

# EFR summary

Seminar Data Analytics

2025 – 2026



Week 1-2

**Deloitte.**

DeNederlandscheBank

EUROSYSTEEM

## **Details**

**Subject:** Seminar Data Analytics (AAC) 2025 - 2026

**Teacher:** M. Merlo, P. Breuer

**Date of publication:** 16.03.2026

---

© This summary is intellectual property of the Economic Faculty association Rotterdam (EFR). All rights reserved. The content of this summary is not in any way a substitute for the lectures or any other study material. We cannot be held liable for any missing or wrong information. Erasmus School of Economics is not involved nor affiliated with the publication of this summary. For questions or comments contact [summaries@efr.nl](mailto:summaries@efr.nl)

# Koffietje doen?

## Start jouw carrière bij BDO

Maak kennis met BDO Accountants & Adviseurs, de beste plek om als toptalent aan de slag te gaan. De koffie staat voor je klaar. Vertellen wij je over wat jij kan bijdragen, en jij ons over je ambities.



*Scan de QR-code  
en plan jouw  
koffiemoment in.*



[werkenbijbdo.nl](https://werkenbijbdo.nl) ▶

**BDO**

# Lecture 1: Anomaly Detection

**Data analytics** = transforming raw data into decision-relevant information. It is an iterative and feedback-driven process. Insights reshape original questions.

Data is a core economic resource.

Organizations use analytics to: Optimize performance, Reduce risk (fraud detection, compliance) and Support strategic decisions (pricing, investments)

Accountants are evolving from 'scorekeeper' (recording past performance) to 'decision facilitator' (providing strategic insights)

→ More data-driven tasks, Risk assessment & predictive analysis, Collaboration with data specialists and Evaluating analytics

## Anomaly Detection

Fraud and scandals often show abnormal patterns before collapse. For example:

- Enron: Stable profits despite rising off-balance-sheet obligations
- WorldCom: Inflated earnings by capitalized operating expenses

**An outlier:** "An observation which deviates so much from the other observations as to arouse suspicions that it was generated by a different mechanism." (Hawkins, 1980), Also referred to as **anomaly**.

**Anomaly Detection:** Identify observations that deviate from expected patterns, making them inconsistent with the rest of the data

- Most observations = normal behavior
- Anomalies = rare and generated by a different process

Examples: Fraudulent transactions, System errors, Data processing mistakes

### Why Anomaly Detection Matters

- **Improves Data Quality:** Identifies errors, duplicates, inconsistencies. Prevents distorted statistical results
- **Enhances ML Performance:** Prevents models from fitting noise. Detects structurally unusual observations
- **Supports Better Decision-Making:** Ensures reliable conclusions, Reduces misleading analysis
- **Reduces Financial & Reputational Risk:** Warning signals for fraud or failure. Protects investors, firms, auditors

Anomaly Detection → Cleaner Data → Better Models → Better Decisions → Lower Risk

### Applications of Anomaly Detection

- **Fraud detection** (unauthorized transactions)
- **Network security** (abnormal logins)
- **Medical diagnosis** (rare disease patterns)
- **Manufacturing** (equipment failure)
- **Climate & environment** (extreme weather)

### Types of Anomalies

#### 1. Point Anomalies :

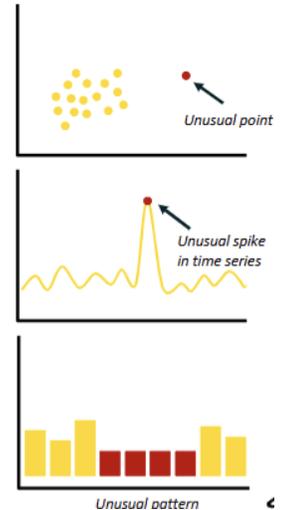
Single unusual observation. Example: One unusually large withdrawal

#### 2. Contextual Anomalies

Unusual given a specific context. Example: Vendor payment at midnight

#### 3. Collective Anomalies

Group of observations jointly unusual. Individual values may appear normal. Example: Sequence of small refunds



### Challenges of Anomaly Detection

#### 1. Conceptual Challenge

- Defining “normal” depends on benchmark (industry, firm size, time)
- Normal behavior evolves over time

#### 2. Statistical Challenge

- Class imbalance (rare events)
- Limited/no labeled data

#### 3. Decision Challenge

- Risk of misclassification

	Actual Normal	Actual Anomaly
Predicted Normal	True Negative	False Negative
Predicted Anomaly	False Positive	True Positive

### Classes of Techniques

#### Statistical Methods

- Use numerical thresholds to detect outliers, examples: Z-Score, IQR

#### Visual Methods

- Spot unusual patterns through plots, example: Histogram

#### Machine Learning Methods

- Detect anomalies using algorithms,
- Approaches: Supervised, Semi-Supervised, Unsupervised

