

Nuages de Points et Modélisation 3D

4 - Machine learning I

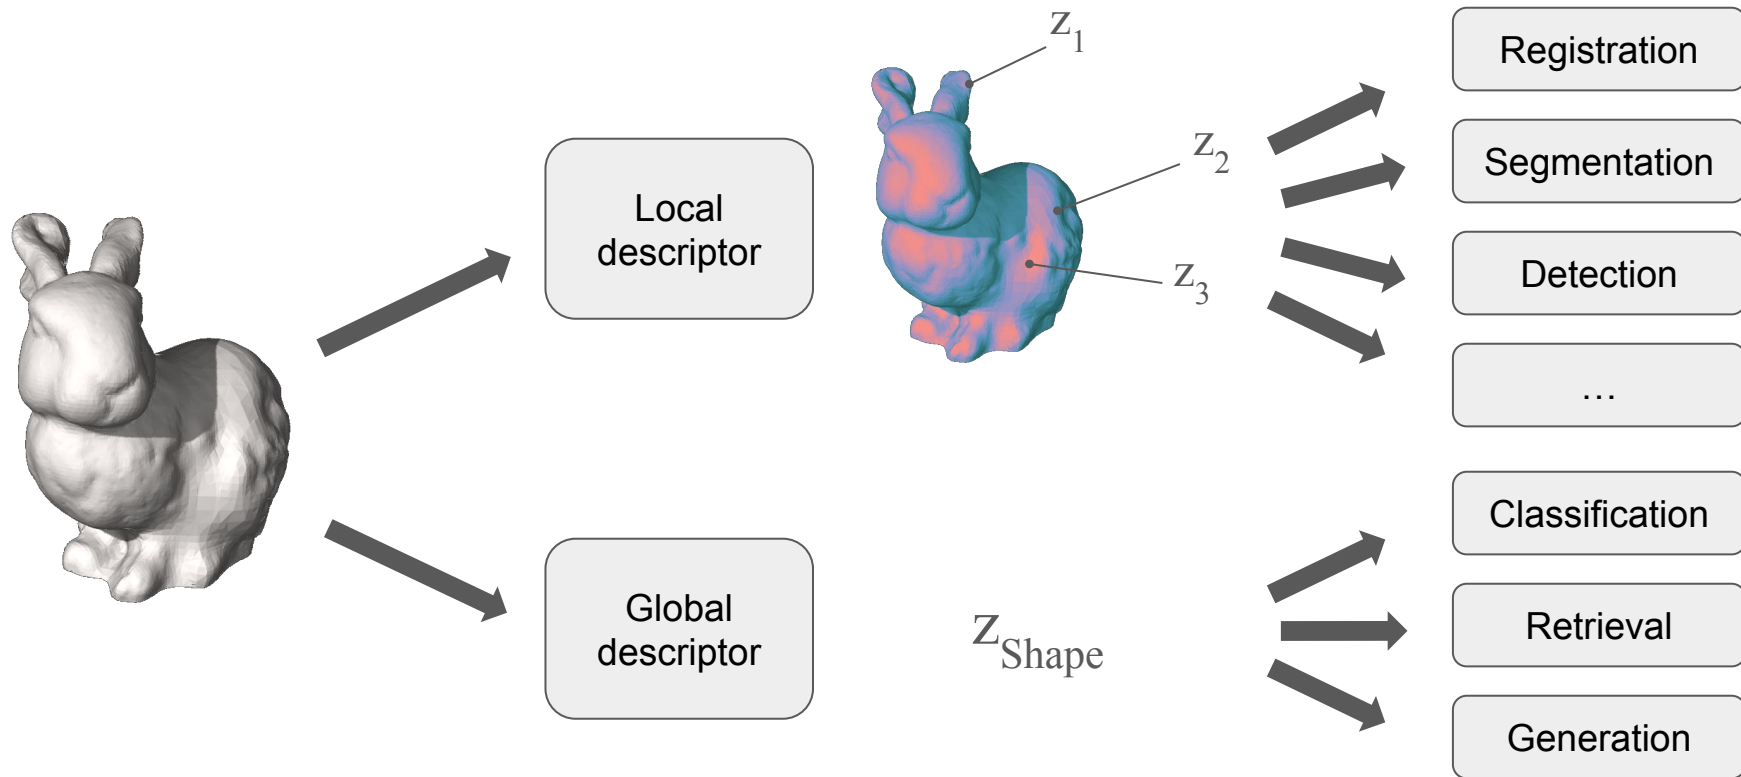
Overview

Machine learning courses

- Surface reconstruction
 - Descriptors and machine learning
 - Image based processing
 - Geometric deep learning
 - Convolutional and Transformer based architectures
 - Tasks and corresponding architectures
- } Today
- } ML course 2
- } ML course 3
- } ML course 4

