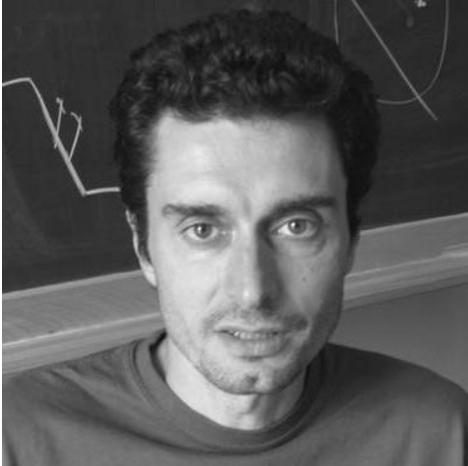


Intrinsic Dimension for continuous data

Antonietta Mira

Università della Svizzera italiana, Lugano



Alessandro
Laio



Elena
Facco



Michele
Allegra



Diego
Doimo



Aldo
Glielmo

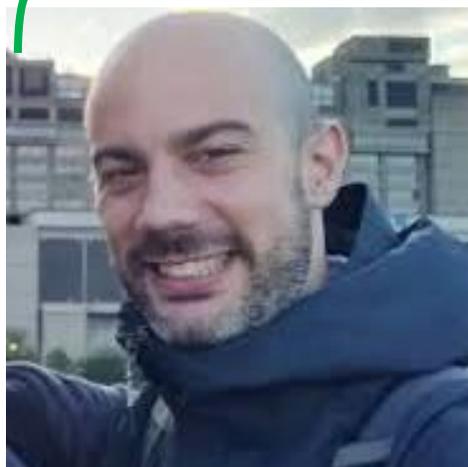


Iuri
Macocco

USI

SISSA

QUT



Francesco
Denti



Antonio
Di Noia



Federico
Ravenda



Edgar
Santos-Fernandez

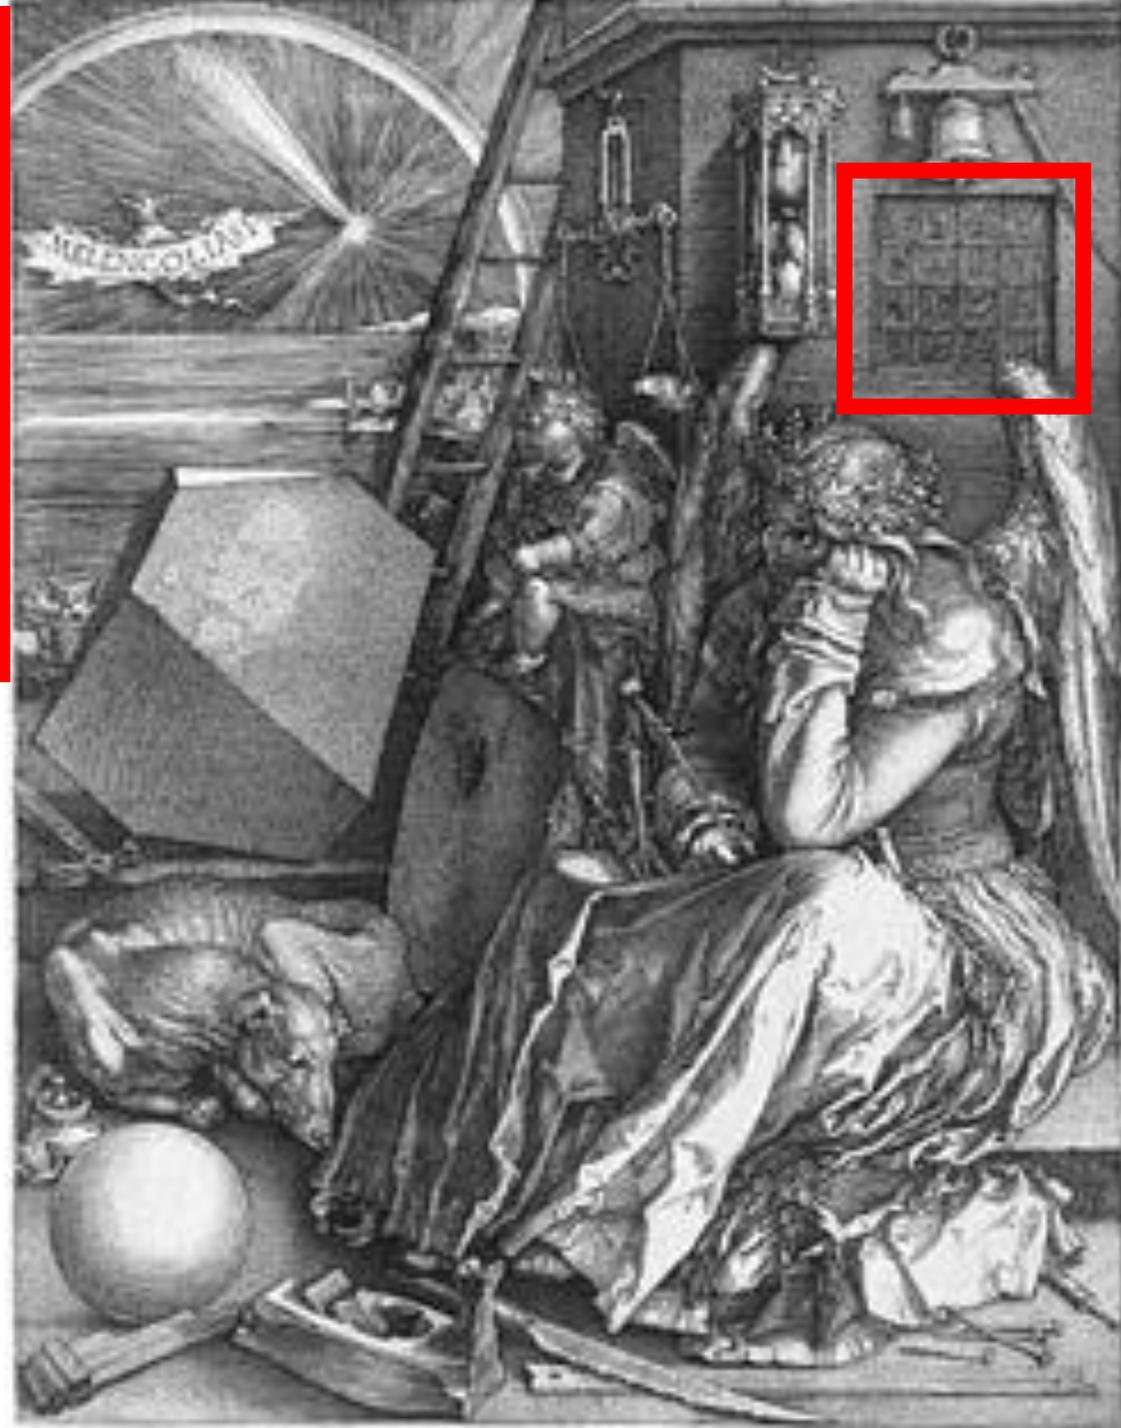


Abhishek
Varghese



Kerrie
Mengersen

16	3	2	13
5	10	11	8
9	6	7	12
4	15	14	1



A. Durer, 1514

synthesis of Earth and Heaven,
 search for the intrinsic
 dimension of
 all things

- **Interdisciplinary:** statistics + physics + domain experts
- **2019 - ongoing:**

- ❖ 5 (SISSA + USI)

- ❖ 2 (QUT + USI)

- ❖ 2 (USI)

Quote / Quiz

”The non-mathematician, is seized by a mysterious shuddering when (s)he hears of ‘four-dimensional’ things, (s)he is seized by a feeling, which is very similar to the thoughts awakened by the occult.”

Albert Einstein, 1920

Classical mechanics:

space and time = separate entities

Special relativity:

space and time = interwoven into a four dimensional

Unified construct known as "space-time"

From physics to statistics / DS / ML

Dimension expansion

- **data**: EM, ABC, Knockoffs, LLM ...
- **parameter**: hybrid MCMC, Slice sampler, mixture models

Dimension reduction

- **data**:
 - Kernel PCA, Canonical Correlation Analysis, Clustering,
 - VAE, t-SNE, Isomap, manifold learning, sketching,
 - random projections, projection pursuit
- **parameter**:
 - Variable / feature selection, Factor analysis

Focus on dimension reduction

Open questions

A subset of the variables or non-linear combinations thereof are often sufficient to describe a real-world data set

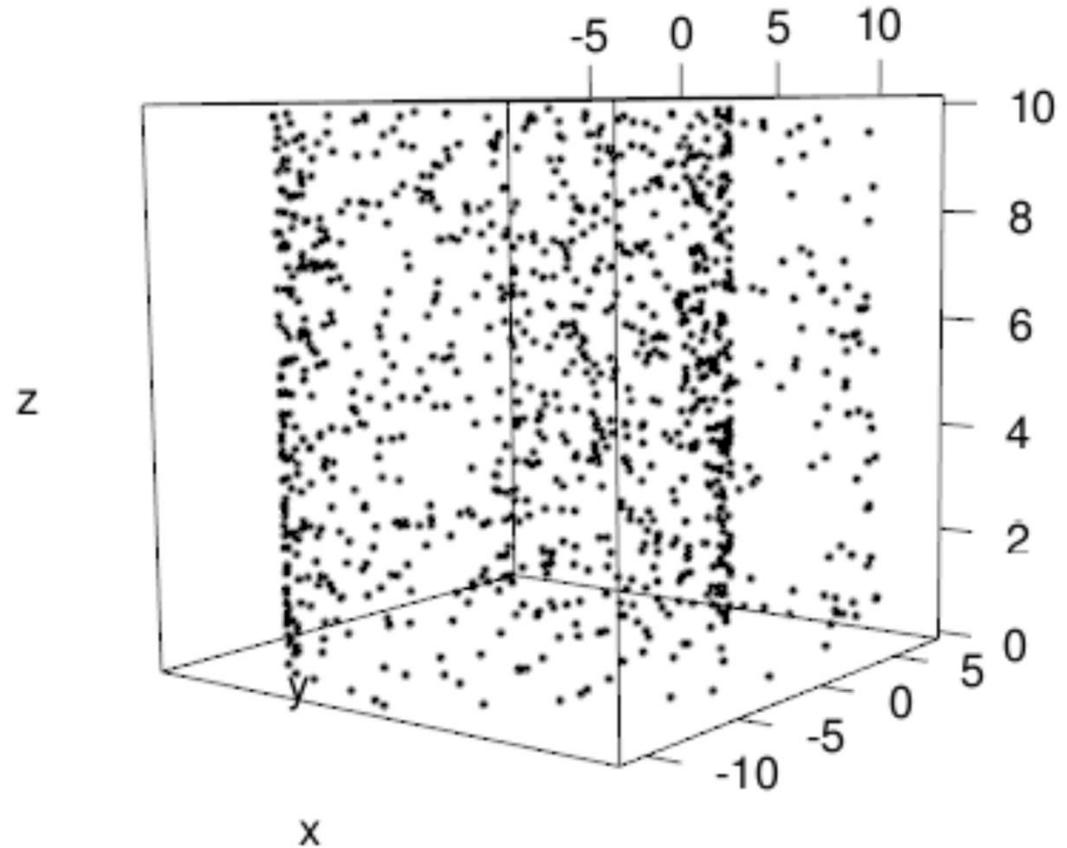
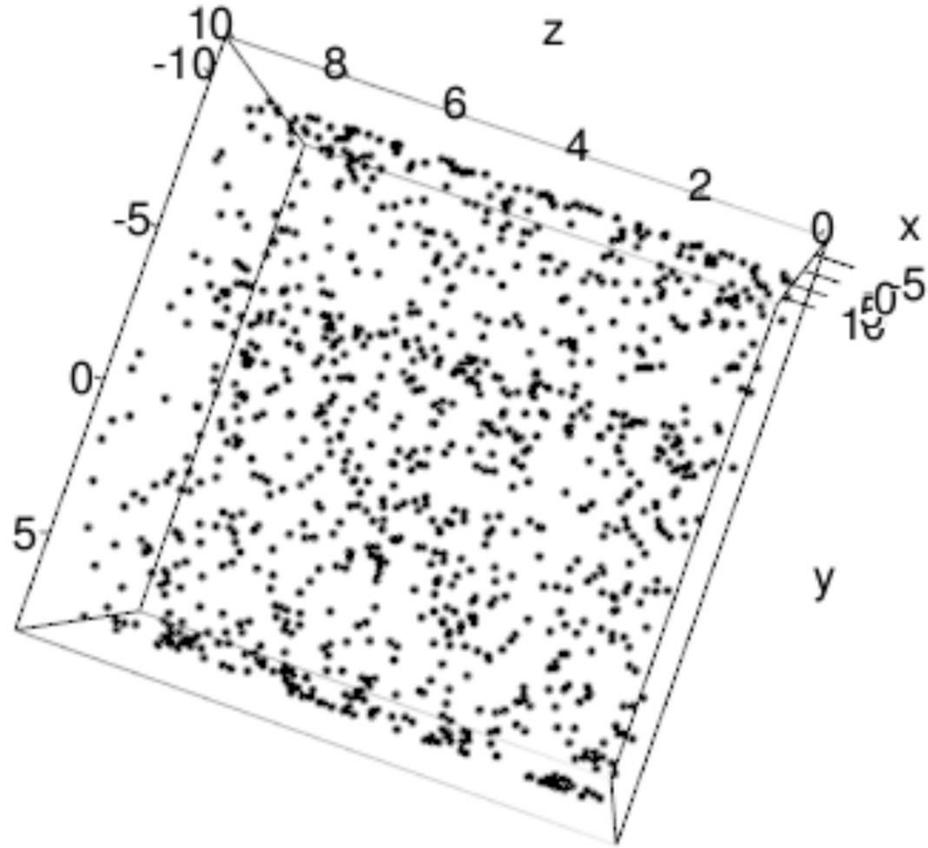
How many (and which) variables are needed to summarize a data set without significant information loss?

What is the appropriate scale at which one should analyze data?

Two questions, often considered unrelated, but strongly entangled, can be addressed within a unified "scale-dimension" framework

We introduce an approach in which the optimal number of variables and the optimal scale are determined self-consistently, bypassing the scale at which data are affected by noise

A matter of perspective



Can you guess the dimension of the support of the data generating process, d ?

Motivation: dimensionality reduction

In the Swissroll example $D = 3$ and $d = 2$

d , is called the intrinsic dimension (ID) of the data

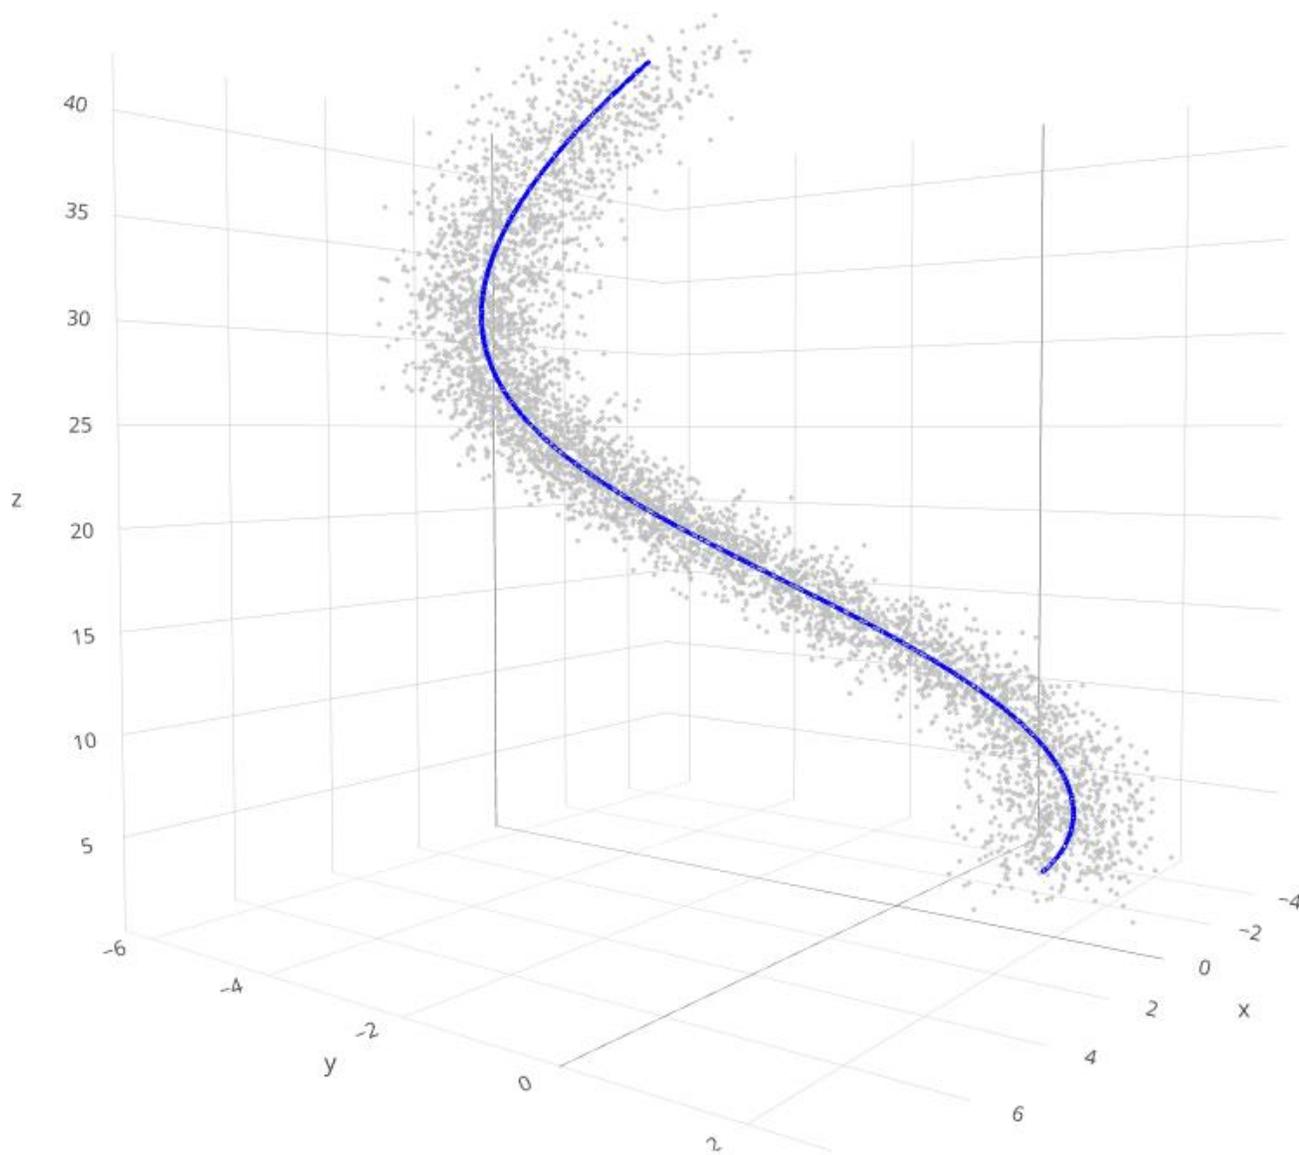
ID = needed for manifold learning (D. Dunson et al.)

ID = regression with covariates on a manifold (J. Rousseau et al.)

ID = key concept in unsupervised learning and feature selection

ID = lower bound to number of variables needed to describe a system

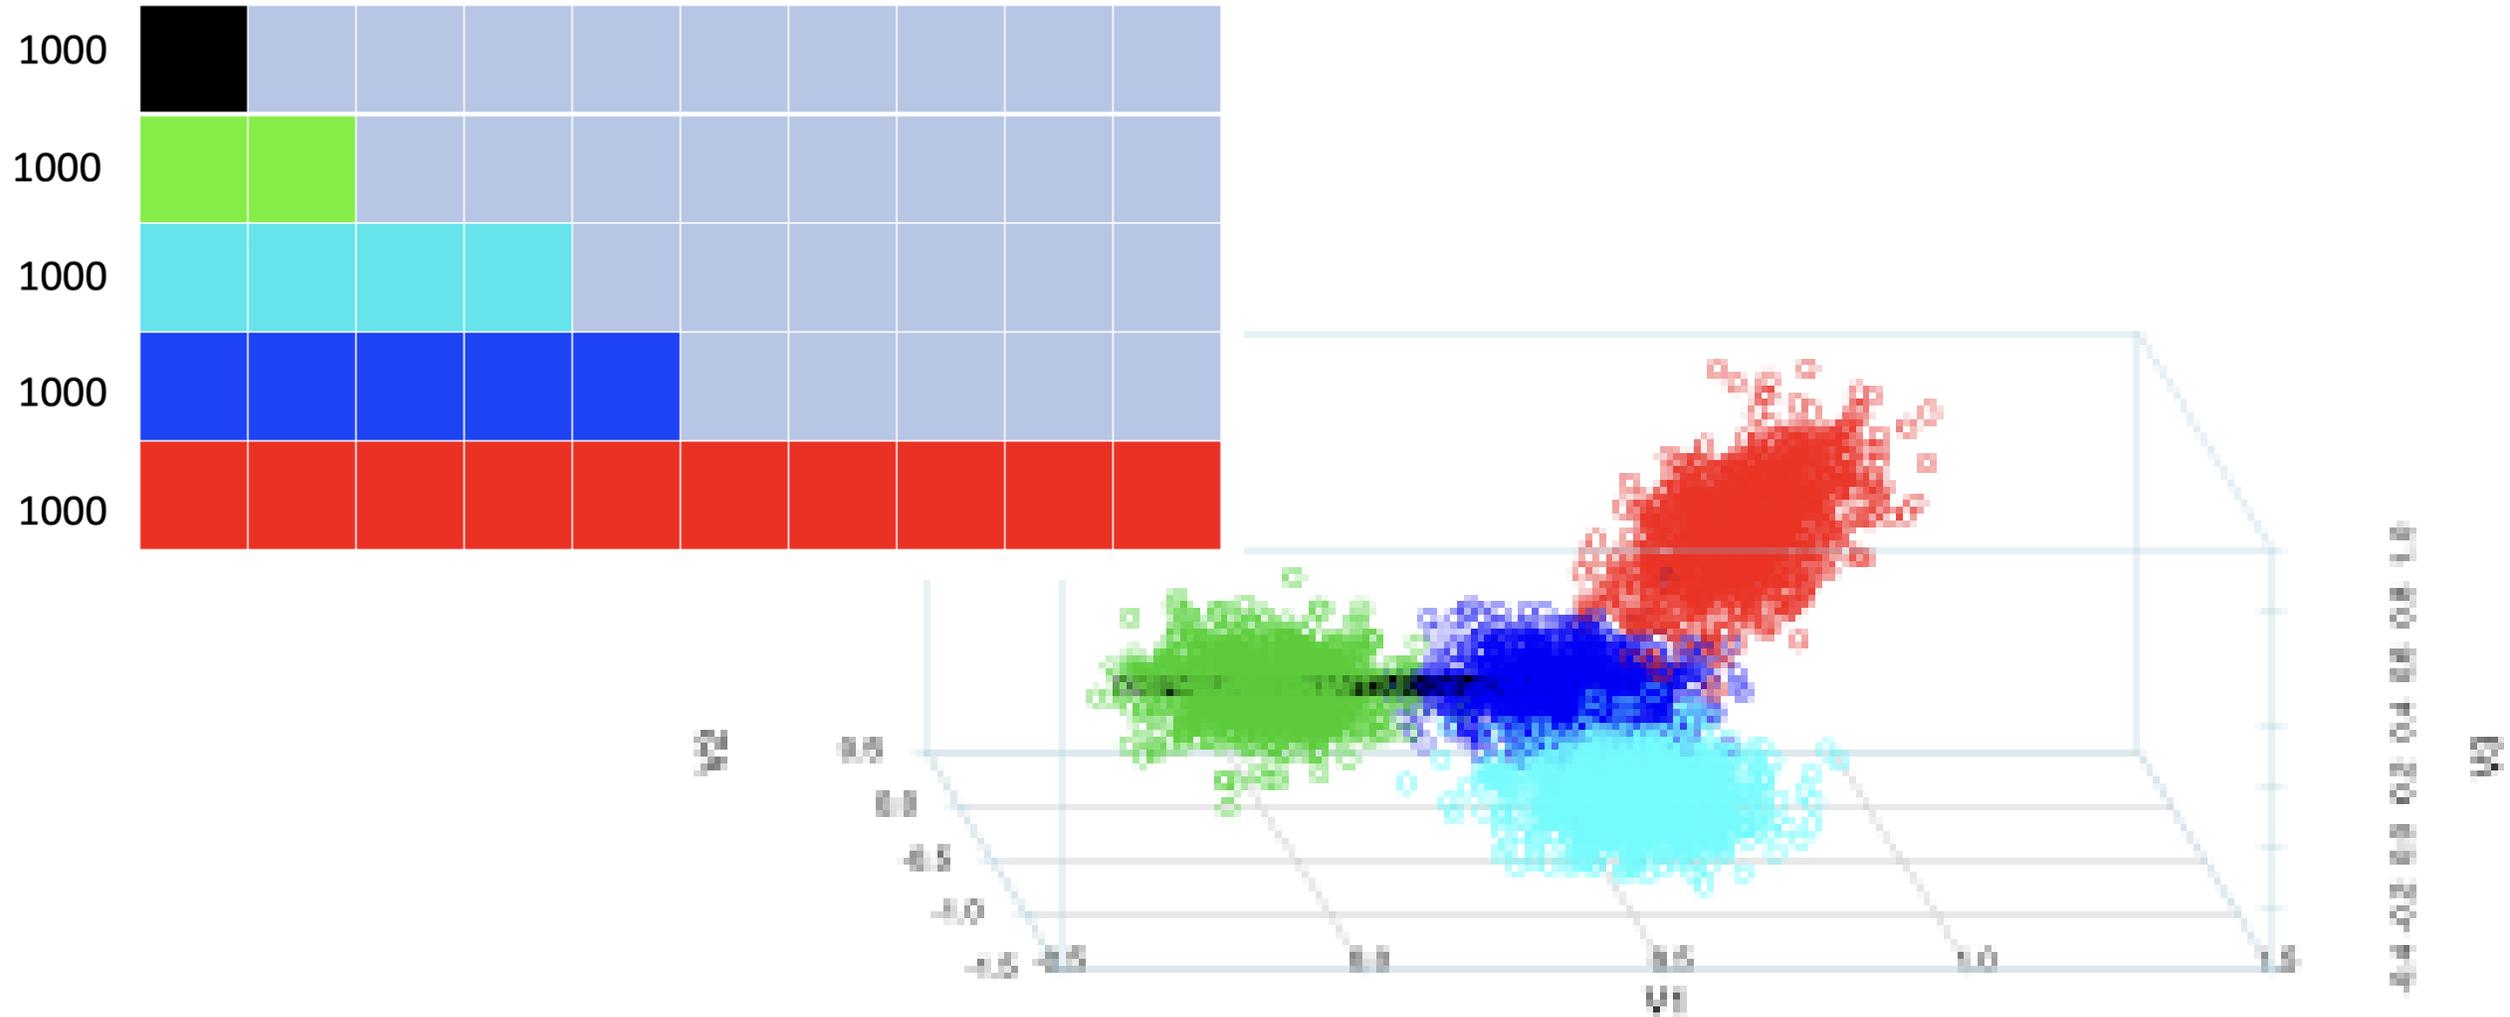
PROBLEM 1: Noise makes the identification of the ID hard



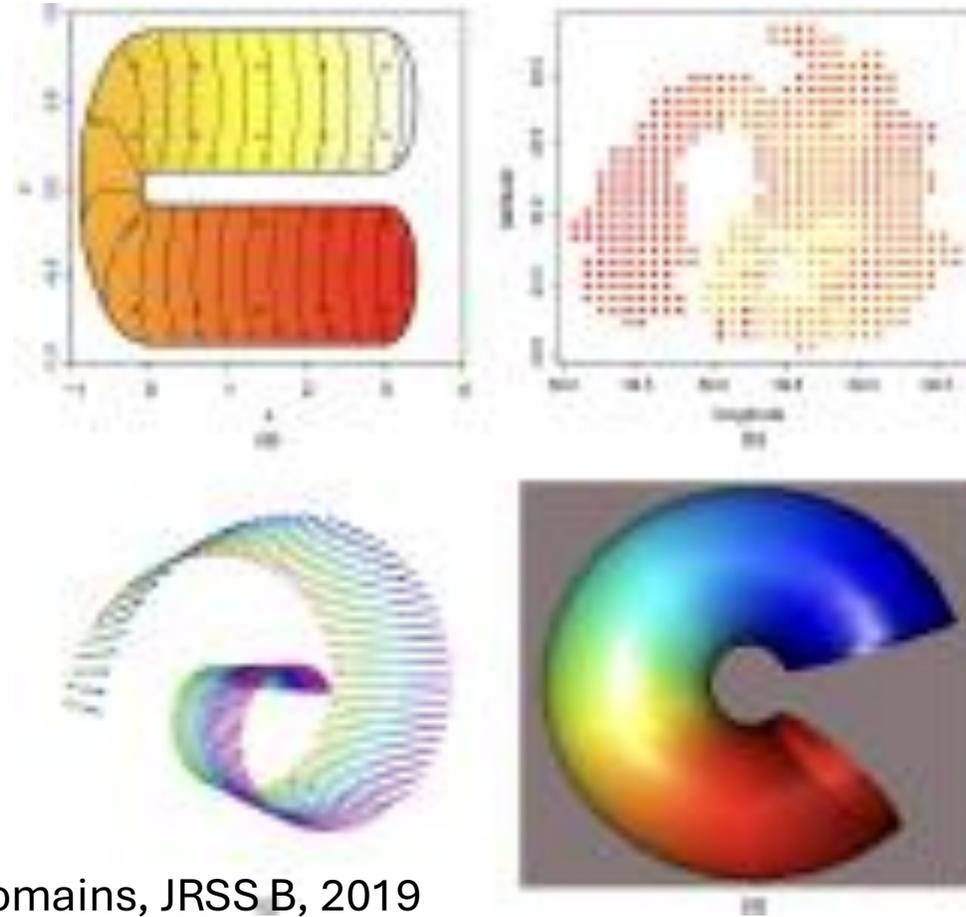
PROBLEM 2: Sometimes the ID can vary within the same dataset

*5 x 1000 observations generated from 5 Gaussians of dimensions 1, 2, 4, 5 and 9, partially overlapping:
3 dim projection*

Data matrix = 5000 x 9



PROBLEM 3: data can be discrete or continuous



D. Dunson et al. GP on constrained domains, JRSS B, 2019

A - a test function increases smoothly within a U-shaped boundary **B - remote sensed chlorophyll data in the Aral sea**
C - a spiralling band in a three-dimensional Euclidean space + data D - Bitten torus + data on the surface

Regions with the same ID host points differing in core properties:

- folded vs unfolded state in **protein** configurations
- active vs non-active regions in **brain** imaging data
- patients vs controls in **gene** expression data
- firms with different **financial risk** in balance sheets data
- winning vs losing teams in **basketball** data (AoAS)
- country specific NPI in **pandemic** evolution (Sci. Rep.)
- ID of an undirected unweighted **network** (Sci. Rep.)
- Identified vs unidentified models via ID in **MCMC**
- ID of layers in a **CNN** and transformes in **LLM**

A simple topological feature uncovers a rich data structure

What is the Intrinsic Dimension (ID)?

Dimension of the support of the data generating distribution

Different estimation approaches:

projective, topological, Nearest-Neighbours (NN) ...

Lack of statistically sound estimation procedures

Aim: directly target the ID as a parameter to estimate

Fundamental issue: scale dependence

Goal: statistical guarantees

- Allow for **uncertainty** quantification
- Rely on exact or **asymptotic** distributional results

NN based ID estimators: 2 classes

For each point i in the data set

- NN order class

fix two NN orders: $n_2 > n_1$

find distance from i to the two NN: $r_{i,n_2} > r_{i,n_1}$

statistics: $\mu_i = r_{i,n_2} / r_{i,n_1}$

- NN distance class

fix two distances: $r_2 > r_1$

count points within that distance from i : $n_{i,r_2} > n_{i,r_1}$

statistics: $n_{i,r_1} | n_{i,r_2}$

We find the distribution of the two statistics = function (d)

CLASS I: ID Estimators based on fixing NN order:

- TWO-NN = ratios of distances 2nd to 1st NN
- GRIDE = Generalized Ratios ID Estimator
- Hidalgo = Finite Mixture of TWO-NN ratios
- BNP Hidalgo = Infinite Mixture of TWO-NN ratios

CLASS II: ID Estimators based on fixing NN distance:

- I3D = ID for Discrete Data
- BIDE = Binomial ID Estimator
- ABIDE = Adaptive BIDE
- BABIDE = Bayesian ABIDE

PROBLEM SOLVING

TWO-NN requires a weaker local homogeneity assumption

GRIDE is more robust to noise in the data

Hidalgo and **BNP Hidalgo** robust if more than one ID exists

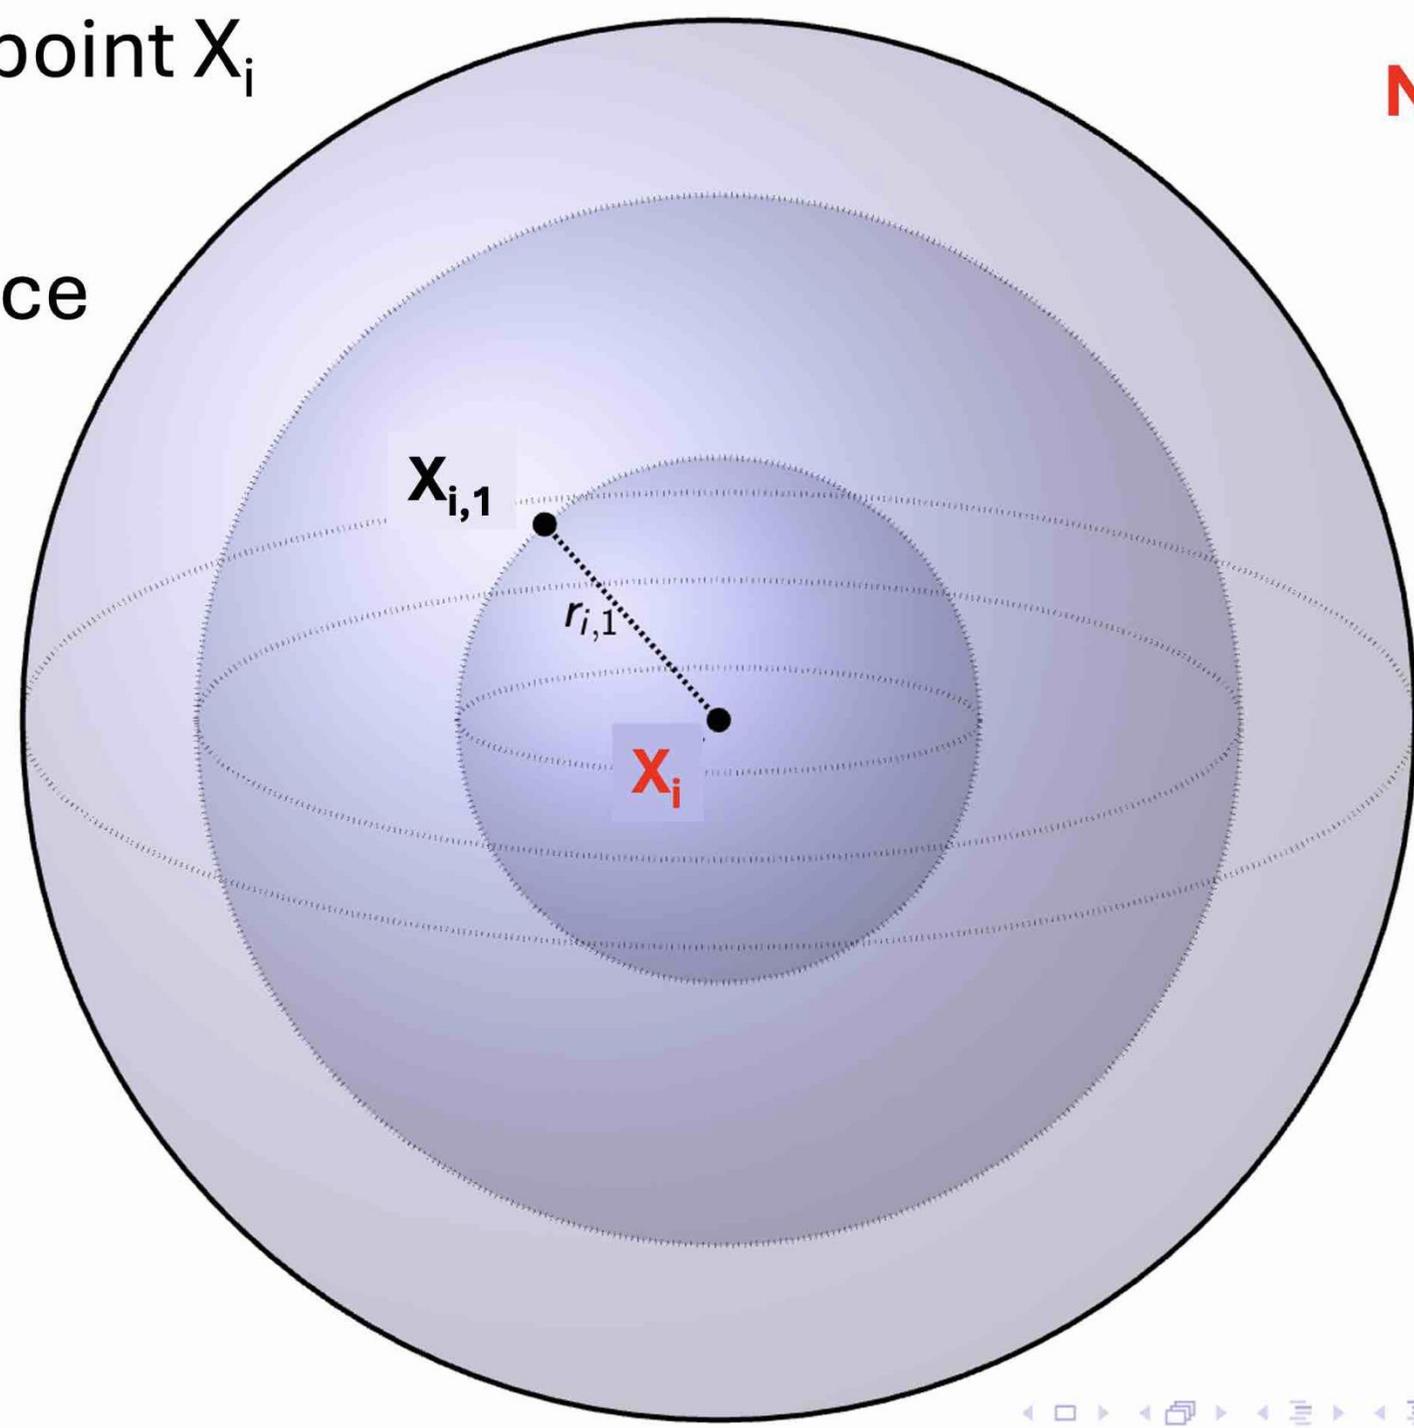
Adaptive ID estimators allow to escape the noise