**Note:** No single method detects all outliers → use multiple approaches.

# Statistical & Visual Methods

## Z-Score

The Z-Score standardizes an observation relative to the sample distribution

$$z = \frac{x - \mu}{\sigma}$$

- $x$  = observation
- $\mu$  = mean
- $\sigma$  = standard deviation

### Decision Rule:

Flag if:  $|z| > k$

Common threshold:  $k = 3$

### Interpretation:

- Measures distance from mean in standard deviations
- Larger  $|z|$  = more extreme

### Appropriate When:

- Numeric data
- Approximately normal distribution
- Mean & SD meaningful
- Detecting extreme values

### Strengths:

- Simple
- Unit-free
- Fast screening tool

**So**, Z-score standardizes observations relative to mean and variance. Z-score works best for symmetric, normally distributed numeric data.

### Z-Score: Robustness Issues

#### Sensitivity:

- Sensitive to skewness
- Sensitive to extreme outliers (mean & SD shift)

#### Masking Effect:

- Extreme outlier prevents other outliers from being detected

## Interquartile Range (IQR)

- Q1 (25th percentile)
- Q3 (75th percentile)
- $IQR = Q3 - Q1$

### Decision Rule:

Flag if:  $x < Q1 - 1.5 \times IQR$  or  $x > Q3 + 1.5 \times IQR$

### Interpretation:

- Uses median as central measure
- Focuses on central 50%
- Less affected by extreme values

**So**, IQR method is more robust than Z-score.

*Point anomaly* can be detected using the IQR method

### Appropriate When:

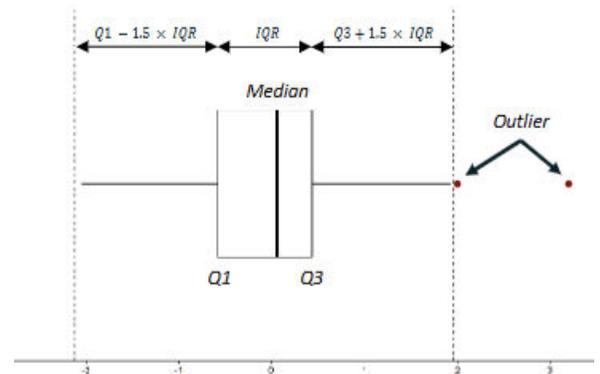
- Numeric data
- Skewed data
- Extreme values present
- Mean/SD unreliable

### Strengths:

- Robust
- No normality assumption
- Simple

### Limitations of Statistical Methods

- **Threshold-based:** Require predefined thresholds, Results depend on benchmark
- **Univariate** (one variable at a time): Cannot detect multivariate anomalies
- **Focus on extreme values:** May miss local anomalies



## Histogram

Divides data into equal-width intervals (bins) and counts the number of observations per bin.

**Anomaly Identification:** Low-density regions may indicate anomalies.

### Appropriate When:

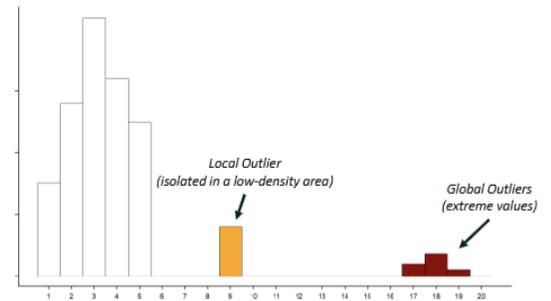
- Numeric data
- Exploring distribution
- Detecting extreme observations
- Visualizing low-density (sparse) regions

### Strengths:

- Intuitive
- Not limited to extreme values

### Limitations:

- Sensitive to bin width
- No formal decision rule
- Univariate only



## When to Use These Methods

### Statistical Methods:

- Univariate
- Extreme outliers
- Quick screening

### Visual Methods:

- Visualizing low-density areas

Both struggle with high-dimensional data.

- Statistical methods on individual columns miss cross-variable patterns
- Visual methods cannot display high-dimensional relationships

Solution: Dimensionality reduction or Multivariate anomaly algorithms (e.g., Isolation Forest)

## Multivariate Anomaly Detection

Anomalies can be unusual combinations of normal values.

Example:

Person	Height	Weight
1	190 cm	80 kg
2	145 cm	45 kg
3	190 cm	45 kg

Individually plausible, jointly unusual (Person 3).

### Limitation:

Z-score & IQR analyze variables separately.

Z-score fails to detect an anomaly in a 2D dataset because the anomaly is only extreme in combination of variables.