I - Descriptors and machine learning

Descriptors




I - Descriptors and machine learning

A - Local descriptors

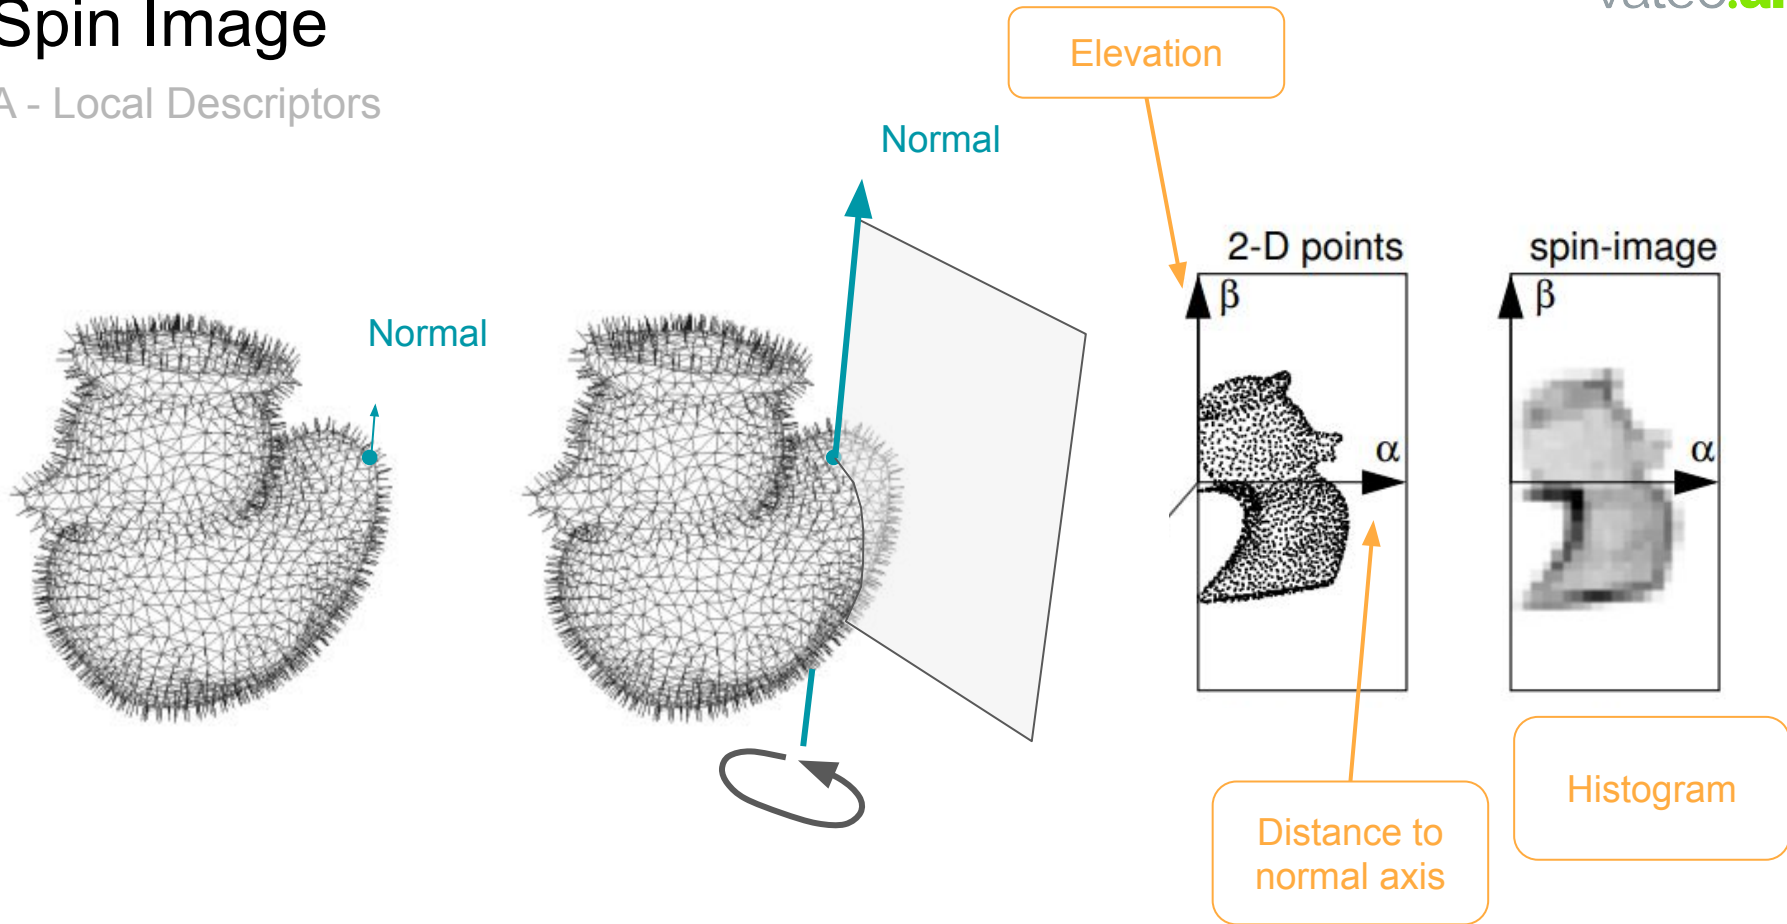
Spin Image

A - Local Descriptors



Spin Image

A - Local Descriptors



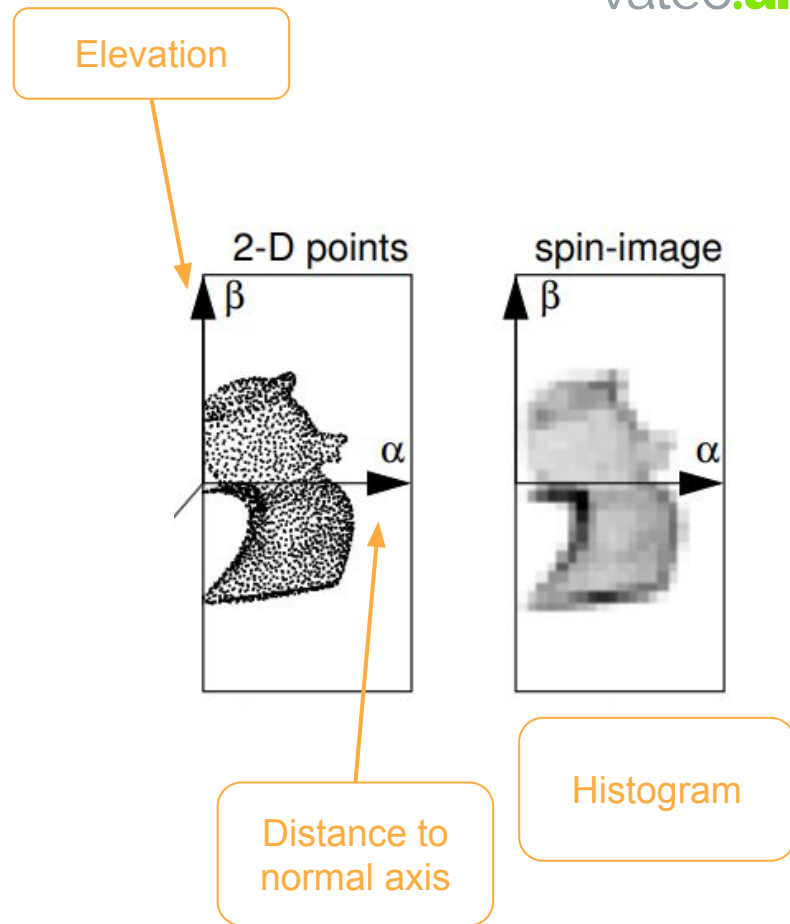
Spin Image

A - Local Descriptors

Given normal \mathbf{n} at point \mathbf{p}

For each point of interest:

- Compute alpha and beta:
 $\beta = \langle \mathbf{n}, \mathbf{q} - \mathbf{p} \rangle$
 $\alpha = \|(\mathbf{q} - \mathbf{p}) - \beta \mathbf{n}\|$
- Accumulate in the histogram image

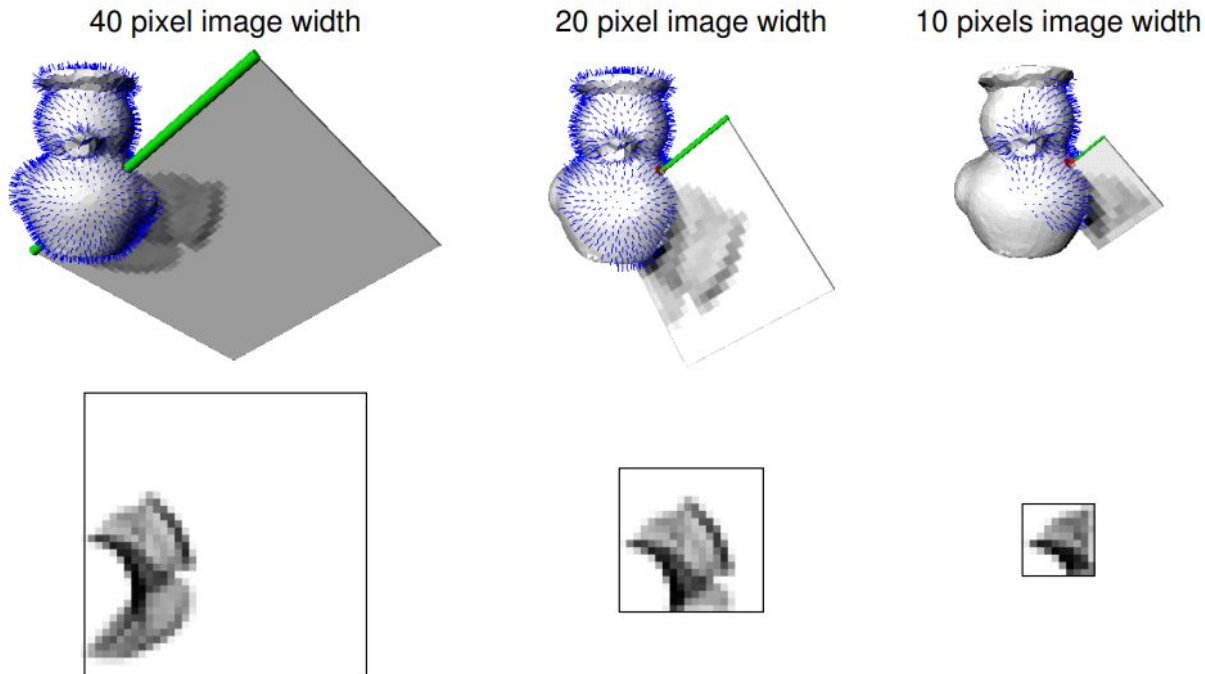


Spin Image

A - Local Descriptors

Influence of the accumulator size

- Defines the size of the observed neighborhood
- More or less context




Spin Image



A - Local Descriptors

Filter on the normal angle



- Retain meaningful points




180° support angle



90° support angle




60° support angle




Spin Image

A - Local Descriptors


Influence of
histogram for
descriptor
matching




scene spin-image




scene points
accumulated




scene spin-image




scene points
accumulated




model spin-image



model points
accumulated




model spin-image




model points
accumulated

Spin Image

A - Local Descriptors



Object retrieval in 3D point clouds




DoN - Difference of Normals

A - Local Descriptors

Observation:

Normals estimated with radius vary with r and the smoothness of the surface

⇒ use this variation as a descriptor



DoN - Difference of Normals

A - Local Descriptors

More formally:


$$\Delta_{\hat{\mathbf{n}}}(\mathbf{p}, r_1, r_2) = \frac{\hat{\mathbf{n}}(\mathbf{p}, r_1) - \hat{\mathbf{n}}(\mathbf{p}, r_2)}{2}$$

r_1 and r_2 are parameters.

More formally, it is an approximation of the **curvature** of the surface at given radius

Curvature

A - Local Descriptors




Curvature:

radius of maximal the sphere that can be fitted on the concave side.


DoN - Difference of Normals

A - Local Descriptors


Practical influence of the radius size




(a) Point cloud (478,377 points).




(b) $|\Delta_{\hat{n}}(0.2 \text{ m}, 2 \text{ m})|$.



(c) $|\Delta_{\hat{n}}(0.8 \text{ m}, 8 \text{ m})|$.



(d) $|\Delta_{\hat{n}}(2 \text{ m}, 20 \text{ m})|$.



DoN - Difference of Normals

A - Local Descriptors

Usage for
unsupervised
segmentation:

- Put a threshold on the norm of the DoN
- Apply **clustering**



(a) Original point cloud (620,820 points).



(b) Clusters found in $|\Delta_{\hat{n}}(0.2 \text{ m}, 2 \text{ m})| \geq 0.25$.



(c) Person cluster.



(d) Traffic light cluster.



(e) Window cluster.



(f) Car cluster.



(g) Tree cluster.


