Discovering, Visualizing and Targeting Customer Segments

Dr. Anestis Fachantidis

Intelligent Systems

Intelligent Systems Group
Computer Science Department
Aristotle University

The Ever Increasing Personalization

Personalization is the top priority of the Marketers today.

According to Forrester:

- *In email only 7% of consumers agree with the statement "Email offers are usually well timed with my needs" (down from 9% in 2014). Meanwhile, marketers ranked "improving the customer experience" first in their list of priorities for the year.* [1]
- *Only 30% of customers today feel that they are getting their desired level of personalization*

Machine Learning (ML) is the key enabling technology for personalization

- *Discovering and exploiting the connection of attributes (customer characteristics) to specific output variables (customer behavior) could (almost) work as a ML definition*

[1] "The Contextual Marketing Imperative: The Evolution Of Personalization From Push Messaging To One-To-One Personal Customer Experiences," Forrester 2015.

The Customer Success Journey

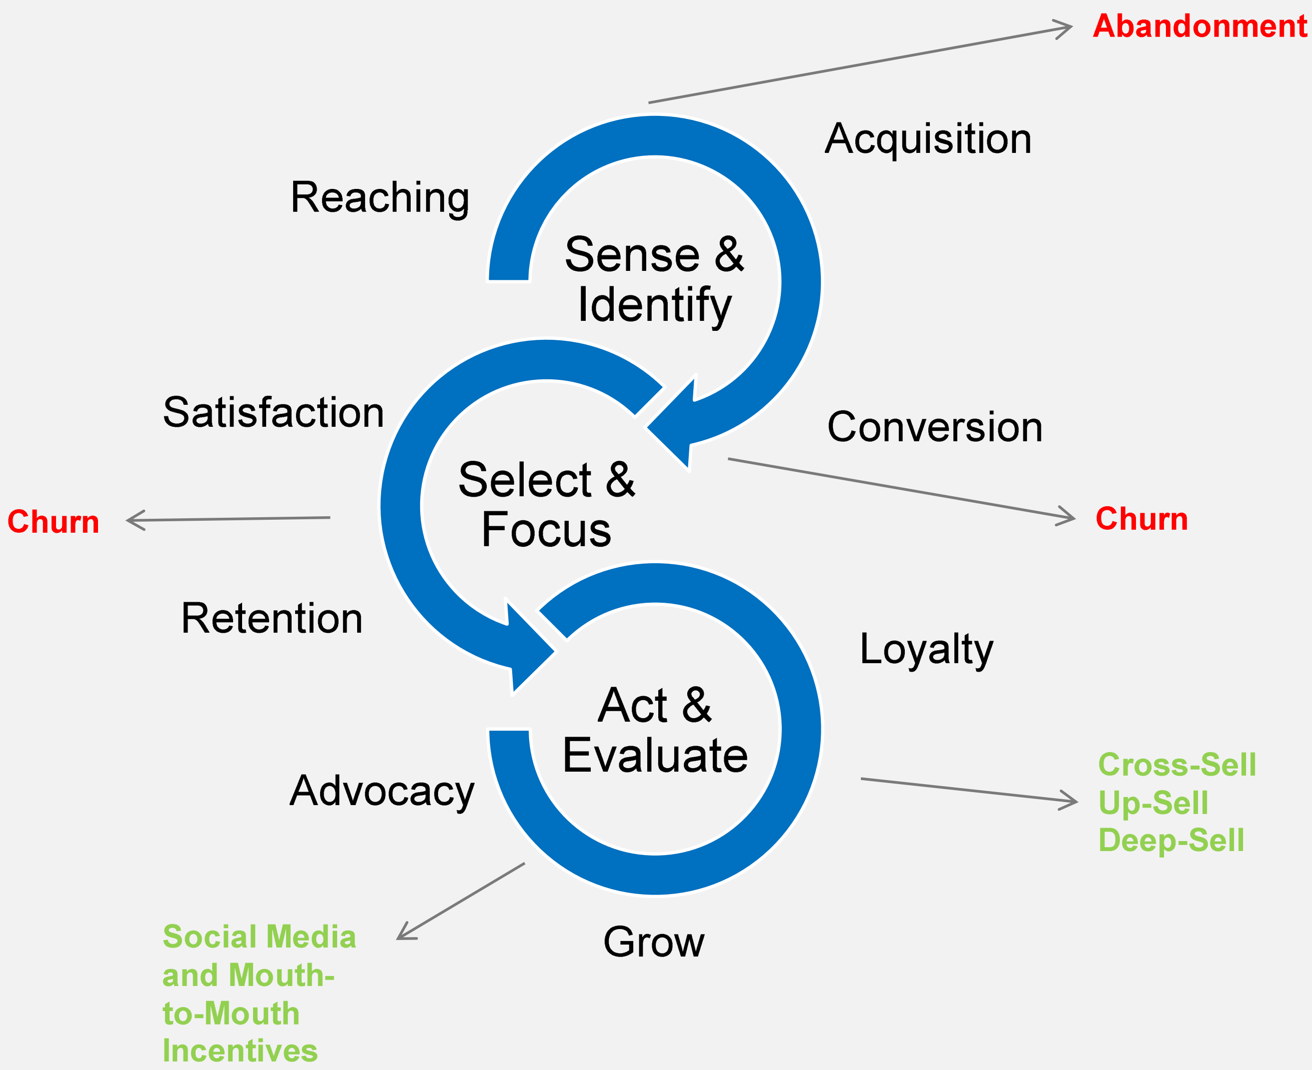

Customer Segmentation

Customer Segmentation

Dividing a company's customer base into smaller, distinct and internally homogeneous groups of customers based on a set of customer characteristics

Example of customer characteristics – the classic RFM model:

- Recency
- Frequency
- Monetary

A finer customer model (more customer characteristics **or** more detailed ones) will enable the formation of *micro-segments*