### Machine Learning Advantage:

- Learn joint patterns
- Detect deviations in high-dimensional space

## Machine Learning Methods

### What is Machine Learning?

**Artificial Intelligence (AI):** Systems mimicking human intelligence.

**Machine Learning (ML):** Subfield of AI. Learns patterns from data. No explicit rule programming

**Applications:** Image recognition, Speech recognition, Recommendation systems

### Types of Machine Learning

#### 1. Supervised Learning

- Uses labeled data to accurately classify data or predict outputs
- Predict output ( Y ) from input variables ( X )

Examples: Classification (spam vs. non-spam) or Regression (predict revenue)

#### 2. Unsupervised Learning

- Uses unlabeled data
- Discovers hidden patterns and relationships in data from input variables ( X )

Examples: Anomaly detection, Clustering or Dimensionality reduction

## Why Anomaly Detection Is Usually Unsupervised

- Anomalies are rare
- Labels often unavailable
- Fraud patterns evolve over time

Supervised methods require labeled anomalies and stable fraud definitions

Therefore, focus on unsupervised algorithms that learn normal behavior and flag deviations

## But Supervised Methods Are Not Useless

Supervised methods are powerful when labels exist and provide clear classification

Difference in questions:

- Supervised → *What predicts fraud?*
- Unsupervised → *What looks different from the rest?*

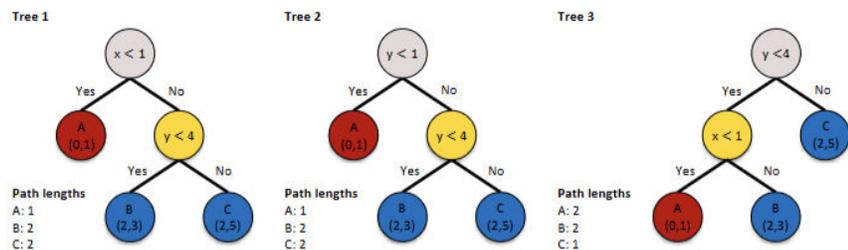
## Isolation Forest

Isolation Forest is an unsupervised, tree-based ML method. It is designed for anomaly detection and works well on numeric tabular data. Goal: anomalies easier to isolate.

**Assumption:** Anomalies are rare & distinct → random partitions isolate them quickly.

“Is feature ( $x$ ) larger/smaller than threshold ( $k$ )?”

Data Point	x	y
A	0	1
B	2	3
C	2	5



Fewer splits needed to isolate a point → more abnormal.

## Isolation Forest: Algorithm

For each tree:

1. Randomly select a feature (e.g.,  $x$  or  $y$ )
2. Randomly select a split value (between min and max)
3. Split data
4. Repeat until:
  - Only one observation in node OR
  - Maximum depth reached
5. Measure path length ( $h(x)$ )
6. Repeat across many trees
7. Compute average path length
8. Convert to anomaly score

## Anomaly Score

Output:

- Continuous score in  $\in(0,1)$
- Not binary

**So**, Short path  $\rightarrow$  high anomaly score.

Two main parameters:

### 1. Number of Trees (t)

- More trees  $\rightarrow$  more stable
- Higher computation cost

### 2. Subsample Size (m)

- Default often 256
- Not necessarily full dataset

Why smaller samples?  $\rightarrow$  Faster, Anomalies stand out more

Why Ensemble?

- Single tree: High variance and Random splits unstable
- Isolation Forest: Many trees, Average path length stabilizes results and Reduces randomness

**So**, Ensemble reduces variance and increases stability.

## Isolation Forest: Strengths & Limitations

### Strengths:

- Scales well to large datasets
- Works in high dimensions
- Robust to irrelevant features
- Fast to compute

### Limitations:

- Limited interpretability (cannot clearly explain which features cause anomaly)
- Produces ranking, not classification (requires professional judgment)
- Stochastic (slightly different results per run)

Formula:

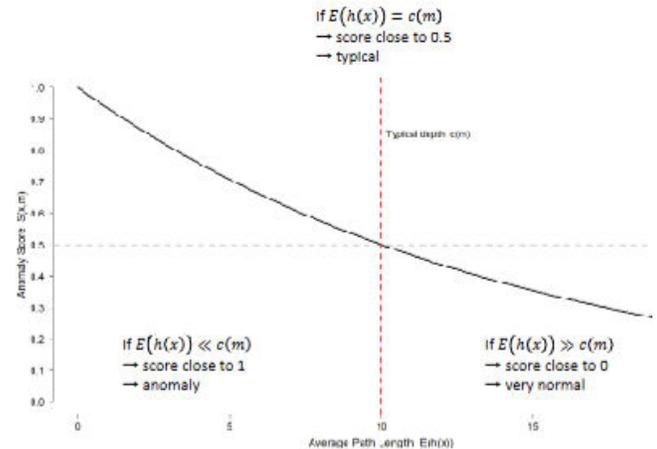
$$S(x, m) = 2^{-E(h(x))/c(m)}$$

, where

$E(h(x))$ : average path length of data point  $x$

$c(m)$ : expected average path length in a tree of size  $m$

$m$ : number of data points used to build each isolation tree



## Case Study: Financial Data

- Goal: Detect unusual firm-years.
- Dataset: 1,500 firm-years
- Variables: Assets, Sales, Net income, Debt, Operating cash flow, R&D, Industry
- Challenge: Financial data messy (missing values, scale differences)

Steps: 1. Prepare data, 2. Build model, 3. Extract scores, 4. Interpret results

**Problem:** Raw levels scale with firm size. Example:

SmallCo (Assets=100, NI=5), BigCo (Assets=10,000, NI=500) → Both ROA = 5%

Solution: Use ratios such as ROA, Leverage (Debt/Assets), Efficiency (Sales/Assets)

Ratios reflect economic behavior, not size.