Apparte on clustering

Euclidean Cluster Extraction

Clustering

for every point, $\mathbf{p}_i \in P$ perform the following steps:

- If already visited, skip
- Else:
 - Initiate cluster
 - add \mathbf{p}_i to the current queue Q ;




Euclidean Cluster Extraction

Clustering

for every point, $p_i \in P$ perform the following steps:

- If already visited, skip
- Else:
 - Initiate cluster
 - add p_i to the current queue Q ;
 - While Q is not empty:
 - Dequeue first element q
 - Add q to cluster C




Euclidean Cluster Extraction

Clustering

for every point, $\mathbf{p}_i \in P$ perform the following steps:

- If already visited, skip
- Else:
 - Initiate cluster
 - add \mathbf{p}_i to the current queue Q ;
 - While Q is not empty:
 - Dequeue first element q
 - Add q to cluster C
 - search for the set neighbors of q in a sphere with radius r ;
 - for every neighbor ,
 - if point not already processed
 - Add to queue Q




Euclidean Cluster Extraction

Clustering

for every point, $p_i \in P$ perform the following steps:

- If already visited, skip
- Else:
 - Initiate cluster
 - add p_i to the current queue Q ;
 - While Q is not empty:
 - Dequeue first element q
 - Add q to cluster C
 - search for the set neighbors of q in a sphere with radius r ;
 - for every neighbor ,
 - if point not already processed
 - Add to queue Q




Euclidean Cluster Extraction

Clustering

for every point, $\mathbf{p}_i \in P$ perform the following steps:

- If already visited, skip
- Else:
 - Initiate cluster
 - add \mathbf{p}_i to the current queue Q ;
 - While Q is not empty:
 - Dequeue first element q
 - Add q to cluster C
 - search for the set neighbors of q in a sphere with radius r ;
 - for every neighbor ,
 - if point not already processed
 - Add to queue Q




Euclidean Cluster Extraction

Clustering

for every point, $\mathbf{p}_i \in P$ perform the following steps:

- If already visited, skip
- Else:
 - Initiate cluster
 - add \mathbf{p}_i to the current queue Q ;
 - While Q is not empty:
 - Dequeue first element q
 - Add q to cluster C
 - search for the set neighbors of q in a sphere with radius r ;
 - for every neighbor ,
 - if point not already processed
 - Add to queue Q




Euclidean Cluster Extraction

Clustering

for every point, $p_i \in P$ perform the following steps:

- If already visited, skip
- Else:
 - Initiate cluster
 - add p_i to the current queue Q ;
 - While Q is not empty:
 - Dequeue first element q
 - Add q to cluster C
 - search for the set neighbors of q in a sphere with radius r ;
 - for every neighbor ,
 - if point not already processed
 - Add to queue Q




Euclidean Cluster Extraction

Clustering

for every point, $\mathbf{p}_i \in P$ perform the following steps:

- If already visited, skip
- Else:
 - Initiate cluster
 - add \mathbf{p}_i to the current queue Q ;
 - While Q is not empty:
 - Dequeue first element q
 - Add q to cluster C
 - search for the set neighbors of q in a sphere with radius r ;
 - for every neighbor ,
 - if point not already processed
 - Add to queue Q




Euclidean Cluster Extraction

Clustering

It is a **region growing** algorithm in n dimensions.

Can be further extended:

- Planarity (reestimate regression plane)
- Normal filtering
- More generally descriptor filtering




DBSCAN: Density-based spatial clustering of applications with noise

Clustering

Algorithm:

- Build the graph of neighborhoods of radius r




DBSCAN: Density-based spatial clustering of applications with noise

Clustering

Algorithm:

- Build the graph of neighborhoods of radius r
- Search the core points: points that have minPts points in their neighborhood (including itself)

$\text{minPts}=4$




DBSCAN: Density-based spatial clustering of applications with noise

Clustering

Algorithm:

- Build the graph of neighborhoods of radius r
- Search the core points: points that have minPts points in their neighborhood (including itself)
- Reachable point q if there is path (p, p_1, \dots, p_n, q) where p_i are core points

$\text{minPts}=4$




DBSCAN: Density-based spatial clustering of applications with noise

Clustering

Algorithm:

- Build the graph of neighborhoods of radius r
- Search the core points: points that have minPts points in their neighborhood (including itself)
- Reachable point q if there is path (p, p_1, \dots, p_n, q) where p_i are core points
- Outliers: isolated points

$\text{minPts}=4$



DBSCAN: Density-based spatial clustering of applications with noise


Clustering

Algorithm:

- Build the graph of neighborhoods of radius r
- Search the core points: points that have minPts points in their neighborhood (including itself)
- Reachable point q if there is path (p, p_1, \dots, p_n, q) where p_i are core points
- Outliers: isolated points

Cluster of p is all the reachable points from p .

$\text{minPts}=4$



K-means


Clustering

One of the most used clustering algorithm.

Parameter: k , number of cluster

Advantage:

- Simple
- Proof of convergence




K-means

Clustering

Initialization:

Select random k points



K-means


Clustering

Initialization:

Select random k points

Step 1:

Assign each point to its closest center



K-means


Clustering

Initialization:

Select random k points

Step 1:

Assign each point to its closest center
(Assignment based on Voronoï cells)



K-means

Clustering

Initialization:


Select random k points

Step 1:

Assign each point to its closest center

Step 2:

Compute new centroids



K-means

Clustering

Initialization:

Select random k points


Step 1:

Assign each point to its closest center

Step 2:

Compute new centroids

Iterate until convergence (no point changes cluster)



K-means

Clustering

Initialization:

Select random k points

Step 1:

Assign each point to its closest center

Step 2:

Compute new centroids

Iterate until convergence (no point changes cluster)



K-means

Clustering

Initialization:

Select random k points


Step 1:

Assign each point to its closest center

Step 2:

Compute new centroids

Iterate until convergence (no point changes cluster)



K-means

Clustering


Advantages:

- Simple
- Any dimension
- Only on parameter
- Will converge

Disadvantages

- Convergence may be slow
→ bad initialization
- Need to know k

Mutiple method for estimating the number of clusters, better initialization...