Adaptive I3D for discrete data

$X_{i,1}$ = 1-st NN to point X_i

NN order class

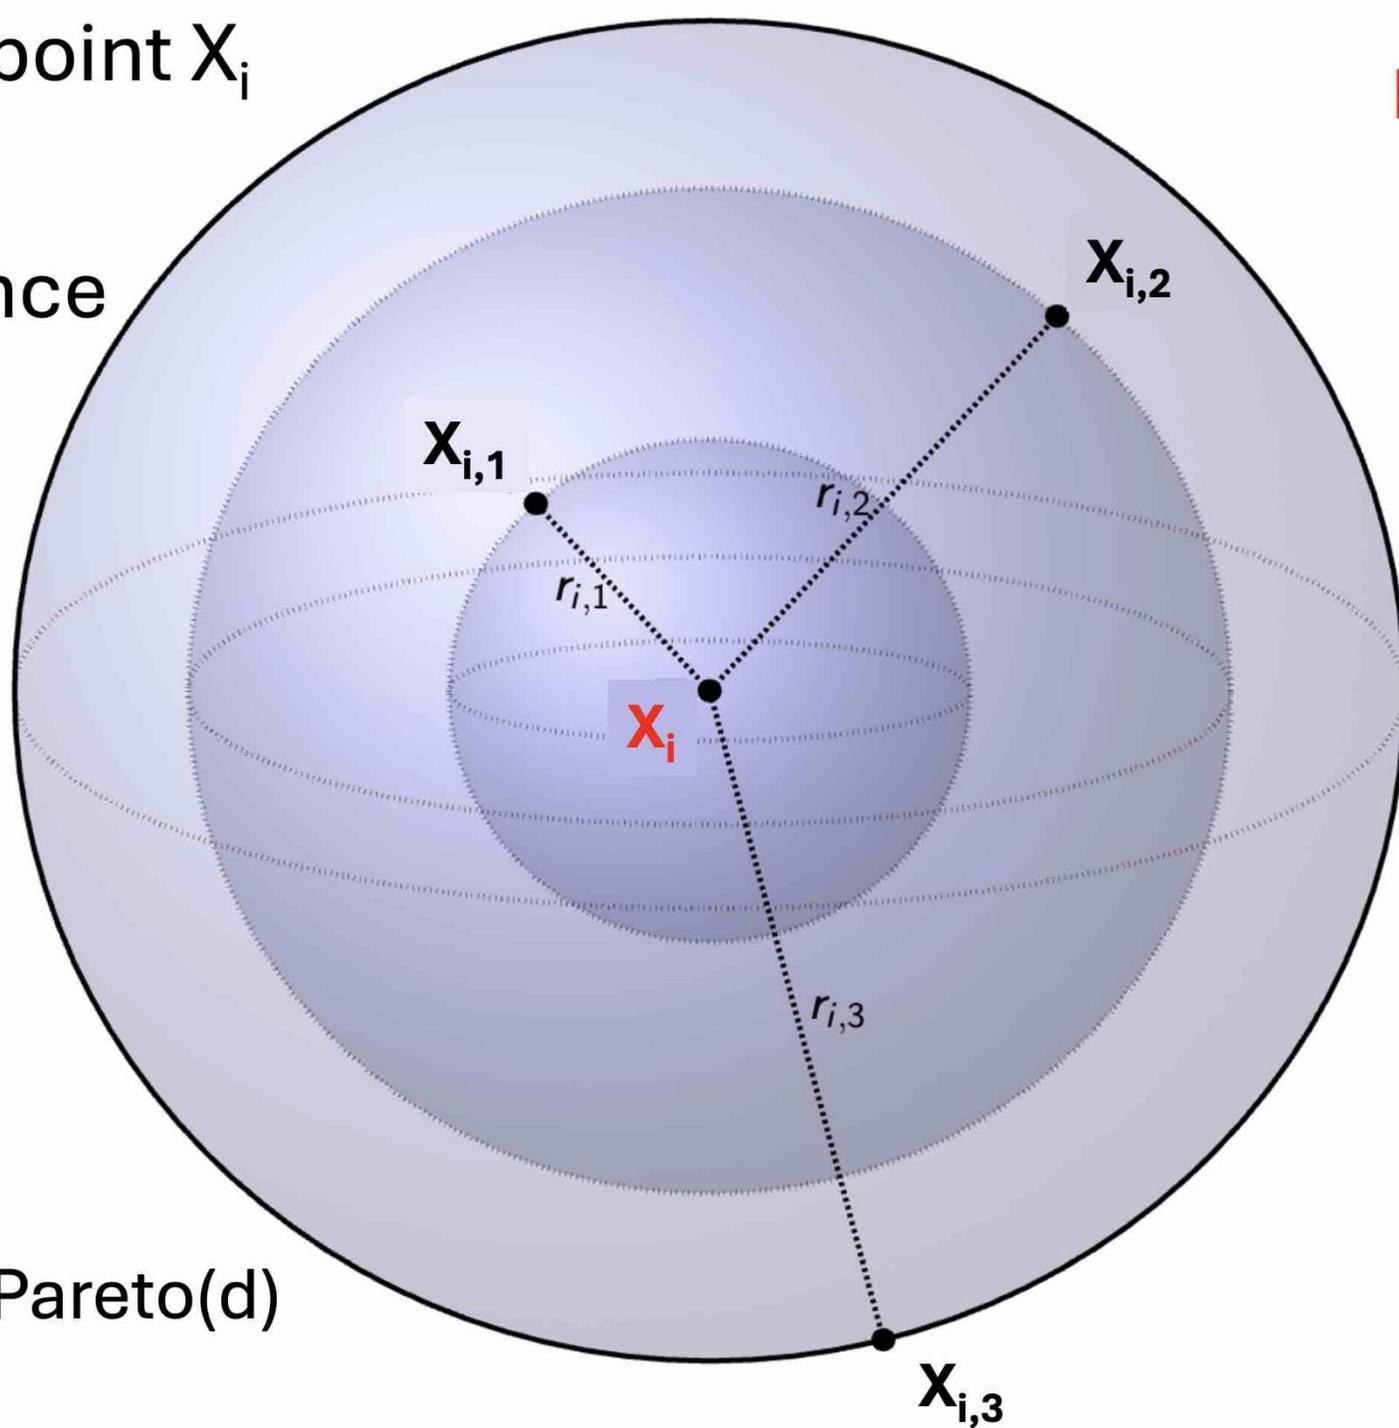
$r_{i,1}$ is their distance



$X_{i,j}$ = j-th NN to point X_i

NN order class

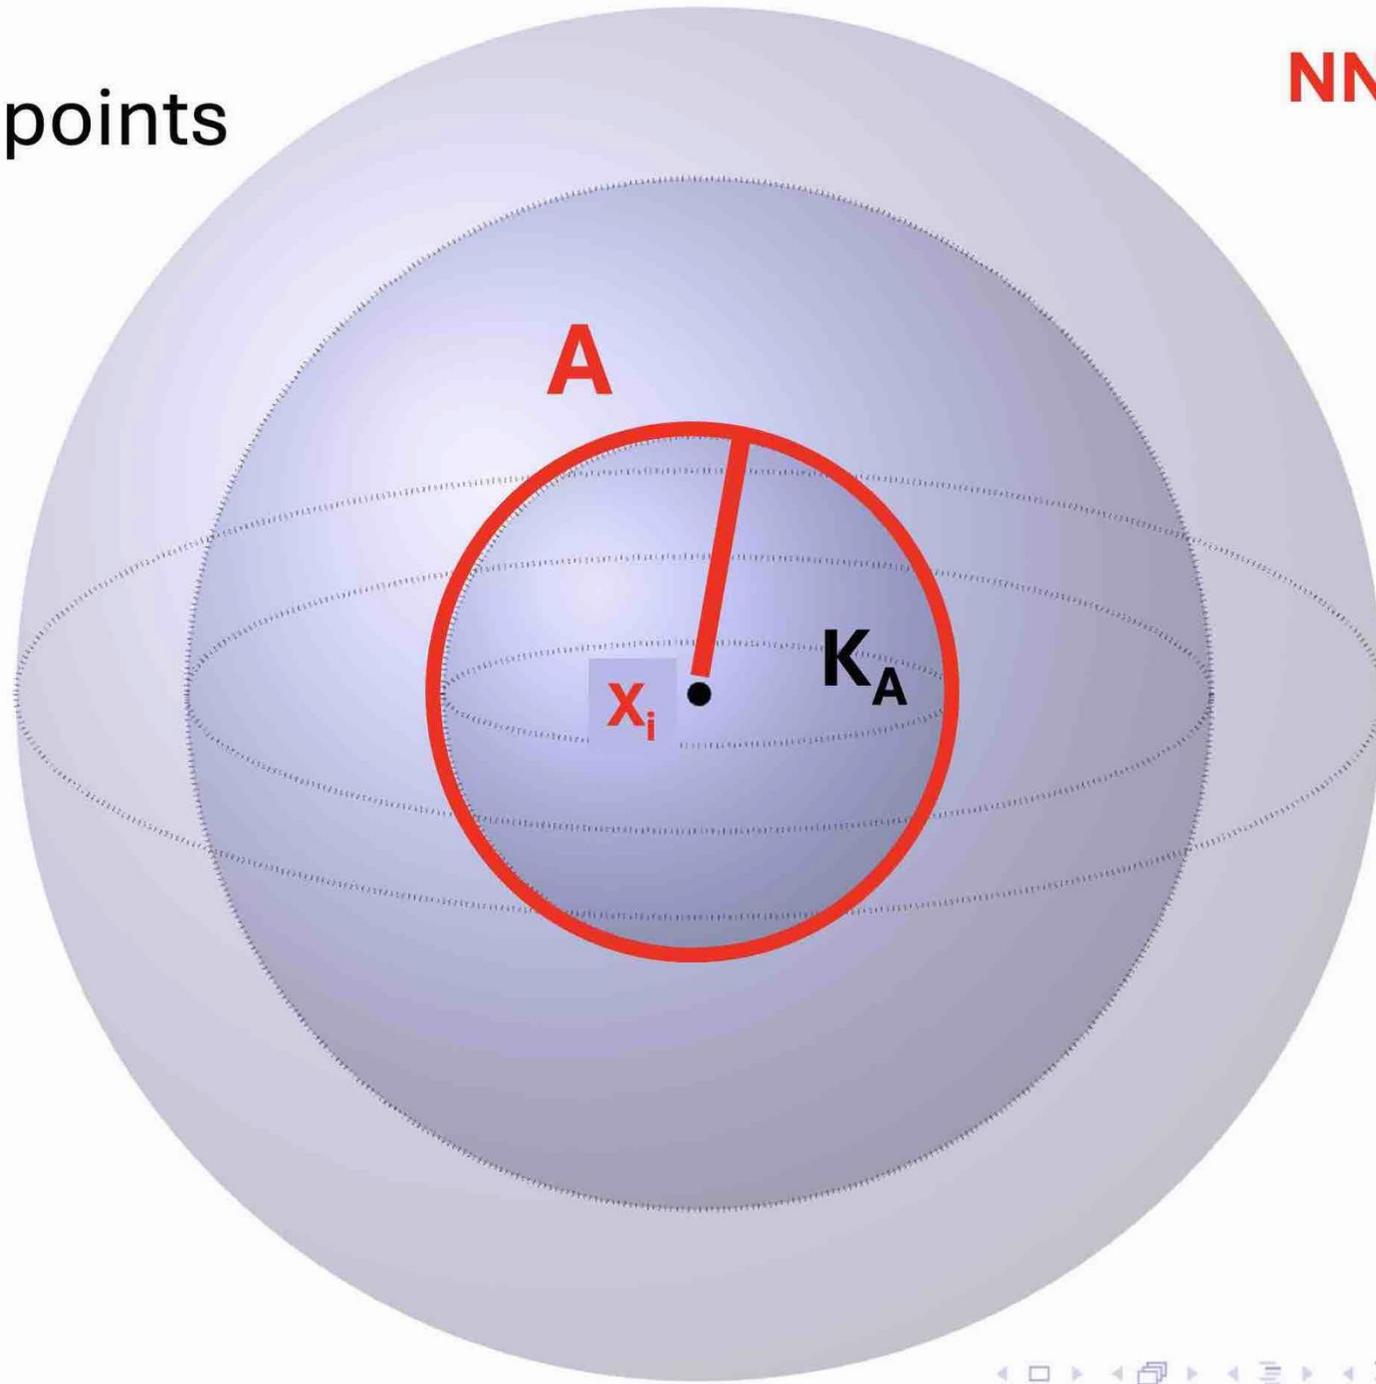
$r_{i,j}$ is their distance



$\mu_i = r_{i,2} / r_{i,1} \sim \text{Pareto}(d)$

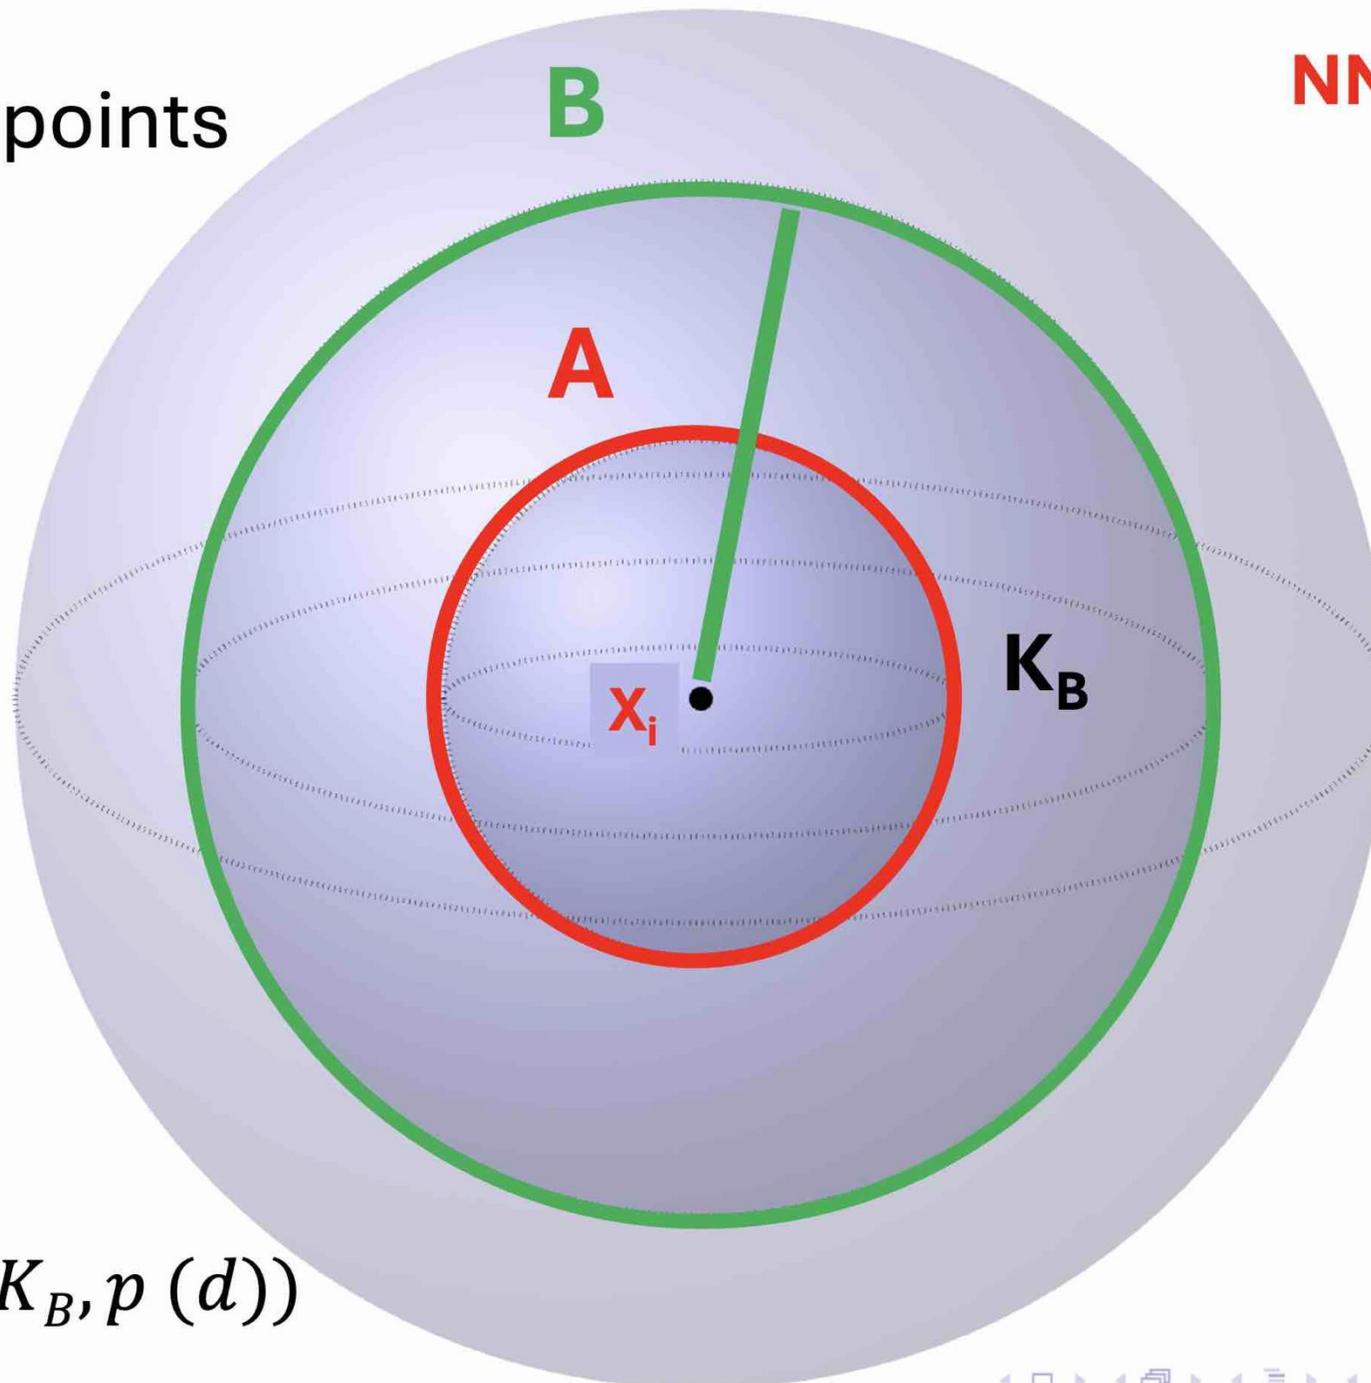
K_A = number of points
in ball A

NN distance class



K_B = number of points
in ball B

NN distance class



$$K_A | K_B \sim \text{Binom}(K_B, p(d))$$

- Let n_1, n_2 be two positive integers with $n_2 > n_1$ and define the ratio $\mu_{i,n_1,n_2} = r_{i,n_2}/r_{i,n_1}$ which is a.s. well defined for continuous data

Theorem

μ_{i,n_1,n_2} has density

$$f_{\mu_{i,n_1,n_2}}(\mu) = \frac{d(\mu^d - 1)^{n_2 - n_1 - 1}}{\mu^{d(n_2 - 1) + 1} B(n_2 - n_1, n_1)}, \quad \mu > 1 \quad (1)$$

where $B(\cdot, \cdot)$ is the Beta function and d is the id.

The proof (F. Denti) follows from the PPP local homegeneity assumption that implies that the volumes of the shell $v_{i,j}$ between the i -th and the j -th NN have an exponential distribution with parameter ρ_i

- Take $n_1 = 1$ and $n_2 = 2$, the density in (1) reduces to a *Pareto*(1, d).
- **TWO-NN MLE estimator** is given by

$$\hat{d} = \frac{n - 1}{\sum_{i=1}^n \log(\mu_i)}, \quad (2)$$

with $1 - \alpha$ confidence interval $\tau_{1-\alpha} = \left[\frac{\hat{d}}{q_{IG_{n,n-1}}^{1-\alpha/2}}, \frac{\hat{d}}{q_{IG_{n,n-1}}^{\alpha/2}} \right]$.

- **TWO-NN Bayesian estimator**: prior $d \sim \text{Gamma}(a, b)$

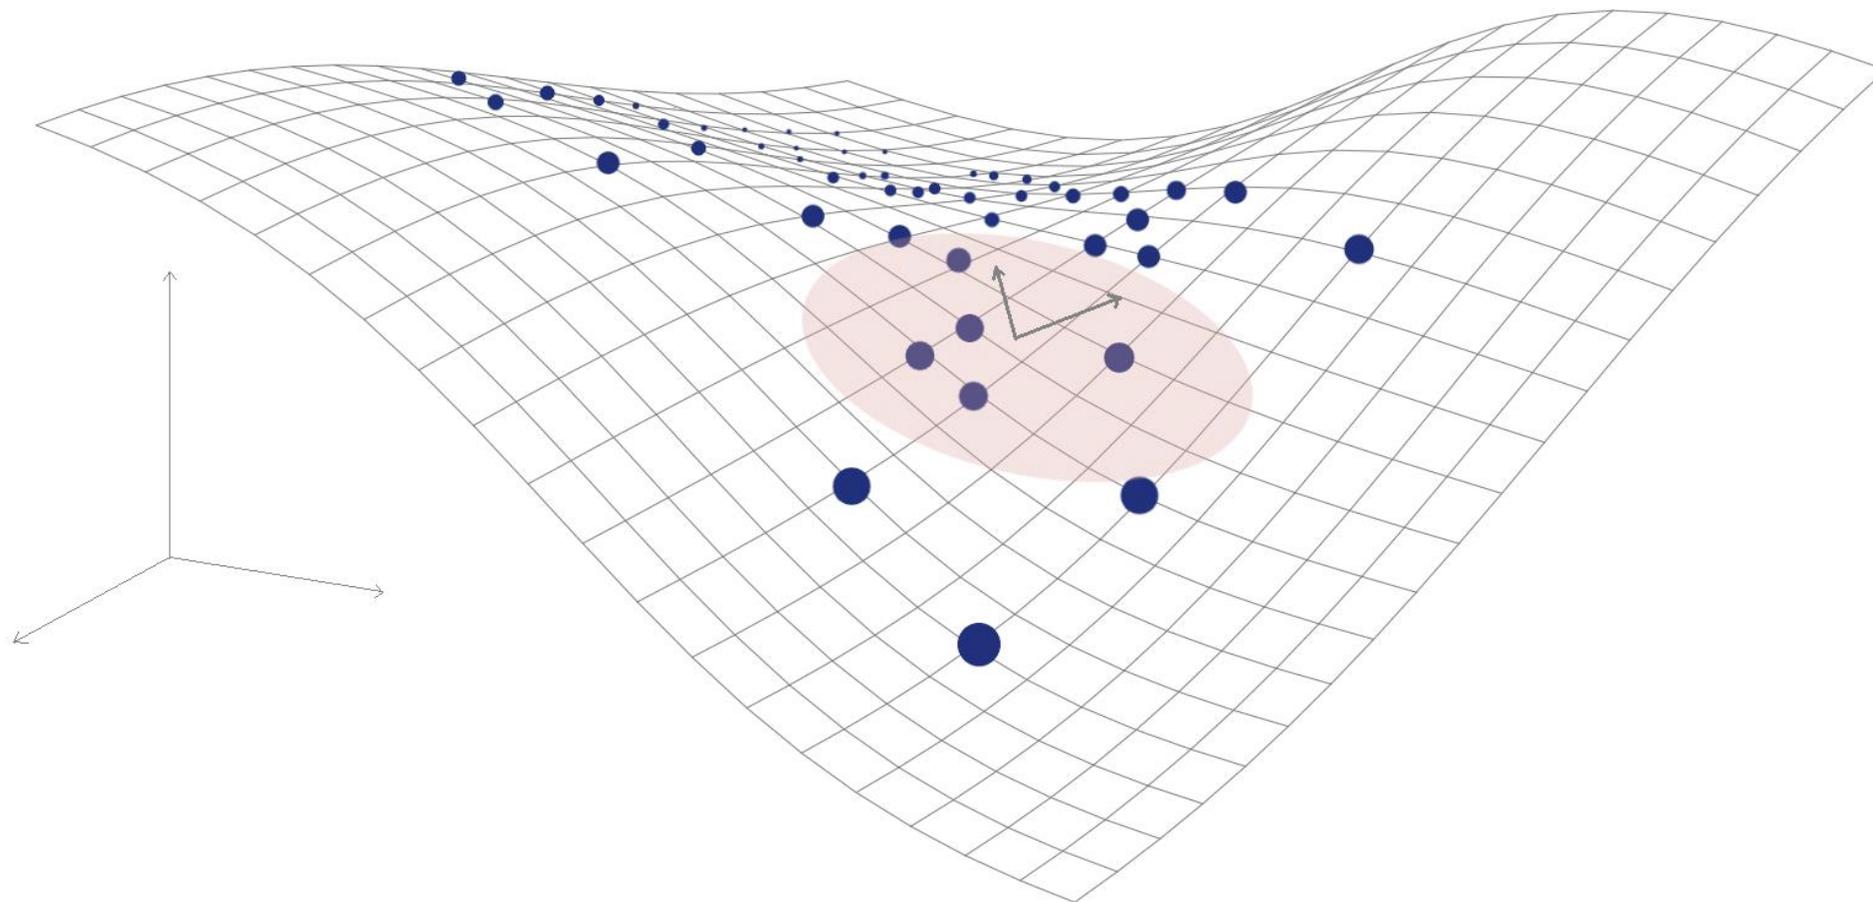
$$d | \mu_1, \dots, \mu_n \sim \text{Gamma}\left(a + n, b + \sum_{i=1}^n \log(\mu_i)\right), \quad (3)$$

whose mean is asymptotically equivalent to (2)

- **GRIDE**: $n_1 = n_2/2$

D = 3 = embedding dimension

d = 2 = intrinsic dimension



Body
temp.

Heart
rate

Blood
pressure

Person 1

Person 2

Person 3

Person 4

Person 5

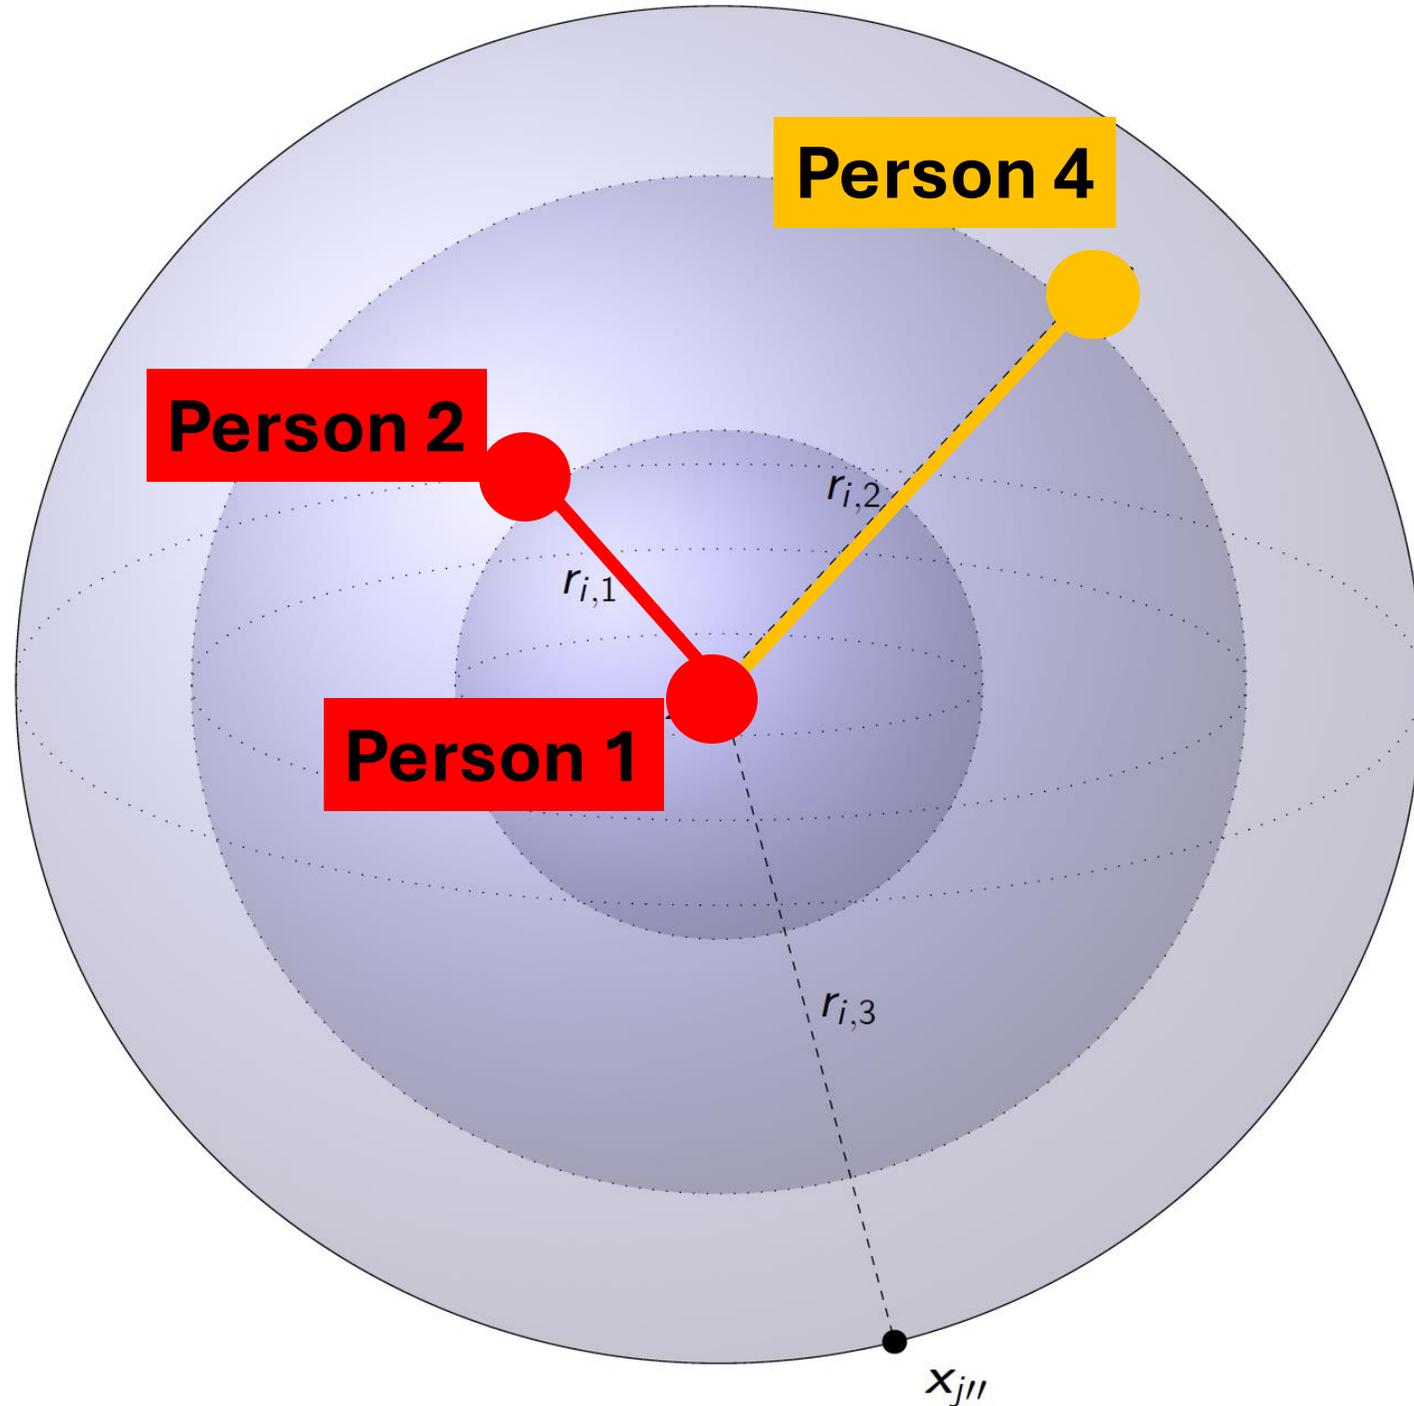
D = 3

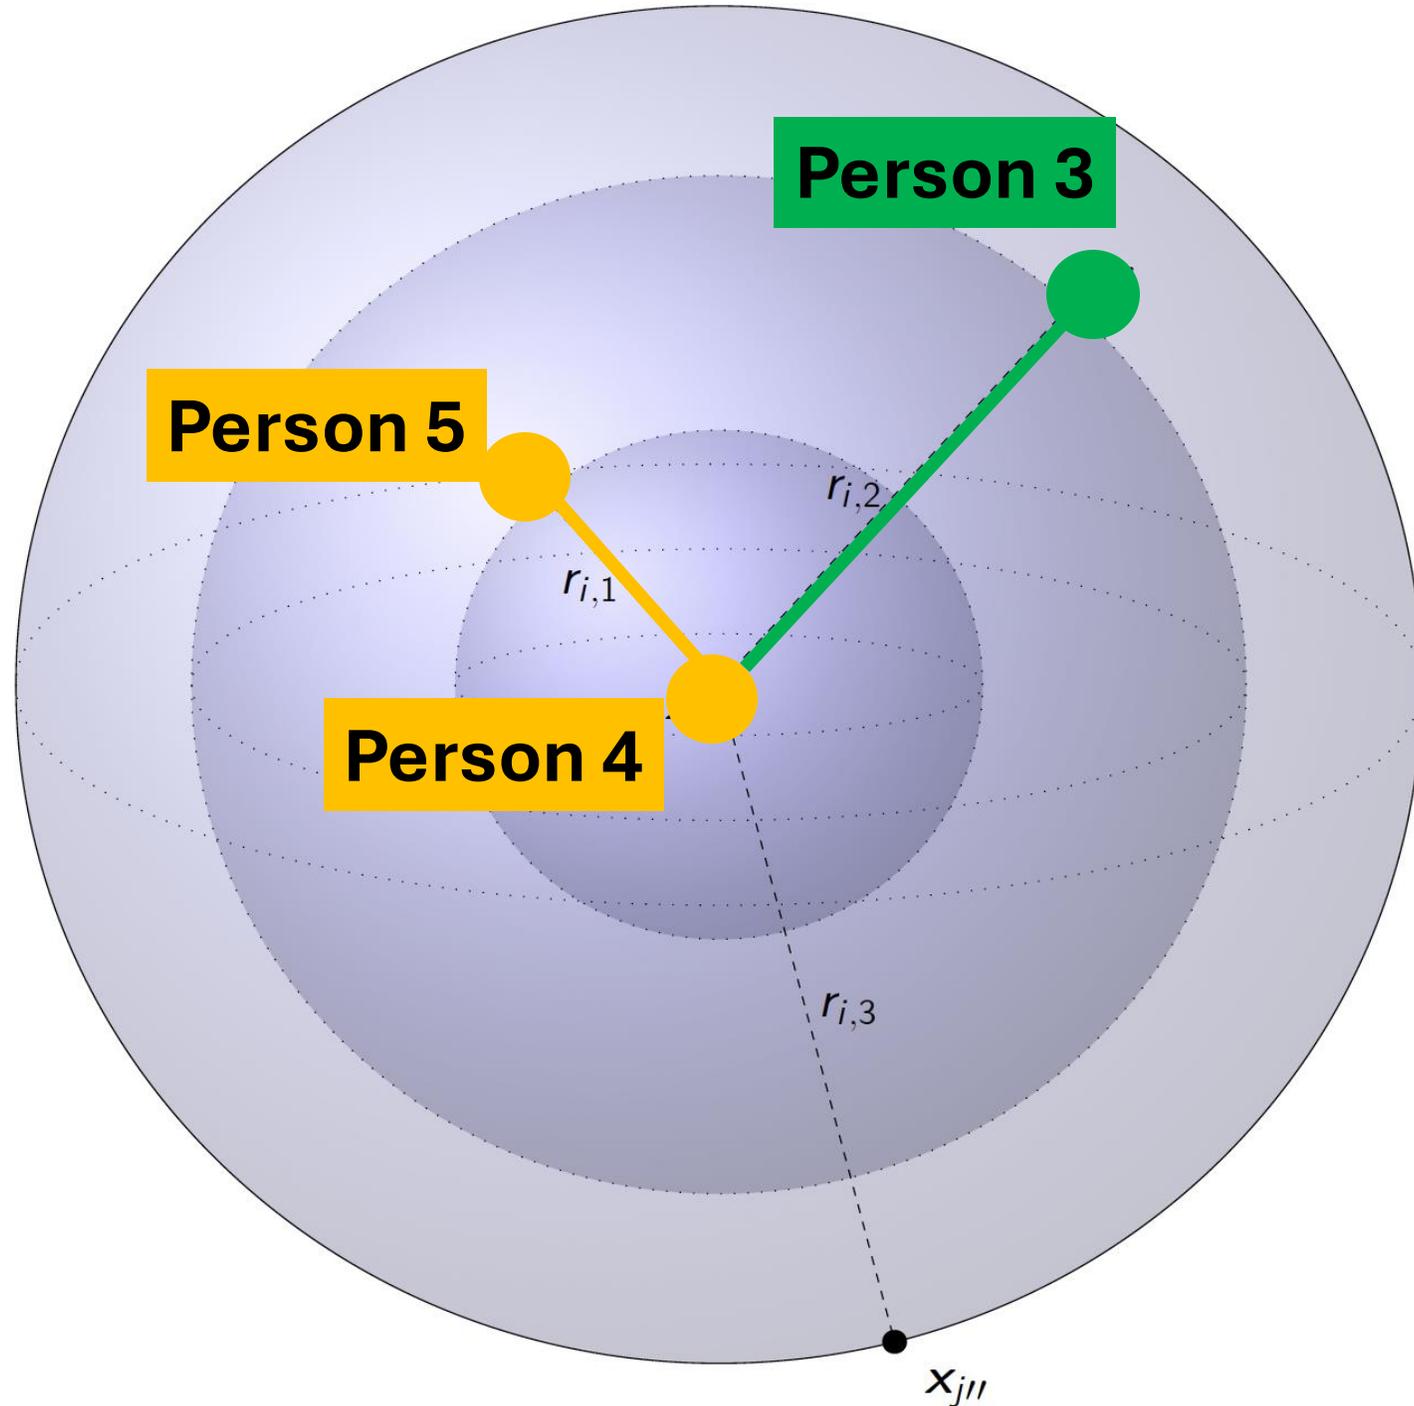
d1 = 0

d2 = 2

d3 = 3

N = 5





Person 3

Person 5

Person 4

x_{jII}

$r_{i,1}$

$r_{i,2}$

$r_{i,3}$

How to deal with heterogeneous ID case? HIDALGO!