**Problem:** Isolation Forest requires complete data.

If missing: Cannot assign branch, Tree inconsistent. Solutions:

- Missing indicator variable
- Impute numeric (mean/median)
- Impute categorical (mode)
- Drop rows only if rare (<1–2%)

Important: Missingness itself can signal anomalies.

**Problem:** Different variables have different ranges. Example:

- ROA: -20% to +20%
- Leverage: 0–90%
- Turnover: 0–5+

Isolation Forest splits between min and max. Large-range variables dominate splits.

Solution: Standardization reduces dominance.

### Interpretation of the Model Output

The Isolation Forest model ranks firm-years by anomaly score. The highest-ranked observation (Firm 423 in 2024) has a score of 0.639, indicating it is relatively unusual.

By examining the standardized financial variables, we see strong negative profitability, negative cash flow, very low leverage, and missing R&D data, which together explain why the model flags this firm-year as anomalous.

Standardized values = statistical abnormality

Economic values = business explanation

Variable	Standardized Value	Economic Value	Interpretation
ROA	-2.59	-11.4%	Strongly negative
Leverage	-2.05	1.0%	Very low
Margin	-2.34	-12.5%	Strongly negative
OCF/A	-1.54	-7.7%	Negative
Turnover	-0.03	0.91	Normal
R&D missing	1	NA	Missing

# Lecture 2 – Unsupervised Learning

## Recap: Lecture 1

Anomaly detection identifies observations that deviate from expected patterns in the data. Assumptions:

- Most observations represent normal behaviour
- Anomalies are rare
- Anomalies are generated by a different process

## Methods for detecting anomalies

Method	Type	Example	Strengths	Limitations
Statistical	Univariate	Z-score, IQR	Simple, fast screening, easy to interpret	Requires thresholds, sensitive to distribution assumptions
Visual	Univariate	Histogram	Intuitive, helps understand distribution	No formal decision rule, sensitive to bin width
Machine Learning	Multivariate	Isolation Forest	Detects anomalies across many variables	Harder to interpret

## Supervised vs Unsupervised Learning

Machine learning can learn from data in two ways.

### Supervised Learning

Uses labeled data

Predicts outcome **Y**

Goal: Prediction

### Unsupervised Learning

Uses unlabeled data

Discovers patterns in **X**

Goal: Structure / pattern

Example Supervised: Fraud prediction, Credit risk prediction, Revenue forecasting

Example Unsupervised: Anomaly detection, Dimensionality reduction, Clustering

# Unsupervised Learning: Structure Discovery

Two main approaches:

## 1. Factor Analysis

Focuses on structure in variables.

Goal: Reduce dimensionality and identify underlying latent factors

Example idea: Many variables → fewer underlying factors.

## 2. Cluster Analysis

Focuses on structure in observations (objects).

Goal: Identify groups of similar observations

# Dimensionality

High-dimensional data: A dataset may contain many variables ( $p$ ).

Problems:

### 1. Correlated and noisy variables

- Ideally variables provide independent information
- In reality many variables are redundant

### 2. Computational complexity

- More variables → more computation time

### 3. Visualization difficulty

Solution: **Dimensionality reduction**

Transform high-dimensional data into fewer dimensions while keeping as much information as possible.

*Example:* Large firms tend to have **high values on variables like assets, sales, employees, etc.** Therefore: these variables may represent a **single underlying concept: Firm size.** Instead of 5 variables → possibly **1 dimension.**

# Factor Analysis

Main idea: Given many interrelated variables, find a smaller number of latent variables (factors) that describe them.

Central concept: Factor analysis is based on correlations between variables.

Uses

- Dimensionality reduction
- Discover underlying factors

Example: Many financial variables → underlying factors like profitability or risk.

## Types of Factor Analysis

### 1. Principal Component Analysis (PCA)

Purpose: Dimension reduction

Example: Combine 10 financial ratios → profitability component

### 2. Exploratory Factor Analysis (EFA)

Purpose: Discover latent constructs

Example: Personality traits from survey questions.

### 3. Confirmatory Factor Analysis (CFA)

Purpose: Test theoretical models

Example: Test if survey questions follow a one-factor model.

Also called: Structural Equation Modeling (SEM).

