




Big Data ACUA6002 Lecture 4

Dimensionality Reduction

*„The more you know, the less you need“
-Yvon Chouinard*


Can we condense information?

**This lesson refers to chapter 4 of the GIDS book*

- Methods for Dimensionality Reduction
 - Principal Component Analysis (PCA)
 - Linear discriminant Analysis (LDA)
 - Multidimensional Scaling (MDS)
 - t-distributed Stochastic Neighbor Embedding (t-SNE)

Datasets


- Datasets used : adult dataset and outliers dataset
- Example Workflows:
 - „Dimensionality Reduction“ <https://kni.me/w/EnmfIBCuLOFpYvc2>
 - PCA
 - t-SNE



PCA: Principal Component Analysis

Principal Component Analysis: Goal

- Method from statistics to construct a **projection** from the high-dimensional space to a lower-dimensional space
- Uses the **variance** in the data as structure preservation criterion
- Projection must be **to a linear subspace** which preserves as much as possible of the original **variance** of the data




First PC (solid line) in the direction of maximum variance; second PC (dashed line) in the direction of second max. variance and orthogonal to first PC.

- (Sample) variance for one numerical attribute:
- Where x_i is the i -th sample in the dataset, n is the number of samples, and $\bar{x}$ is the sample mean.
- Variance of a multidimensional data set = Sum of the variances of the attributes
- Covariance Matrix:

Where $\mathbf{x}_i$ is the i -th sample vector in the dataset, n is the number of samples, and $\bar{\mathbf{x}}$ is the mean vector.

Principal Component Analysis: Projection



- The projection can be represented by a matrix M mapping the data points to the plane by:
- Where $\bar{x}$ is the sample mean of x , i.e. the vector of mean values, n the number of data, d the data dimensionality, and ...

- ... and
- is the j -th ***principal component***, that is the vector in the direction of the j -th max. variance of the dataset
- With constraints:

Solution:

- are the **eigenvectors** of the covariance matrix C of the dataset
- are the **eigenvalues** of the covariance matrix C of the dataset and the variance associated with each eigenvector

The projection matrix is given by
where the principal components are the normalized eigenvectors of the covariance matrix C of the data

Sorted by the corresponding eigenvalues

Note.

- is called an eigenvalue of a matrix A , if there is a non-zero vector $\mathbf{v}$ such that $A\mathbf{v} = \lambda\mathbf{v}$ holds.
- The vector $\mathbf{v}$ is called eigenvector to the eigenvalue (as you learned in your maths lectures).

Principal Component Analysis: Dimensionality Reduction

- If moves the data from to , there is no reduction in dimensionality.
- We look for with .
- Total variance in the dataset with .
- Let's preserve just a fraction of the total variance:
- If we use the top we lose only a bit of the original variance
- We project the data onto the first principal components corresponding to the eigenvalues , thus preserving the given fraction of the data variance

Principal Component Analysis

- PCA is a statistical procedure that **orthogonally** transforms the original n coordinates of a data set into a new set of n coordinates, called principal components
- We calculate the largest eigenvalue and the corresponding eigenvector of the covariance matrix C , as the direction of the largest variance in the dataset
- Each succeeding component must follow the direction of the **next largest possible variance** under the constraint that it is orthogonal to (i.e., uncorrelated with) the preceding components
- describes most of the variability in the data, adds the next big contribution, and so on. In the end, the last do not bring much more information to describe the data.
- Thus, to describe the data we use only the top (i.e.,) components with little - if any - loss of information

Principal Component Analysis



	Principal Components			
	PC1	PC2	PC3	PC4
Proportion of variance	0.73	0.229	0.0367	0.00518
Cumulative Proportion	0.73	0.958	0.9948	1.0000

Preservation of the variance of the Iris data set depending on the number of principal components.