Heterogeneous ID algorithm - Hidalgo model

allows for the possibility that the ID may not be uniform in the dataset.

Under some assumptions the distribution of $\mu_i = r_{i2}/r_{i1}$ is a **mixture of Pareto** distributions

$$f(\mu_i) = \sum_{k=1}^K p_k d_k \mu_i^{-d_k-1}$$

The **likelihood** of the data is

$$\mathcal{L}(\boldsymbol{\mu}|\mathbf{d}, \mathbf{p}) = \prod_{i=1}^N \sum_{k=1}^K p_k d_k \mu_i^{-d_k-1}$$

where $\boldsymbol{\mu} = (\mu_1 \dots \mu_N)$

Then we can again estimate

$$\mathbf{d} = (d_1 \dots d_K), \quad \mathbf{p} = (p_1 \dots p_K)$$

The **likelihood** of the data is

$$\mathcal{L}(\boldsymbol{\mu}|\mathbf{d}, \mathbf{p}) = \prod_{i=1}^N \sum_{k=1}^K p_k d_k \mu_i^{-d_k-1}$$

where $\boldsymbol{\mu} = (\mu_1 \dots \mu_N)$

Then we can again estimate

$$\mathbf{d} = (d_1 \dots d_K), \quad \mathbf{p} = (p_1 \dots p_K)$$

Fix $P_{prior}(\mathbf{d}, \mathbf{p})$ and compute the posterior means

$$P_{post}(\mathbf{d}, \mathbf{p}) \propto \mathcal{L}(\boldsymbol{\mu}|\mathbf{d}, \mathbf{p})P_{prior}(\mathbf{d}, \mathbf{p})$$

Independent priors on \mathbf{d} and \mathbf{p}

Prior on \mathbf{d} : $d_k \sim \text{Gamma}(a_0, b_0), \quad k = 1, \dots, K$

Prior on $\mathbf{p} \sim \text{Dir}(\alpha_1, \dots, \alpha_K)$

Prior on $\mathbf{Z}|\mathbf{p} \sim$ discrete distribution on $(1, \dots, K)$ w.p. \mathbf{p}

- To adopt a full Bayesian approach, we need to address the uncertainty on the number of mixture components K
- Instead of making K **stochastic**, we adopt a Bayesian nonparametric approach, letting $K \rightarrow \infty$

We now model μ_i as a infinite mixture of Pareto distributions:

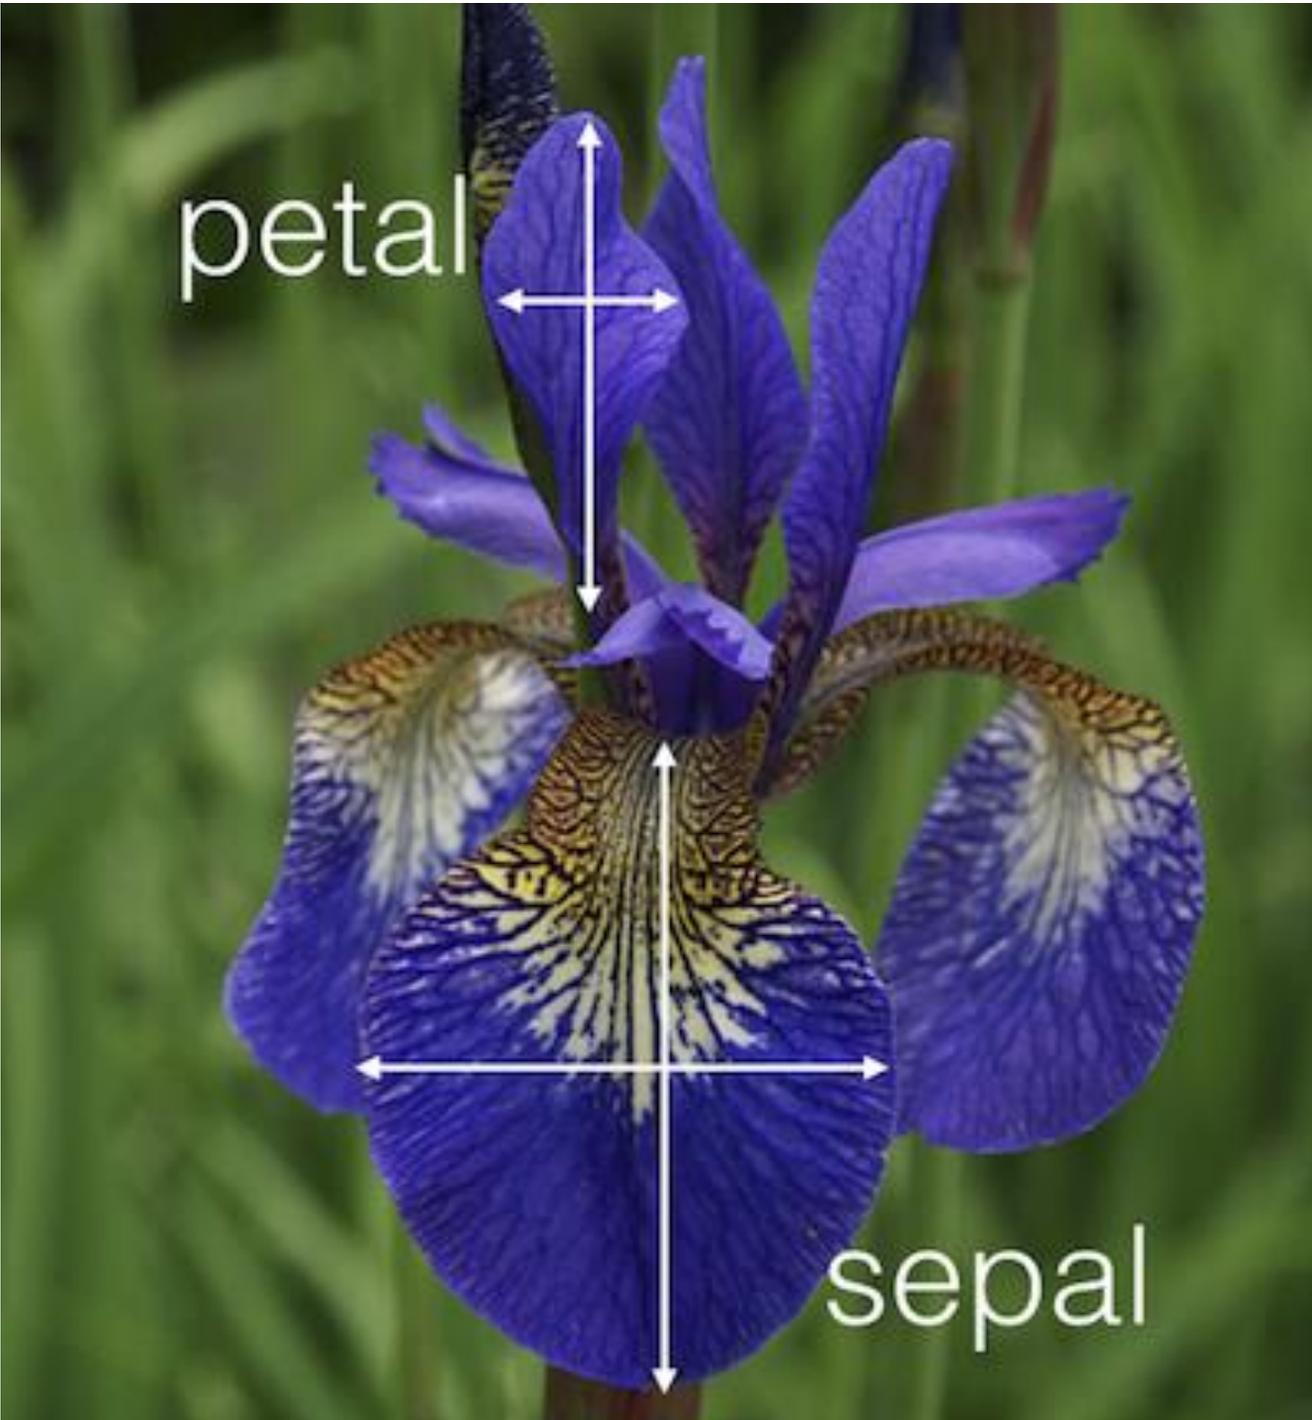
$$\sum_{i=1}^{+\infty} p_i \cdot \mathcal{P}(\mu_i | d_i)$$

We adopt a Dirichlet process prior for the parameters that model the ID



Three types of Iris flowers

$N = 50$



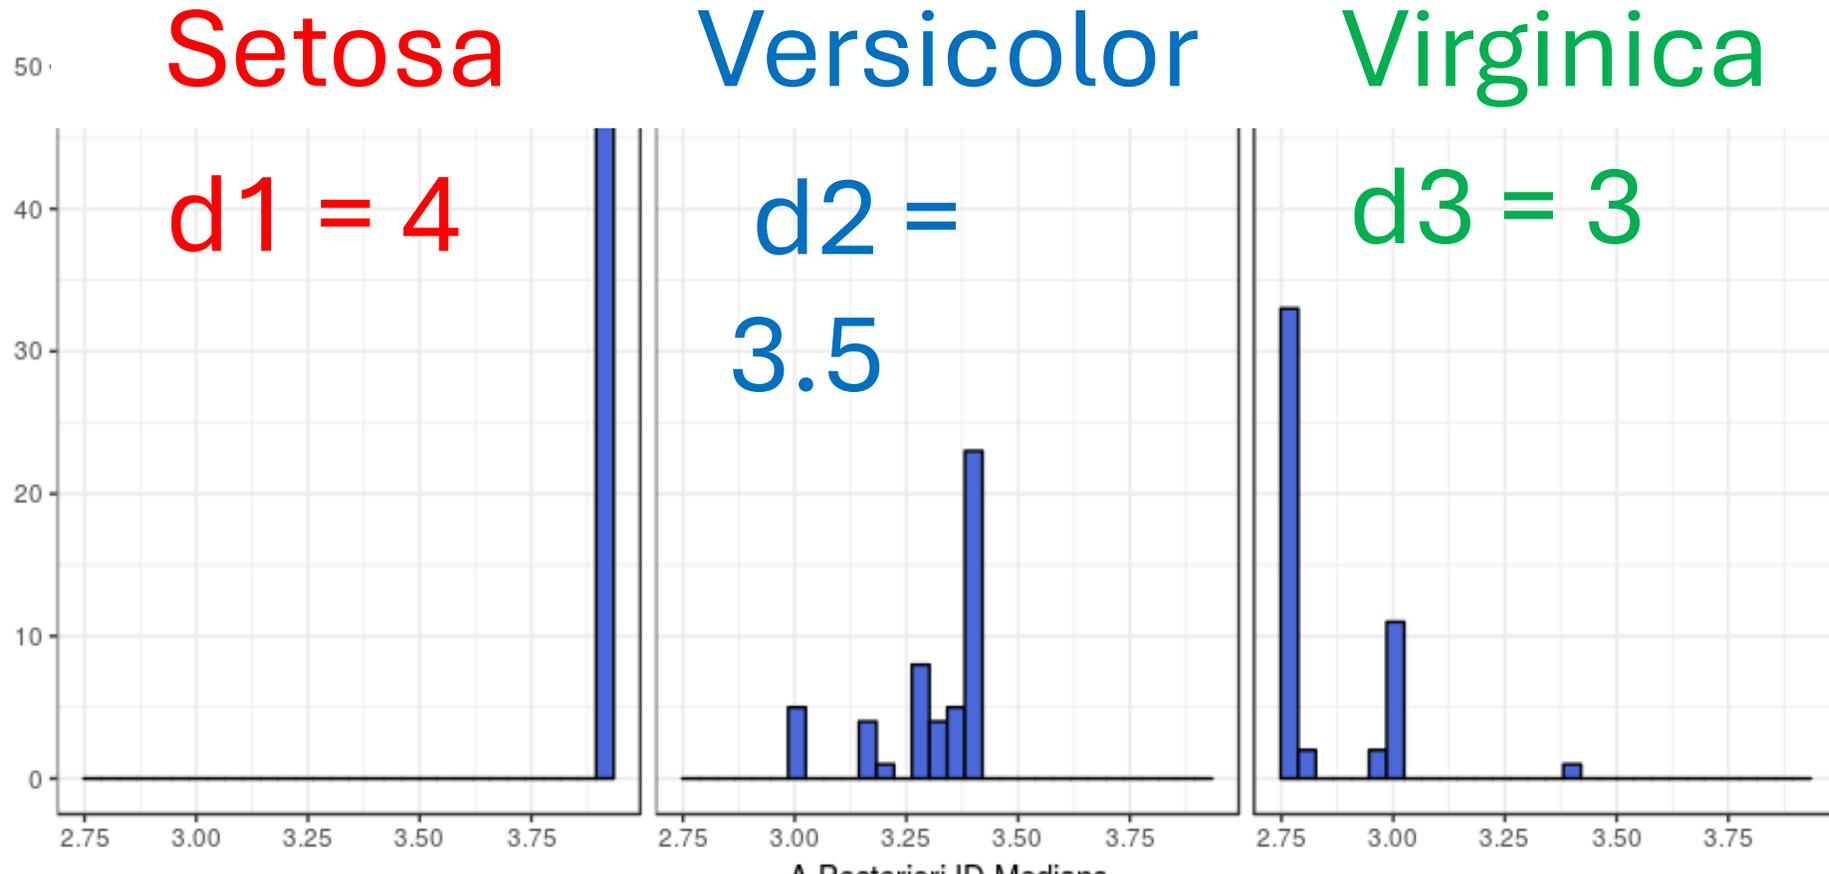
D = 4

	Petal length	Petal width	Sepal length	Sepal width
Flower 1				
Flower 2				
Flower 3				
Flower 4				
⋮				
Flower 50				

D = 4

N = 50

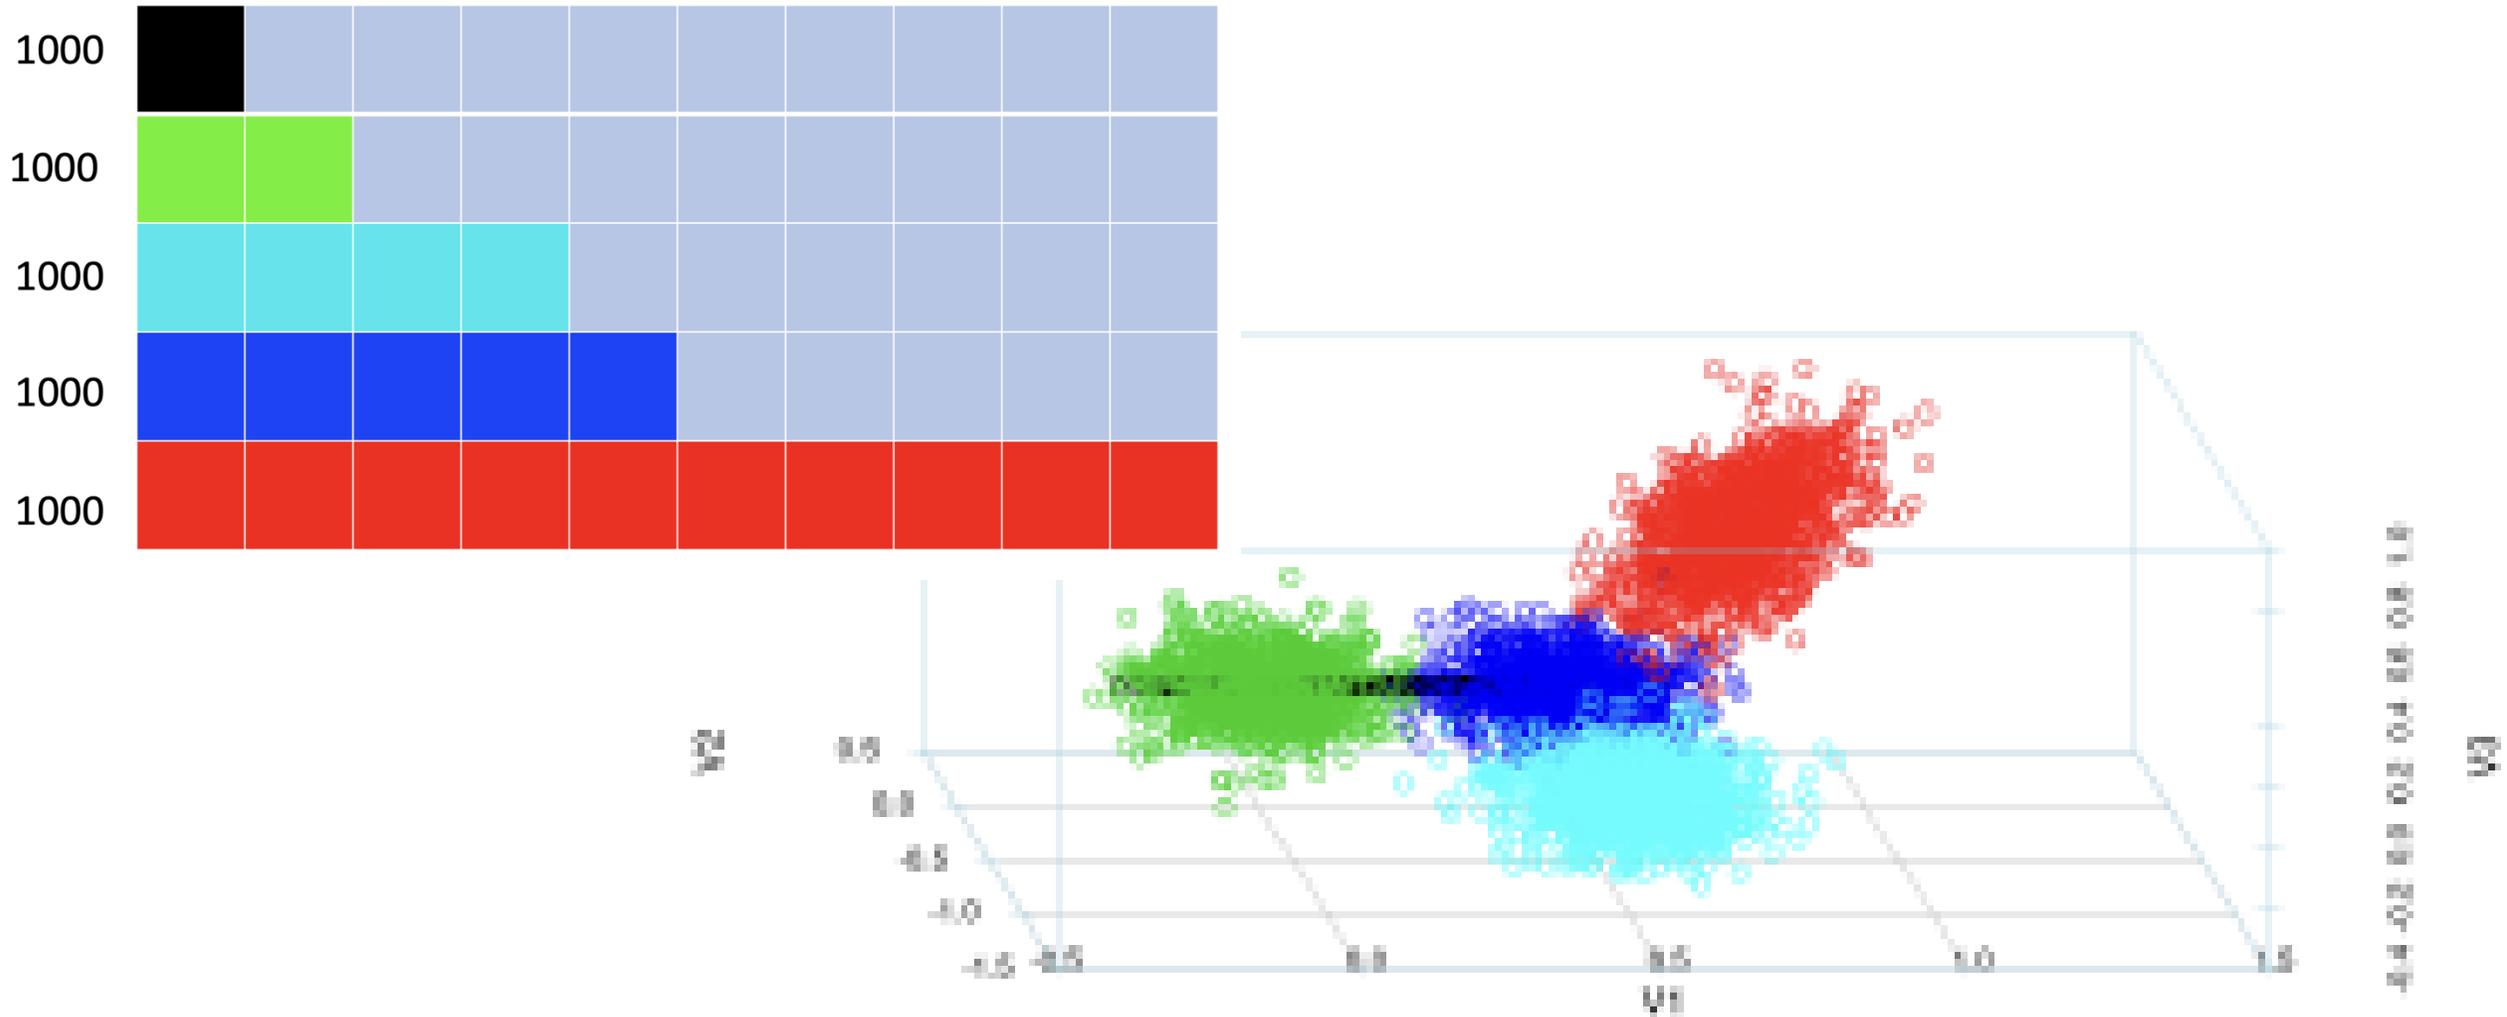
*3 clusters
almost coincident
with the flower species*



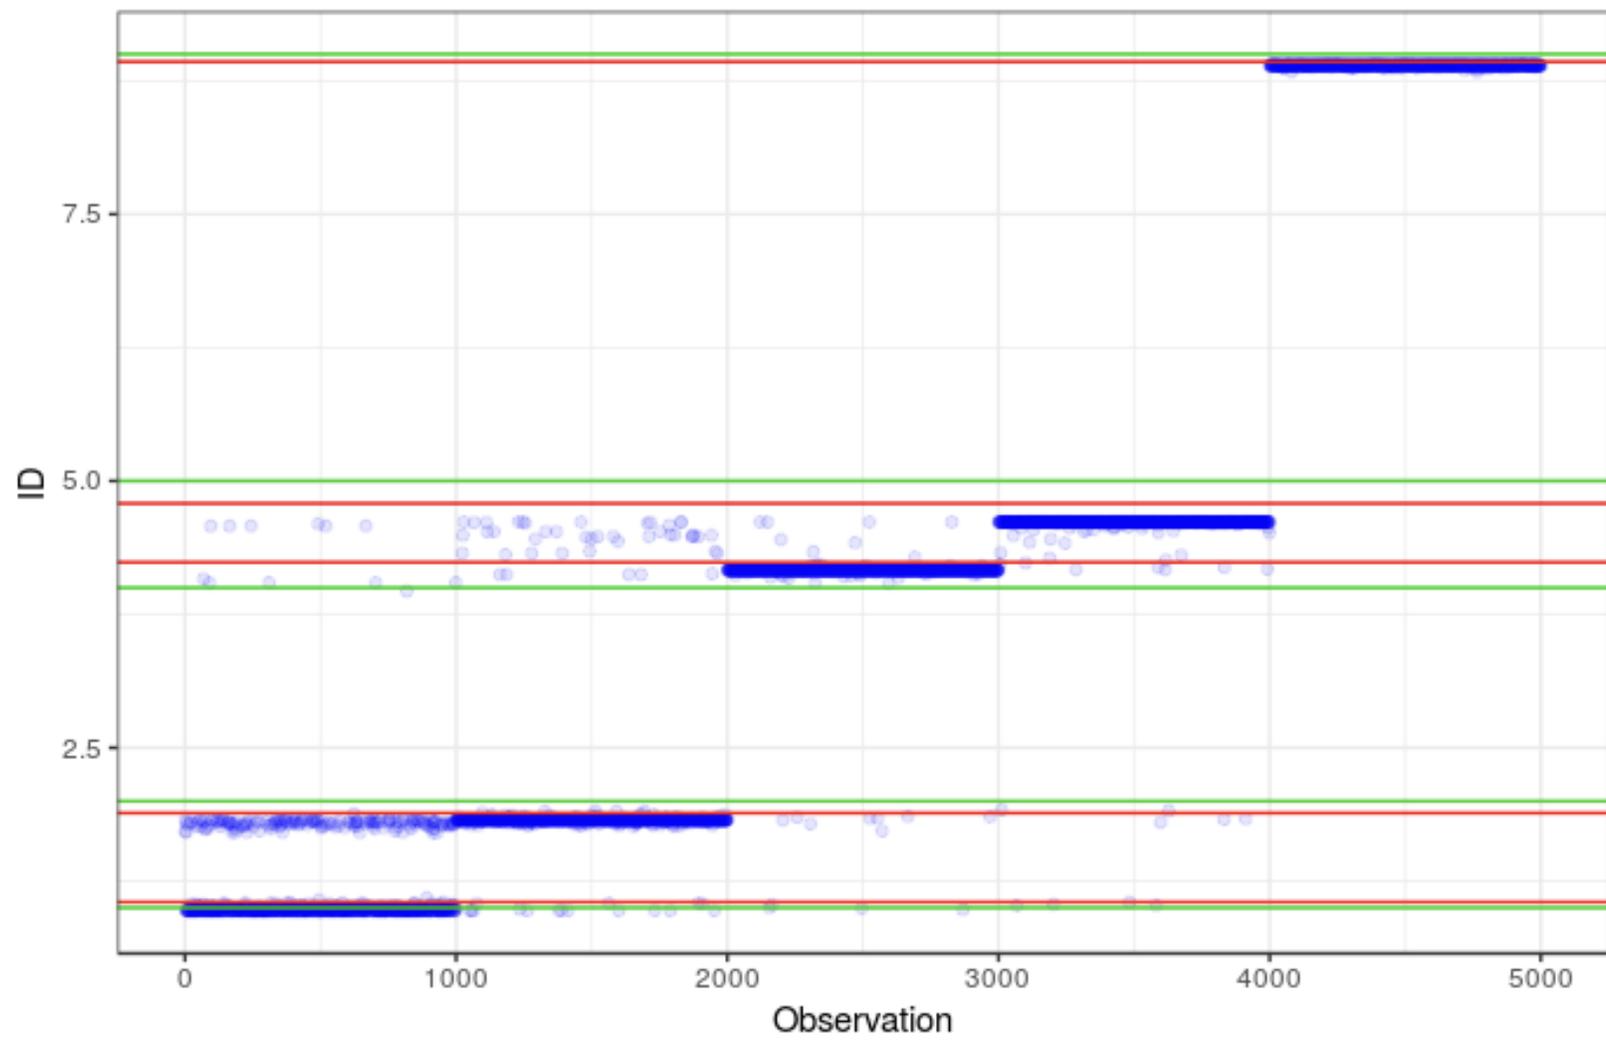
Simulated data

Data matrix = 5000 x 9

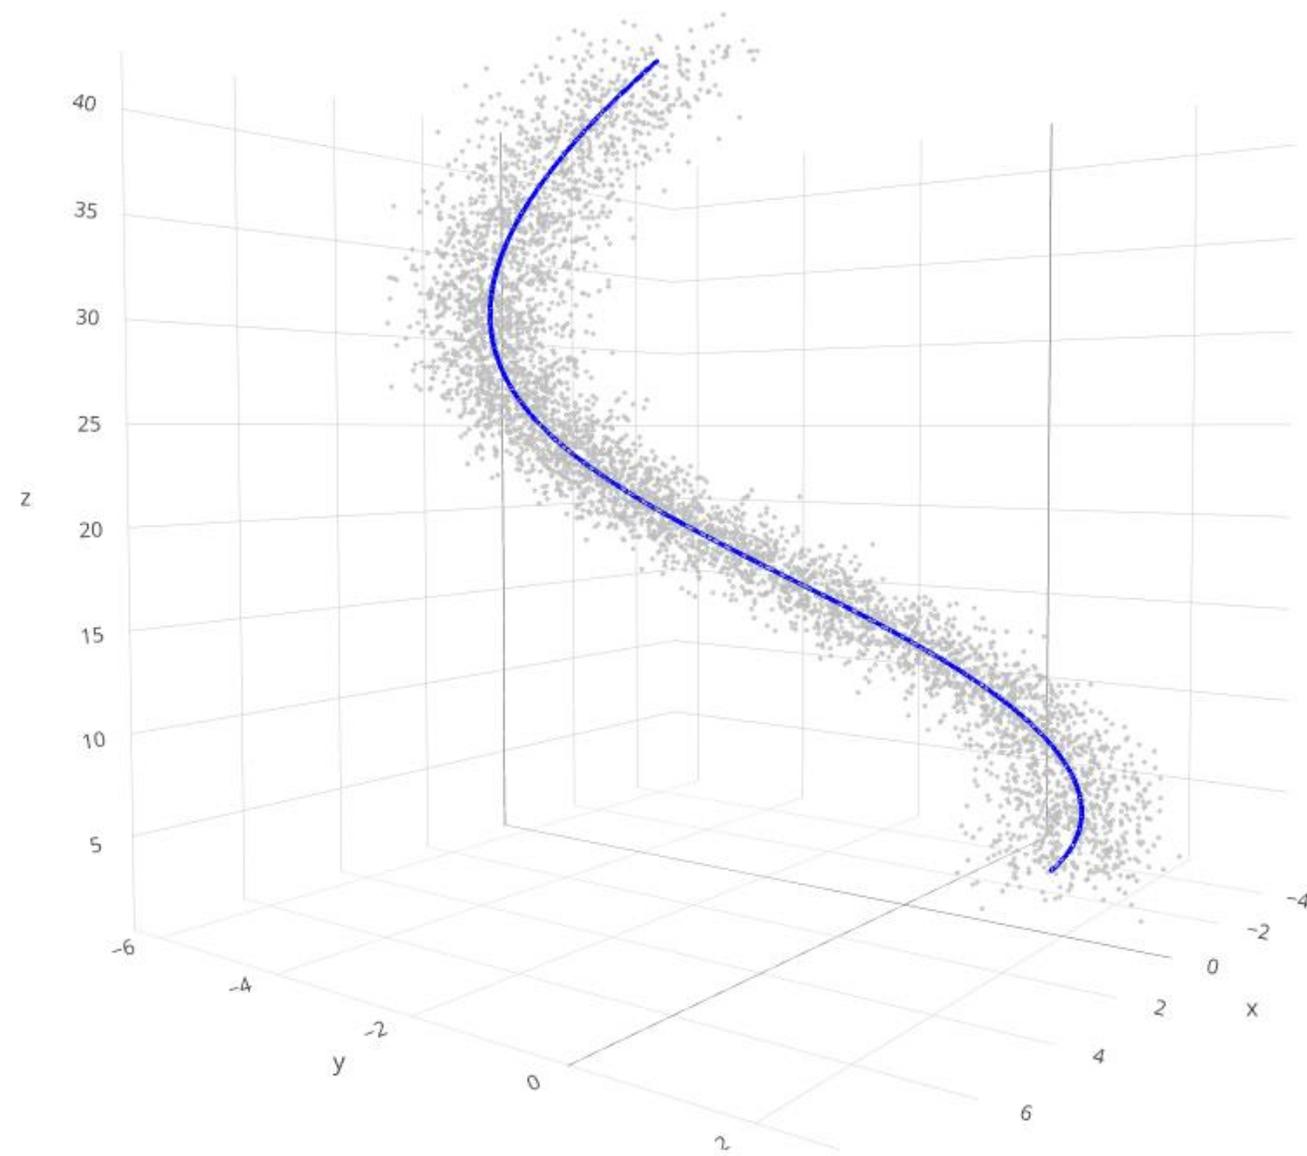
*5 x 1000 observations generated from 5 Gaussians of dimensions 1, 2, 4, 5 and 9, partially overlapping:
3 dim projection*