## Principal Component Analysis

Tries to represent data with fewer variables.

Idea: Create new variables that are linear combinations of the original variables

The new variables are called **principal components**

- The components are uncorrelated
- Number of components  $\leq$  number of variables

Usage

- Dimensionality reduction
- Preprocessing step before further analysis

Each principal component is a **linear combination** of variables:

$$Z_s = u_1X_1 + u_2X_2 + u_3X_3 + \dots + u_pX_p$$

Where:

- $(X)$  = original variables
- $(u)$  = weights
- $(Z)$  = principal component

Goal: Each component should explain as much variance as possible.

**Eigenvalues:** Measure the amount of variance explained.

- Larger eigenvalue → more information captured

Important: Variables are usually standardized before PCA.

### Why standardize variables?

PCA gives more weight to variables with higher variance.

If variables use different units: Results become **biased**

Example: Height vs weight measured in different units.

Standardization formula:

$$X_{ij}^s = \frac{X_{ij} - \bar{x}_j}{s_j}$$

Where:

- subtract mean
- divide by standard deviation

Result: mean = 0, variance = 1

**Important:** Correlation  $\neq$  causation.

When interpreting PCA components: Be careful not to overinterpret correlations.

### PCA Strengths

- Reduces dimensionality: Summarizes correlated variables
- Simplifies datasets: Makes visualization easier
- Useful for feature engineering: Components can be used in prediction models

### PCA Limitations

- Components may be hard to interpret: New variables are linear combinations of many variables
- Sensitive to scaling: Requires standardization
- Only captures linear relationships

### Conducting PCA Steps:

1. Formulate the problem
2. Determine number of dimensions
3. Retain variables
4. Interpret dimensions
5. Use components in further analysis

## PCA Example Dataset

Dataset: job satisfaction

Variables (rated 1–4): Supervisor, Salary, Career possibilities, Customers, Colleagues, Work content, Company, Work-life balance

Scale: 1 = very unsatisfied, 4 = very satisfied

### PCA Step 1: Problem Definition

Goal: **Dimensionality reduction**

Question: Can we combine **8 job satisfaction variables** into **fewer components**?

Variables: Work content, Colleagues, Supervisor, Salary, Career possibilities, Customers, Company, Work-life balance. Each rated on a 4-point scale.

Factor analysis relies on **correlations between variables**.

Expectation: Variables that are **highly correlated** likely belong to the **same dimension**. Correlation matrix →

Example correlations:

- Supervisor & colleagues → moderate correlation
- Career possibilities & salary → stronger correlation

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Work content (1)	1						
Colleagues (2)	0.210	1					
Supervisor (3)	0.243	0.389	1				
Salary (4)	0.153	0.127	0.235	1			
Career possibilities (5)	0.350	0.184	0.288	0.417	1		
Customers (6)	0.219	0.165	0.139	0.045	0.125	1	
Company (7)	0.239	0.213	0.317	0.151	0.334	0.242	1
Work-life balance (8)	0.101	0.098	0.172	0.075	0.158	0.114	0.309

Interpretation: Highly correlated variables may form the **same component**.

### PCA Step 2: Determining the Number of Dimensions

When conducting PCA, we must decide how many components (dimensions) to keep. Important note: There is no universally accepted optimal rule. Approaches:

#### 1. A priori determination

- The researcher decides the number of dimensions beforehand
- Based on **theory or prior knowledge**

#### 2. Kaiser criterion

- Keep components with **eigenvalues > 1**
- Reason: each retained component should explain at least as much variance as one variable

#### 3. Scree plot

- Plot **eigenvalues vs number of components**
- Look for the **“elbow” point** where the curve flattens

#### 4. Proportion of variance

- Choose number of components that explain enough **cumulative variance**

## Kaiser Criterion

Key rule: Keep components where: **Eigenvalue > 1**

Why? Because each retained component should explain **at least the variance of one original variable**.

## Example Table

Dimension	Eigenvalue	Variance %	Cumulative Variance %
Dim 1	2.51	31.39	31.39
Dim 2	1.10	13.73	45.12
Dim 3	0.99	12.32	57.44
Dim 4	0.93	11.61	69.05
Dim 5	0.75	9.42	78.47

Interpretation: Only **Dim 1 and Dim 2** have eigenvalues > 1 → keep **2 dimensions**.

## Scree Plot

A scree plot is a graphical way to determine the number of components.

- **X-axis:** number of dimensions
- **Y-axis:** eigenvalues

Eigenvalues are ordered from largest to smallest.

**Interpretation:** Look for the **“elbow” point** where the curve suddenly flattens.

Before elbow → important dimensions

After elbow → less useful dimensions

**Limitation:** The elbow is sometimes **hard to identify clearly**.



## Proportion of Variance

Another method is looking at **cumulative explained variance**.

Idea: Keep enough dimensions to explain a large portion of total variance.

