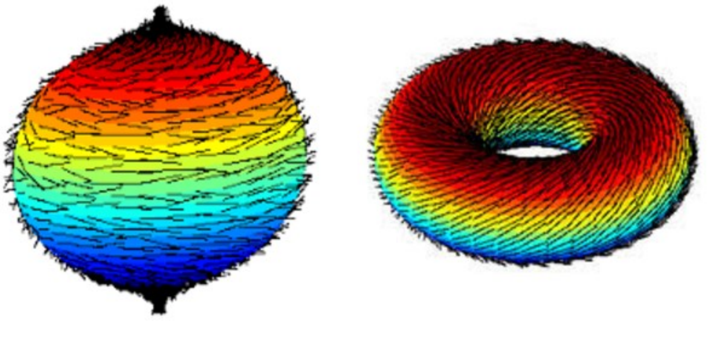

Heuristically said: "it is not possible to completely comb a hairy ball"

50

INTRINSIC CONVOLUTION ON SURFACES

The torus is the only closed orientable surface admitting a translational group.

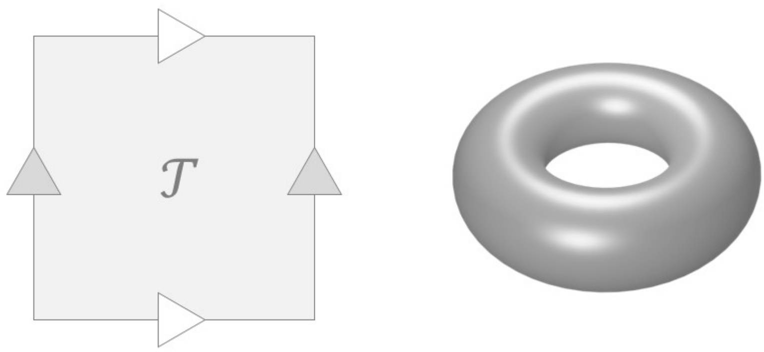

Maron et al, "Convolutional Neural Networks on Surfaces via Seamless Toric Covers", SIGGRAPH 2017

51

51

CONVOLUTION ON SURFACES

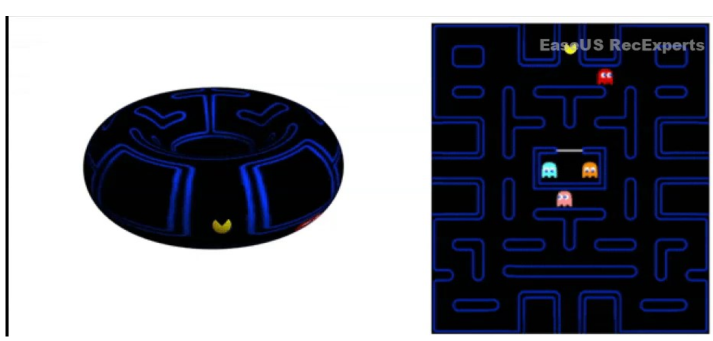

52

52

PRELIMINARIES: DIFFERENTIAL GEOMETRY

53

Prototypical non-Euclidean objects

Manifolds

Graphs

54

54

DIFFERENTIAL GEOMETRY

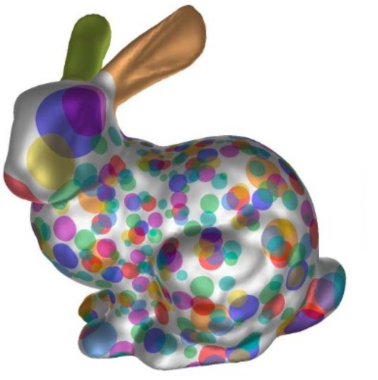

The study of local properties of curves and surfaces.
Each neighborhood has a well-behaved mapping to some subset $\mathcal{U} \subset \mathbb{R}^2$.