- **Caveats:**
- Results of PCA are quite difficult to interpret
- Data normalization is required (z-score normalization)
- Only effective on numeric columns

Principal Component Analysis: the „Cube“ data

PCA




- Impression that the data are uniformly distributed over a grid
- Data are visibly not uniformly distributed
- Outlier not represented
- Outlier is visible

LDA: Linear Discriminant Analysis

Linear Discriminant Analysis: Goal

- LDA is a **classification algorithm** based on class statistics
- It can also be used as a method to construct a **projection** from the high-dimensional space to a lower-dimensional space
- Uses the class **separation** as structure preservation criterion
- Projection must be **to a linear subspace** which preserves as much as possible of the original **separation** of the classes in the data




First LD (red line) in the direction of maximum discriminability between classes; second LD (blue line) in the direction of second max. discriminability and orthogonal to first LD.

- describes best the class separation in the data, adds the next big contribution, and so on. In the end, the last do not bring much more information to separate the classes.
- Thus, for our classification problem, we could use only the top $\frac{1}{2} \leq \frac{1}{2}$ (i.e., ,) discriminants with little - if any - loss of information for classification
- **Caveats (as for PCA):**
 - Results of LDA are quite difficult to interpret
 - Normalization required
 - Only effective on numeric columns

Linear Discriminant Analysis vs. Principal Component Analysis

- LDA and PCA are both statistical procedures that **orthogonally** transform the original n coordinates of a data set into a new set of n coordinates, called linear discriminants (LD) or Principal Components (PC).
- Discriminants maximize the separation between classes, components maximize the variance in the data
- PCA: unsupervised
- LDA: supervised



MDS: Multidimensional Scaling

- Is not constructing an explicit mapping from the high-dimensional space into the low-dimensional space
- Only positions data points in the low-dimensional space
- Uses the **distance** between the high dimensional data points as structure preservation criterion (not the variance, not the class separability)

Multi-Dimensional Scaling: Goal

- Input: $(x_1, \dots, x_n)$ with $(input\ space)$
- Output: $(y_1, \dots, y_n)$ with $(output\ space, usually\ q=2\ or\ q=3)$

Goal:

- Define a data point y_i for each data object x_i such that all distances between point y_i and all other points in the output space are **roughly** the same as the corresponding distances between the original data objects x_i and in the input space .

and

- Distance function requirements:
 - Non negative
 - Symmetric
 - Positive definite
- Usually Euclidean distance
- Starting Point: Distance matrix in input space
- Where is the distance between data point and data point in the input space.
- Usually Euclidean distance in after normalization

- **Objective functions (error measures)** to define the quality of a solution
- **Sum of squared errors**
- **Normalised sum of squared errors**
- In contrast to E_0 , the value of E_1 does neither depend on the number of data objects nor of the original distances

(The normalization factor does not influence the location of the minimum of the objective function)


- **Relative Squared Error**
- Mixed absolute and relative error (also called **Stress**)
- MDS based on is called **Sammon Mapping**

- **How to minimize the error measures?**
- No closed solution
- Non-linear optimization problem with (for) parameters to be optimized
- Even for a small dataset like the Iris dataset (), we need to optimize 300 parameters!
- Because we are searching for the best position of in a two-dimensional space (x,y) for each one of the 150 data points
- Heuristic strategy needed Typically a **gradient descent** method

Reminder: gradient descent based solutions

The gradient, i.e., the vector of partial derivatives of the objective function with respect to the model parameters, points in the direction of steepest ascend.

- The gradient is the vector of partial derivatives:
- The gradient points in the direction of steepest ascend.



Reminder: gradient descent based solutions

The gradient, i.e., the vector of partial derivatives of the error function with respect to the model parameters, points in the direction of steepest ascend.


1. Start at a random point , i.e. an arbitrary choice of the model parameters
2. Calculate the gradient of the objective function
3. From move a certain step in the opposite direction of the gradient (for an error function to minimize) and reach point
4. Calculate the new value of the objective function in
5. Repeat from 2.

Procedure continues until no more improvements on the objective function can be achieved (step 4), a fixed number of gradient steps has been carried out, or is too close to .