Typical threshold in **social sciences**: **≈ 60% cumulative variance**

## Example

From the table:

- Dim 1 → 31.39%
- Dim 1 + 2 → 45.12%
- Dim 1 + 2 + 3 → 57.44%
- Dim 1 + 2 + 3 + 4 → 69.05%

Thus **4 dimensions** would pass the 60% threshold.

	Eigenvalue	Variance %	Cumulative Variance %
Dim. 1	2.51	31.39	31.39
Dim. 2	1.10	13.73	45.12
Dim. 3	0.99	12.32	57.44
Dim. 4	0.93	11.61	69.05
Dim. 5	0.75	9.42	78.47
Dim. 6	0.65	8.13	86.60
Dim. 7	0.58	7.23	93.82
Dim. 8	0.49	6.18	100.00

### PCA Step 3: Retain the variables

Next step: determine **which variables contribute most to each component.**

Contribution = percentage of variance of a dimension explained by a variable.

Interpretation: Some variables contribute strongly to **one specific dimension.**

	Dim. 1	Dim. 2
Work content	13.23	0.05
Colleagues	11.43	1.82
Supervisor	17.19	0.00
Salary	9.58	39.38
Career possibilities	18.39	15.48
Customers	6.63	22.29
Company	16.93	6.98
Work-Life Balance	6.61	14.00

### Contribution Bar Plot (Dimension 1)

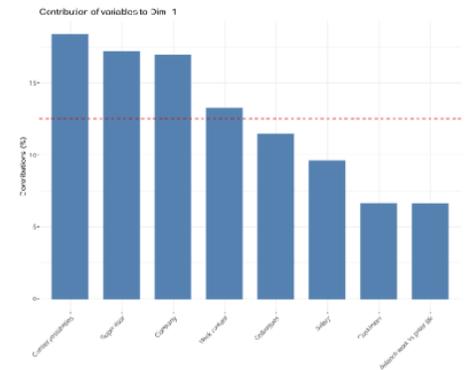
- Each bar = contribution of a variable
- Dashed red line = **expected average contribution**

Expected value:  $1/p$

Where:  $p$  = number of variables

Here:  $1/8 = 12.5\%$

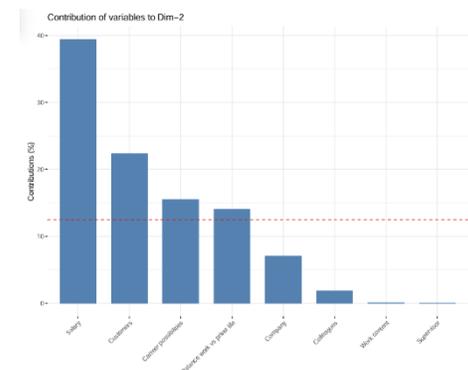
**Interpretation:** Variables above the dashed line contribute **more than average** to the component.



### Contribution Bar Plot (Dimension 2)

**Interpretation:** Variables with the highest contributions define the **meaning of that dimension.**

Example: Dimension 2 may relate more to: Salary, Career possibilities, Customers, Work-life balance



### PCA Step 4: Interpret the dimensions

Interpretation is based on **variable loadings**. Distance from origin indicates **strength of relationship**.

1) Identify variables with high loadings

- **Variables at the end of an axis** → strongly related to that dimension
- **Variables near the origin** → not strongly related to any dimension
- **Variables between axes** → related to multiple dimensions

2) Look for common theme

Moderate positive correlation between:

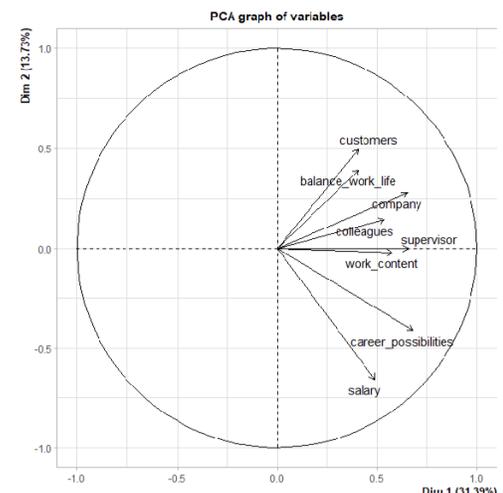
→ customers, work-life balance and dimension 2.

High positive correlations between:

→ supervisor, career possibilities, company and dimension 1.

High to moderate negative correlation between:

→ salary, career possibilities and dimension 2



## PCA Step 5: Use in subsequent analysis

After PCA, we compute **component scores**.