DIFFERENTIAL GEOMETRY

Differential geometry gives us powerful tools to compute lengths, areas, integrals, gradients, etc. on surfaces.

MANIFOLDS

Manifolds are unions of charts:

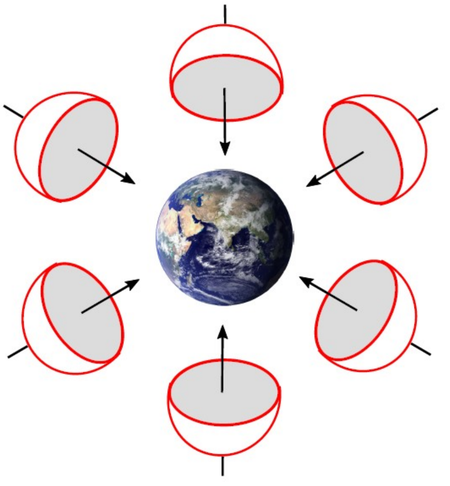

57

2D MANIFOLDS (SURFACES)

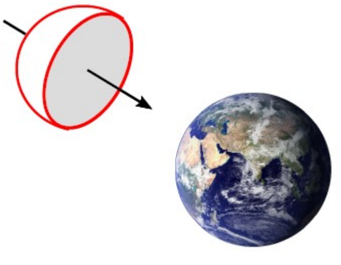

chart

Each chart can be seen as a mapping $\phi: \mathbb{R}^2 \to S \subset \mathbb{R}^3$.

We require ϕ to be smooth and invertible (diffeomorphism).

- The domain of ϕ is the parametric space and is Euclidean.
- The image of ϕ is the embedding and is a surface.

58

2D MANIFOLDS (SURFACES)

Recipe for a regular surface in $\mathbb{R}^3$:

- Cut pieces of a plane
- Deform these pieces
- Glue them together in a shape so that there are no sharp points, edges, or self-intersections (regularity)

59

59

2D MANIFOLDS (SURFACES)

Regularity ensures that we can talk about tangent planes at each point

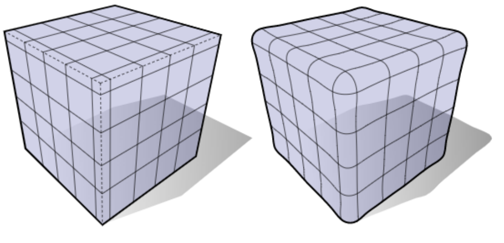

In the language of differential geometry:

"L2-dimensional Riemannian sub-manifold"

60

60

MANIFOLDS

Manifolds can be k-dimensional, meaning that we have charts:

$$\Phi: \mathbb{R}^k \to \mathcal{M} \subset \mathbb{R}^d \quad with \quad k < d$$

The parametrization is not unique:

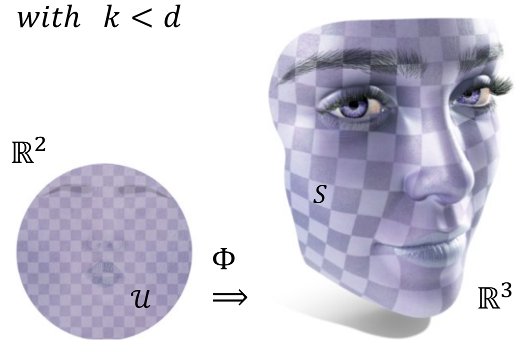

61

61

MANIFOLDS

Manifolds can be k-dimensional, meaning that we have charts:

$$\Phi: \mathbb{R}^k \to \mathcal{M} \subset \mathbb{R}^d \quad with \quad k < d$$

The parametrization is not unique:

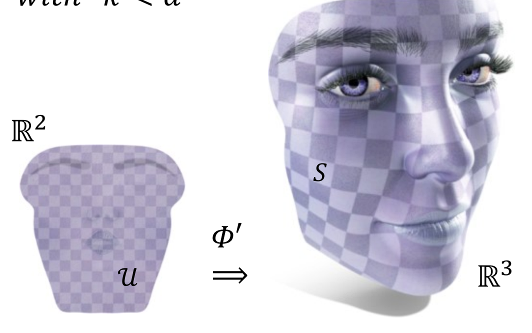

However, all encode exactly the same geometric information.