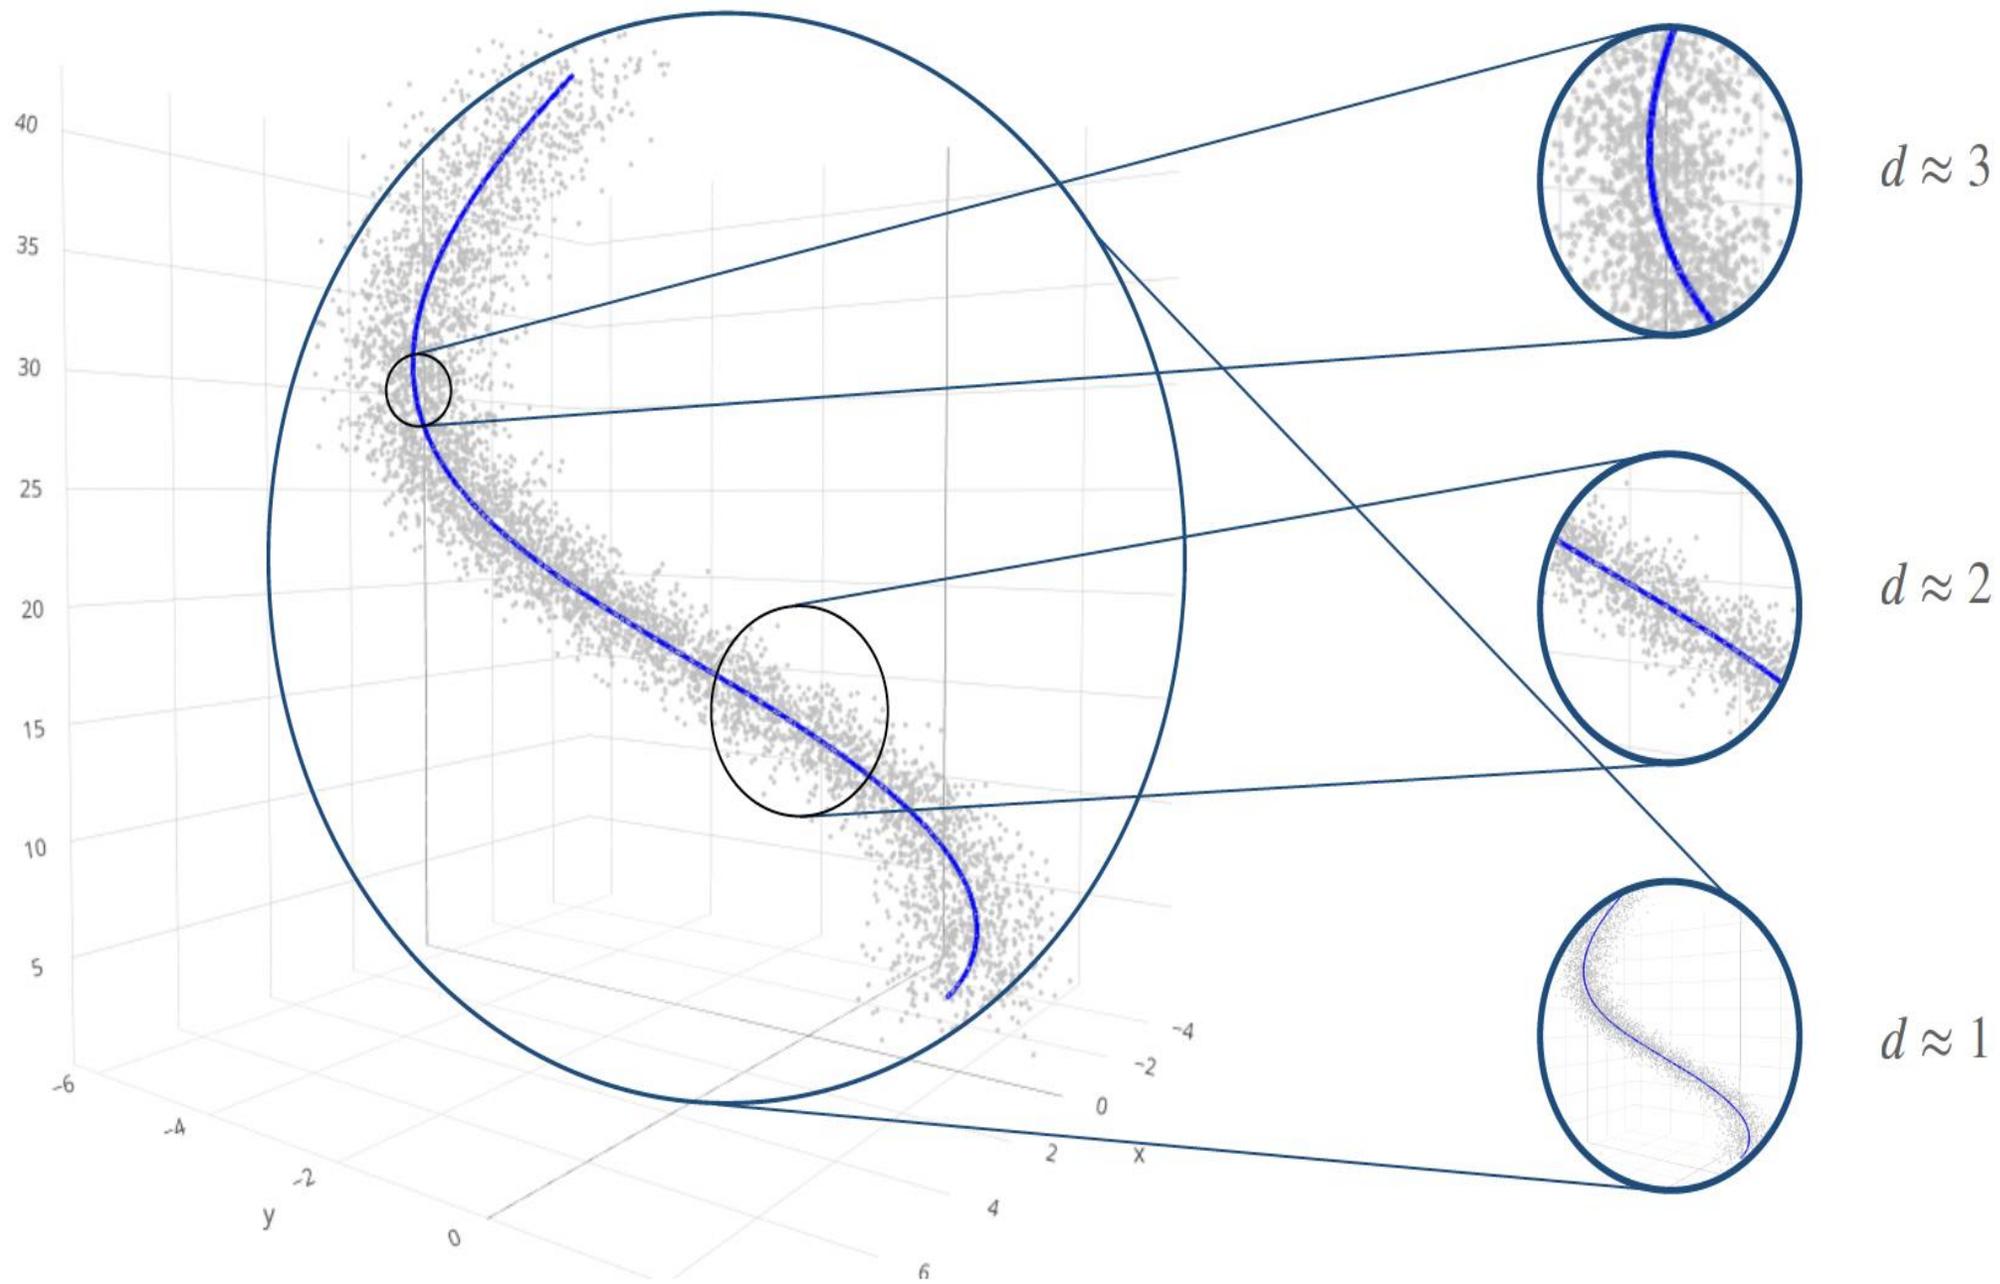


Posterior medians of d_i



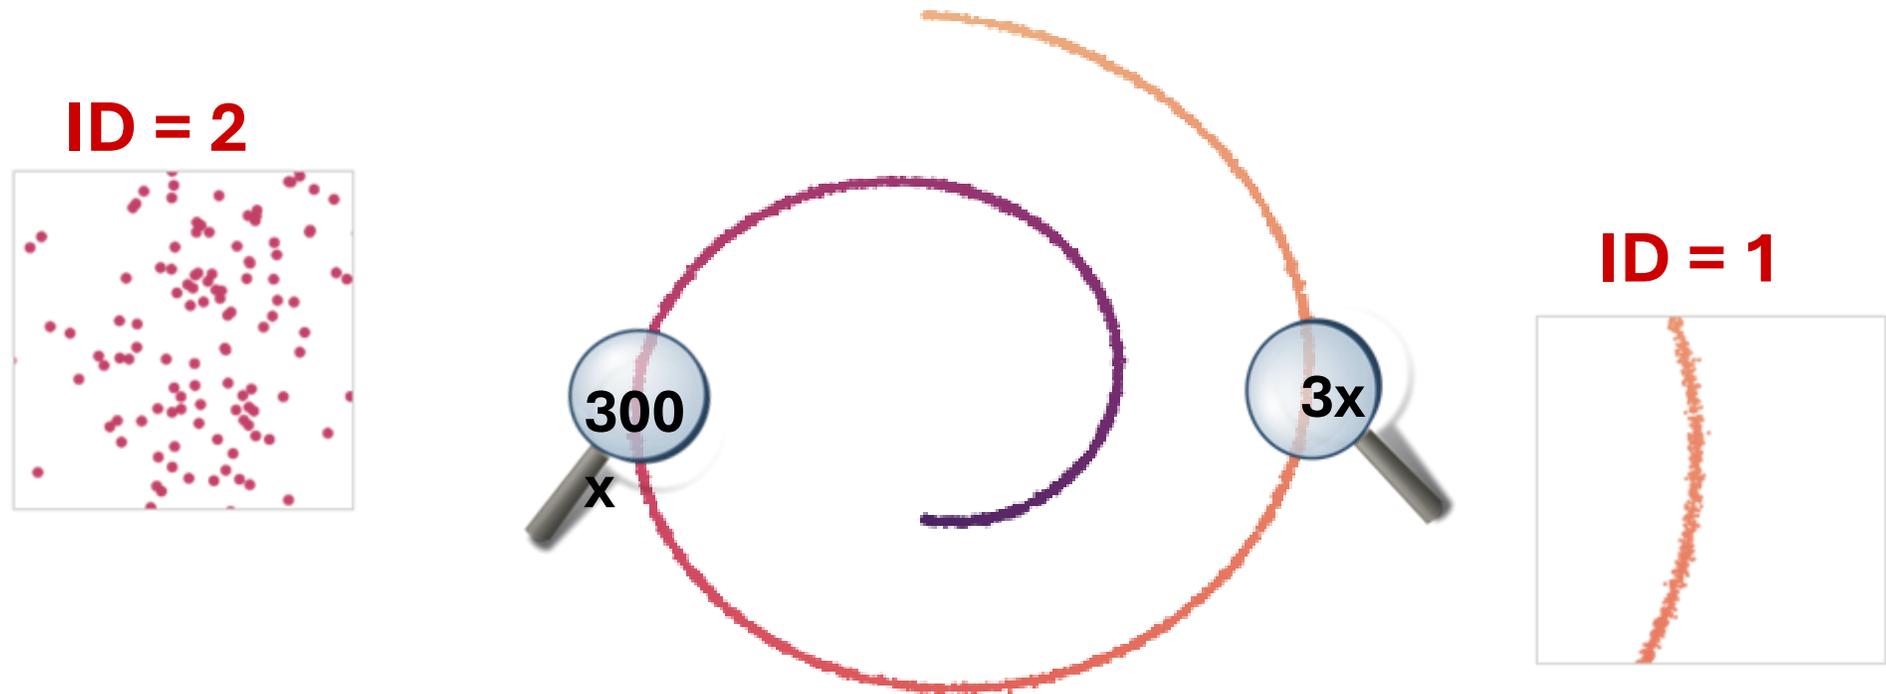
NOISE





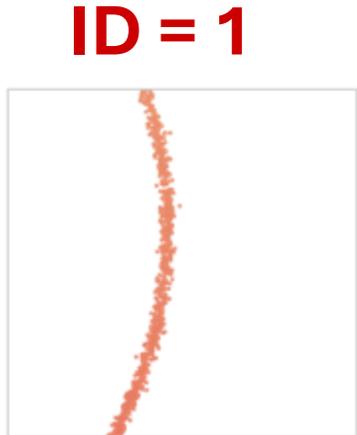
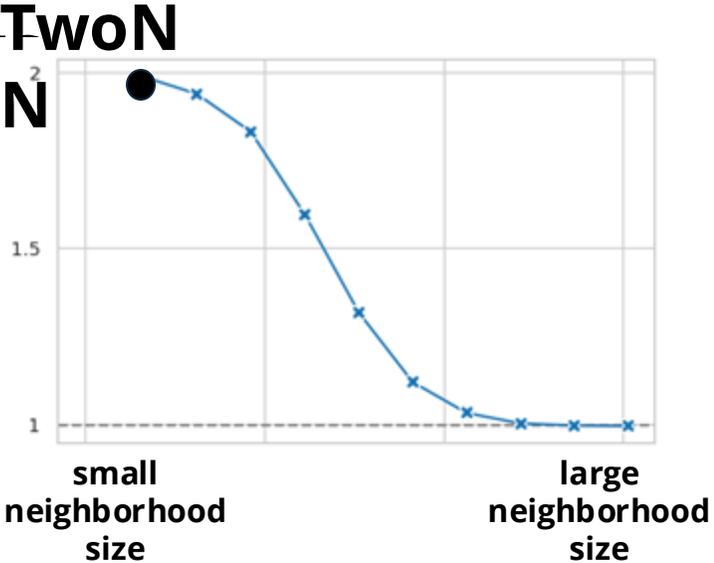
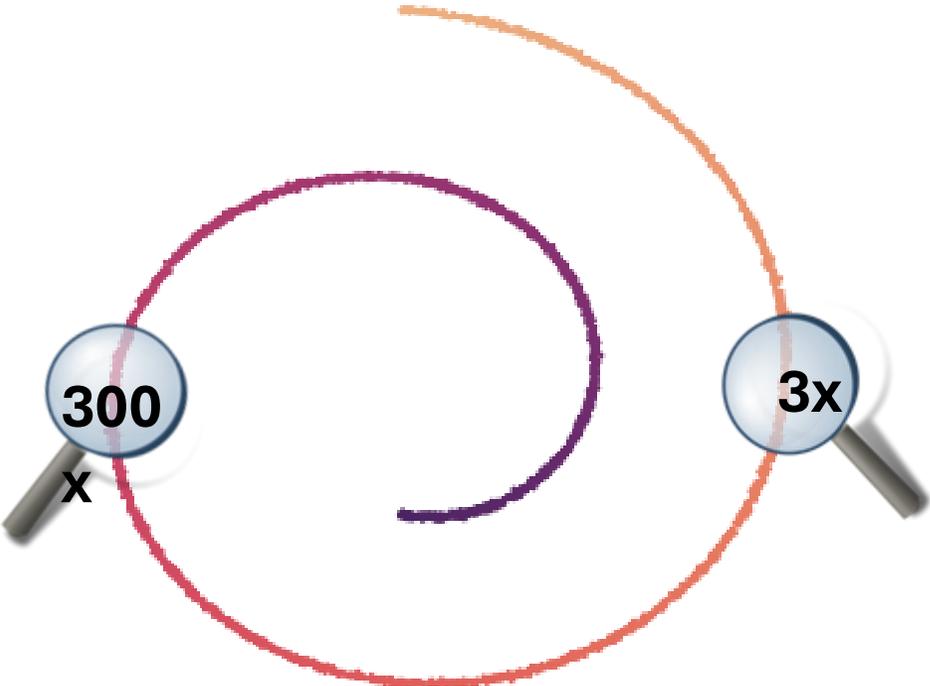
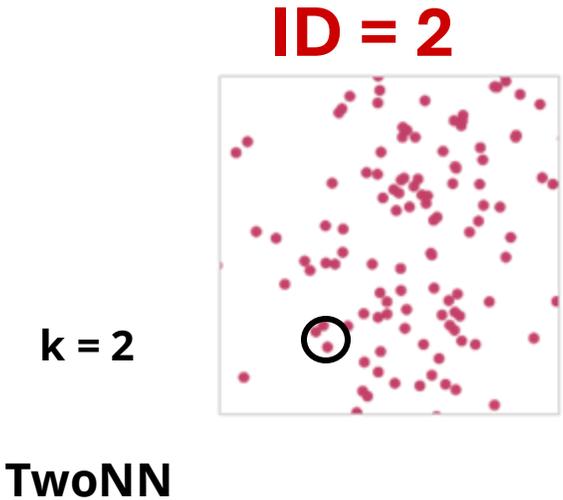
Intrinsic dimension estimation on noisy datasets

- The estimated intrinsic dimension is a **scale dependent** quantity



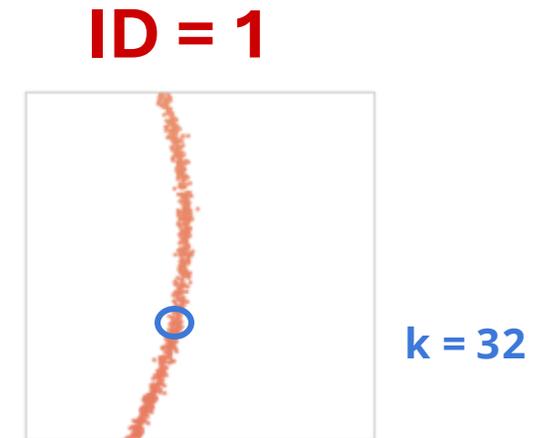
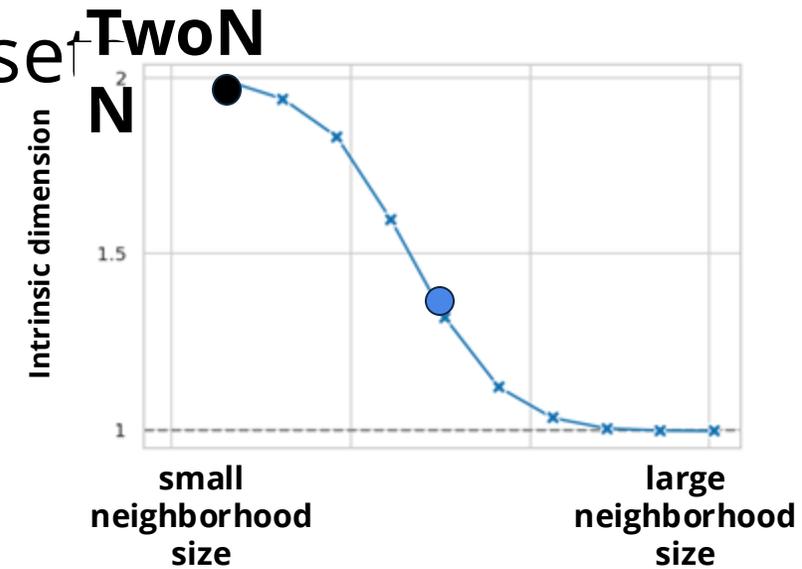
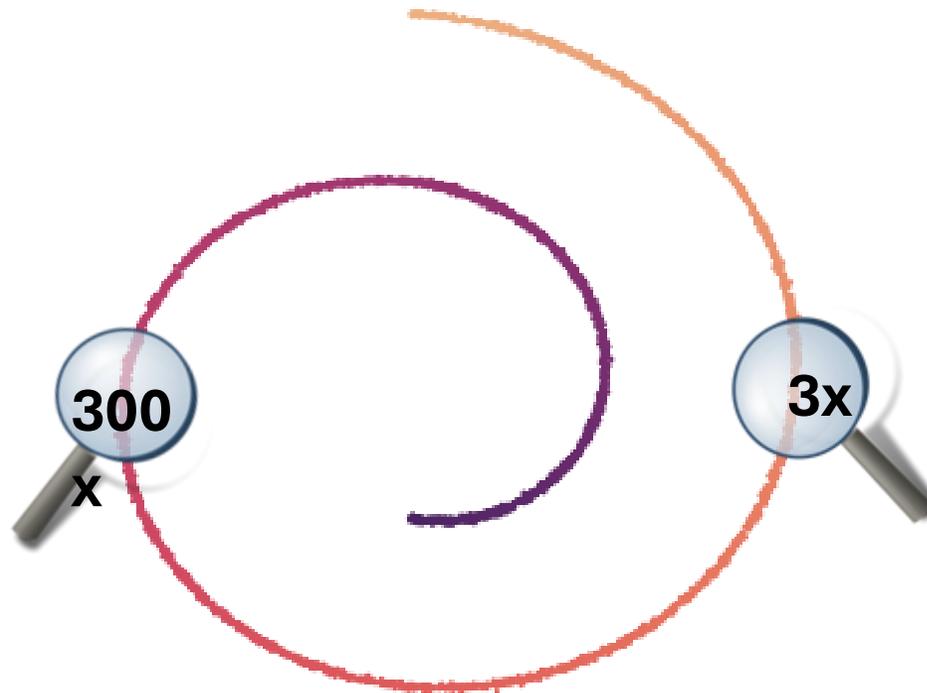
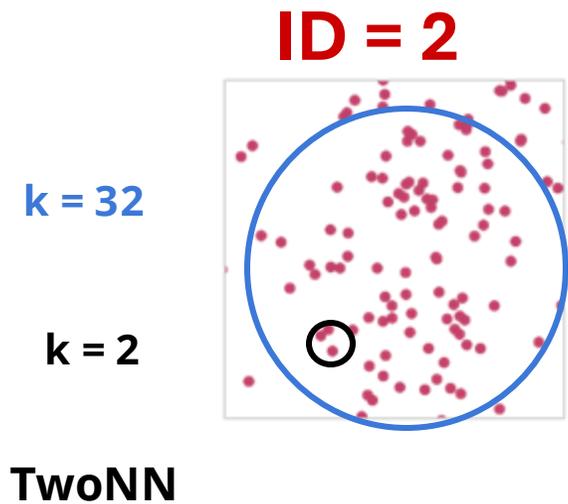
Intrinsic dimension estimation on noisy dataset

- The intrinsic dimension is a **scale dependent** quantity



Intrinsic dimension estimation on noisy dataset

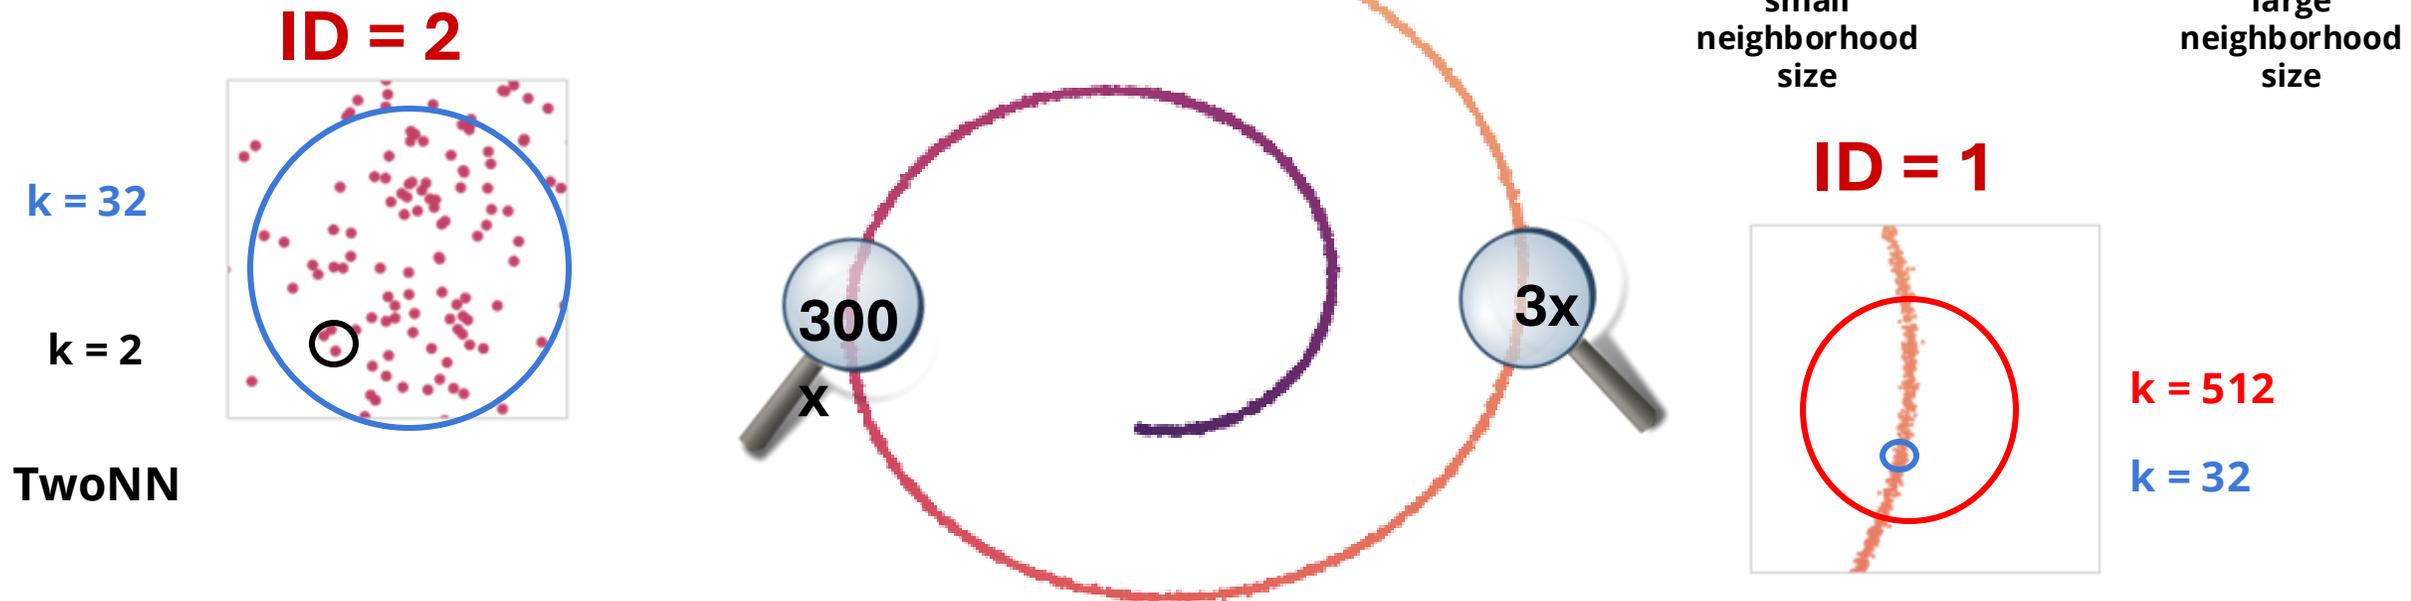
- The intrinsic dimension is a **scale dependent** quantity



- The scale dependence of the intrinsic dimension can be probed with an **higher order ratio approach**

Intrinsic dimension estimation on noisy dataset

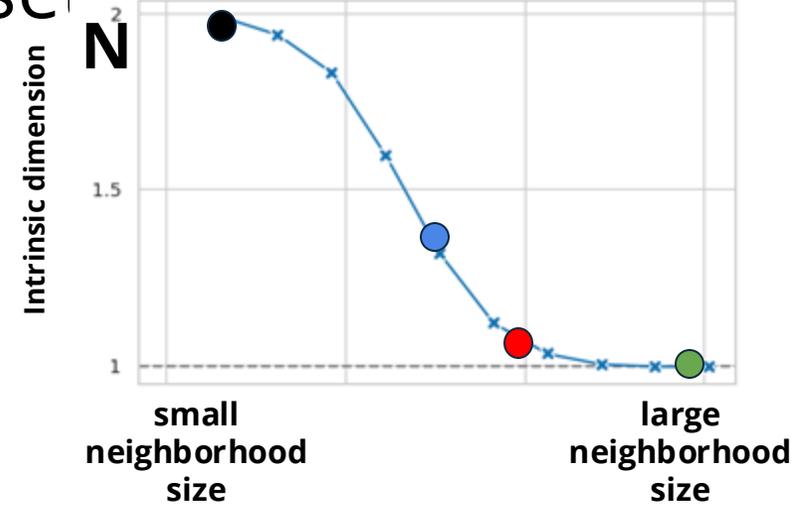
- The intrinsic dimension is a **scale dependent** quantity



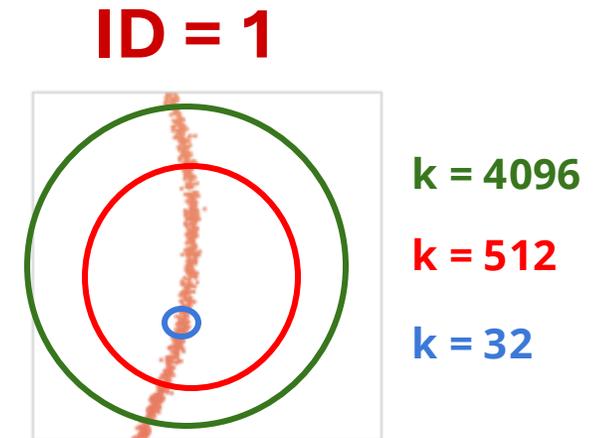
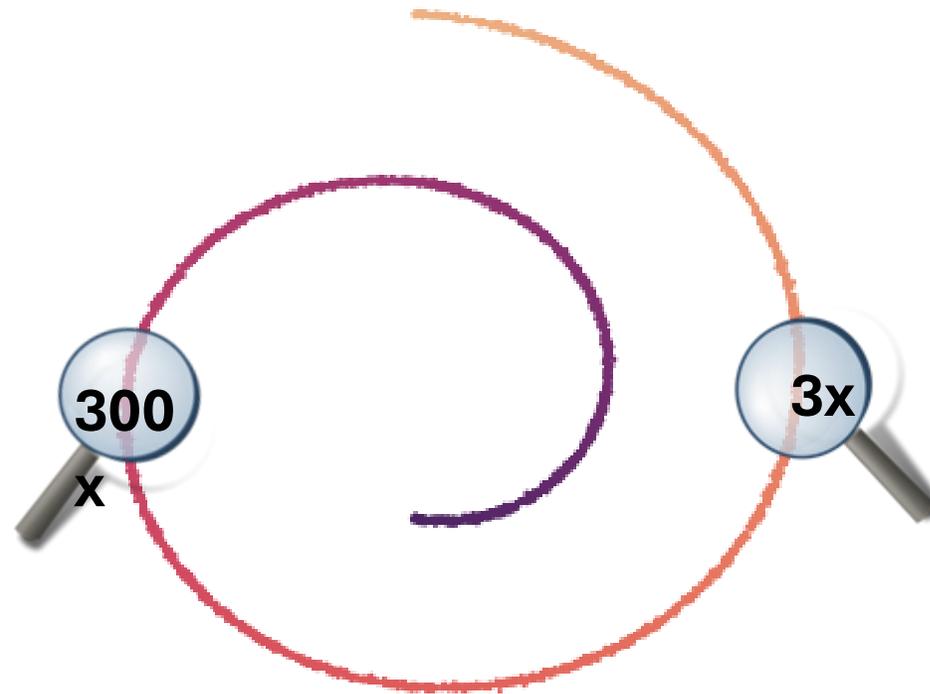
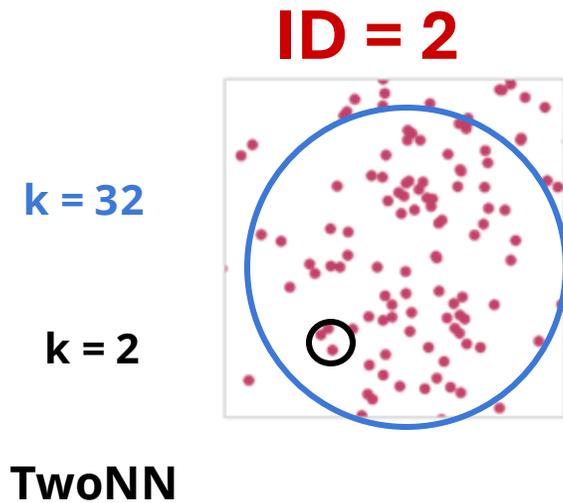
- The scale dependence of the intrinsic dimension can be probed with a **higher order ratio approach**

Intrinsic dimension estimation on noisy dataset

TwoN



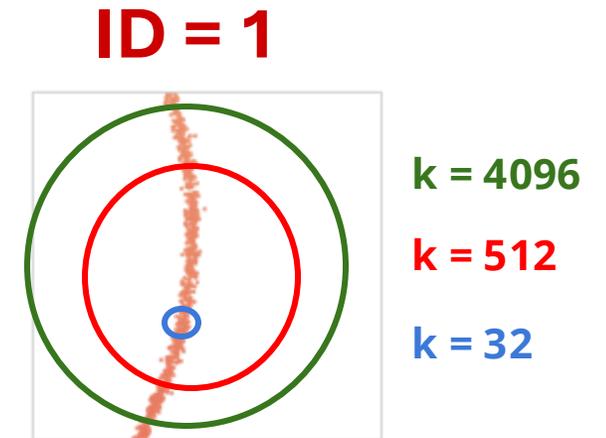
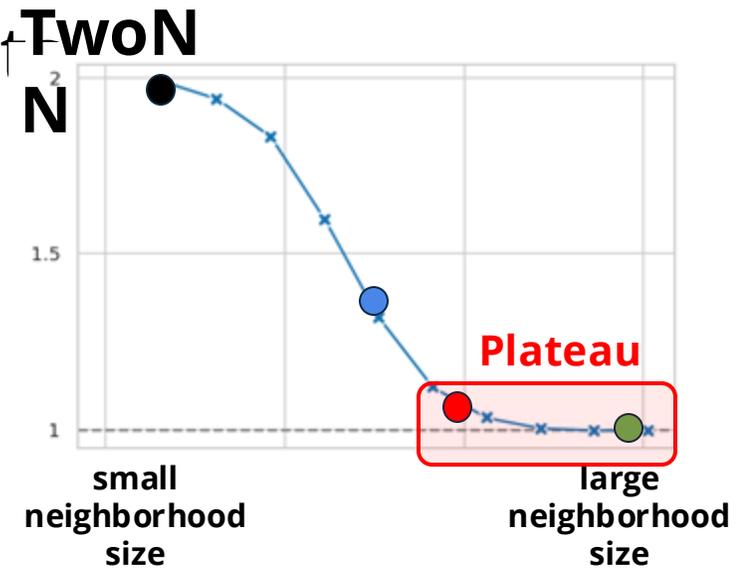
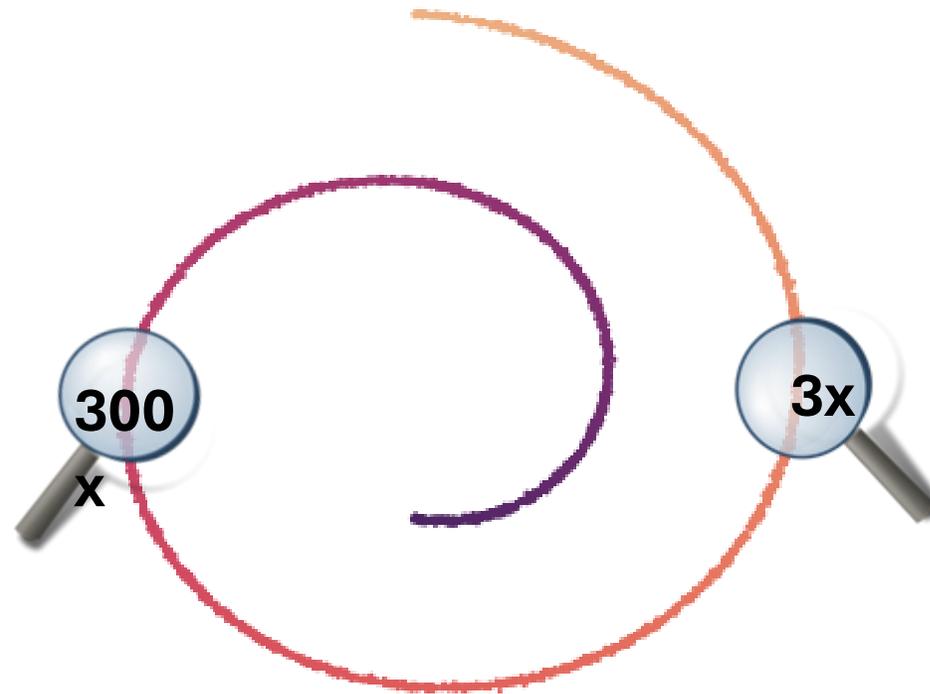
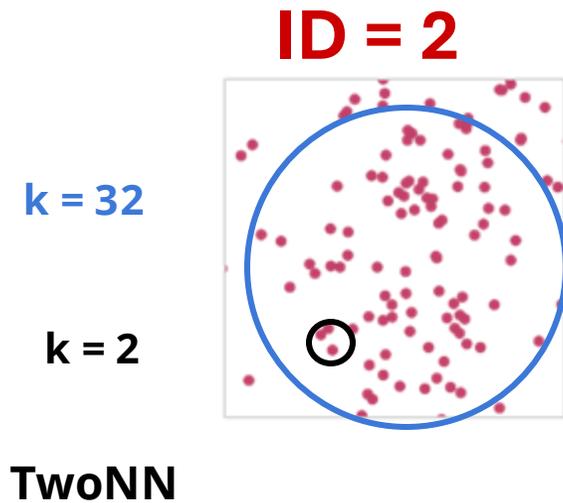
- The estimated intrinsic dimension is a **scale dependent** quantity



- The scale dependence of the intrinsic dimension can be probed with an **higher order ratio approach**

Intrinsic dimension estimation on noisy dataset

- The intrinsic dimension is a **scale dependent** quantity



- The scale dependence of the intrinsic dimension can be probed with a **higher order ratio approach**

Generalized ratios ID estimator (GRIDE)

Homogeneous Poisson process assumption:

1. The **number of points** $k(A_1)$, $k(A_2)$ falling in two non overlapping regions A_1 , A_2 are **independent random variables**
1. The **number of points** k in a region of volume V is a **Poisson random variable**:

$$P(k, V) = \frac{(\rho V)^k}{k!} e^{-\rho V}$$

Generalized ratios ID estimator (GRIDE)

**$k^{\text{th}}, 2k^{\text{th}}$
neighbors**

$$\mu_{k_i} = \frac{r_{2k_i}}{r_{k_i}}$$



$$p(\mu_{k_i}|d) = \frac{d(\mu_{k_i}^d - 1)^{k-1}}{\mu_{k_i}^{(2k-1)d+1} \text{Beta}(k, k)}$$