Formula:  $Z_s = u_1X_1 + u_2X_2 + \dots + u_pX_p$

Each observation gets values for Dimension 1 and Dimension 2

Respondent	Dim 1	Dim 2
1	3.99	2.19
2	-0.32	-1.18
3	2.52	-1.79

Standardized realizations of X for respondent 1

Respondent	Work content	Colleagues	Supervisor	Salary	Career possibilities	Customers	Company	Work-life balance
1	1.57	1.44	1.52	0.47	0.40	2.61	1.99	1.93

Variable weights for dimension 1

Dimension	Work content	Colleagues	Supervisor	Salary	Career possibilities	Customers	Company	Work-life balance
1	0.36	0.34	0.41	0.31	0.43	0.26	0.41	0.26

$$Z_1 = 0.36 \times 1.57 + 0.34 \times 1.44 + 0.41 \times 1.52 + 0.31 \times 0.47 + 0.43 \times 0.40 + 0.26 \times 2.61 + 0.41 \times 1.99 + 0.26 \times 1.93 = 3.99$$

## PCA in Supervised Learning

- PCA can be used for **feature engineering**.  
Idea: Replace several correlated variables with one principal component.
- Example: Bankruptcy prediction  
Variables: Profit margin, EBIT / Assets, Net income / Assets  
These are highly correlated.
- Instead create: **Profitability component**  
Prediction model: Bankruptcy = f(leverage, liquidity, profitability component)

## Benefits

- Reduces multicollinearity
- Simplifies models

# Cluster Analysis

**Main idea:** Group observations into clusters so that:

- objects in the **same cluster are similar**
- objects in **different clusters are different**

Based on similarity between observations

## Uses

- Market segmentation
- Fraud detection
- Data reduction

For example: Playing Cards What are the natural groupings?

Possible clusters: By suit (hearts, clubs, etc.), By number

→ Clusters depend on the features used.

## Data Exploration

- Cluster analysis is **exploratory**.
- Quote from Hastie et al. (2009): There is no direct measure of success in unsupervised learning.
- Therefore: Results must be interpreted carefully, Researchers rely on heuristics and judgment

## Types of Clustering Methods

### 1. Hierarchical Clustering

Creates a **tree structure of clusters**.

Two approaches:

- **Agglomerative**
  - Start with individual observations
  - Merge clusters
- **Divisive**
  - Start with one cluster
  - Split into smaller clusters, each observation is one cluster

### 2. Non-hierarchical Clustering

Example: **k-Means clustering**

- Number of clusters **k chosen beforehand**
- Observations assigned to nearest cluster center

## k-Means Clustering

k-Means is a **partitioning method**. It divides objects into **k clusters**.

Clusters are: distinct and non-overlapping

**Why popular?** → Simple, Efficient, Works well with large datasets

**Step 1:** Generate **k random cluster centers**. These are initial means.

**Step 2:** Assign each observation to the **nearest cluster center**.

Distance is usually measured using **Euclidean distance**.

Thus each point belongs to the **closest centroid**.

**Step 3:** Update cluster centers.

Each cluster center becomes: the mean of the observations assigned to that cluster.

Thus cluster centers move toward the **center of their assigned points**.

**Step 4** Repeat steps:

1. Assign observations to nearest center
2. Update cluster means

Continue until **convergence**.

Convergence occurs when: cluster assignments **no longer change**

## k-Means: Strengths

### 1. Identifies groups of similar observations

- Useful for discovering patterns in data.

### 2. Good for segmentation

- Frequently used in marketing, finance, and customer analysis.

### 3. Simple and computationally efficient

- Easy to implement.
- Works well on **large datasets**.

### 4. Reveals heterogeneity

- Different clusters may represent **different behavioral patterns**.

## k-Means: Limitations

### 1. Number of clusters must be chosen beforehand

- Selecting the correct **k** can be difficult.

### 2. Sensitive to initialization

- Different starting cluster centers can produce **different results**.

### 3. Sensitive to outliers

- Extreme values can **distort cluster centers**.

## Conducting k-Means Clustering

The clustering process follows **five main steps**:

1. Formulate the problem
2. Select a distance measure
3. Decide on the number of clusters
4. Interpret the clusters
5. Use clusters in subsequent analysis

## Example Dataset: Loan Applications

Dataset: **500 rejected loan applications** from a US online bank.

After removing non-numeric variables (date, loan title, ZIP code, state), four numeric variables remain:

Variable	Scale	Score	Category
Amount requested	\$1,000 – \$35,000	0–299	Very Poor
Credit risk score	0 – 850	300–579	Poor
Debt-to-income ratio	0 – 3.16	580–669	Fair
Employment length	0 – 10 years	670–739	Good
		740–799	Very Good
		800–850	Excellent

### Step 1: Formulate the Problem

Problem definition: Goal is segmentation

Research question: Are there distinct clusters of rejected loan applicants based on their characteristics?

Important: Each variable uses

**different measurement scales.**

Loan application	Amount requested	Credit risk score	Debt-to-income ratio	Employment length
1	25,000	690	0.49	0
2	3,000	0	0	0
3	5,000	654	0.09	0

### Step 2: Select a Distance Measure

Clustering is based on similarity between observations.