62

62

PRELIMINARIES: METRIC GEOMETRY

63

MEASURING DISTANCE

Working with curved surfaces rather than flat domains requires us to reconsider all the basic notions that we took for granted in high school geometry.

How do you measure distance between x and y in this picture?

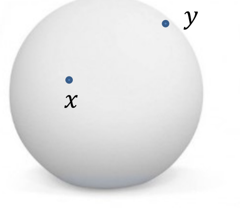

There is not a unique way!

- You can pass through the sphere with a straight line (Euclidean)
- You can walk on the surface in a “straight" path (non-Euclidean)

64

64

EUCLIDEAN DISTANCE

The Euclidean distance measures the length of a straight line connecting two points:

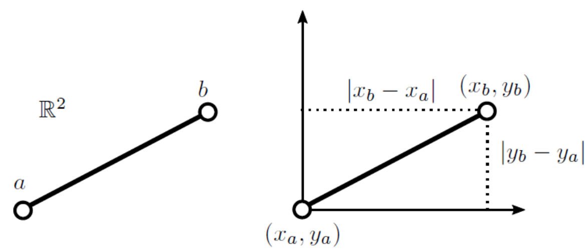

Apply Pythagoras' theorem: $d(a, b) = (|x_b - x_a|^2 + |y_b - y_a|^2)^{1/2}$

In vector notation: $d(\boldsymbol{a}, \boldsymbol{b}) = \|\mathbf{a} - \mathbf{b}\|_2$

Where $\mathbf{a} = \binom{x_a}{y_a}$ and $\mathbf{b} = \binom{x_b}{y_b}$

65

65

L_p DISTANCE IN $\mathbb{R}^k$

One can generalize to different power coefficients $p \geq 1$

$$\|\mathbf{x} - \mathbf{y}\|_2 = (|x_1 - y_1|^2 + |x_2 - y_2|^2)^{1/2}$$
$$\|\mathbf{x} - \mathbf{y}\|_p = (|x_1 - y_1|^p + |x_2 - y_2|^p)^{1/p}$$

As well as generalize from $\mathbb{R}^2$ to $\mathbb{R}^k$:

$$\|\mathbf{x} - \mathbf{y}\|_p = \left(\sum_{i=1}^{k} |x_i - y_i|^p\right)^{1/p}$$

This definition gives us the L_p distance between vectors in $\mathbb{R}^k$.

Examples:

- Euclidean (L_2) distance between 3D points
- Manhattan (L_1) distance between cities in a map

66

66

L_p UNIT BALLS

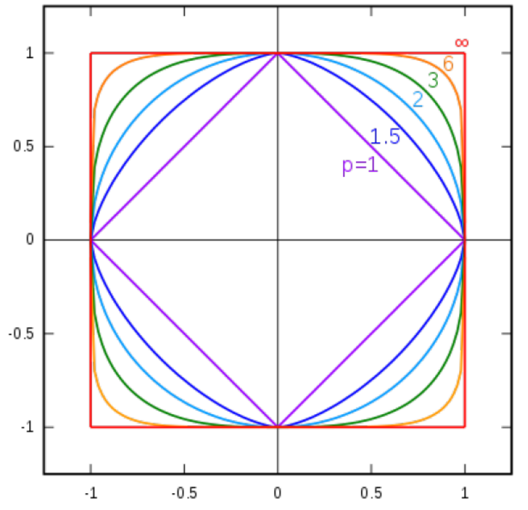

67

67

METRIC SPACES

The pair $(object, distance)$ forms a metric space. More formally:

A set $\mathcal{M}$ is a metric space if for every pair of points $x, y \in \mathcal{M}$ there is a metric (or distance) function $d_{\mathcal{M}}: \mathcal{M} \times \mathcal{M} \to \mathbb{R}_+$ such that:

- $d_{\mathcal{M}}(x, y) = 0 \Leftrightarrow x = y$ (identity of indiscernibles)
- $d_{\mathcal{M}}(x, y) = d_{\mathcal{M}}(y, x)$ (symmetry)
- $d_{\mathcal{M}}(x, y) \leq d_{\mathcal{M}}(x, z) + d_{\mathcal{M}}(z, y)$ for any $x, y, z \in \mathcal{M}$ (triangle inequality)

We will specify a metric space as the pair $(\mathcal{M}, d_{\mathcal{M}})$.

Example:

- The sphere with Euclidean distance is $(\mathbb{S}^2, d_{L_2})$
- The sphere with geodesic distance is $(\mathbb{S}^2, d_g)$

68

68

EXAMPLE: GEODESIC ISOLINES

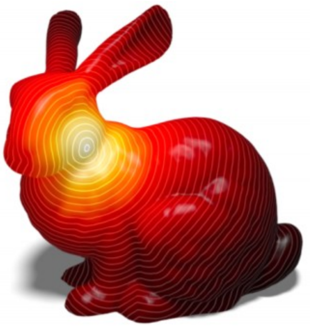

Each isoline identifies a set of points $x \in \mathcal{X}$ at the same distance (according to d_g) from some reference $y \in \mathcal{X}$.

AMBIENT SPACE AND RESTRICTION

If $\mathcal{A}$ is a metric space and $\mathcal{X} \subset \mathcal{A}$, then $\mathcal{A}$ is called ambient space for $\mathcal{X}$.

$\mathcal{A} = (\mathbb{R}^3, d_{L_2})$

A metric on $\mathcal{X}$ can be obtained by the restriction $d_{\mathcal{X}} = d_{\mathcal{A}|\mathcal{X}}$, such that: $d_{\mathcal{X}}(x, y) = d_{\mathcal{A}}(x, y)$ for all $x, y \in \mathcal{X}$.

ISOMETRIES

Let $(\mathcal{M}, d_{\mathcal{M}})$ and $(\mathcal{N}, d_{\mathcal{N}})$ be two metric spaces.
A bijective map $f: \mathcal{M} \to \mathcal{N}$ is called an isometry if:

$$d_{\mathcal{M}}(x, y) = d_{\mathcal{N}}(f(x), f(y))$$

for any $x, y \in \mathcal{M}$.

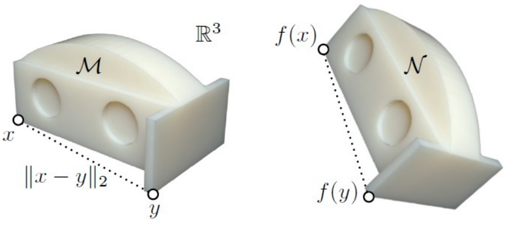

If $d_{\mathcal{M}} = \|.\|_2$ and $d_{\mathcal{N}} = \|.\|_2$ we say "rigid isometry".

71

71

EXAMPLE: NON-RIGID "QUASI"-ISOMETRIES

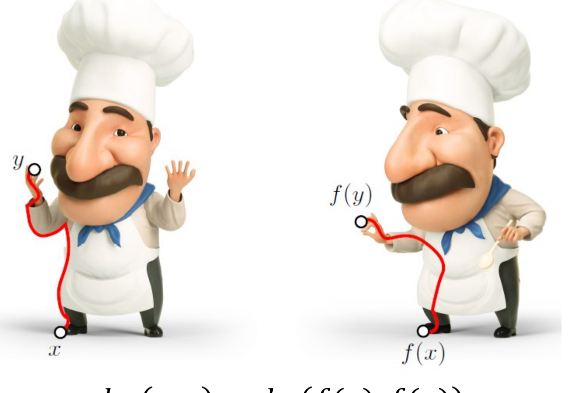

$$d_{\mathcal{M}}(x, y) \approx d_{\mathcal{N}}(f(x), f(y))$$

(here $d_{\mathcal{M}}, d_{\mathcal{N}}$ are geodesic distance functions)