Marketers create targeting strategies tailor-made for each customer segment

The Business Perspective

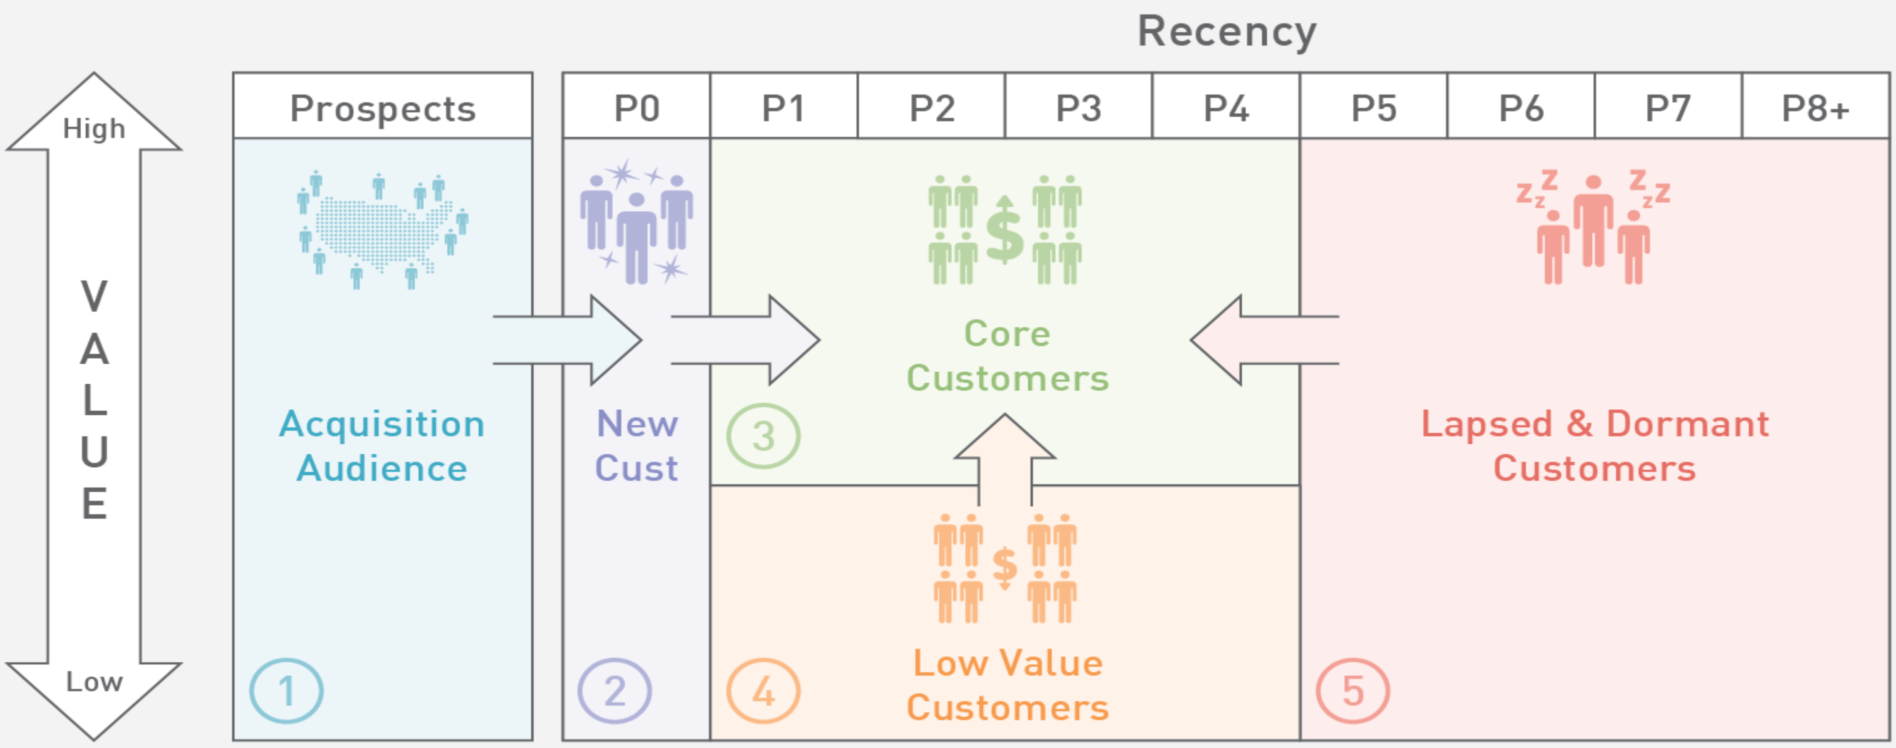

source: wiland.com

The Data Perspective

Categories of any given Merchandise Node

Extracted RFM values for each customer

Customer ID	Category 1	Category 2	Category 3	Category 4	...	Recency	Frequency	Monetary
xxxxxxx1	84.00	13.00	2.00	0.00	...	40	1.34	5.34
xxxxxxx2	40.00	39.00	18.00	2.00	...	13	0.78	2.1
xxxxxxx3	28.00	64.00	3.00	0.00	...	5	0.5	0.9

Values can be absolute monetary dispersion or percentage of

Linking Business & Data Perspectives

The Machine Learning Advantage:

- Non–automated segmentation requires extensive and costly experimentation
- **Data-driven** automated analysis can be far more reliable and efficient
- Clustering algorithms (from unsupervised machine learning) can **discover** the natural-underlying groupings of customers
- Over 100 clustering **algorithms** in literature. Three well known families are:
  - Centroid-based (e.g. K-Means)
  - Hierarchical (named after their linkage criterion)
  - Density based (e.g. DBSCAN)