Data point likelihood

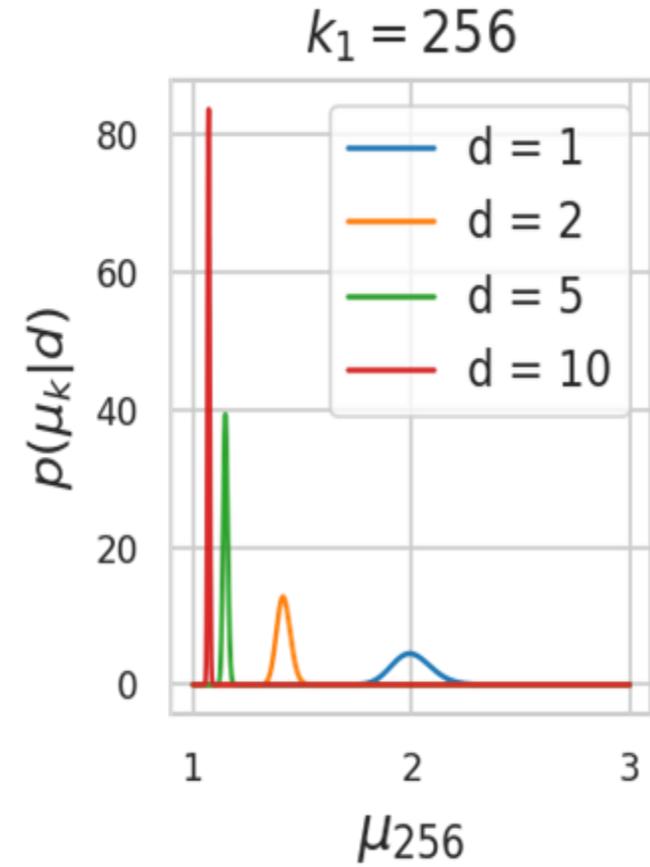
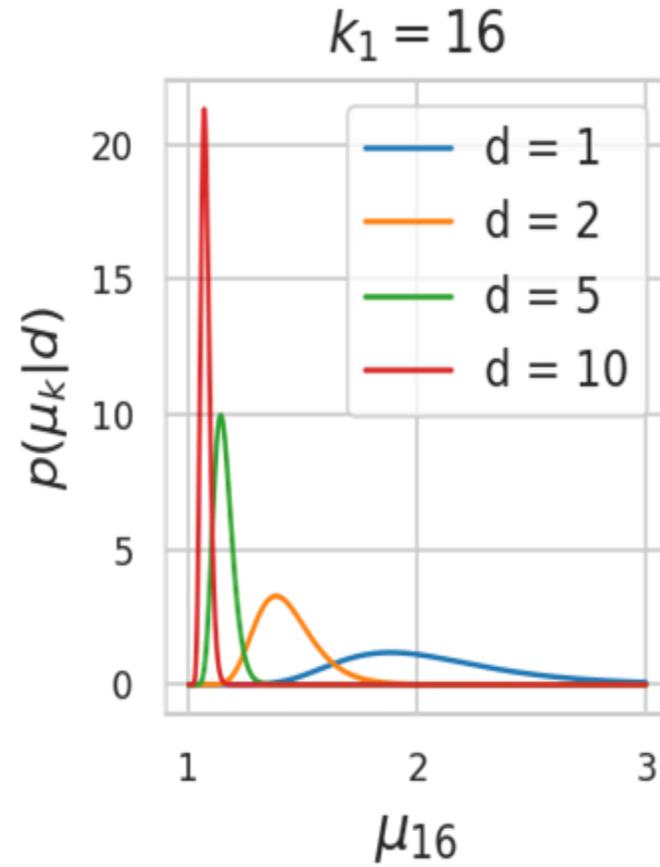
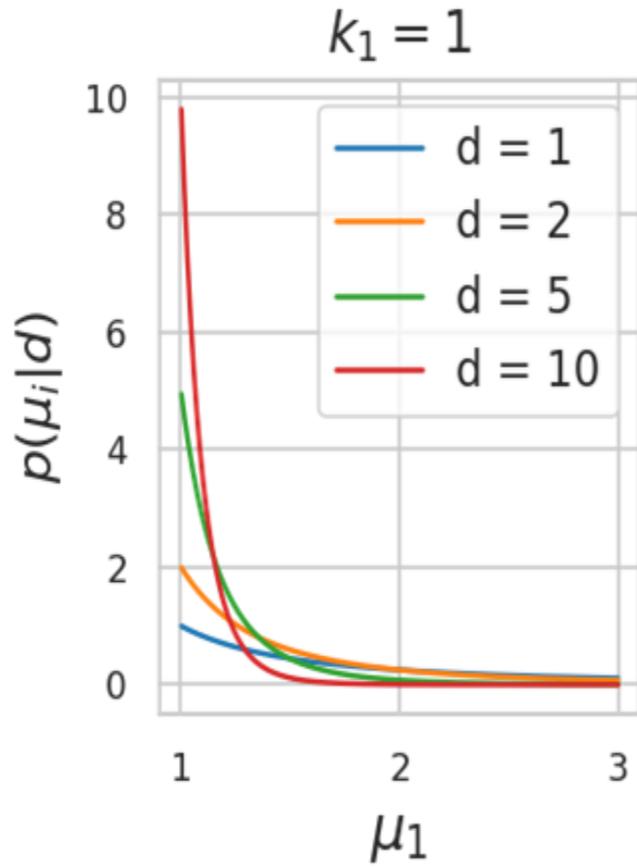
$$p(\boldsymbol{\mu}|d) = \prod_i^N p(\mu_i|d)$$

Data set likelihood

Generalized ratios ID estimator (GRIDE)

TwoNN

Likelihood variance decreases increasing k

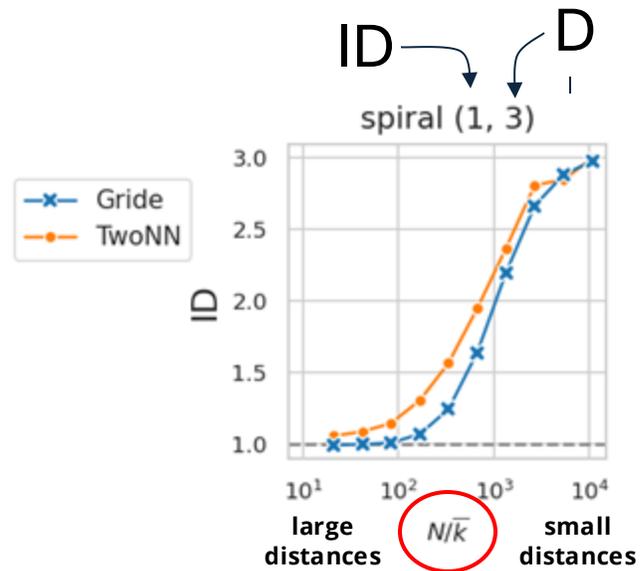
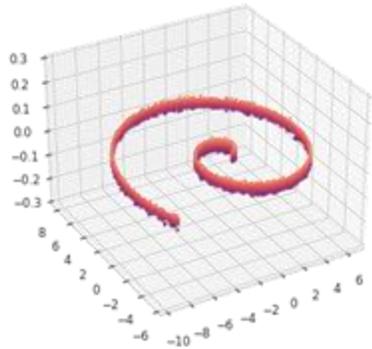


Scale analysis of the ID on synthetic datasets

n° data points = $N = 16000$

$$\sigma = \frac{0.01}{\sqrt{D}}$$

spiral 1d

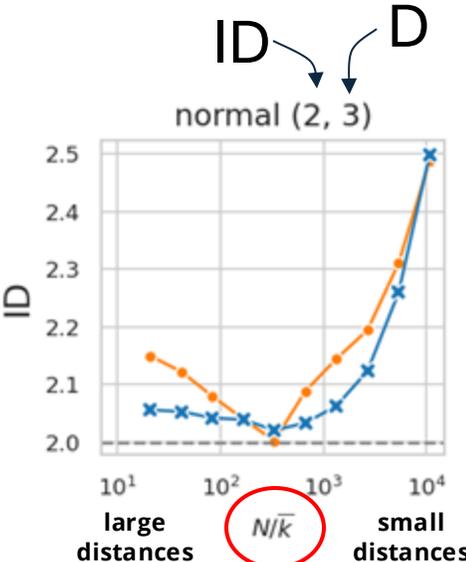
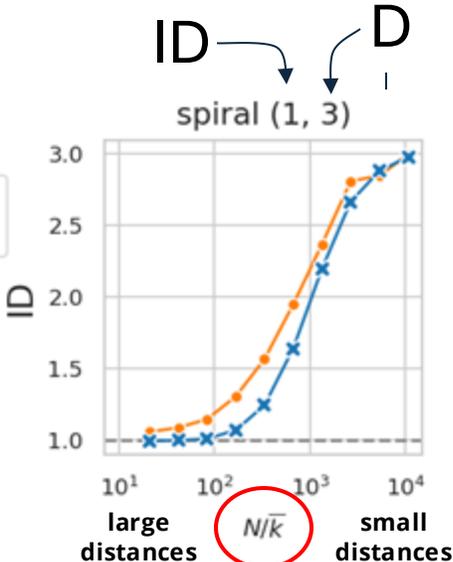
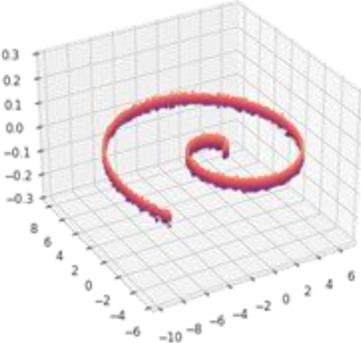


Scale analysis of the ID on synthetic datasets

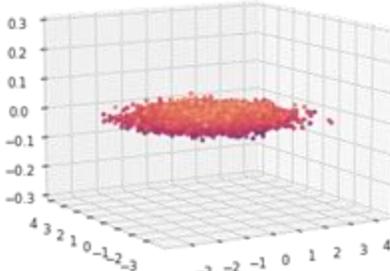
n° data points = N = 16000

$$\sigma = \frac{0.01}{\sqrt{D}}$$

spiral 1d



gaussian 2d

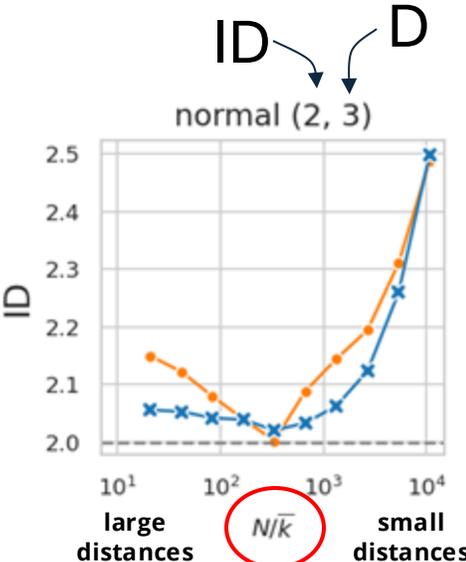
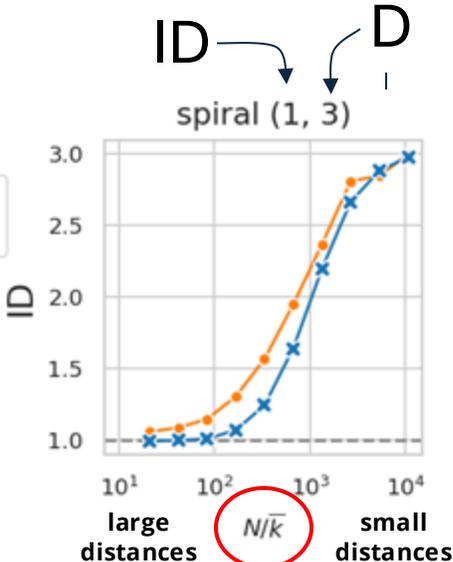
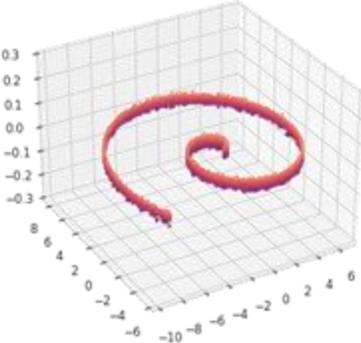


Scale analysis of the ID on synthetic datasets

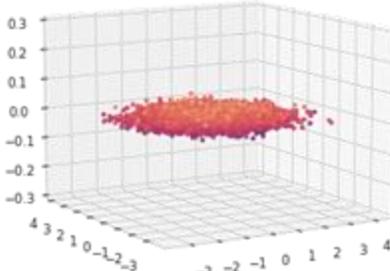
n° data points = 16000

$$\sigma = \frac{0.01}{\sqrt{D}}$$

spiral 1d



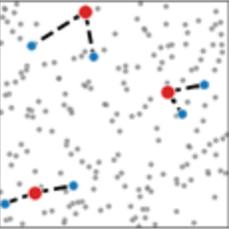
gaussian 2d



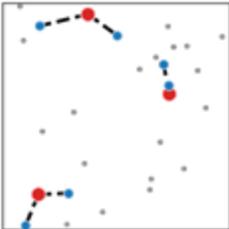
We define the scale with **N/k**

→ It works better when the data are high dimensional

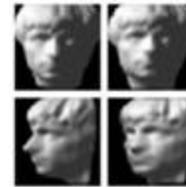
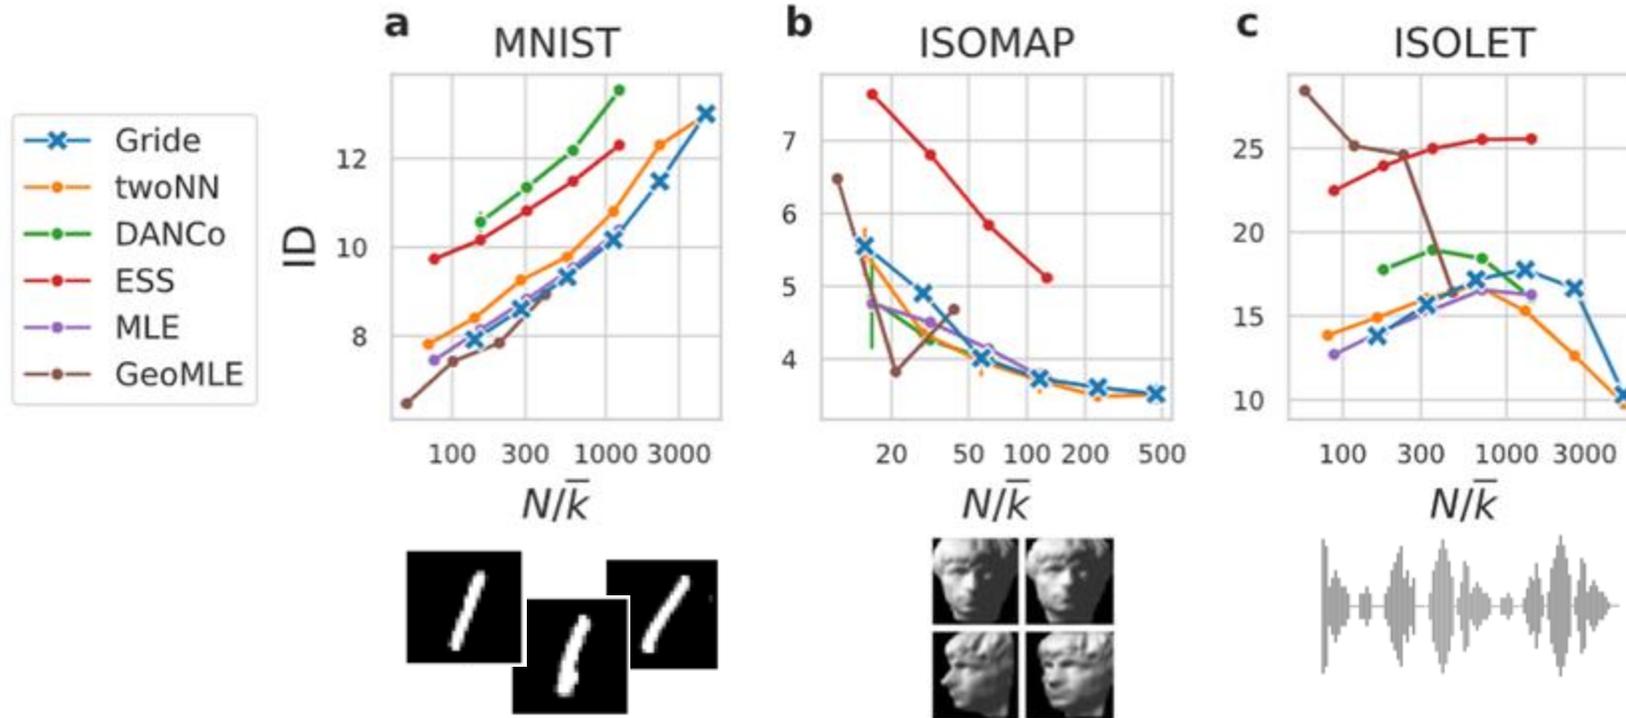
Gride



Decimation



Scale analysis of the ID on real data sets: unknown intrinsic dimension



number of data (N) -
number of features (P)

$N_{\text{tot}} = 6742$
 $P = 784$

$N_{\text{tot}} = 698$
 $P = 4096$

$N_{\text{tot}} = 7797$
 $P = 617$

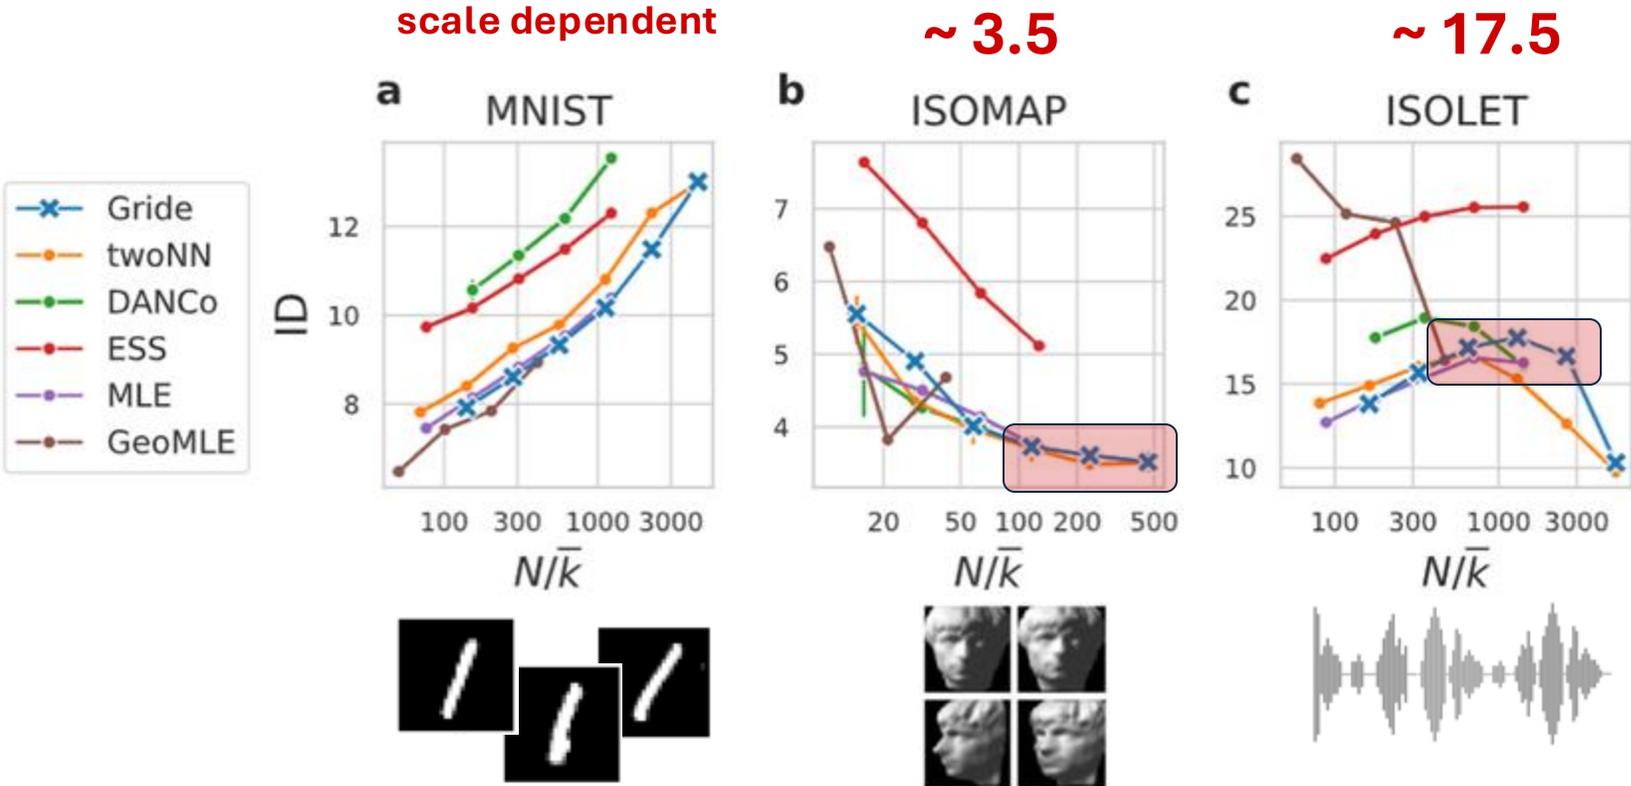
Consensus around
plausible ID ranges:

9-14

3

17-22

Scale analysis of the ID on real data sets: unknown intrinsic dimension



number of data (N) -
number of features D

$N_{tot} = 6742$
 $D = 784$

$N_{tot} = 698$
 $D = 4096$

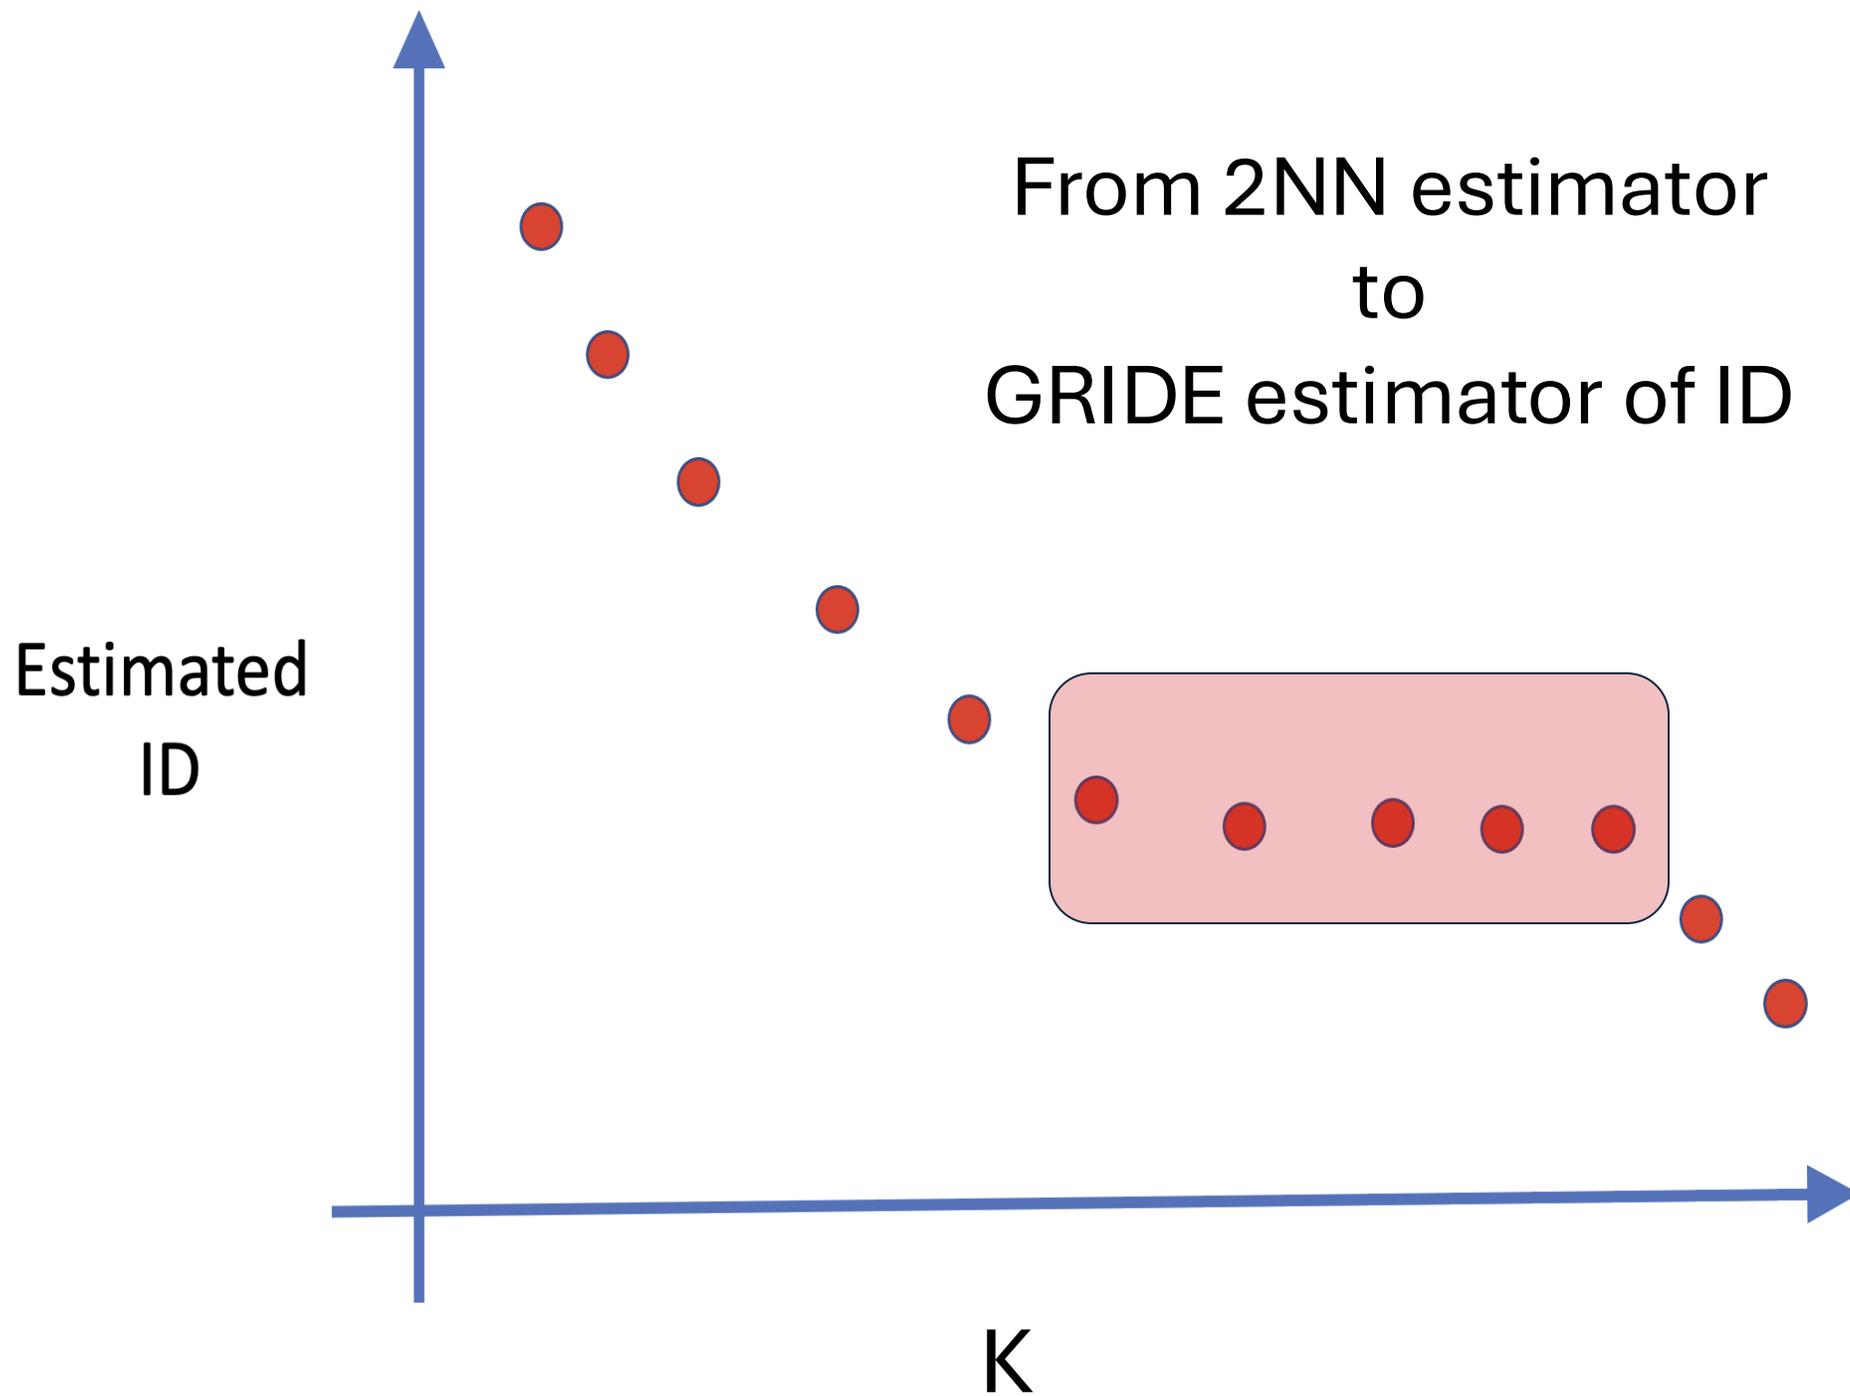
$N_{tot} = 7797$
 $D = 617$

Consensus around
plausible ID ranges:

9-14

3

17-22



Computational efficiency of GRIDE

Dataset CIFAR10: 32 x 32 color images



ID ~ 30

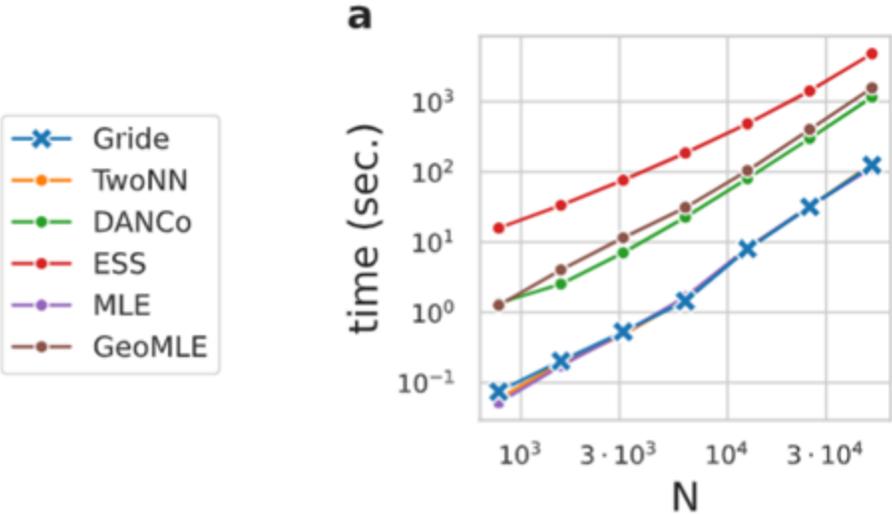
Computational efficiency of GRIDE

Dataset CIFAR10: 32 x 32 color images



ID ~ 30

D = 3072 = 32x32x3



Computational efficiency of GRIDE

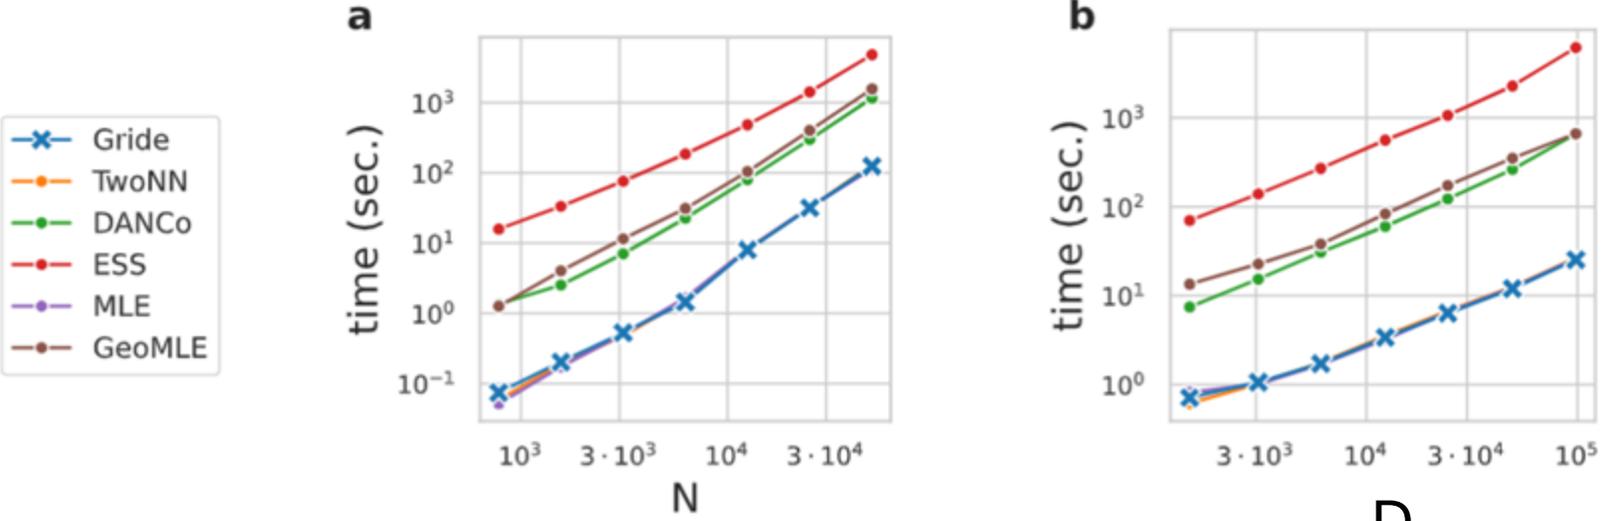
Dataset CIFAR10: 32 x 32 color images



ID ~ 30

D = 3072 = 32x32x3

N = 5000



*A simple topological feature
uncovers a rich data structure*

Folded vs unfolded state in protein configurations

Active vs non-active regions in brain imaging data

Patients vs controls in gene expression data

Firms with different financial risk in balance sheets data

Other applications:

Winning vs losing teams in basketball data

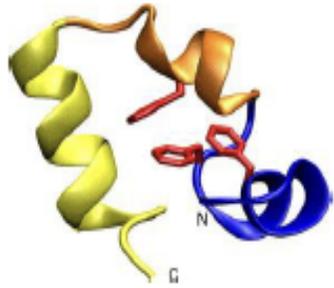
Country specific NPI in Covid-19 pandemic evolution

Identified vs unidentified models in MCMC simulation

Layers in a Deep Neural Network

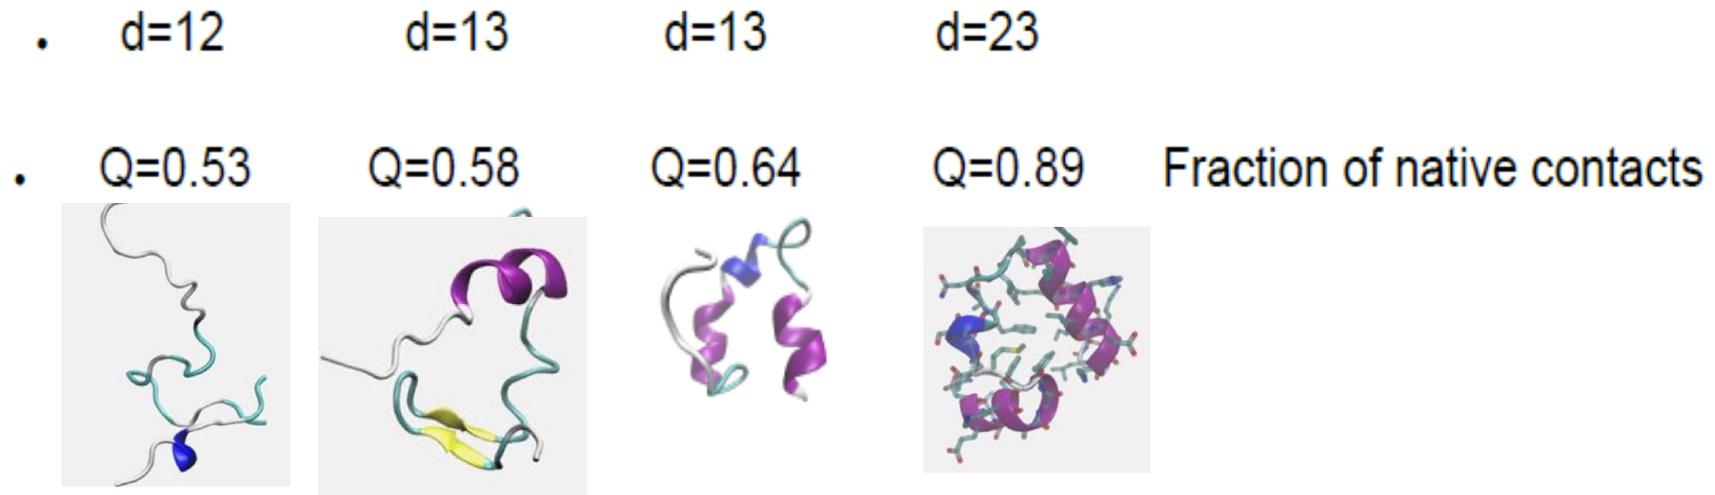
Communities in Network data

Molecular dynamics



- consider a MD of unfolding/refolding villin headpiece
- for each of the $N \sim 32000$ configurations, $D=32$ dihedral angles.

We find four manifolds



The folded state is recognized from its higher ID!