72

72

Isometry as equivalence

"Being isometric" is an equivalence relation, since it is:

- reflective $(a = a)$
- symmetric $(a = b \Rightarrow b = a)$
- transitive $(a = b \wedge b = c \Rightarrow a = c)$

In this sense, we think of isometric shapes as being the same shape:

73

Isometric embeddings

How to compare the metric spaces themselves?

General idea: Embed $(\mathcal{X}, d_{\mathcal{X}})$ and (Y, d_Y) into a new metric space (Z, d_Z), and use a classical distance (e.g., Hausdorff, Chamfer) there.

An isometric embedding is a transformation $f: \mathcal{X} \rightarrow Z$ which preserves the metric for all pairs $x, y \in \mathcal{X}$, i.e.,

$$d_Z(f(x), f(y)) = d_{\mathcal{X}}(x, y)$$

For example, take $d_{\mathcal{X}} = d_g$ and $d_Z = \|.\|_2$:

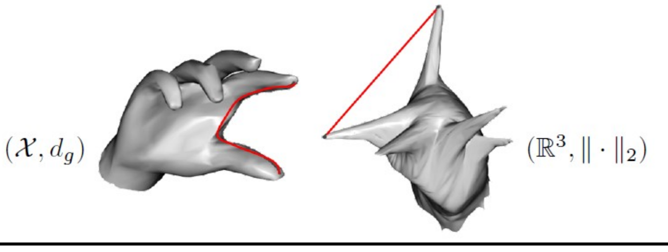

74

DISTANCE BETWEEN ISOMETRIC EMBEDDINGS

It will be a metric on the space of isometry classes.

An isometry class is a set of shapes which are equal up to isometry.
Therefore:

$d_{iso}($  ,  $) = d_{iso}($ 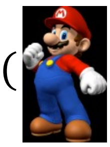 ,  $)$

Question: What is the isometry class for the sphere $(\mathbb{S}^2, d_g)$?

75

75

A CARTOGRAPHER'S PROBLEM

Are isometric embeddings always possible?
Consider the following:

An isometric embedding of $\mathbb{S}^2$ into $\mathbb{R}^2$ is not possible!
Any approximate solution introduces metric distortion.

76

76

NON-EMBEDDABILITY OF THE SPHERE

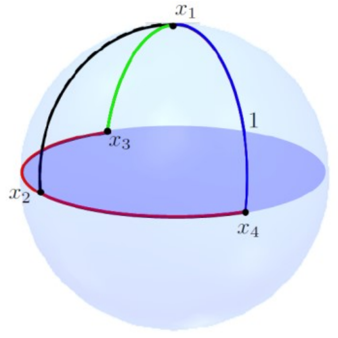

- Consider the triangle $\triangle\,(x_1, x_3, x_4) \Rightarrow$ collinear!
- Consider the triangle $\triangle\,(x_2, x_3, x_4) \Rightarrow$ collinear!
- Then $x_1 = x_2$, which contradicts $d_g(x_1, x_2) = 1$

  $\Rightarrow$ This metric space cannot be embedded into $\mathbb{R}^k$ for any k

77

77

A CARTOGRAPHER'S SOLUTION

78

78

SUGGESTED READING

Bronstein et al, “Geometric deep learning: going beyond Euclidean data“, 2016
https://arxiv.org/abs/1611.08097

Bronstein et al, “Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges”, 2021
https://arxiv.org/abs/2104.13478