Coming back to descriptors

Covariance-based descriptors

A - Local Descriptors

A single point → poor information (only 3 coordinates)

Local descriptors → neighborhood

Previous course

Normal estimation using neighborhood (K-nearest neighbors)

Construction of the covariance matrix

Covariance-based descriptors

A - Local Descriptors

Covariance matrix (again)

Build the covariance matrix for the neighborhood N of a point q :

- Average:
$$\bar{x} = \frac{1}{n} \sum_{x \in P} x$$

Covariance:
$$Cov \in \mathbb{R}^3 \times \mathbb{R}^3$$

$$Cov(i, j) = \frac{1}{n} \sum_{x \in P} (x_i - \bar{x}_i)(x_j - \bar{x}_j) = \frac{1}{n} X^T X$$

Covariance-based descriptors

A - Local Descriptors

Covariance matrix (again)

- Compute covariance matrix
- Diagonalize the Matrix, with P orthonormal (Cov is positive, real, symmetric)

$$Cov = PDP^{-1}$$

$$D = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix}$$

$$P = [e_1 \quad e_2 \quad e_3]$$

$$\lambda_1 \geq \lambda_2 \geq \lambda_3$$
$$e_1, e_2, e_3 \in \mathbb{R}^3$$

Covariance-based descriptors

A - Local Descriptors

Covariance matrix (again)

- Compute covariance matrix
- Diagonalize the Matrix, with P orthonormal (Cov is positive, real, symmetric)

$$Cov = PDP^{-1}$$

$$D = \begin{pmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{pmatrix} \quad P = [e_1 \quad e_2 \quad e_3] \quad \begin{array}{l} \lambda_1 \geq \lambda_2 \geq \lambda_3 \\ e_1, e_2, e_3 \in \mathbb{R}^3 \end{array}$$

Covariance-based descriptors

A - Local Descriptors

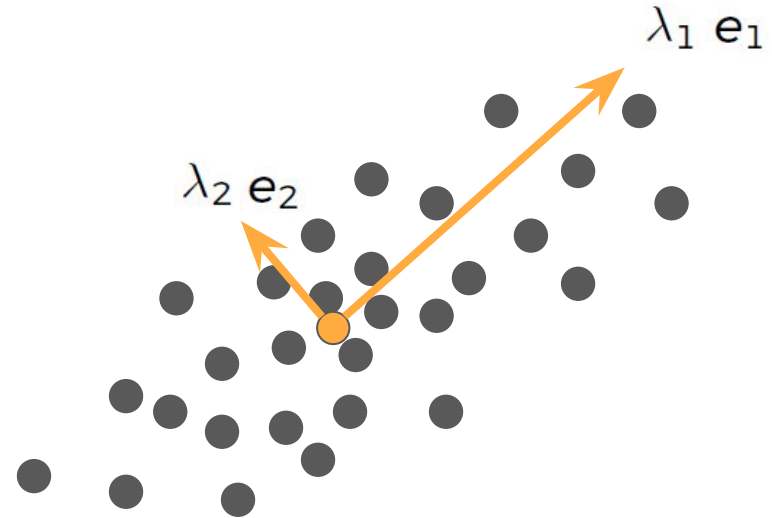
What are $\lambda_1 \geq \lambda_2 \geq \lambda_3$?