**Key idea:** Objects with smaller distances between them are more similar.

**Euclidean Distance** For two observations (x) and (y):  
This measures the straight-line distance between two points in multidimensional space.

$$d_{euc}(x, y) = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}$$

**Important issue:** The scale of variables matters.

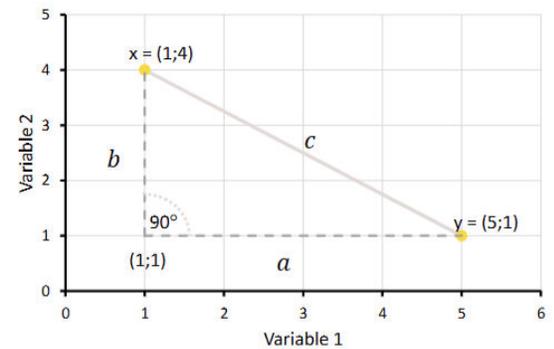
If one variable has much larger values: It dominates the distance measure.

**Solution:** Standardize variables before clustering.

**Euclidean Distance Illustration** using the Pythagorean theorem.

Two points:  $x = (1,4)$   $y = (5,1)$

Distance is calculated as the straight-line distance between them.



### Step 3: Decide the Number of Clusters

In k-means, k must be chosen before running the algorithm. However, selecting k is difficult. Methods to determine k:

#### 1. Elbow Method

Look for the point where the improvement slows.

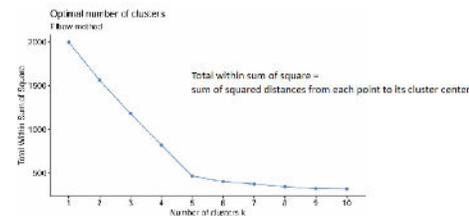
#### 2. Average Silhouette Method

Choose the k that maximizes the silhouette score.

Both methods rely on **visual interpretation**.

#### Elbow Method

- The elbow method uses a plot of: **Within-cluster variance vs number of clusters**
- Within-cluster variance measures how tightly observations are grouped.
- Goal: Minimize **within-cluster variation**.
- Formula idea: Total within sum of squares = sum of squared distances between points and their cluster center.
- **Interpretation:** Choose k where the curve forms an **elbow**.
- After this point: Adding clusters provides little improvement.



#### Silhouette Method

Another approach: **Average silhouette score**.

**Silhouette value:** Measures how well an observation fits its cluster.

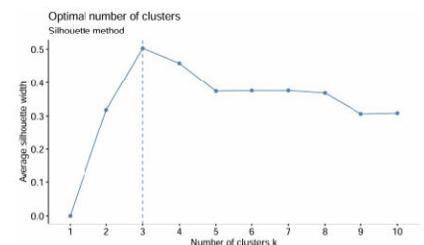
High silhouette score means:

- observation is **similar to its own cluster**
- observation is **different from other clusters**

Range: -1 to +1

Interpretation: Higher average silhouette → **better clustering**.

Choose k with **maximum silhouette score**.



#### NbClust Package

In practice, software can help determine k. Example: **NbClust package (R)**

It evaluates clustering using **30 different indices**.

The package: tests multiple clustering schemes, compares different numbers of clusters and suggests the **best clustering solution**

### Step 4: Interpret Clusters

Once clustering is completed, analyze cluster statistics.

Interpretation: Cluster 2 is the largest group of rejected applicants. →

Cluster size indicates importance of clusters.

#### Cluster Size

1	9
2	275
3	51
4	133
5	32

### Cluster Visualization

Clusters can be visualized in 2D space.

→ `fviz_cluster()`: This function first applies PCA to reduce dimensions. Then clusters are plotted in two dimensions. This allows us to visualize clusters clearly.

### Interpreting Cluster Centroids

Cluster centroids represent average values of variables within each cluster.

Example centroid table:

Cluster	Amount Requested	Credit Score	Debt-to-Income	Employment Length
1	9,888	480	3.03	1.22
2	6,811	615	0.18	0.38
3	14,191	665	0.26	8.37
4	26,376	665	0.28	0.32
5	6,528	0	0.01	0.78

Interpretation: Clusters describe different borrower types.

Example: Cluster 3: higher employment length, moderate credit score

### Step 5: Use in Subsequent Analysis

- Clusters can be used in further analysis.
- Clustering can be used for feature engineering.
- Clusters can also become new variables.

Example: Borrower types. Clusters may represent:

→ Young borrowers: low income, short credit history, high income growth

→ Established borrowers; stable jobs, long credit history

Prediction model example: Default risk =  $f(\text{income, debt ratio, cluster})$

### Benefits

- Captures **heterogeneity**
- Different groups may follow **different behavioral patterns**

Example: Income predicts default for young borrowers,  
Debt ratio predicts default for established borrowers.