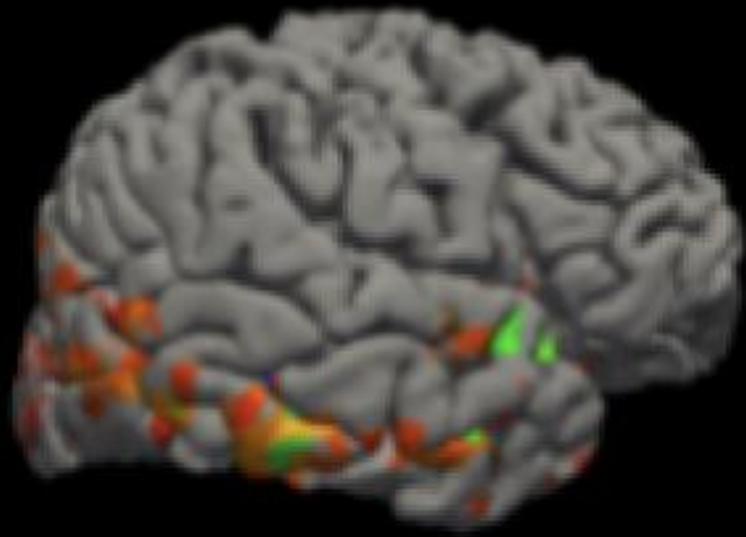
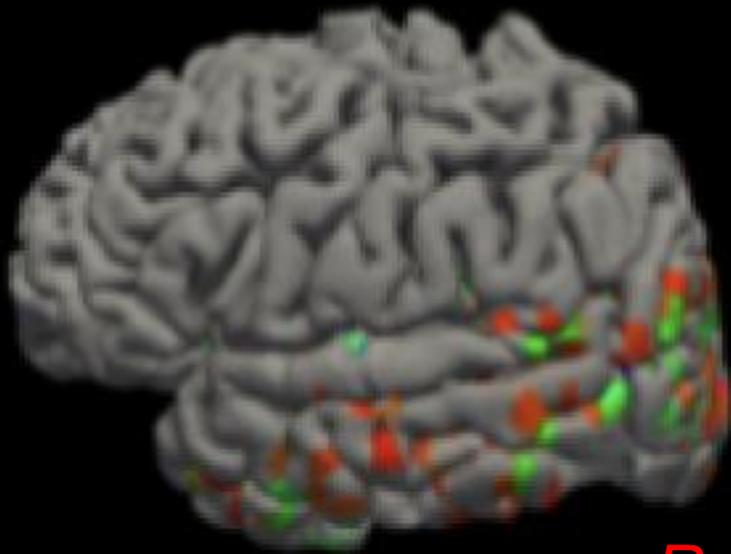
fMRI time series of BOLD signal

$N \sim 30'000$ voxels with
 $D = 202$ scans

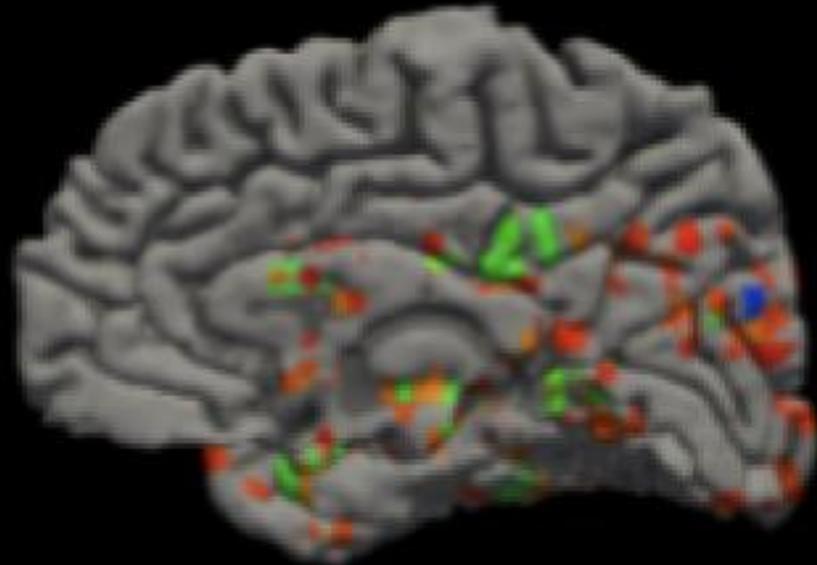
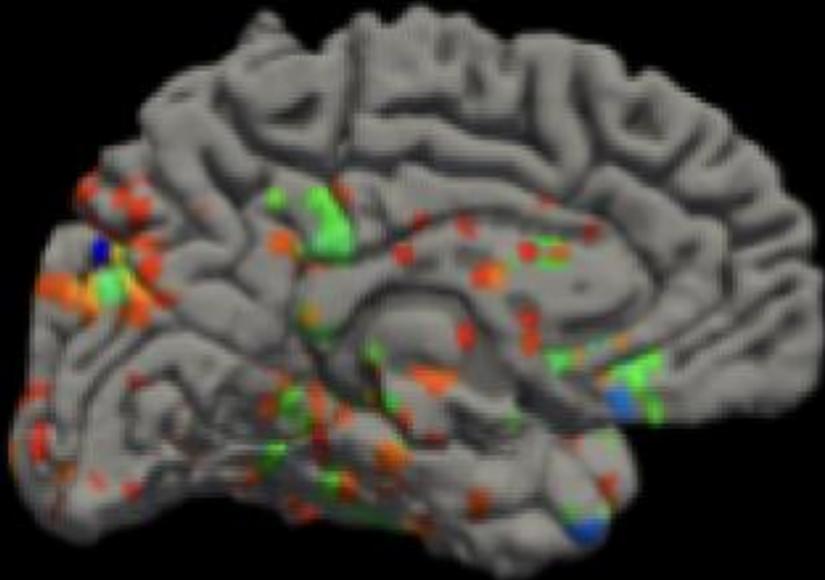
We find two manifolds $d1 = 16$ and
 $d2 = 32$

Task-relevant voxels are in the manifold with higher ID

Low-dimensional manifold mostly includes “noise” voxels



Red: high-ID voxels
Blue: task relevant voxels
Green: intersections

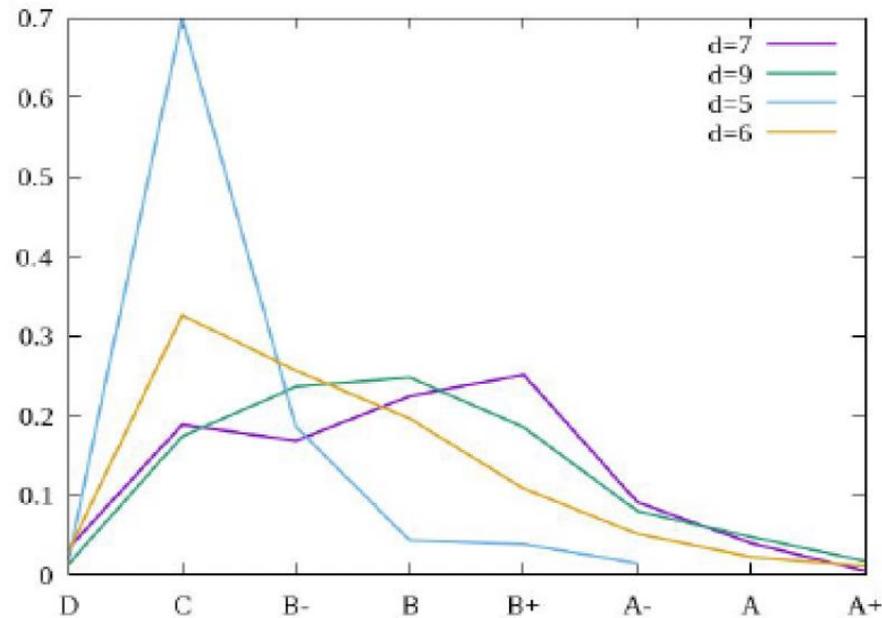


Firms from Compustat

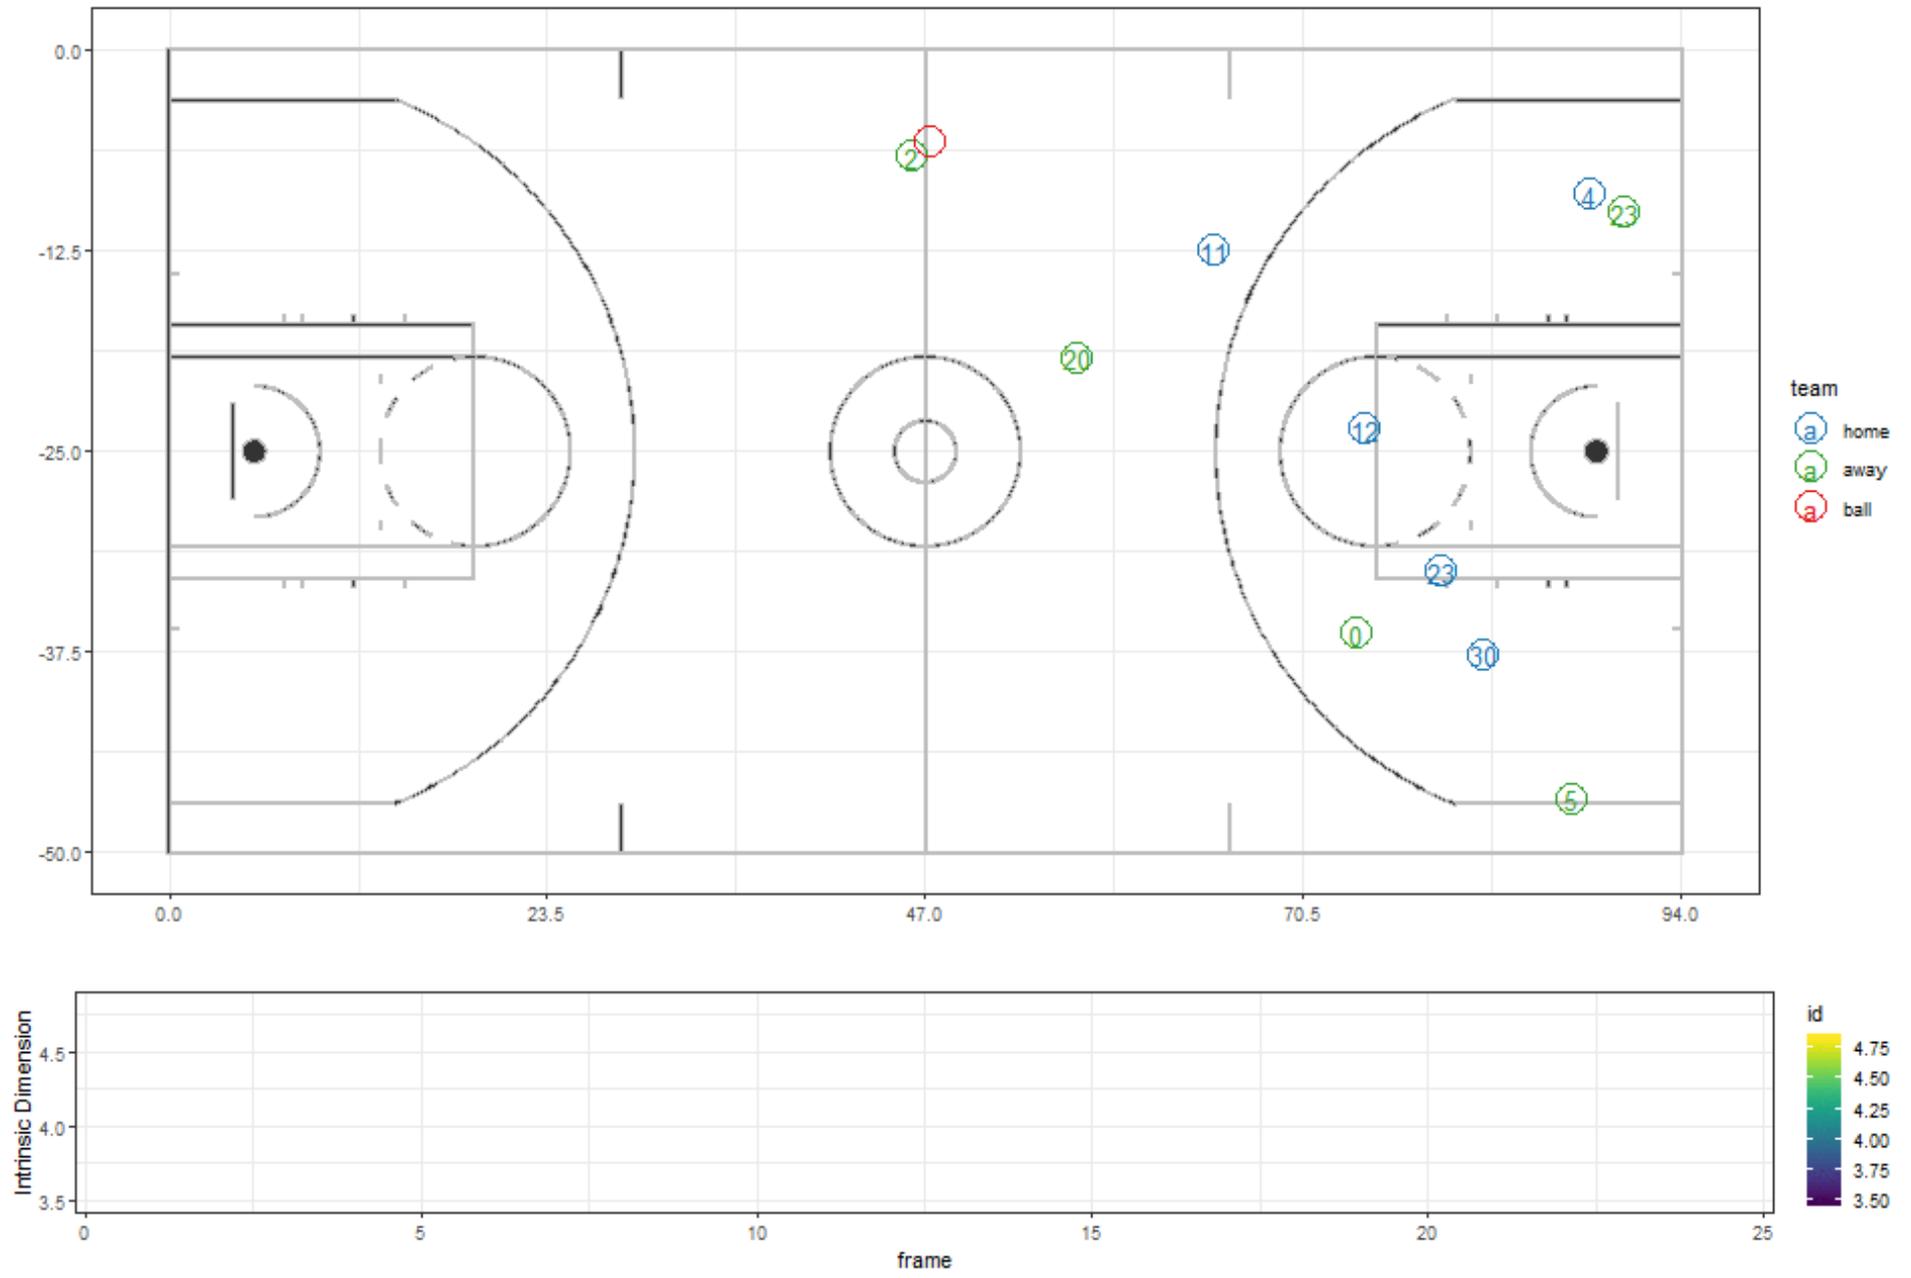
- consider ~8000 firms in the Compustat Database
- for each of the firms, D=31 balance sheet variables

We find four manifolds: d=5, d=6, d=7, d=9

We compute S&P ratings for the different manifolds



Lower dimension tends to have lower ratings!



Courtesy of Edgar Santos-Fernando

Gene expression data

Joint work with Luciano Cascione, Institute of Oncology Research, USI

$D \approx 16.900$ genes expressions

$N = 69$ tissue samples

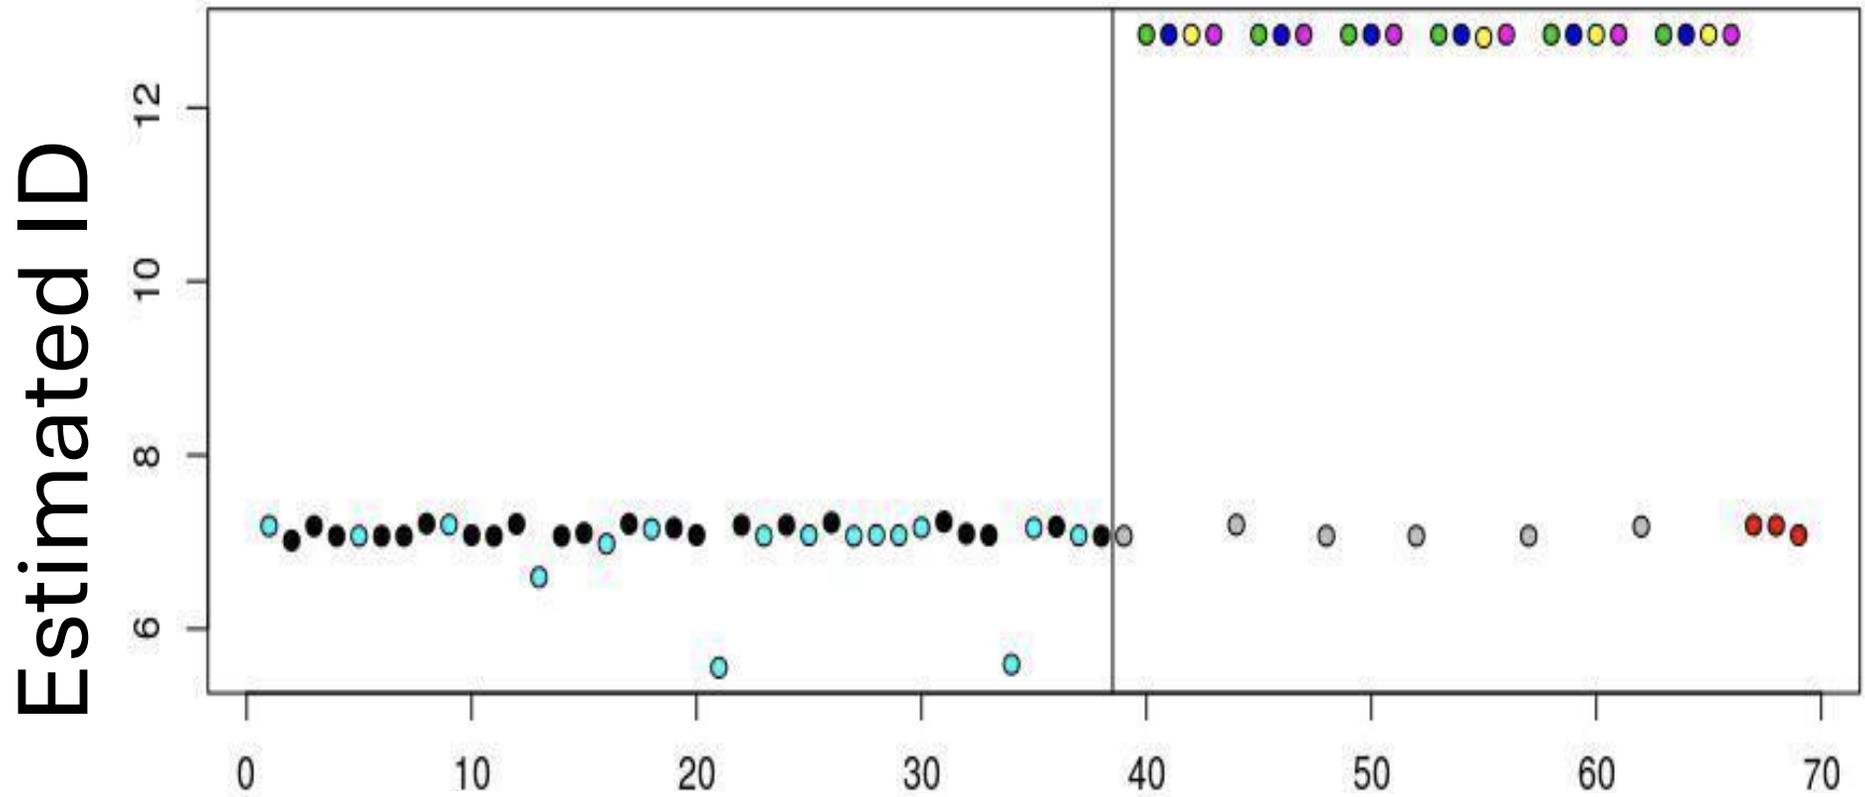
*The **first 38 samples** are CASES affected by
Diffuse large B-cell lymphoma - DLBCL*

*The **last 31** constitute the CONTROL group*

69 tissues colored by cell typology

The vertical line divides case / control groups

The y-axis: posteriori medians of the ID for each tissue sample



38 CASES

Index

31 CONTROLS

Binomial Intrinsic Dimension Estimator (BIDE) ... and its Bayesian counterpart

BIDE (MLE):

$$\hat{d} = \frac{\log\left(\frac{1}{n} \sum_{i=1}^n k_{A,i} / \frac{1}{n} \sum_{i=1}^n k_{B,i}\right)}{\log(\tau)}$$

BBIDE Bayesian estimator:

- Prior: $p = \tau^d \sim \text{Beta}(\alpha_0, \beta_0)$
- Posterior:

$$p | k_{B,1}, \dots, k_{B,n} \sim \text{Beta}(\alpha, \beta)$$

where $\alpha = \alpha_0 + \sum_{i=1}^n k_{A,i}$ and $\beta = \beta_0 + \sum_{i=1}^n (k_{B,i} - k_{A,i})$.

- Posterior expectation and variance:

$$E[d] = \frac{\psi_0(\alpha) - \psi_0(\alpha + \beta)}{\log(\tau)}, \quad \text{Var}[d] = \frac{\psi_1(\alpha) - \psi_1(\alpha + \beta)}{\log(\tau)^2}$$

Adaptive BIDE (ABIDE) and its Bayesian twin

Algorithm Adaptive-BIDE

```
1:  $d_{\text{current}} \leftarrow$  ID from Two-NN
2:  $d_{\text{next}} \leftarrow 0$ 
3: for  $it < \text{max\_iter}$  do
4:    $\tau = 0.2032^{1/d_{\text{current}}}$ 
5:   for  $i < n$  do
6:     compute  $k_i^*$  (using  $d_{\text{current}}$ ) and set  $k_{B,i}^* = k_i^*$ 
7:      $t_{B,i}(k_i^*) = r_{ik_i^*}$ 
8:      $t_{A,i}(k_i^*) = \tau t_{B,i}(k_i^*)$ 
9:      $k_{A,i}^* = \sum_{j=1}^n \mathbf{1}\{t_{A,i}(k_i^*) - r_{i,j} > 0\}$ 
10:  end for
11:   $d_{\text{next}} = \frac{\log\left(\frac{1}{n} \sum_{i=1}^n k_{A,i}^* / \frac{1}{n} \sum_{i=1}^n k_{B,i}^*\right)}{\log(\tau)}$  or  $d_{\text{next}} = \frac{\psi_0(\alpha^*) - \psi_0(\alpha^* + \beta^*)}{\log(\tau)}$ ,
12:  if then  $|d_{\text{current}} - d_{\text{next}}| < \delta$  break
13:  end if
14:   $d_{\text{current}} = d_{\text{next}}$ 
15: end for
16:  $d^* = d_{\text{next}}$ 
17: for  $i < n$  do
18:   compute  $k_i^*$  (using  $d^*$ )
19: end for
20: return  $d^*, k_i^*$ 
```

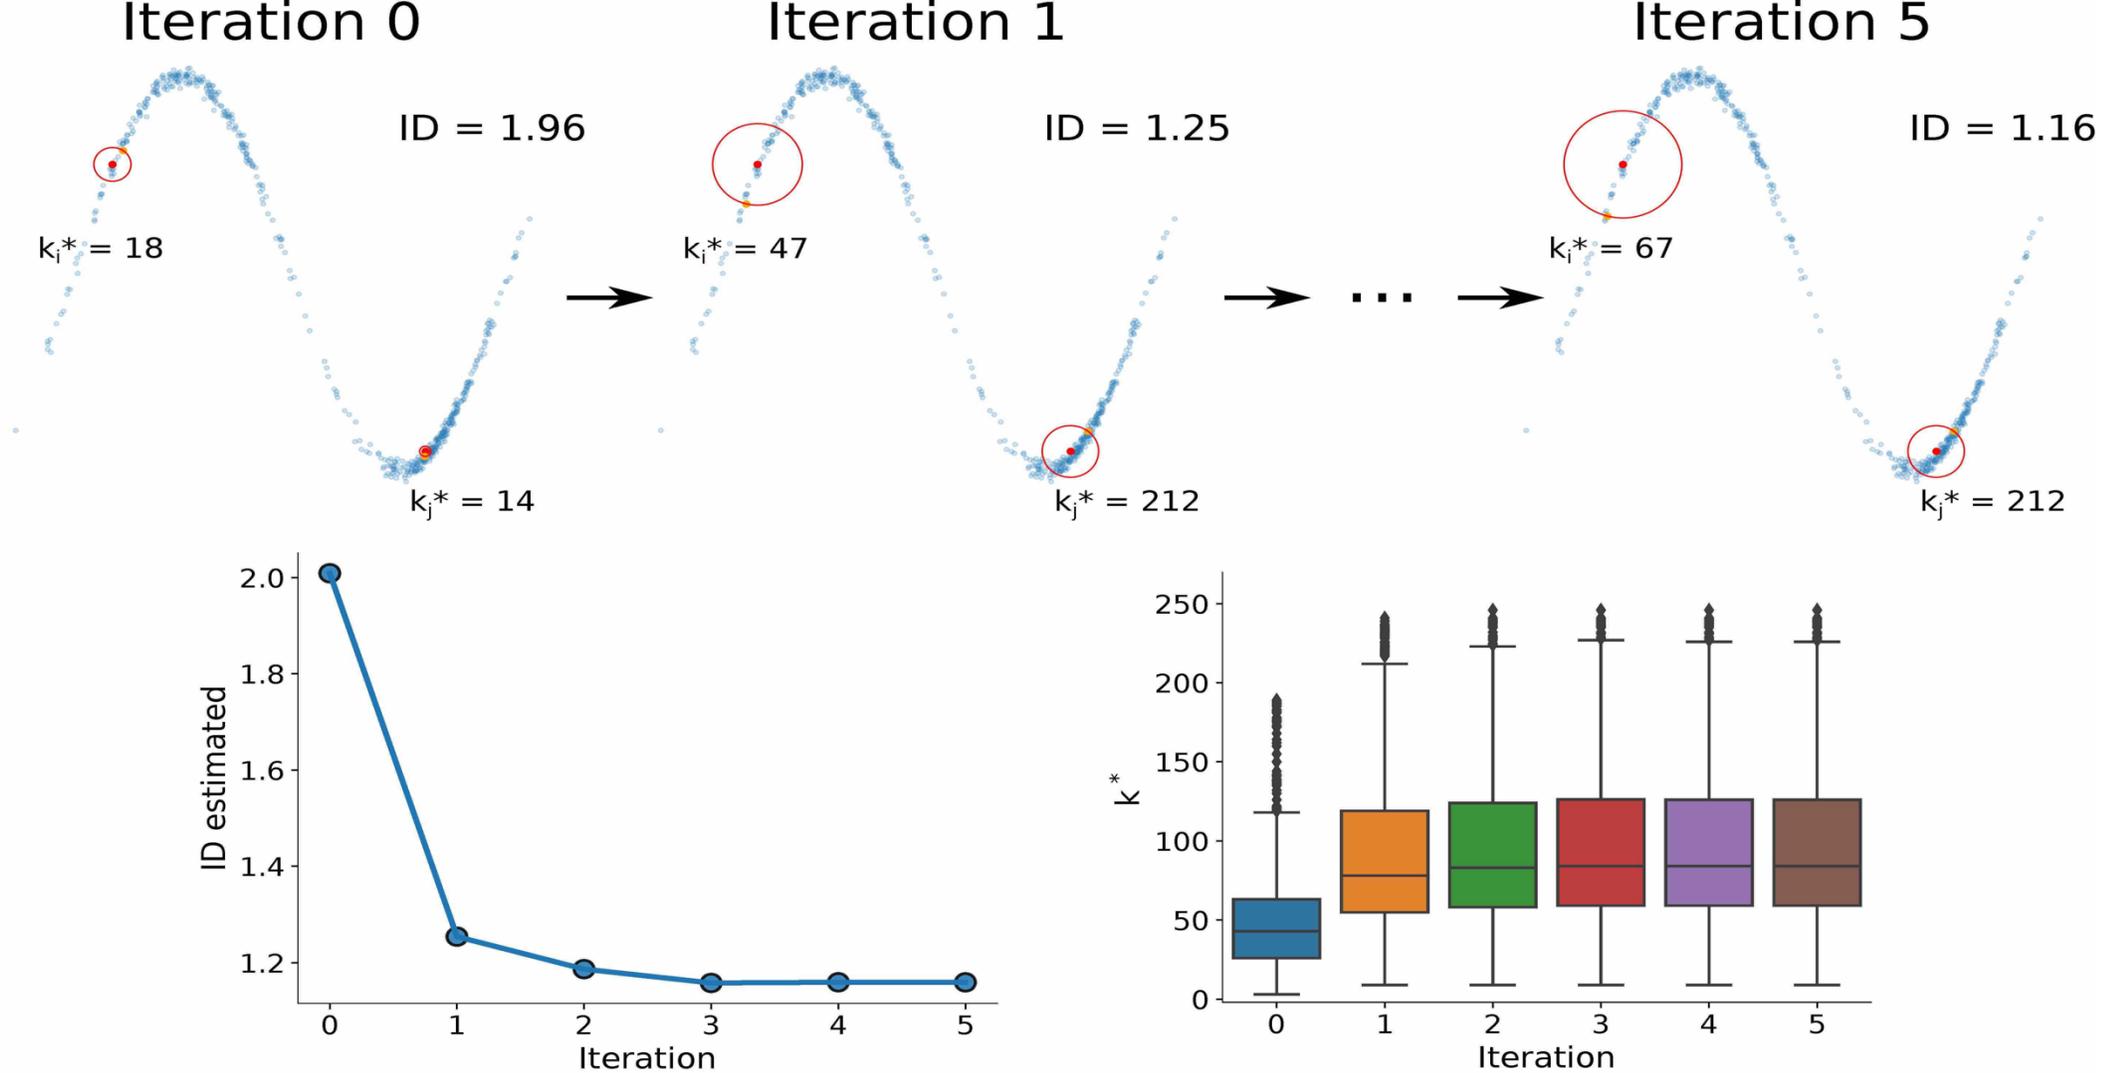


Figure: Consider $\mathcal{N}(\frac{\pi}{2}, 1)$ and $\mathcal{N}(\frac{5}{3}\pi, 0.5^2)$, sample X_1, \dots, X_{1000} points, 500 from each one of the two. Then map the points on a curved manifold by adding a second coordinate given by $Y_i = \sin(X_i)$.

Locally Linear Embedding - LLE

LLE: unsupervised learning algorithm that computes low dimensional, neighborhood preserving embeddings

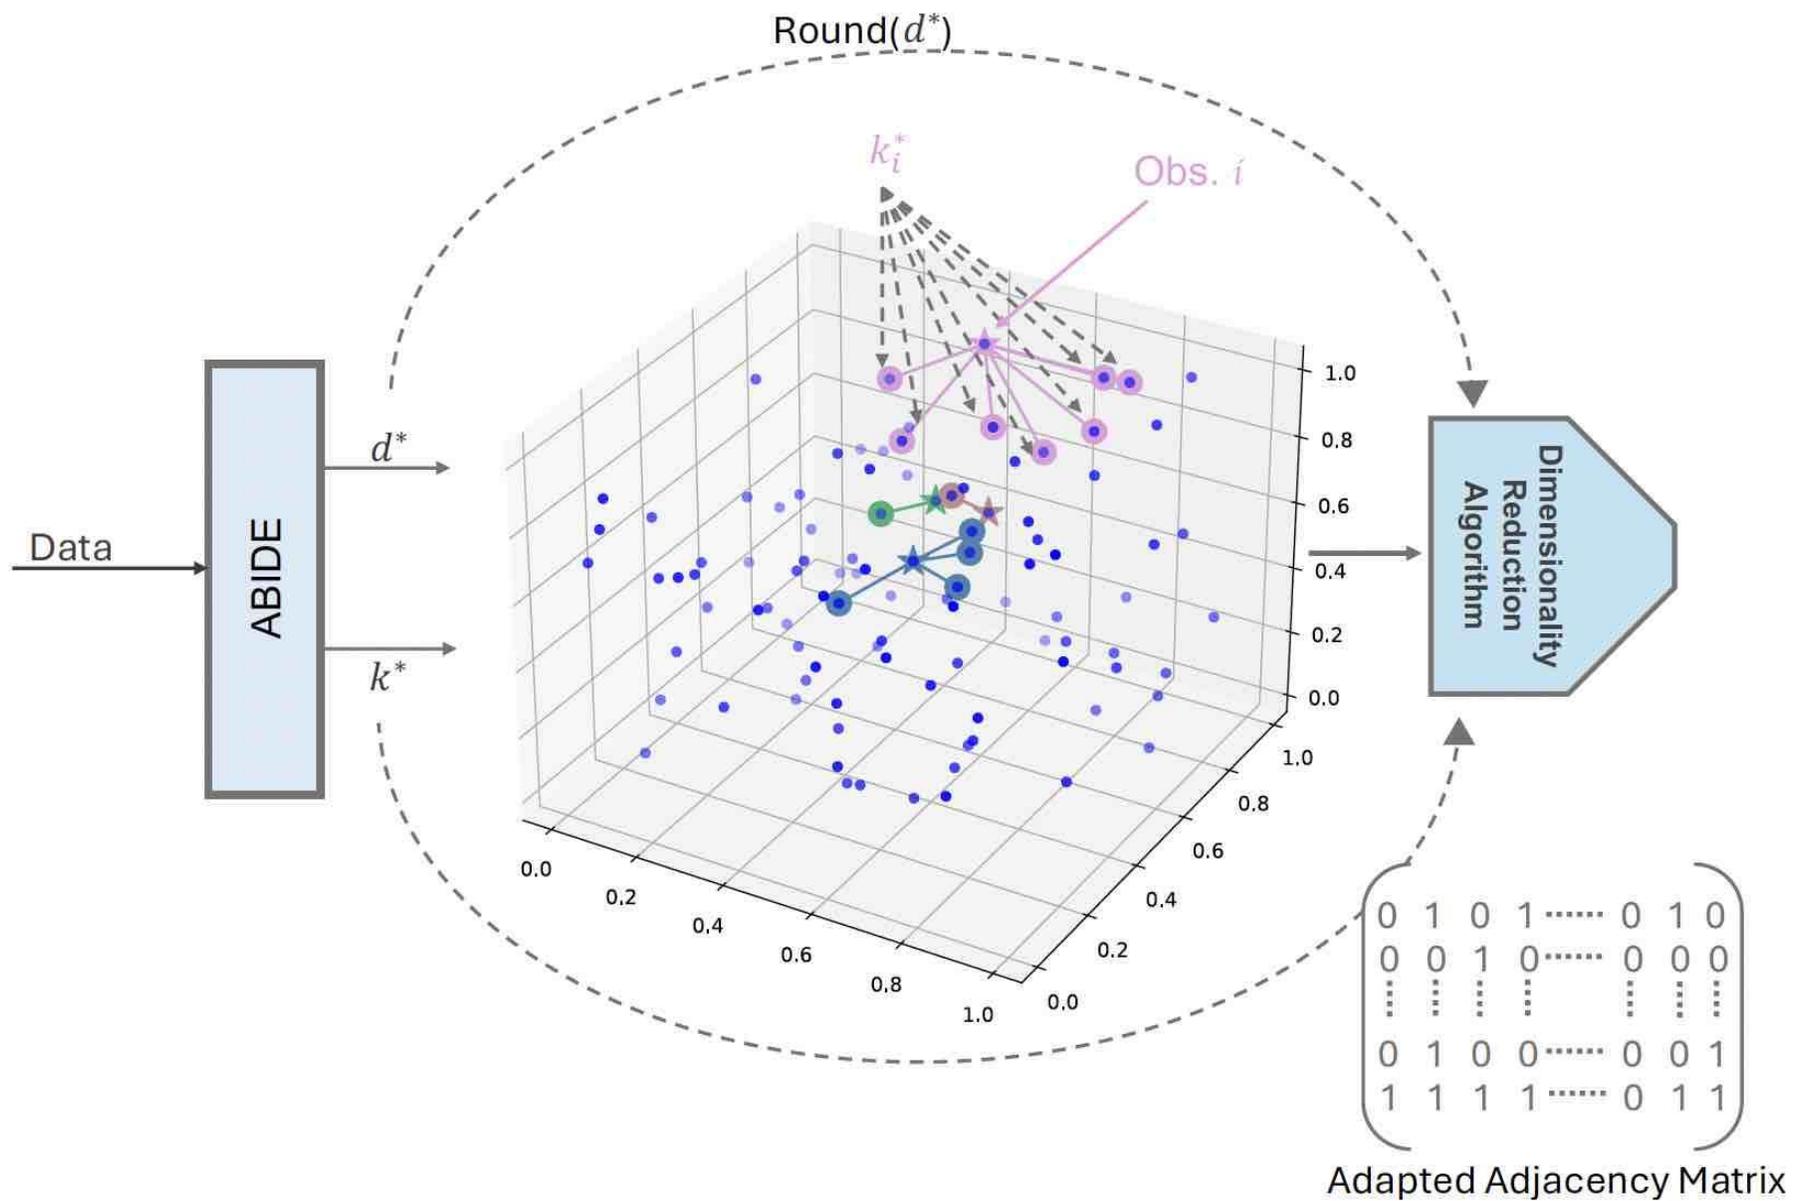
TS Roweis and LK Saul, Science, 2000, \sim 20.000 citations

LLE: eigenvector method for nonlinear dimensionality reduction

Given a dataset $\{\mathbf{x}_i\}_{i=1}^n \subset \mathbb{R}^D$, LLE finds low dimensional embedding vectors $\{\mathbf{y}_i\}_{i=1}^n \subset \mathbb{R}^d$