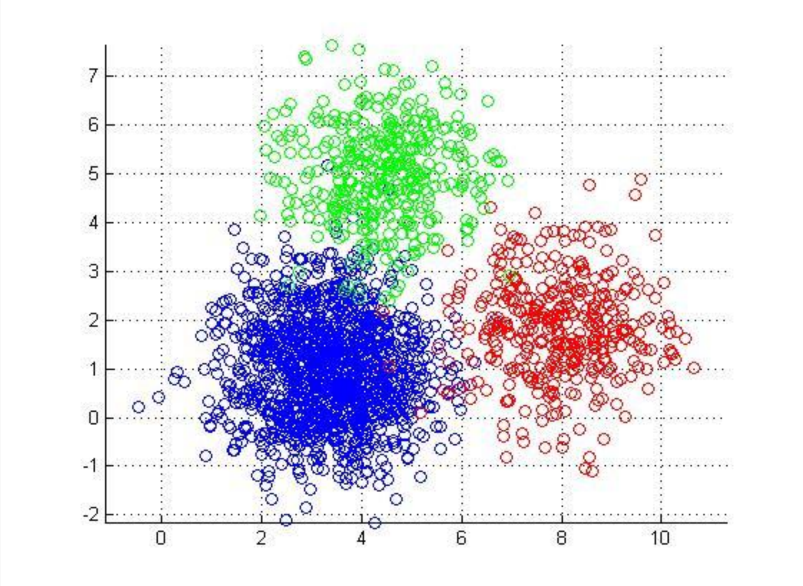

The Machine Learning Perspective

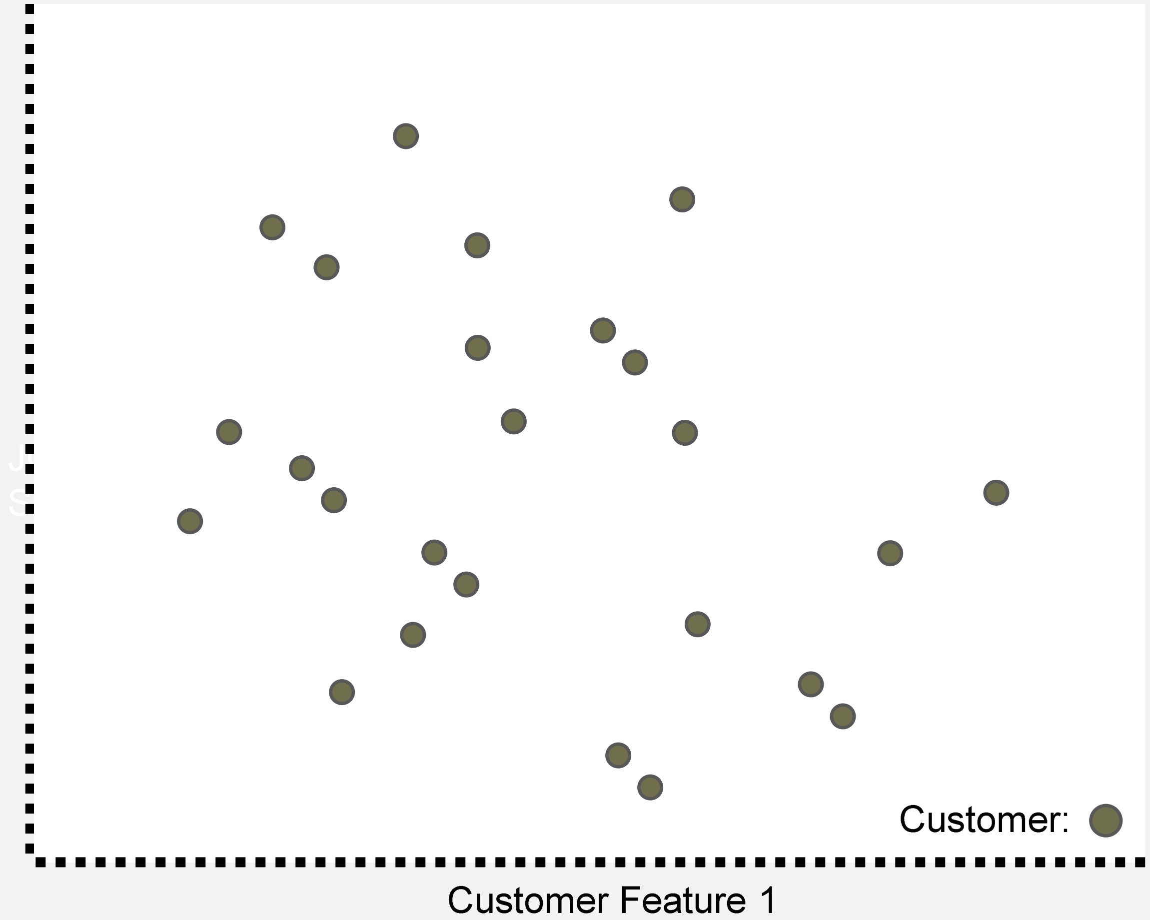

Key Considerations:
1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

The Machine Learning Perspective

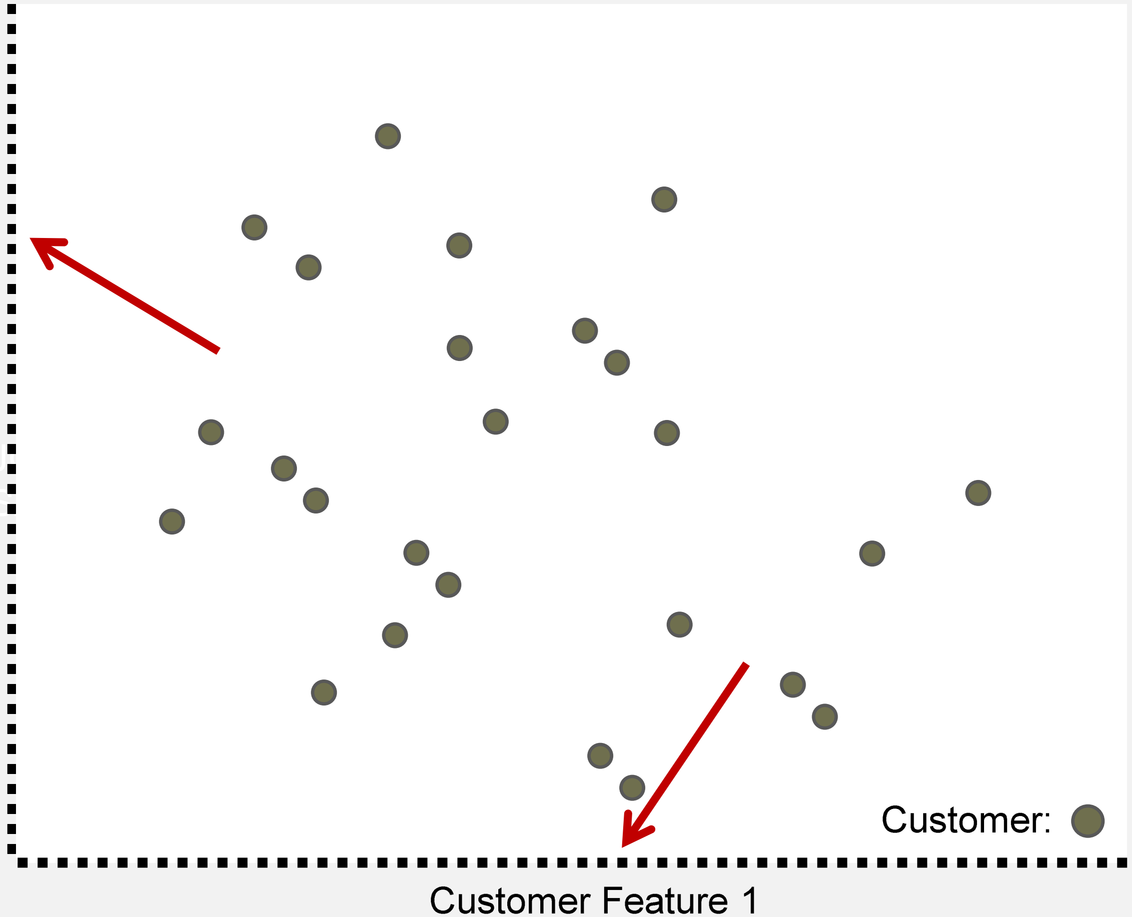

Key Considerations:

1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

- What do these features represent? (customer value?)
- How will they be produced? (aggregations, summaries, etc.)
- What types of features? (qualitative/quantitative etc)