Notes:

- must be differentiable
- The landscape of the objective function (vs.) cannot be plotted already for
- The gradient descent procedure moves in steps along the objective function
- The step size: constant or adaptive?
- Problem of the local minima (maxima)

It is recommended to run a gradient method repeatedly, starting with different initial points to increase the chance to find the global or at least a good local optimum.



Multi-Dimensional Scaling: gradient descent example

- Error measure
- Gradient of error measure with respect to one data point in the output space

using

Multidimensional Scaling: Algorithm

Algorithm MDS($\mathcal{D}$)



input:	data set $\mathcal{D} \subset \mathbb{R}^m$ with $ \mathcal{D} = n$ or distance matrix $[d_{i,j}^{(X)}]_{1 \leq i,j \leq n}$
parameter:	dimension q for the representation, stepwidth $\alpha > 0$, stop criterion SC
output:	set Y of n points in $\mathbb{R}^q$

1	Initialize $Y = \{\mathbf{y}_1, \dots, \mathbf{y}_n\} \subset \mathbb{R}^q$ randomly or better with a PCA projection
2	compute the distances $d_{ij}^{(X)}$
3	do
4	Compute $d_{i,j}^{(Y)} = \ \mathbf{y}_i - \mathbf{y}_j\ $ (for all $i, j = 1, \dots, n$)
5	Compute $\partial E_1 / \partial \mathbf{y}_k$ (for all $k = 1, \dots, n$)
6	update $\mathbf{y}_k^{\text{new}} = \mathbf{y}_k^{\text{old}} - \alpha \cdot \partial E_1 / \partial \mathbf{y}_k$ (for all $k = 1, \dots, n$)
7	while SC is not satisfied

– Caveats:

- Normalization before computing the distances
- Assure that no distance between two data objects is 0 before applying MDS algorithm

Multidimensional Scaling



MDS (Sammon Mapping) for the Iris and the „cube“ datasets

PCA

- Preserves the variance of the dataset
- Provides explicit mapping to the lower-dimensional space
 - easier insertion of new data
- Lower computational complexity
 - Covariance matrix can be calculated in linear time

MDS

- Preserves the distances between data objects
- Does not provide explicit mapping
 - New data cannot be projected
- Higher computational complexity
 - Quadratic for the pairwise distance
 - Complexity problems partially overcome by sampling or variations of the error measures

LDA

- Preserves the separation among classes
- Provides explicit mapping to the lower-dimensional space
 - easier insertion of new data
- Lower computational complexity

t-SNE: t-distributed Stochastic Neighbor Embedding

- Nonlinear dimensionality reduction technique based on nonlinear relationships among the data points
- Primarily used for visualization from very high-dimensional data to 2 or 3 coordinates
- Works with probabilities of two data points being neighbours
- As for MDS, new data cannot be projected. All data must be used to find the new points.

In the input space:


- For any two data points x_i and x_j , probabilities are defined as:
- where p_{ij} is the probability that object x_i picks object x_j as its neighbor
- Parameter α becomes smaller if density of points around x_i is high and vice versa

In the output space:

- The corresponding probabilities for points x_i and x_j are defined using t- or Cauchy distribution-like functions
- The points x_i are chosen in a way that the probabilities p_{ij} and p_{ji} are as similar as possible according to the **Kullback-Leibler divergence**

t-SNE is mainly used for visualization because:


- Aggressive dimensionality reduction
- Transformation of likelihood of pairs into visual proximity
- Capability to represent strange shapes of data



t-SNE applied to the „cube“ dataset. The four original clusters are well preserved in the two-dimensional representation and the outlier is also clearly visible.

Practical Examples with KNIME Analytics Platform

- PCA and t-SNE techniques for data reduction to 2D space



Berthold, M.R., Borgelt, C., Höppner, F., Klawonn, F. and Silipo, R. (2020). *Guide to Intelligent Data Science: How to Intelligently Make Use of Real Data*. 2nd ed. [online] www.springer.com. Springer International Publishing.

Available at:

<https://www.springer.com/gp/book/9783030455736> [Accessed 17 Jun. 2021].

Questions?