PCA spread

Orthogonal basis $e_1, e_2, e_3 \in \mathbb{R}^3$

e_3 is the normal

\Rightarrow there is more information
than just the normal



Covariance-based descriptors

A - Local Descriptors

Sum of eigenvalues

$$\sum \lambda_i$$

Omnivariance

$$\left(\prod \lambda_i\right)^{\frac{1}{3}}$$

Eigenentropy

$$-\sum \lambda_i \ln(\lambda_i)$$

Linearity

$$(\lambda_1 - \lambda_2)/\lambda_1$$

Planarity

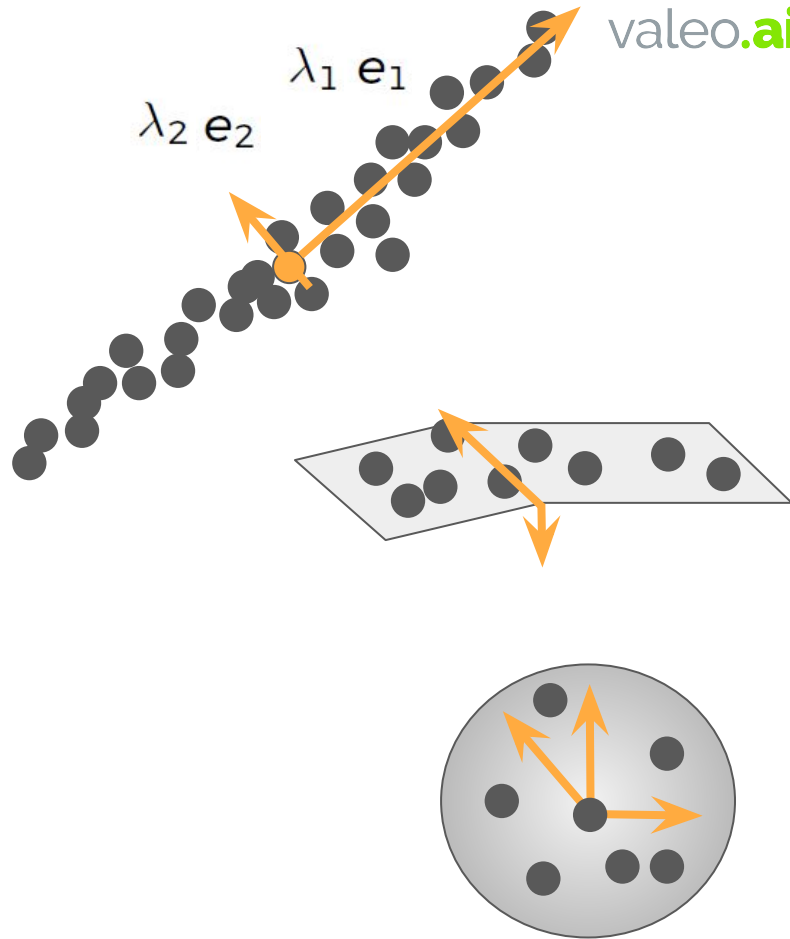
$$(\lambda_2 - \lambda_3)/\lambda_1$$

Sphericity

$$\lambda_3/\lambda_1$$

Change of curvature

$$\lambda_3/(\lambda_1 + \lambda_2 + \lambda_3)$$



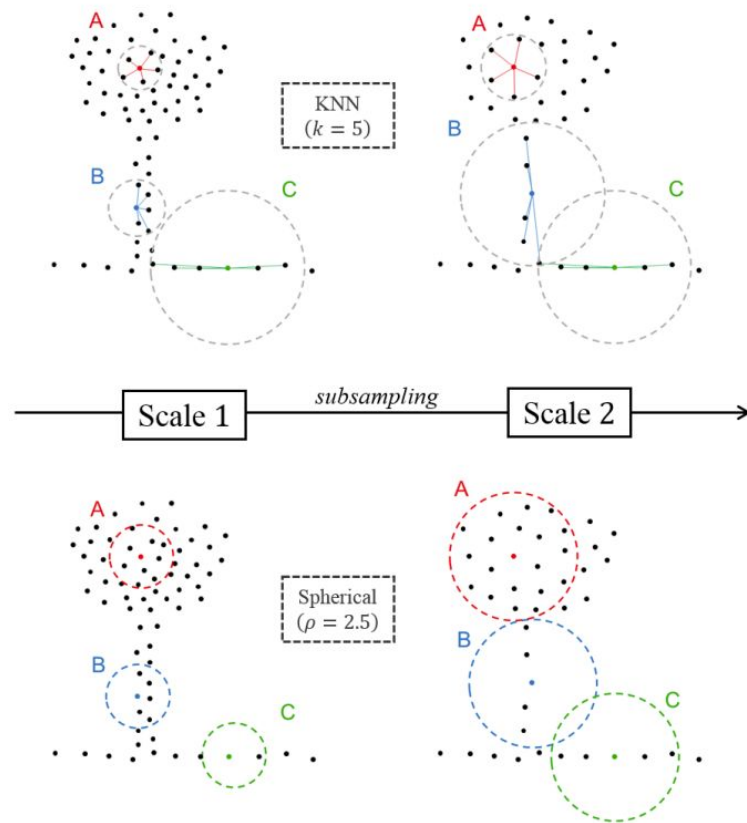

Covariance-based descriptors

A - Local Descriptors

Sum of eigenvalues	$\sum \lambda_i$	Verticality (x2)	$\left \frac{\pi}{2} - \text{angle}(\mathbf{e}_i, \mathbf{e}_z) \right _{i \in (0,2)}$
Omnivariance	$\left(\prod \lambda_i \right)^{\frac{1}{3}}$	Absolute moment (x6)	$\frac{1}{ \mathcal{N} } \left \sum \langle \mathbf{p} - \mathbf{p}_0, \mathbf{e}_i \rangle^k \right _{i \in (0,1,2)}$
Eigenentropy	$-\sum \lambda_i \ln(\lambda_i)$	Vertical moment (x2)	$\frac{1}{ \mathcal{N} } \sum \langle \mathbf{p} - \mathbf{p}_0, \mathbf{e}_z \rangle^k$
Linearity	$(\lambda_1 - \lambda_2) / \lambda_1$	Number of points	$ \mathcal{N} $
Planarity	$(\lambda_2 - \lambda_3) / \lambda_1$	Average color (x3)	$\frac{1}{ \mathcal{N} } \sum c$
Sphericity	λ_3 / λ_1	Color variance (x3)	$\frac{1}{ \mathcal{N} - 1} \sum (c - \bar{c})^2$
Change of curvature	$\lambda_3 / (\lambda_1 + \lambda_2 + \lambda_3)$		

Covariance-based descriptors

A - Local Descriptors



Covariance-based descriptors

A - Local Descriptors

Pros

- Fast to compute
- Memory efficient
- Simple classifier
(SVM, random forests, MLP...)
- Requires limited amount of data for training

Cons

- No long distance relations
- Limited by the descriptors:
 - The system cannot infer the most suited descriptors
- Difficulty to create dataset specific descriptors

I - Descriptors and machine learning

B - Global descriptors


GASD - Globally Aligned Spatial Distribution

B - Global descriptors

Descriptor for a complete shape

→ shape classification

Step 1: compute global orientation
(PCA)



GASD - Globally Aligned Spatial Distribution

B - Global descriptors

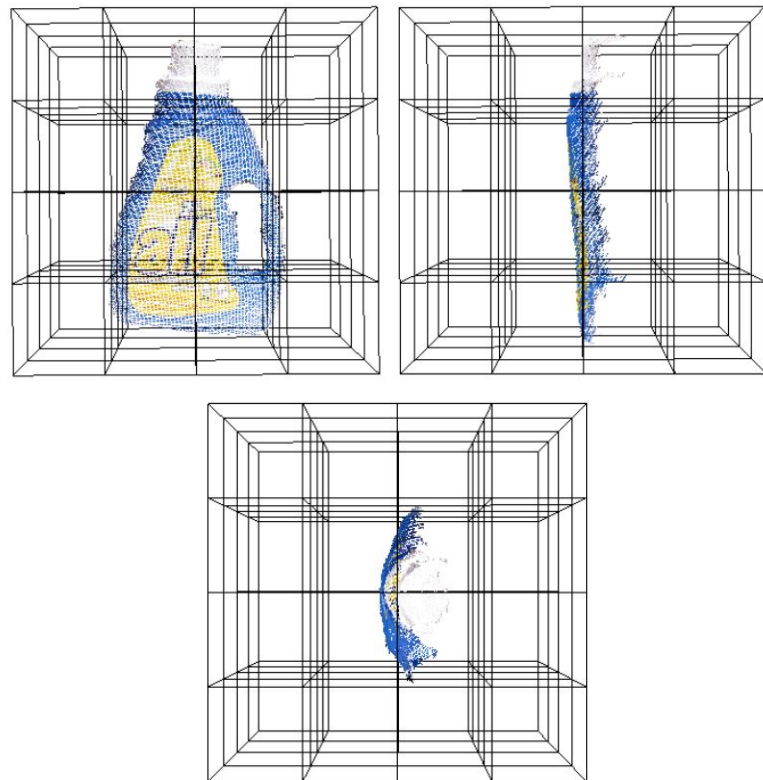
Descriptor for a complete shape

→ shape classification

Step 1: compute global orientation
(PCA)

Step 2:

- Normalized histogram of the count of points in each cell
- Average color per cell



GASD - Globally Aligned Spatial Distribution

B - Global descriptors

Descriptor for a complete shape
→ shape classification



Fig. 14. GASD fails to recognize a partially occluded blue detergent bottle.