The Machine Learning Perspective

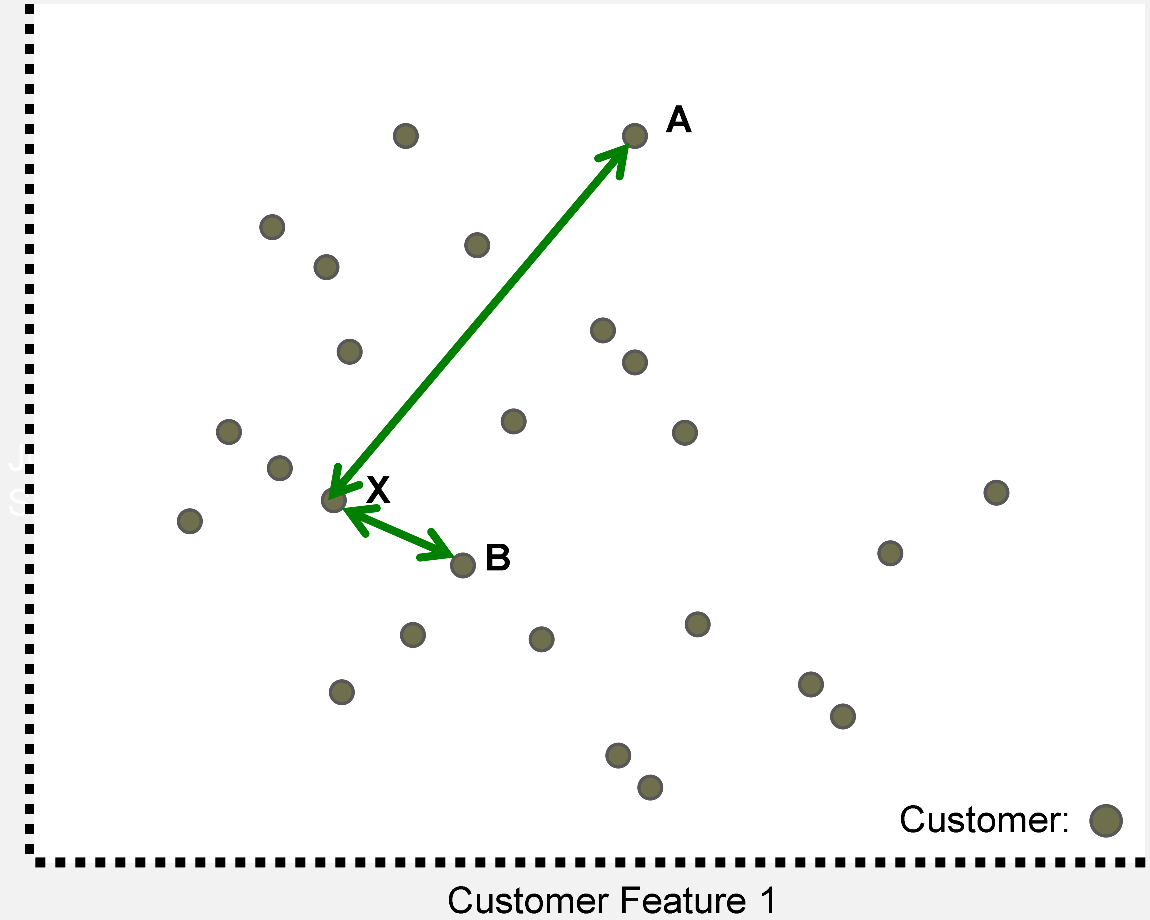

Key Considerations:

1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

Different distance measures will express different types of **similarity between customers**. Example:

- Euclidean distance, usually for real-valued data (customer value), **X is closer to B**
- Cosine distance, suitable for count data (customer basket), **X is closer to A**