Given inputs: k and d

1. Compute the k neighbors of each data point x_i
2. Compute the weights that best reconstruct each data point from its neighbors, minimizing the reconstruction error
3. Compute the vectors y_i best reconstructed by the weights by minimizing a cost function



Diffuse Large B-Cell Lymphoma - DLBCL

$D \approx 16.900$ genes expressions recorded on $N = 69$ tissues

38 patients with DLBCL with different variants of lymphoma

Activated B-Cell-like - ABC

Germinal Center B-Cell-like - GCB

31 healthy donors of B-cell samples at various stages of maturation

Naive B Cell - NB

CentroBlast - CB

CentroCyte - CC

Memory B Cell - MEM

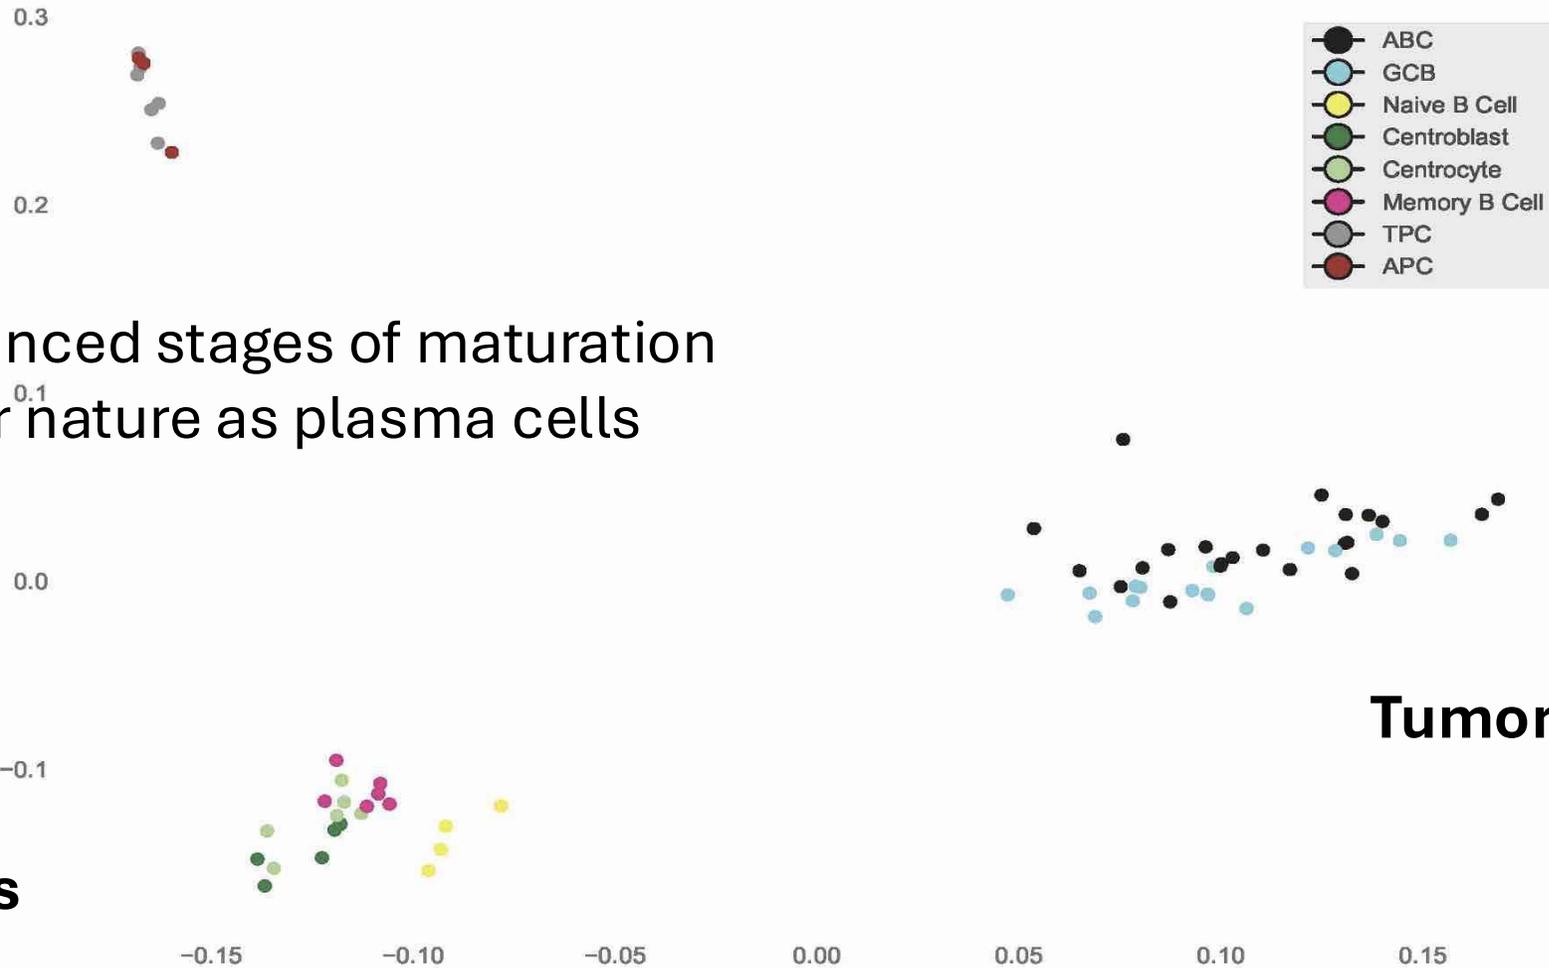
Transitional Plasma Cells - TCP

Advanced-stage Plasma Cells in Bone Marrow - APC / BMPC

TPC and APC, being plasma cells, represent the most advanced stage of B-cell maturation

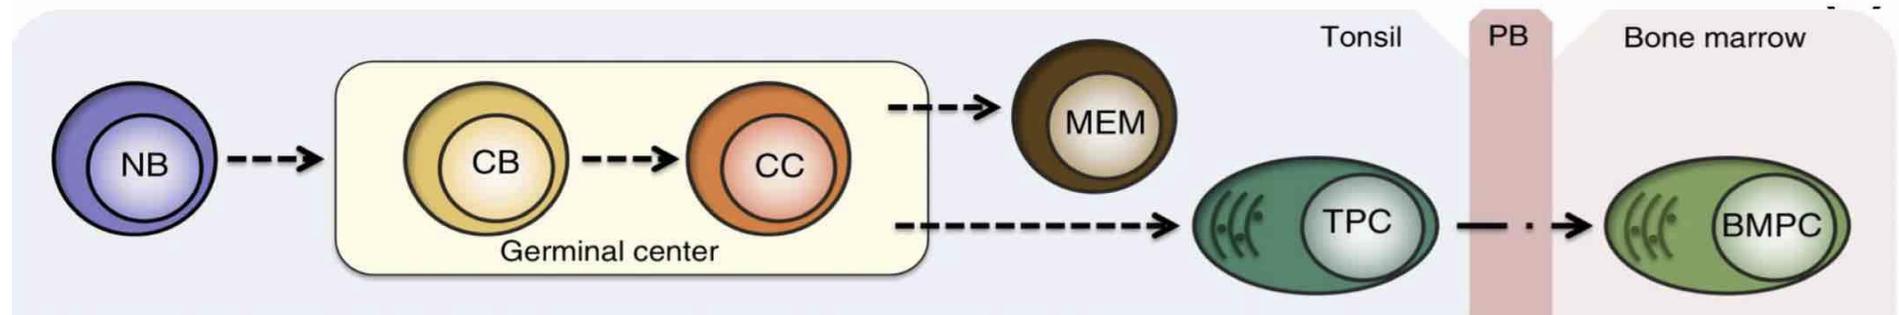
TPC and APC cells

are in the most advanced stages of maturation
consistent with their nature as plasma cells



Healthy cells

Tumor cells



ID of Covid-19

Joint with

A. Varghese, E. Santos-Fernandez, F. Denti, Kerrie Mengersen

	CSI 1.3.20 – 29.5.21	New Cases pmp 1.3.20 – 29.5.21	New Deaths pmp 1.3.20 – 29.5.21
Country 1			
Country 2			
Country 3			
Country 4 , , ,			
Country 115			

D =
 $454 \times 3 =$
 1362

N = 115

excluded: > than 20% missing
 < 1 million population

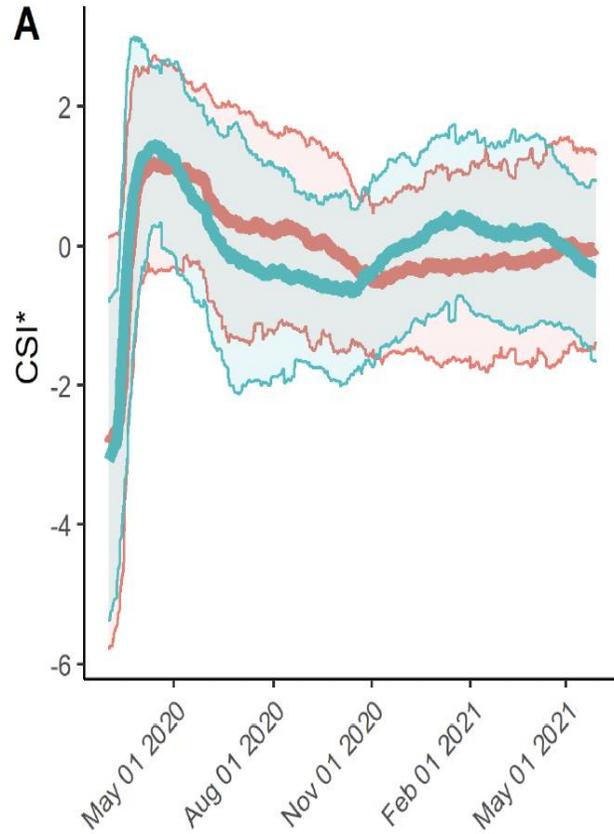
MAIN CONCLUSION

high-income countries are more likely to lie on low-dimensional manifolds,

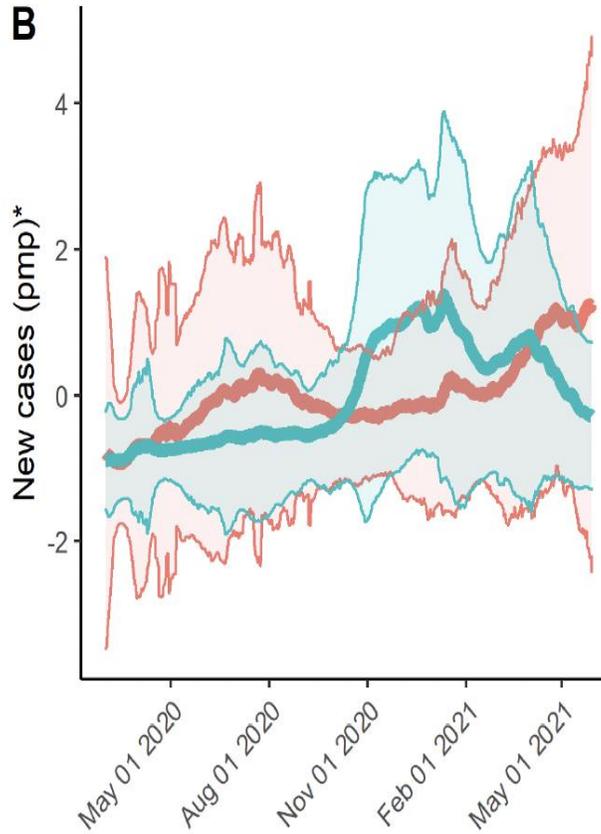
likely arising from aging populations and comorbidities,

causing increased per capita mortality from COVID-19

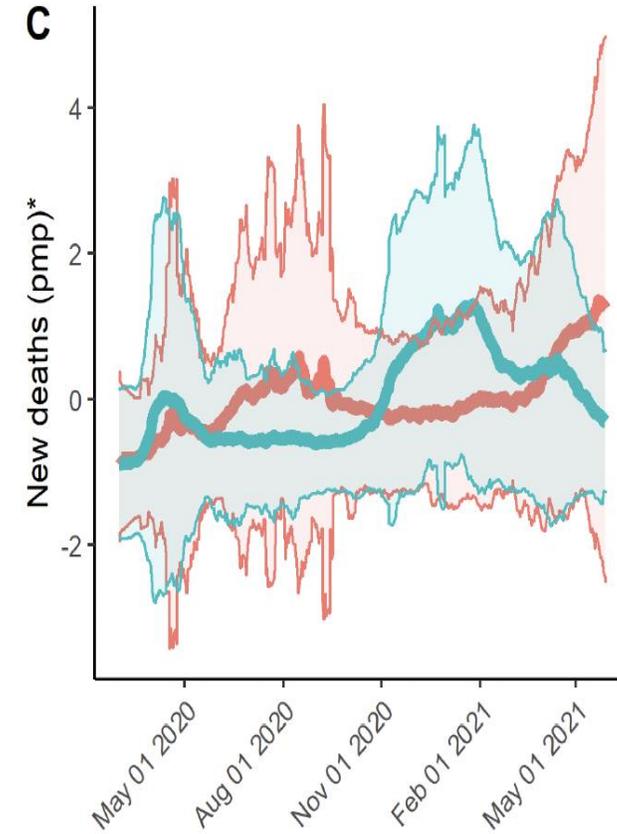
1st Mar 2020 to 29th May 2021



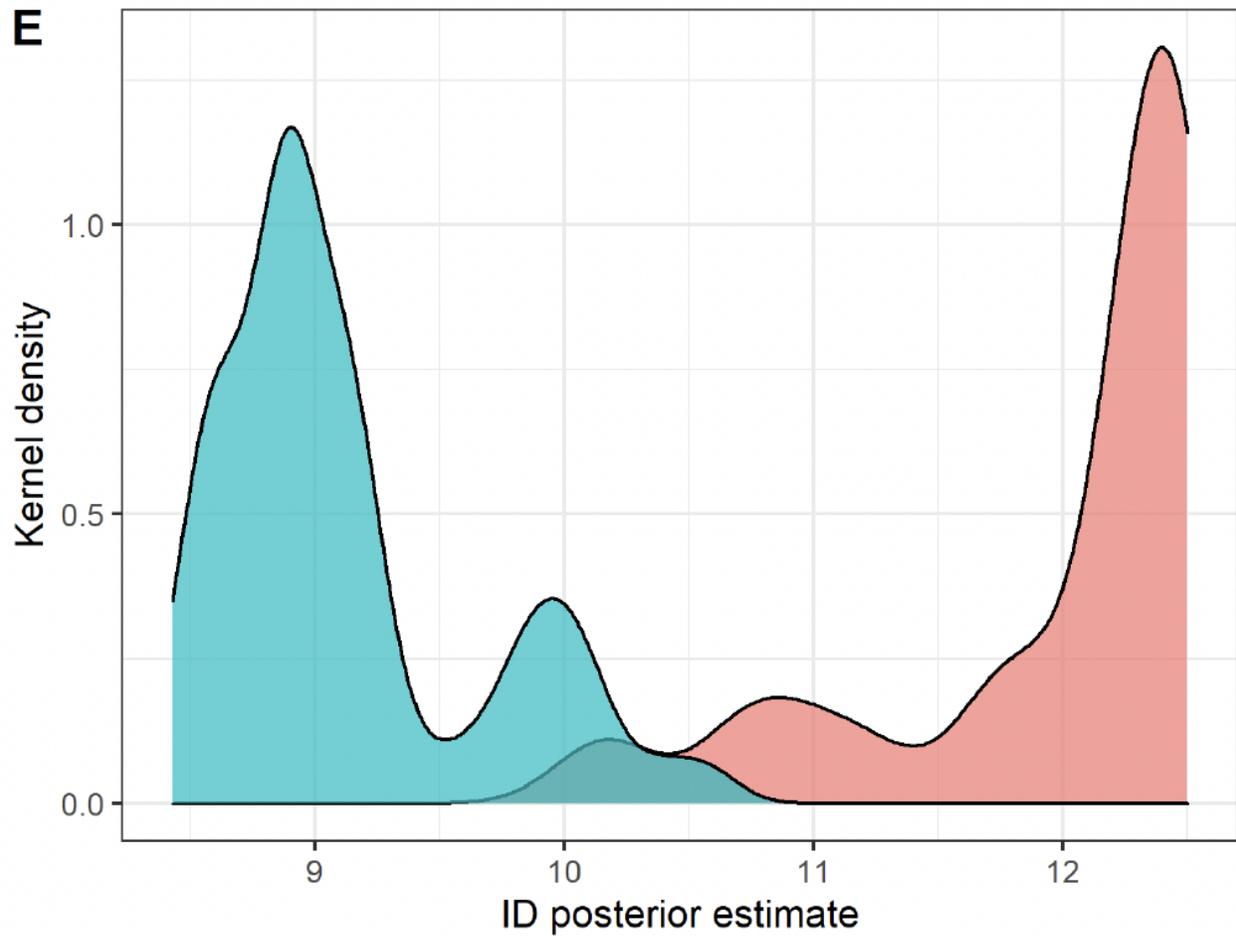
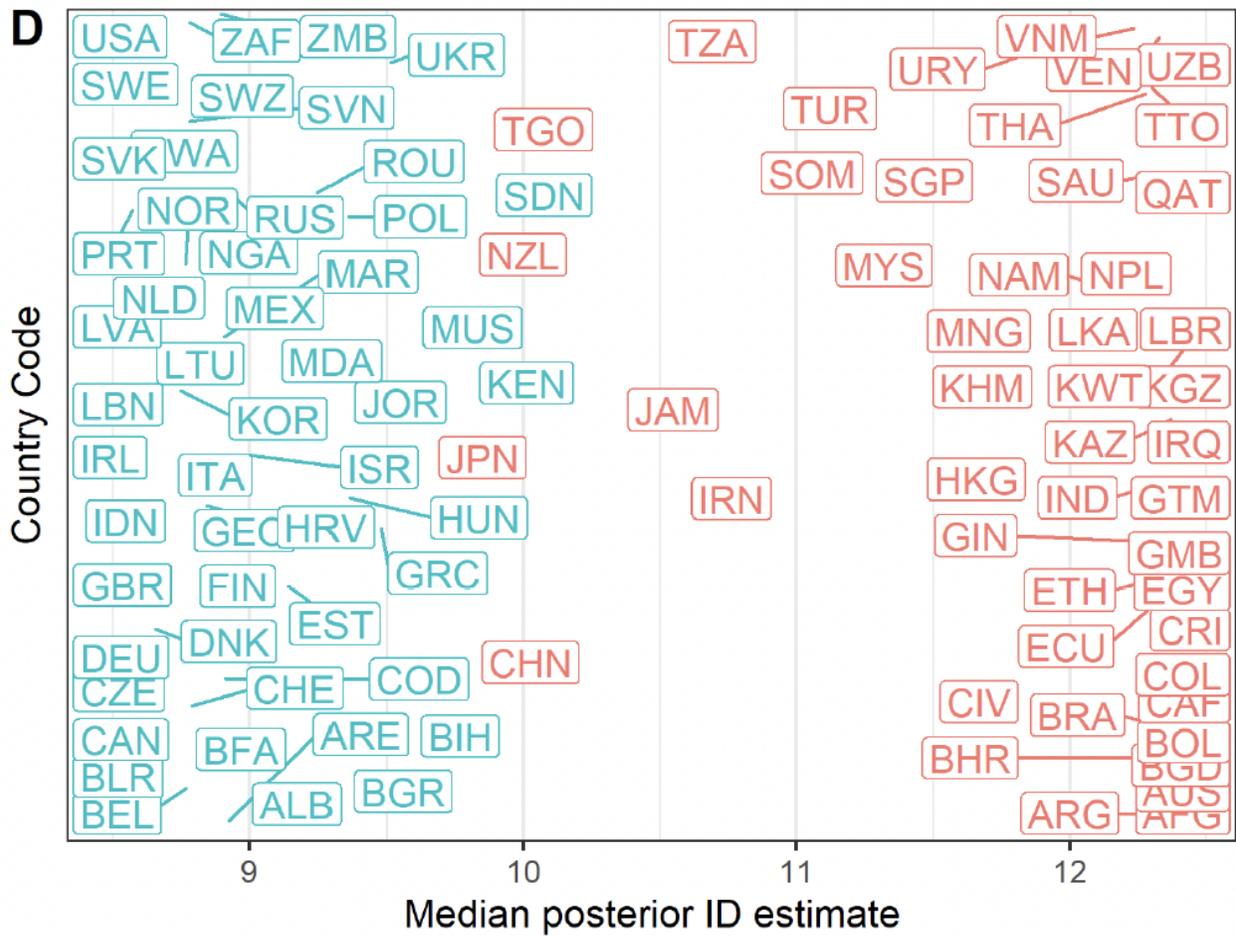
CSI

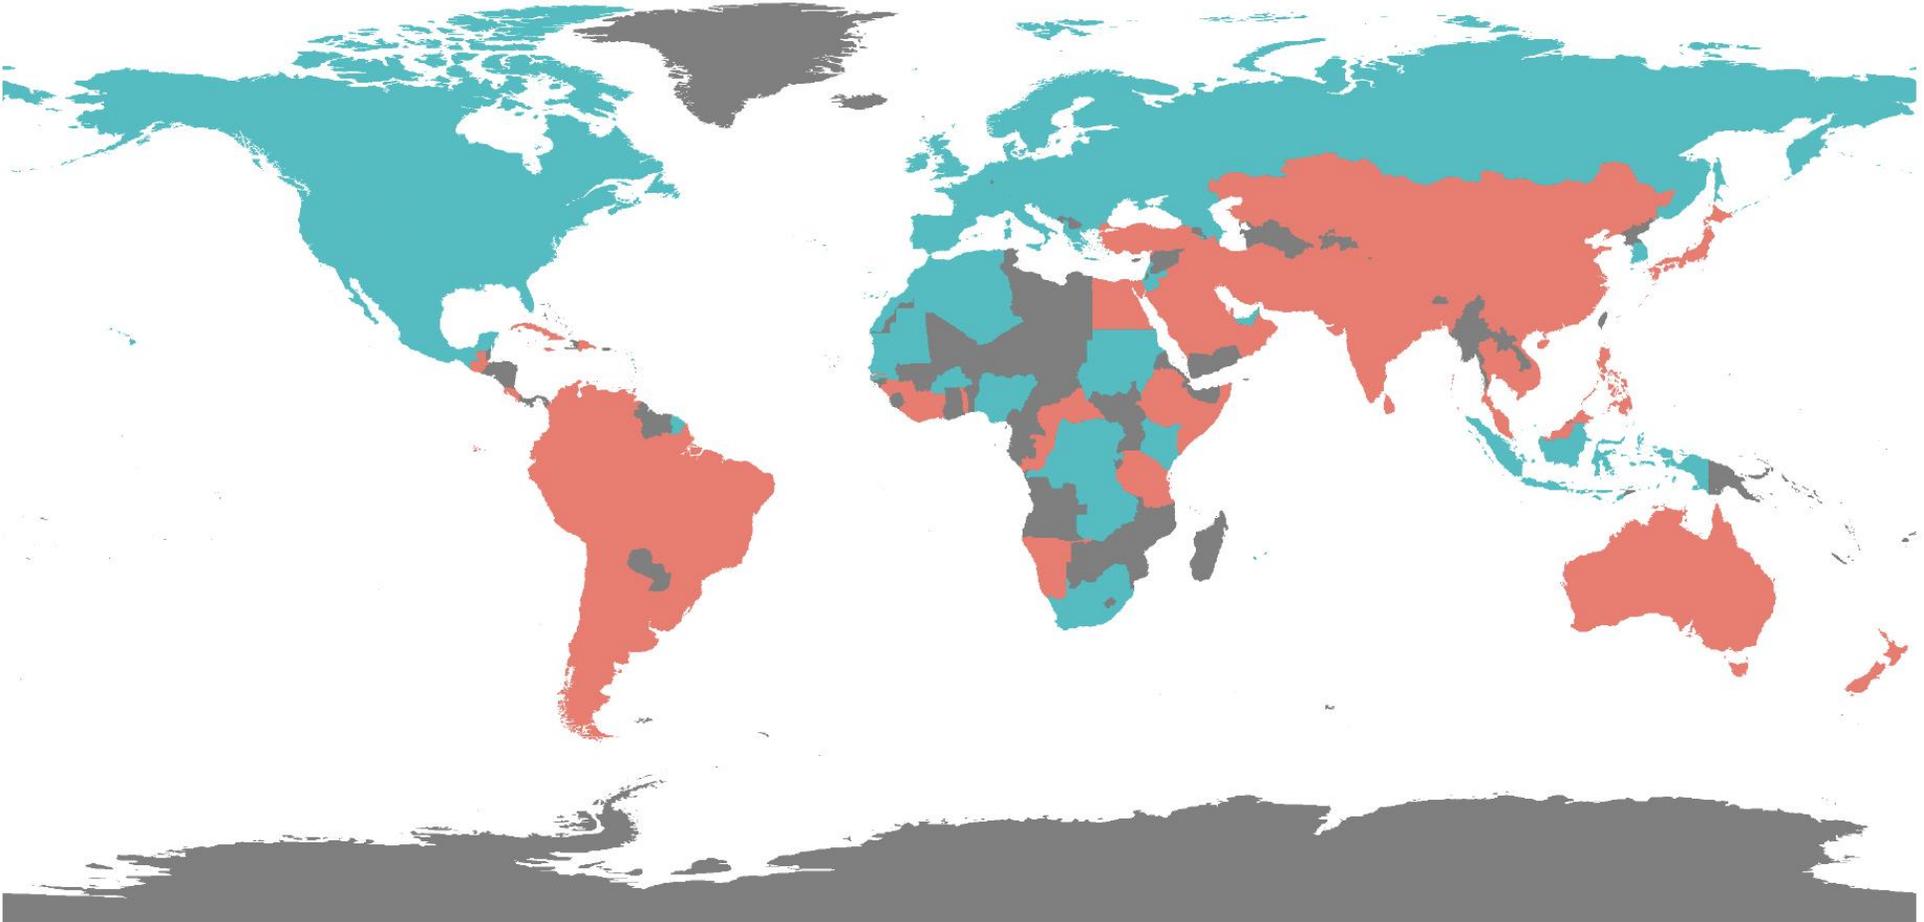


new cases
pmp

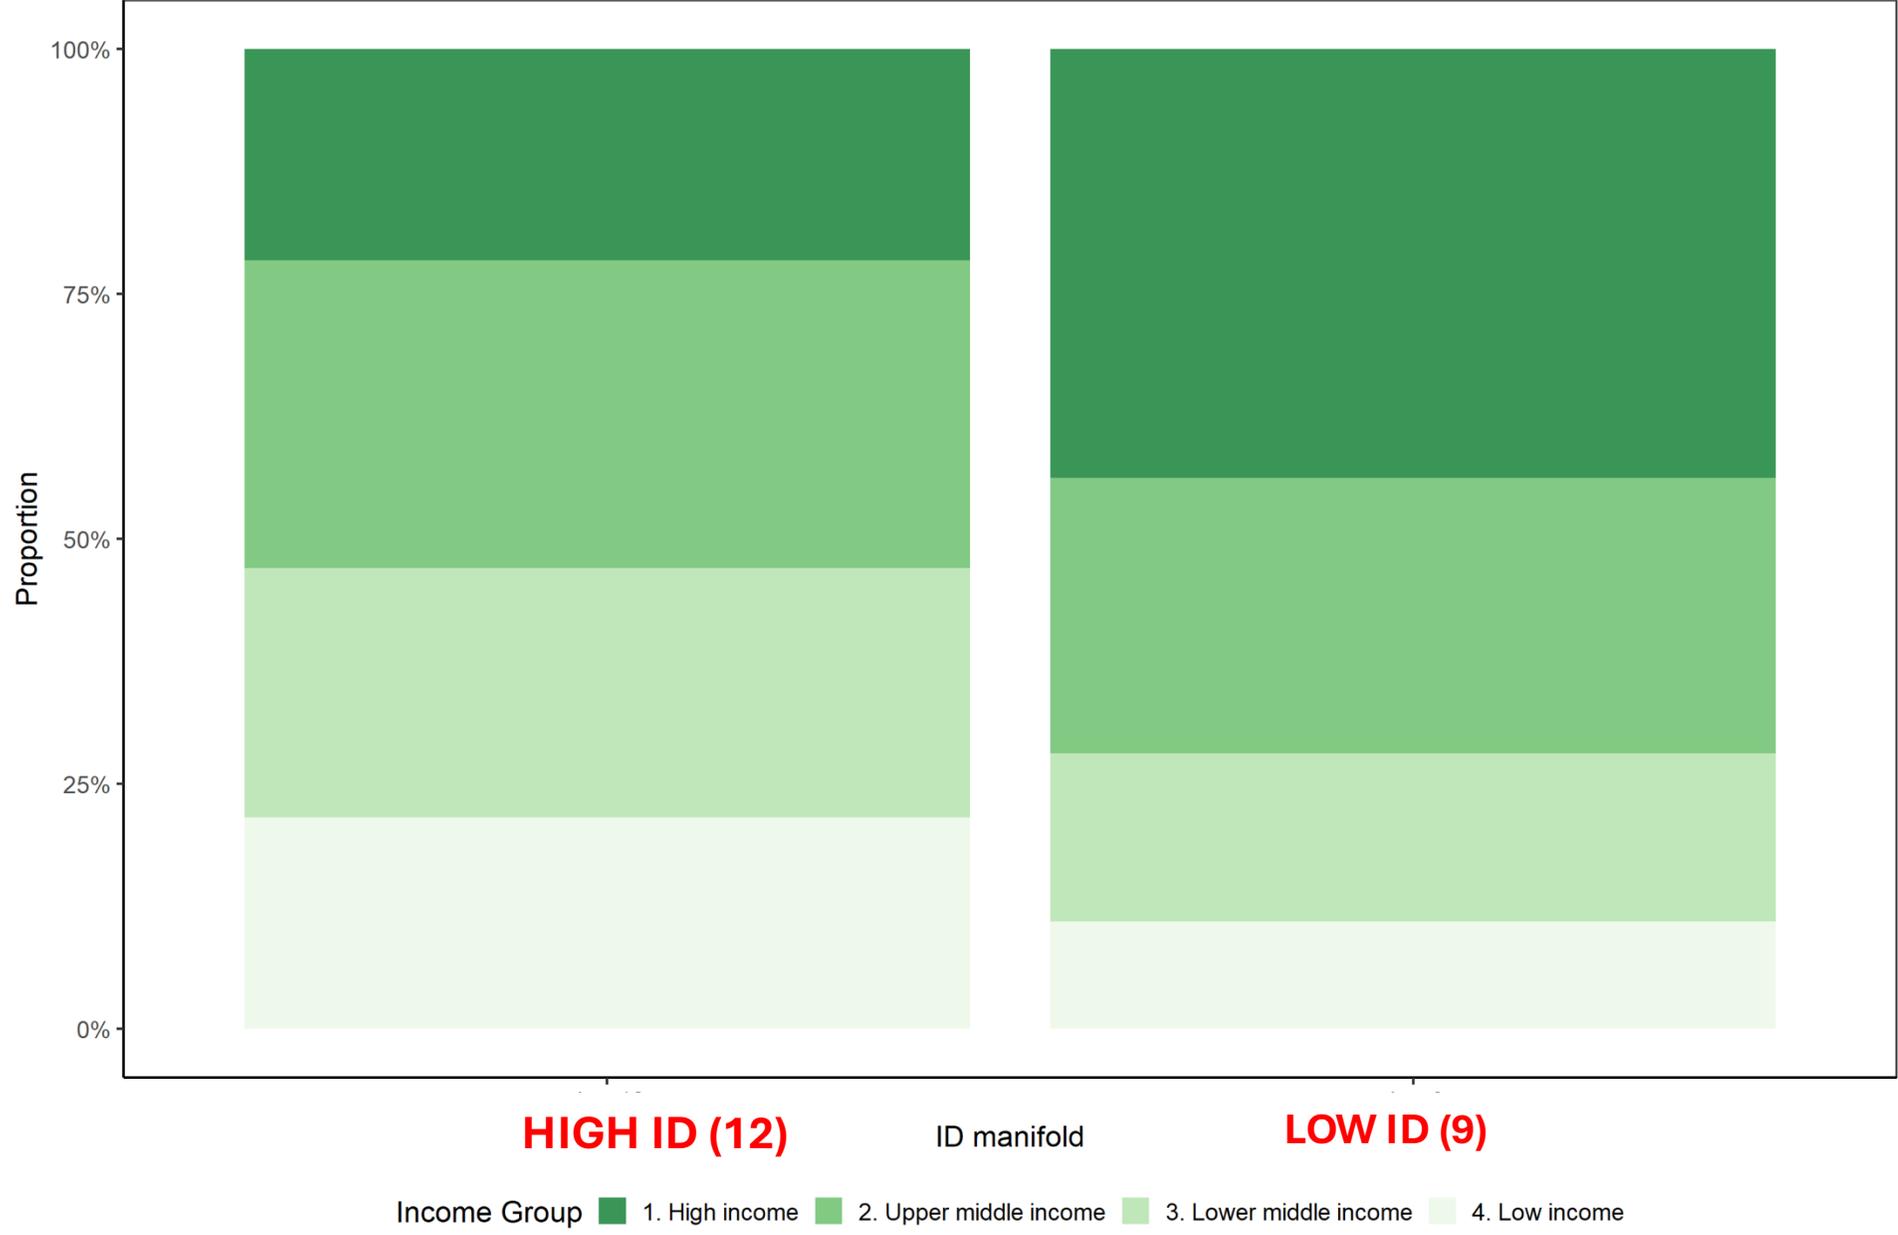


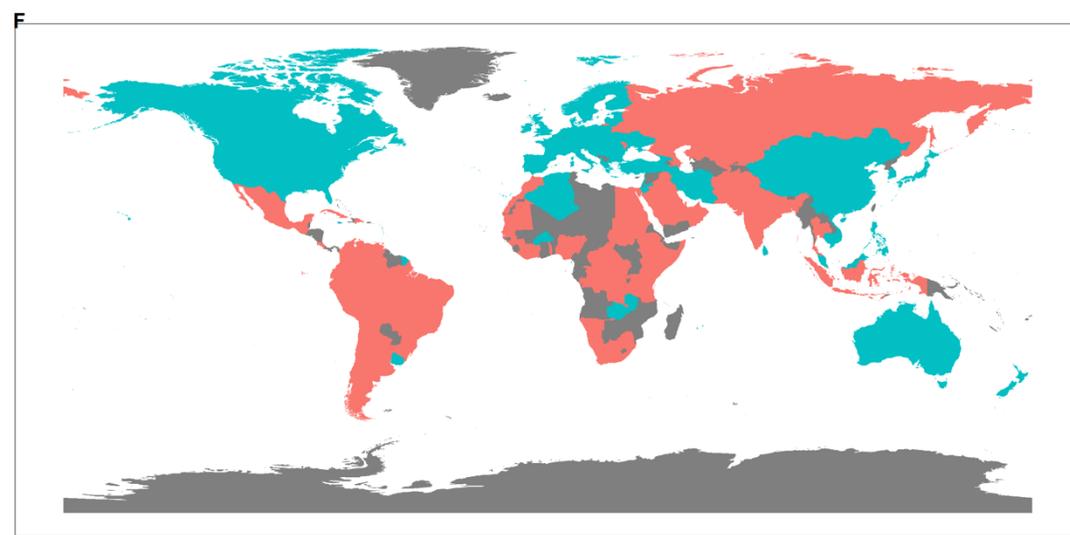
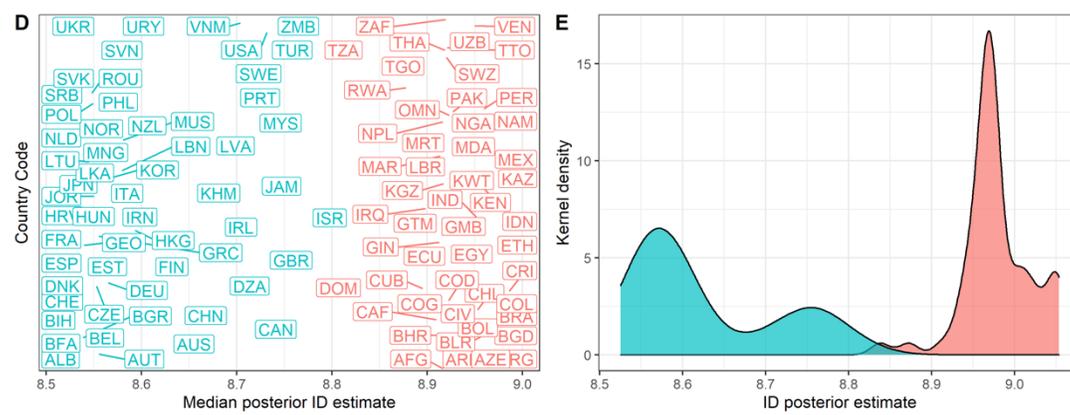
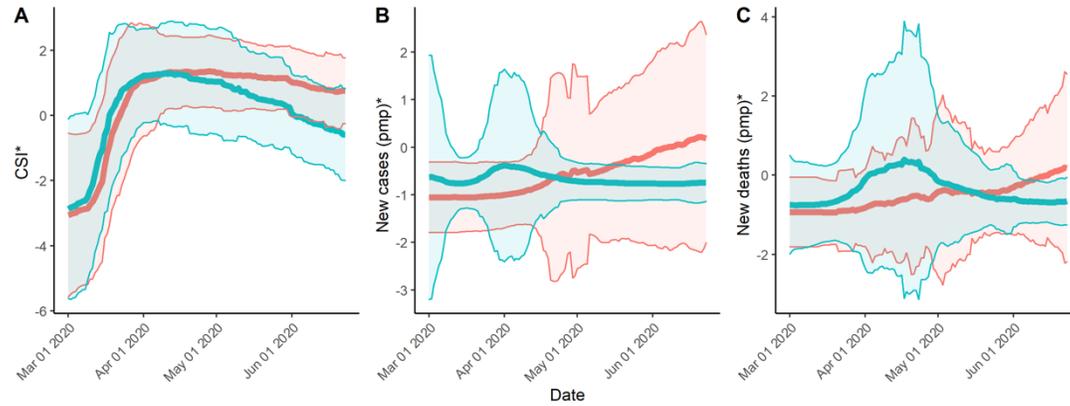
new deaths
pmp