Fig. 7. Object recognition and pose estimation example: input scene (left), results obtained using ESF (center) and GASD-SI (right).

Global descriptors

B - Global descriptors

In practice:

This approach can be used with any descriptor (use accumulator)

Warning:

- Orientation is a hard problem (missing elements, partial view...)

Current trend would be to use augmentations

- Use multiple rotation
- Masking

I - Descriptors and machine learning

C - Downstream tasks

Object retrieval

C - Downstream tasks

Given a descriptor look for the k-closest in a database



Classification

C - Downstream tasks

Associate each shape with a label

- Expert decision (by hand)
- ML: SVMs, RandomForest, MLP...

Requires a **labeled train set**


Classification

C - Downstream tasks

Associate each shape with a label

- Expert decision (by hand)
- ML: SVMs, RandomForest, MLP...

Requires a **labeled train set**



Semantic segmentation

C - Downstream tasks

Associate each point with a label

- Expert decision (by hand)
- ML: SVMs, RandomForest, MLP...

NPM3D

Requires a **labeled train set**




Point matching

C - Downstream tasks

Match local descriptor from one point cloud to another

Estimate transformation (e.g., RANSAC)



II - Image-based approaches

Idea

Image-based approaches



Image
processing




Image processing is a well studied problem

Idea


Image-based approaches



Image processing




?



Idea

Image-based approaches



Regular grid projections

Image-based approaches

Images are pixels arrays

Implicit neighborhoods


Information is in the color and relative position of the pixels

Thanks to this grid structure

Optimized network architectures

Fast (hardware optimization)

Relatively low memory cost



2D projections

Image-based approaches

2D convolution for an image patch centered on pixel n :

$$\mathbf{h}[n] = \sum_{f \in \{1, \dots, C\}} \sum_{m \in \{-M/2, \dots, M/2\}^2} \mathbf{K}_f[m] \mathbf{f}_f[n + m]$$

With \mathbf{f} : input features and \mathbf{K} : convolution kernel

And new architectures for images:

Vision transformers, MLP Mixers, ...

Regular grid projections

Image-based approaches

Images are pixels arrays

Implicit neighborhoods


Information is in the color and relative position of the pixels

Thanks to this grid structure

Optimized network architectures

Fast (hardware optimization)


Relatively low memory cost



Point clouds:

~~Implicit neighborhoods~~

~~Information is in the color and relative position of the pixels~~



Idea

Find a way to create grid data from point cloud

2D projections

Image-based approaches

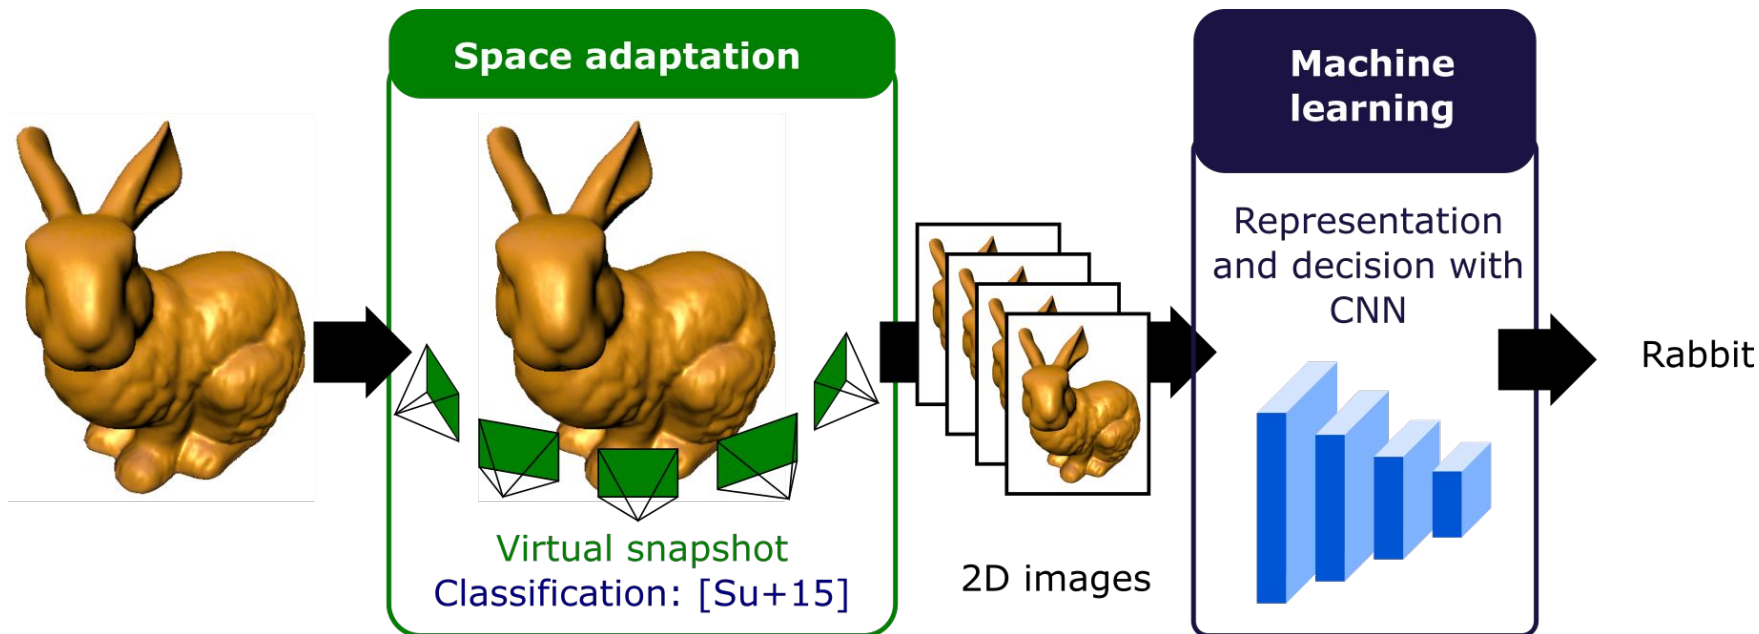
Generate images representing the scene

- Use a 3D renderer
- Take virtual snapshots of the scene
- Work in the image



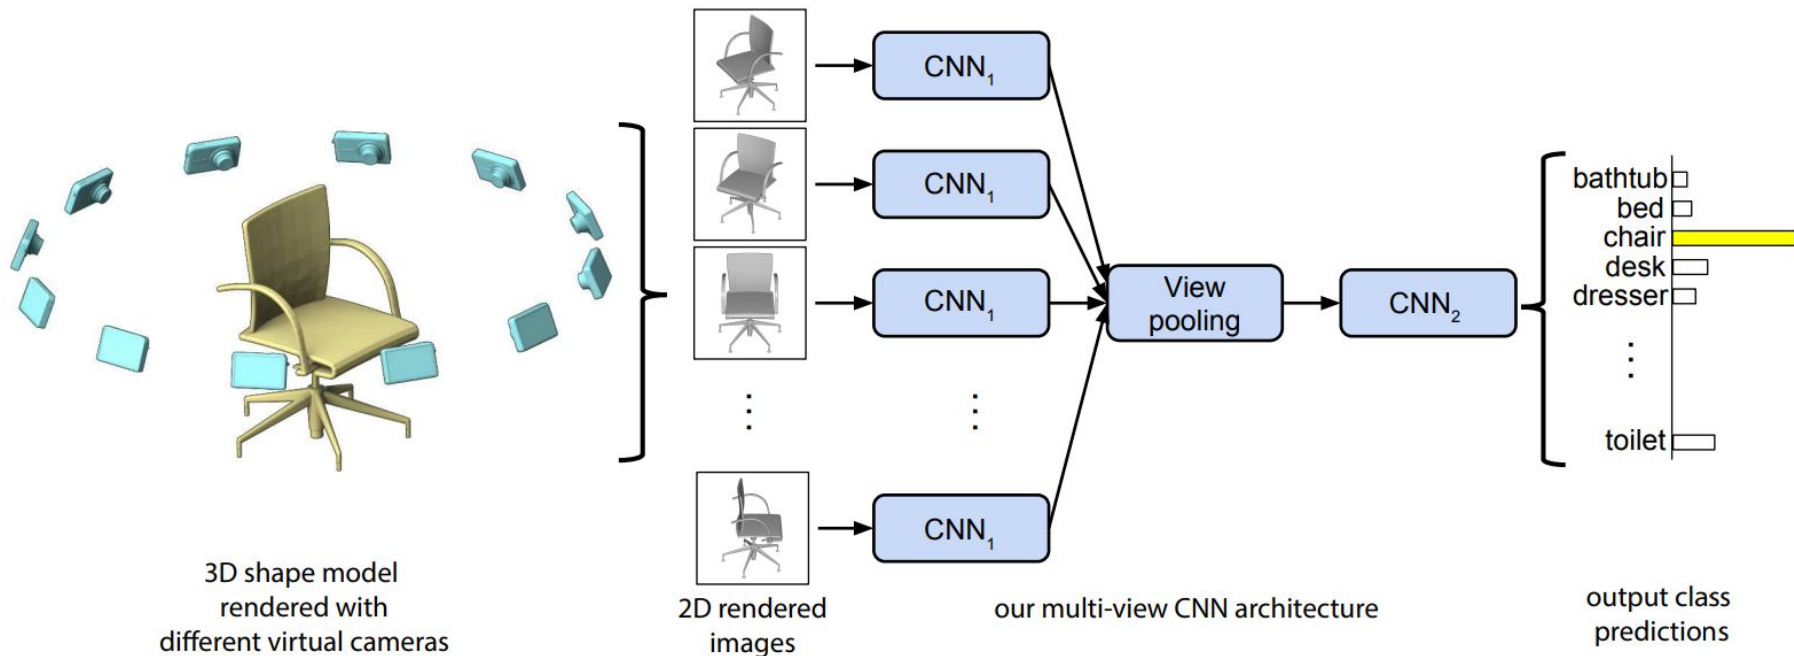
Classification pipeline

Image-based approaches




Classification pipeline

Image-based approaches



Semantic segmentation pipeline

Image-based approaches




SnapNet

Image-based approaches

Reprojection trick

Generate a snapshot of the scene with fake colors corresponding to point ids

- Allow to generate different snapshots (w / wo colors, geometric features, ground truth at training...)
- Easy reprojection of the results on the original points



SnapNet - advantages and limitations

Image-based approaches

Pros

- Benefit from architectures from image processing
- Use of pre-trained models from large image datasets
- Unlimited number of snapshot a given scene (straightforward data augmentation)

Cons


- Good snapshot strategy, which can vary from a dataset to another
- Requires a mesh



Range projection

Image-based approaches


Exploit sensor information to produce images



SalsaNext

Image-based approaches

Use image backbone (U-Net) for semantic segmentation




RangeViT

Image-based approaches

LiDAR segmentation based on range images & vision transformers (ViTs)

- Unify architectures in LiDAR and image domain
- Leverage image pre-trained ViTs for LiDAR segmentation



Practical session

Machine learning 1

Practical session

- Implement covariance based local descriptors.
- Train a classifier
- Evaluate on validation point cloud