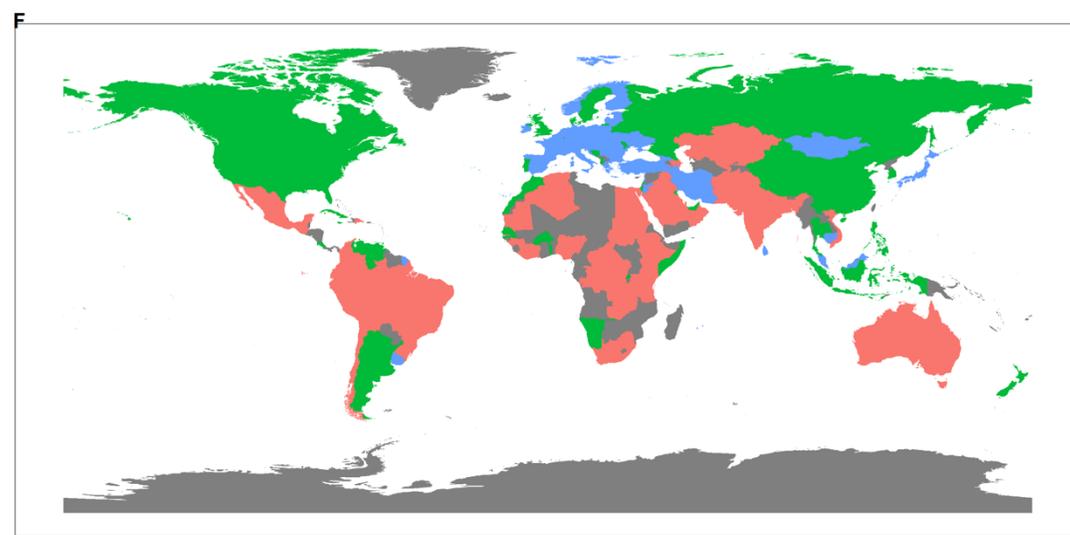
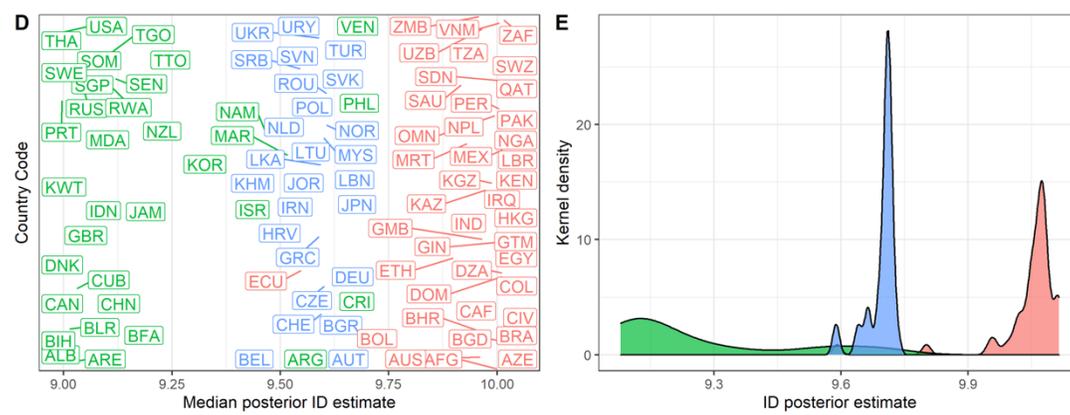
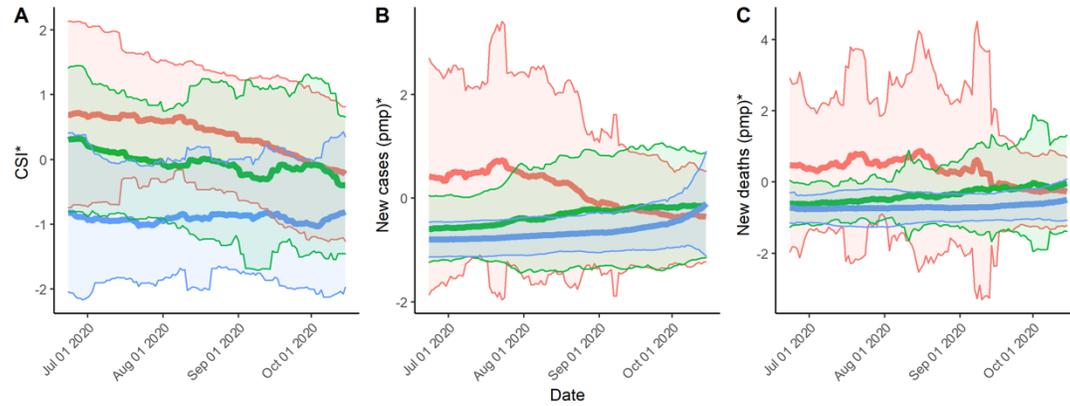


ID manifold ■ $d_1 = 12$ ■ $d_2 = 9$ ■ NA





ID manifold 1 2 NA



ID manifold 1 2 3 NA

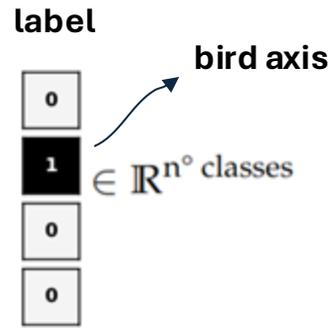
Ansuini, Laio, Macke, Zoccolan, *Advances in NIPS* , 2019

- Study the ID of data representations in CNN
- In a trained CNN, the ID is orders of magnitude smaller than the number of units in each layer
- Across layers, the ID first increases then decreases
- The ID of the last hidden layer predicts classification accuracy on the test set in CNN to classify images
- Not true for untrained networks
Not true for networks trained on randomized labels

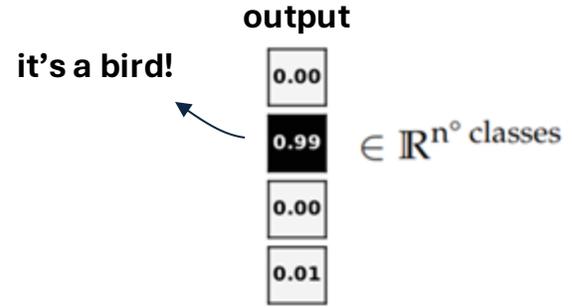
Conclusion: NN that can generalize are those that transform the data into low-dimensional, but not necessarily flat manifolds

ID exploitation for selecting CNN architectures and training procedures

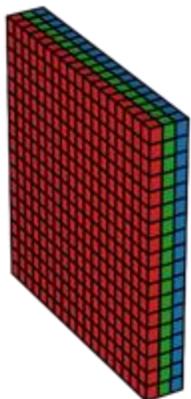
Hidden representations of convolutional networks



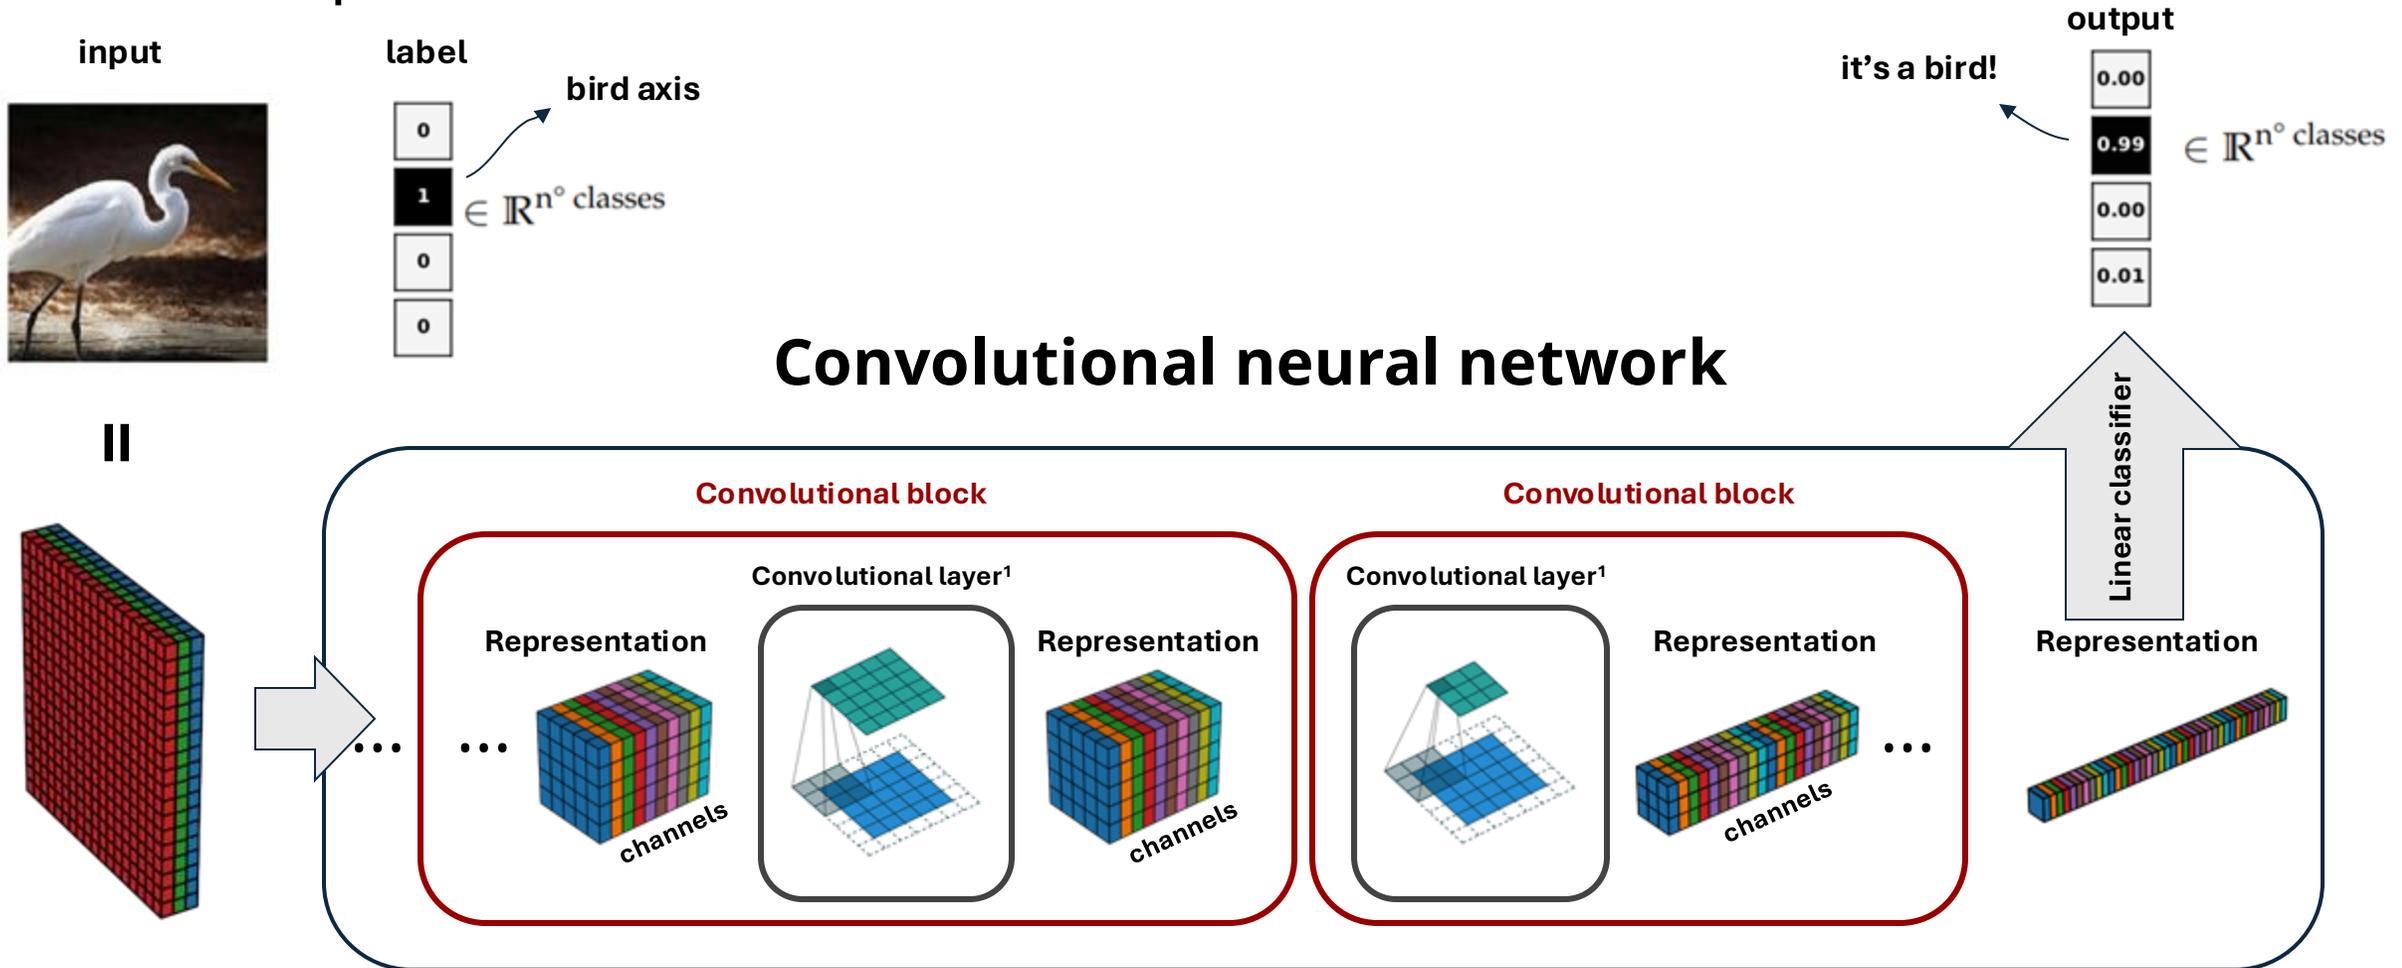
Convolutional neural network



||



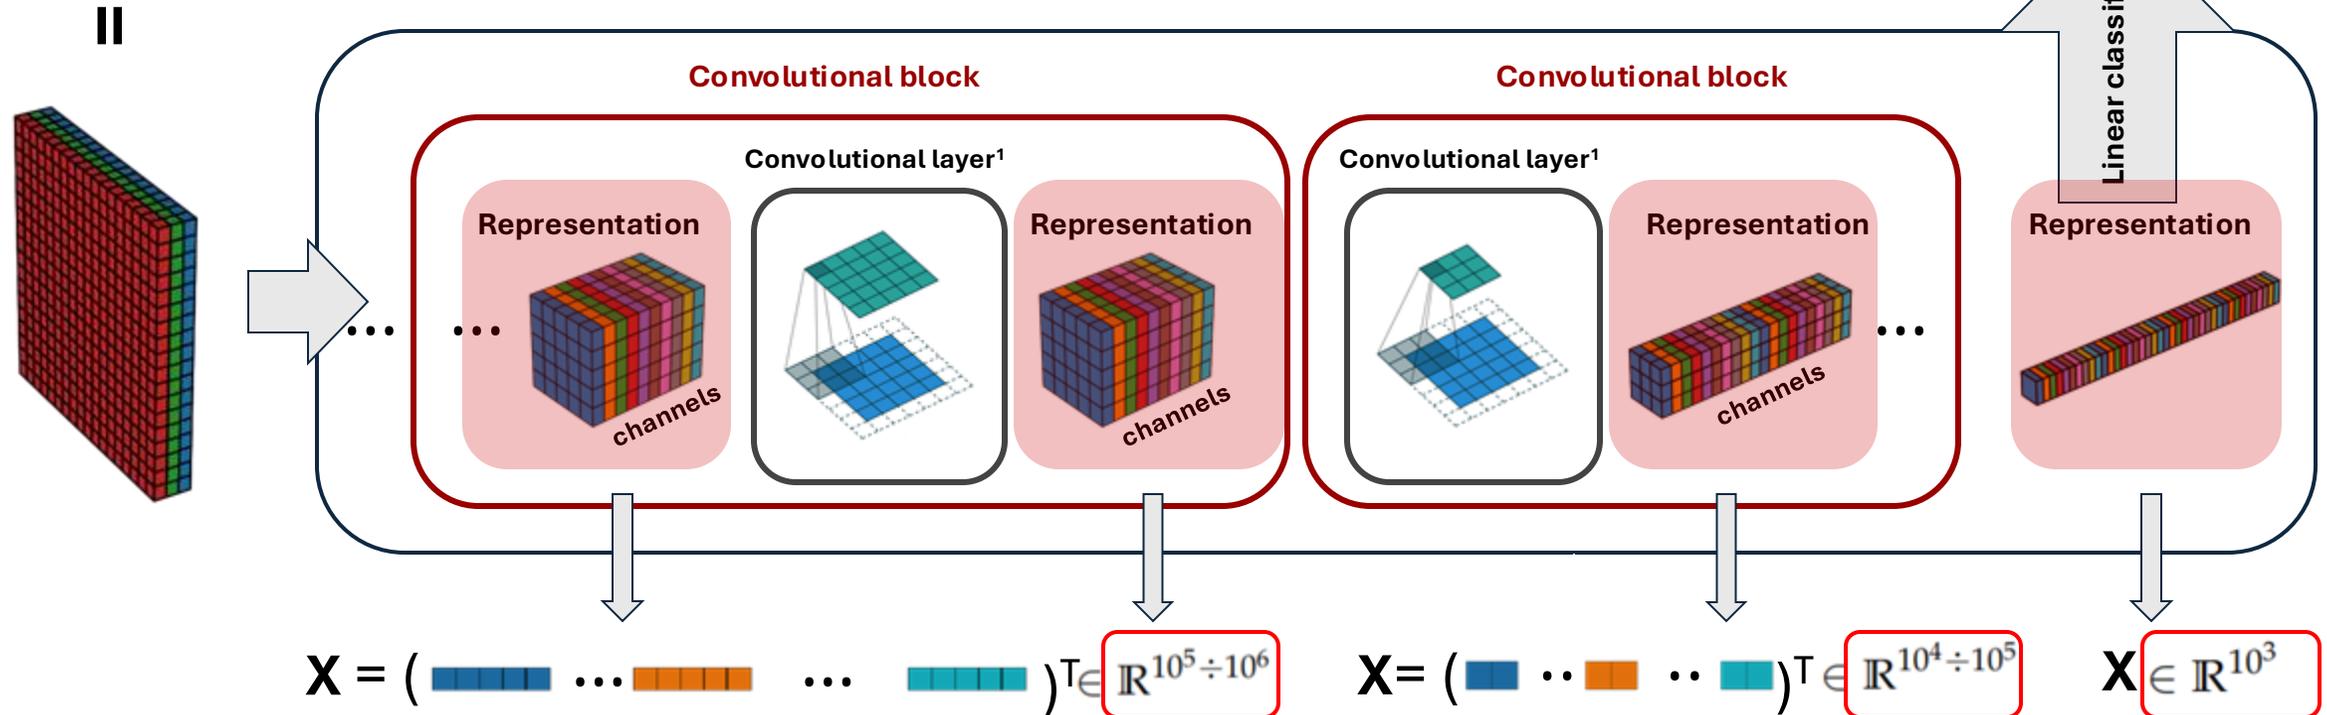
Hidden representations of convolutional networks



Hidden representations of convolutional networks



Convolutional neural network



ID of hidden representations of ResNet152

Dataset ImageNet: 1.2 million images, 1000 classes

We consider a subset of **300 classes with 300 images per class** for a total of **90 000**

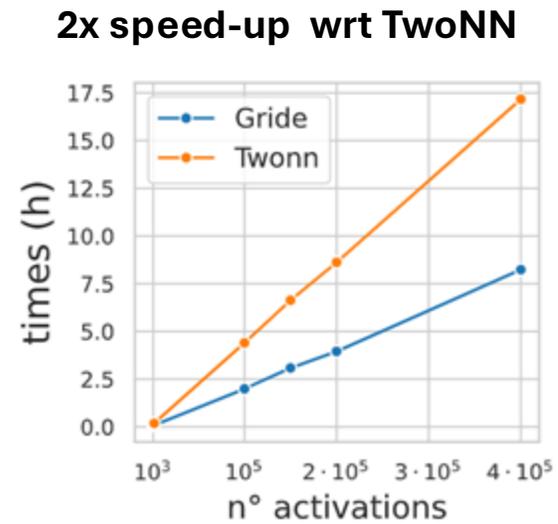
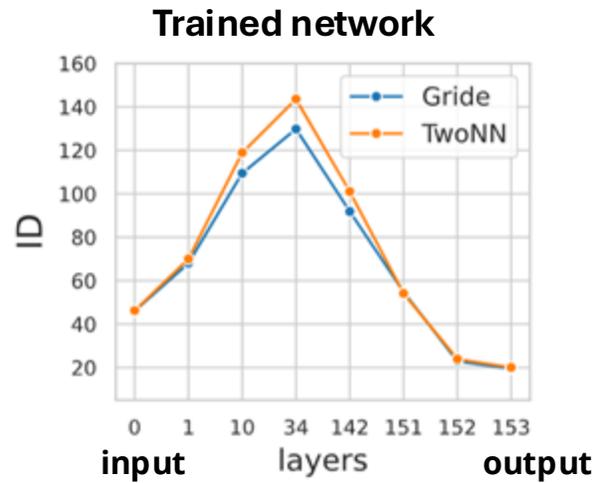


ges



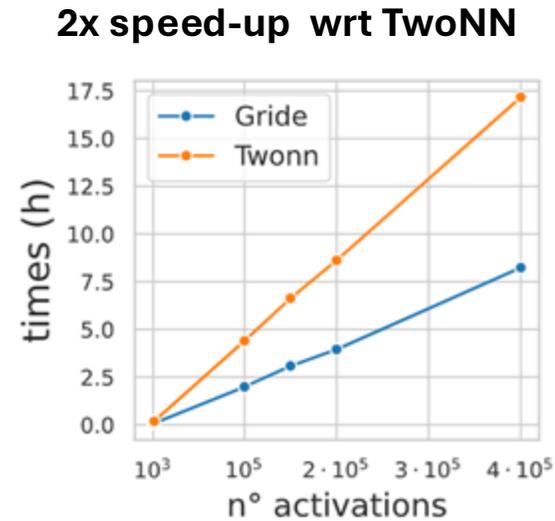
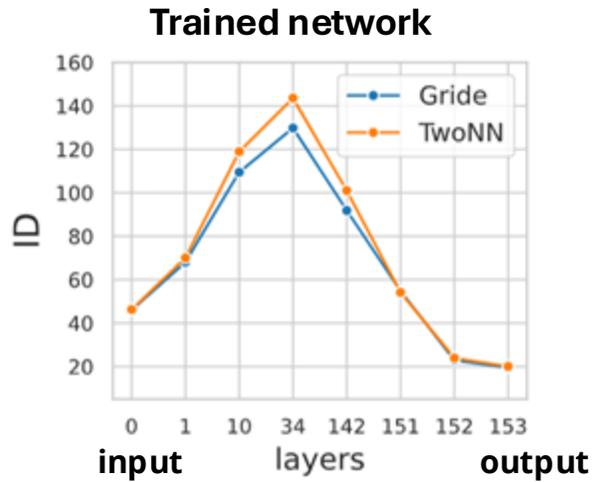
ID of hidden representations of ResNet152

Ansuini et al.,
NeurIPS, 2019

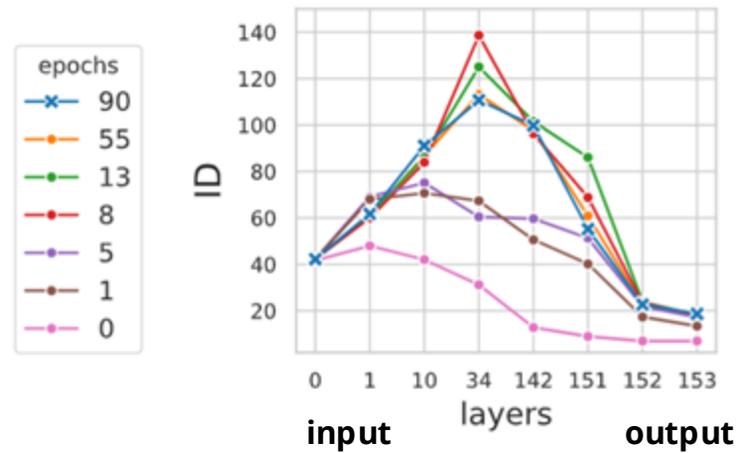


ID of hidden representations of ResNet152

Ansuini et al.,
NeurIPS, 2019



Evolution of the ID during training

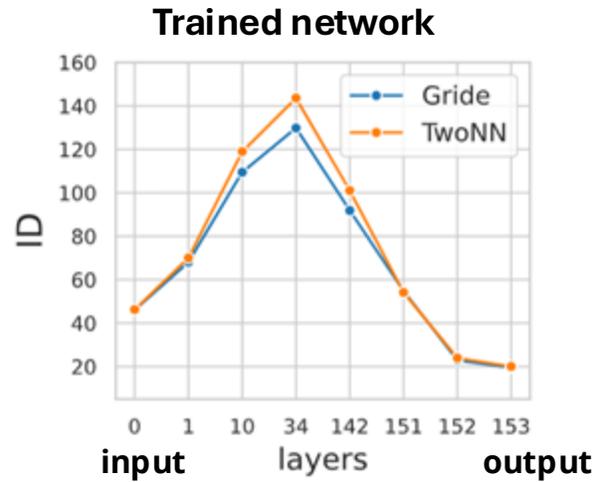


90 epochs

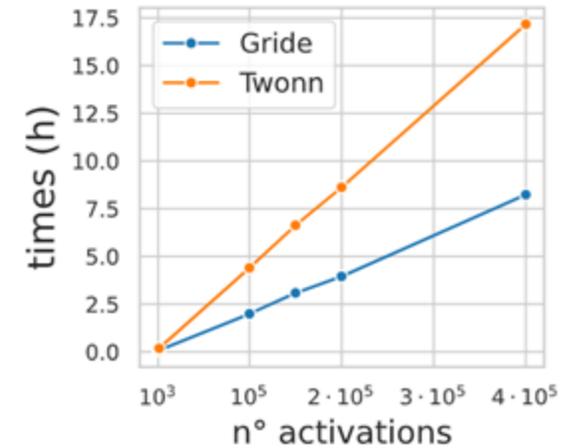
Greatest change of the ID
in the first 10 epochs

ID of hidden representations of ResNet152

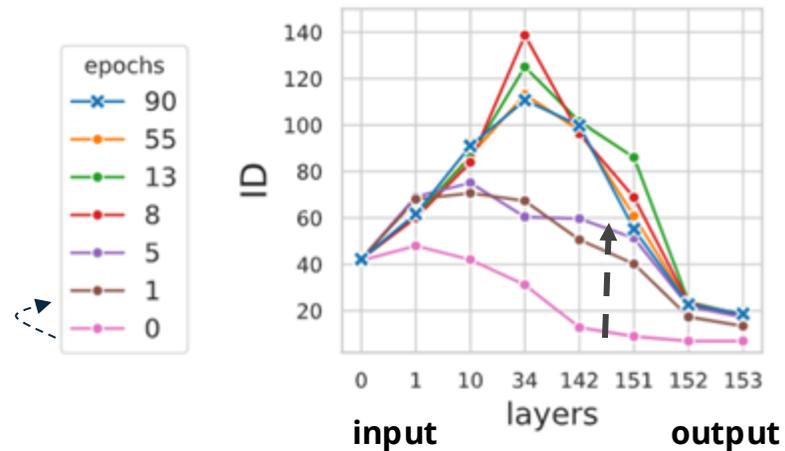
Ansuini et al.,
NeurIPS, 2019



2x speed-up wrt TwoNN



Evolution of the ID during training

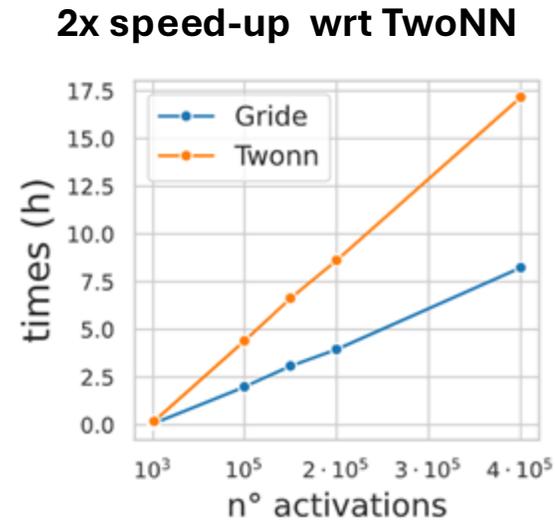
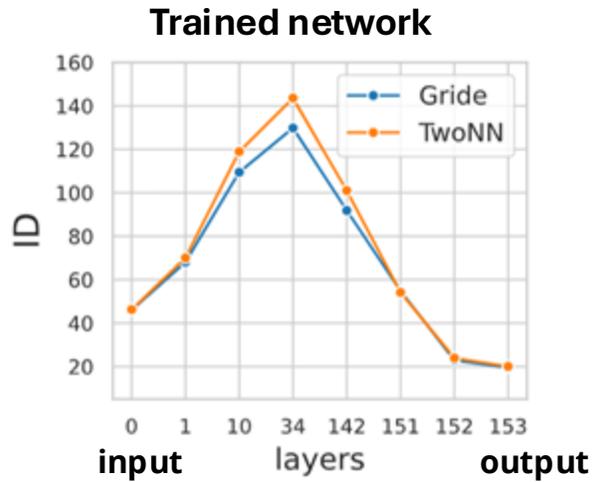


90 epochs

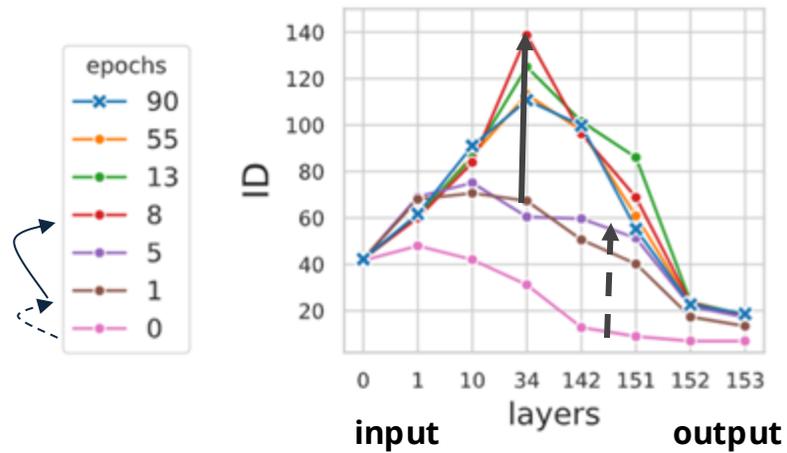
Greatest change of the ID
in the first 10 epochs

ID of hidden representations of ResNet152

Ansuini et al.,
NeurIPS, 2019



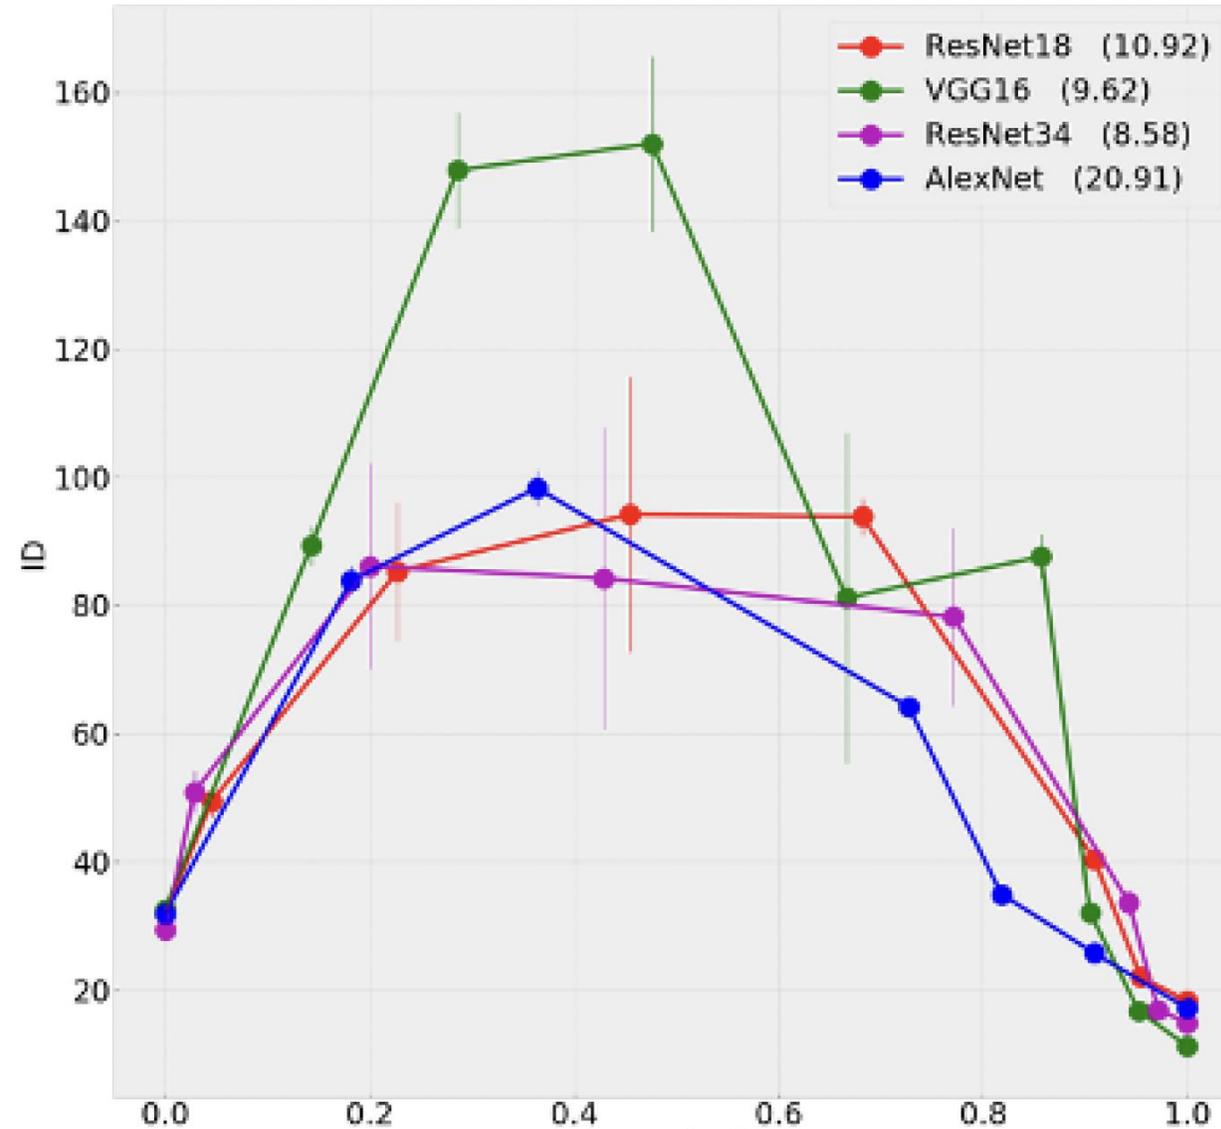
Evolution of the ID during training



90 epochs

Greatest change of the ID
in the first 10 epochs

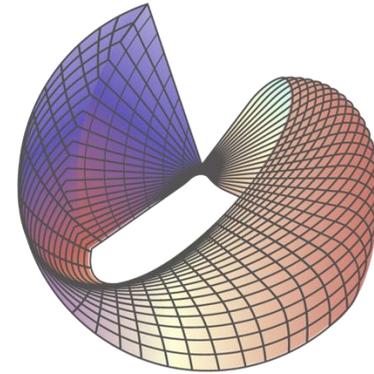
CNN architectures: AlexNet, Vgg, Resnet pretrained on ImageNet
50 samples of ~ 2000 images
ID computed as function of the NN layer depth



R package by F. Denti
Cattolica University, Italy



Python suite by SISSA,
Trieste, Italy



DADApy

References

Facco, d'Errico, Rodriguez, Laio;

Estimating the intrinsic dimension of datasets by a minimal neighborhood information.

Scientific Reports 2017

Allegra, Facco, Denti, Laio, Mira

Data segmentation based on the local ID,

Scientific Reports 2020

Denti, Doimo, Laio, Mira,

Distributional Results for Model-Based ID Estimators,

Scientific Reports 2022

Vargehese, Santos-Fernandez, Denti, Mira, Mengersen,

A global perspective on the intrinsic dimensionality of COVID-19 data

Scientific Reports 2023

Denti,

intRinsic: an R package for model-based estimation of the ID of a dataset,

J Stat Software 2023