The Machine Learning Perspective

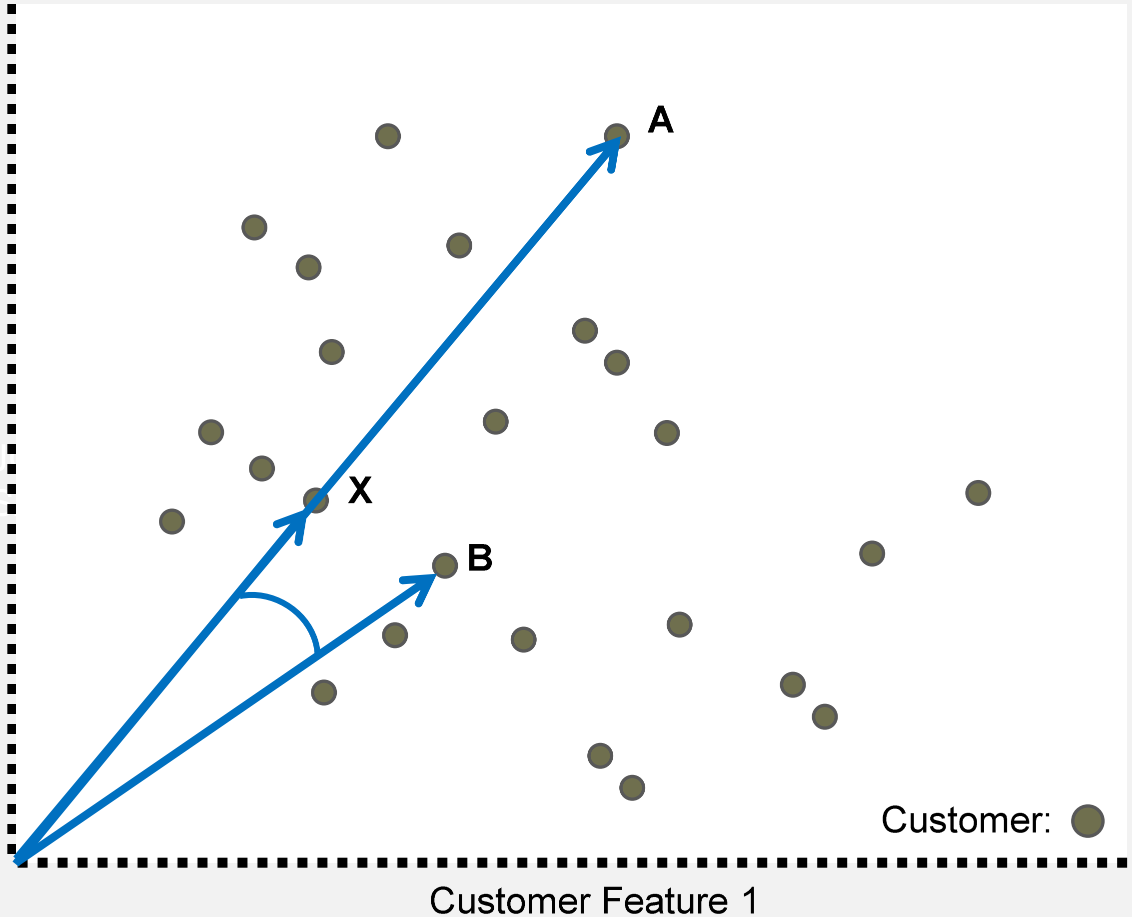

Key Considerations:

1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

Different distance measures will express different types of **similarity between customers**. Example:

- Euclidean distance, usually for real-valued data (customer value), **X is closer to B**
- Cosine distance, suitable for count data (customer basket), **X is closer to A**

The Machine Learning Perspective

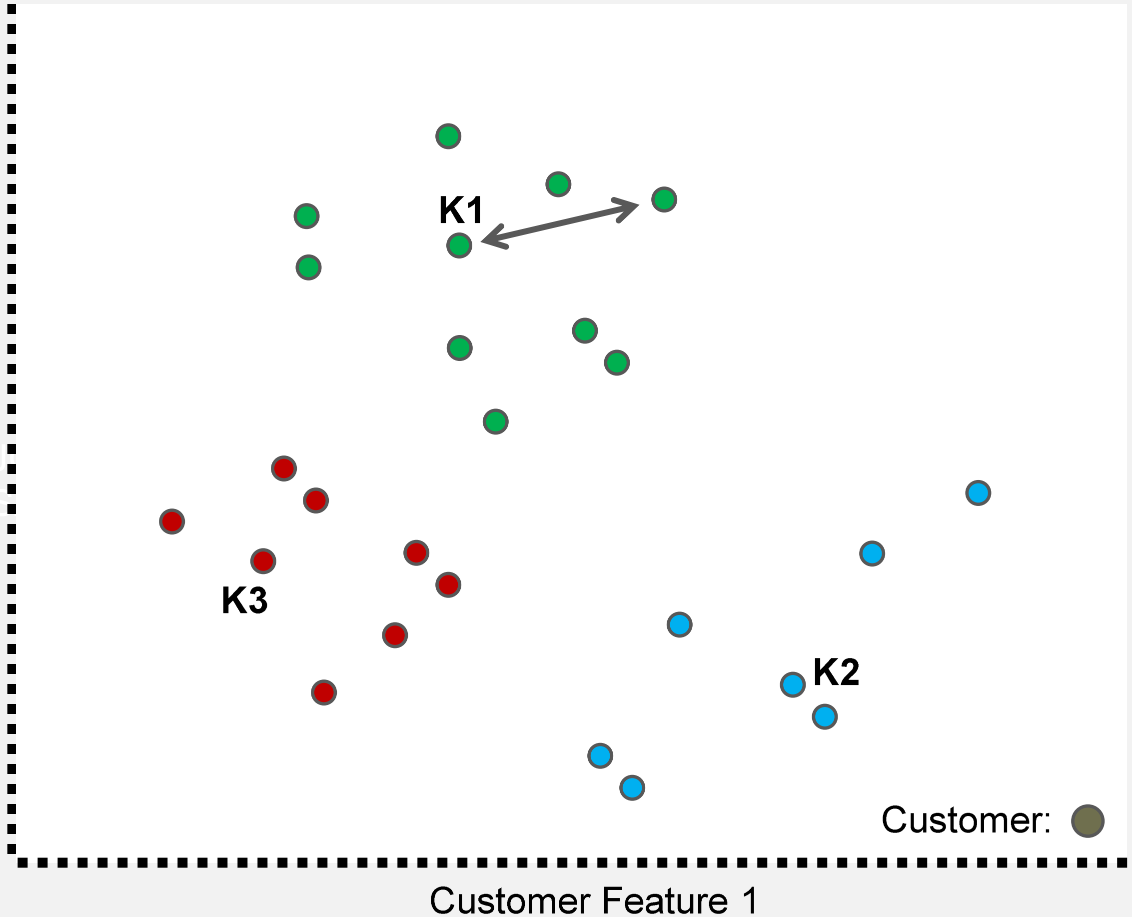

Key Considerations:

1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

Centroid-Based Clustering
(e.g., K-Means)

- Set number of cluster representatives (# of clusters)
- An optimization process of continually setting new cluster centers (customer archetypes) and forming clusters with the data points (customers) closer to them

The Machine Learning Perspective

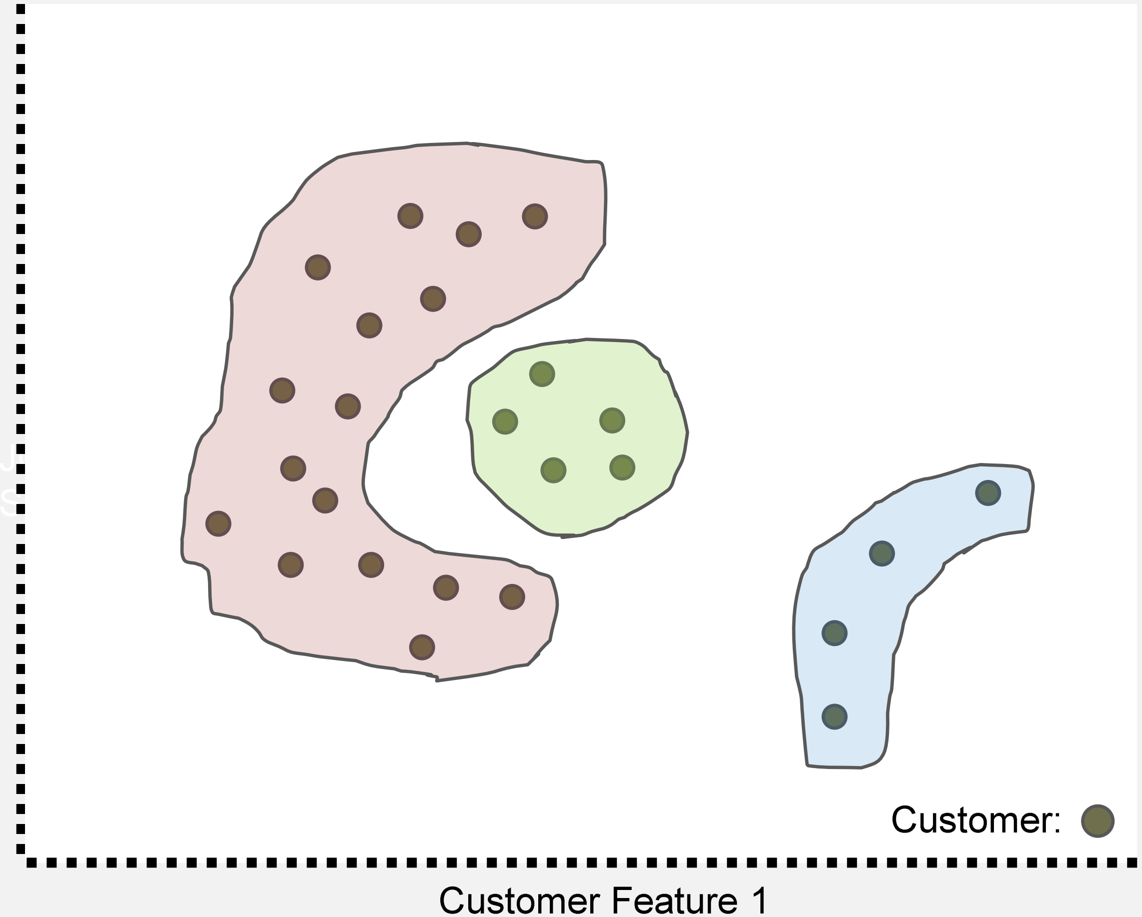

Key Considerations:

1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

Density-Based Clustering (e.g., DBSCAN)

Clusters are defined as areas of **higher density** than the remainder of the data set.

The Machine Learning Perspective

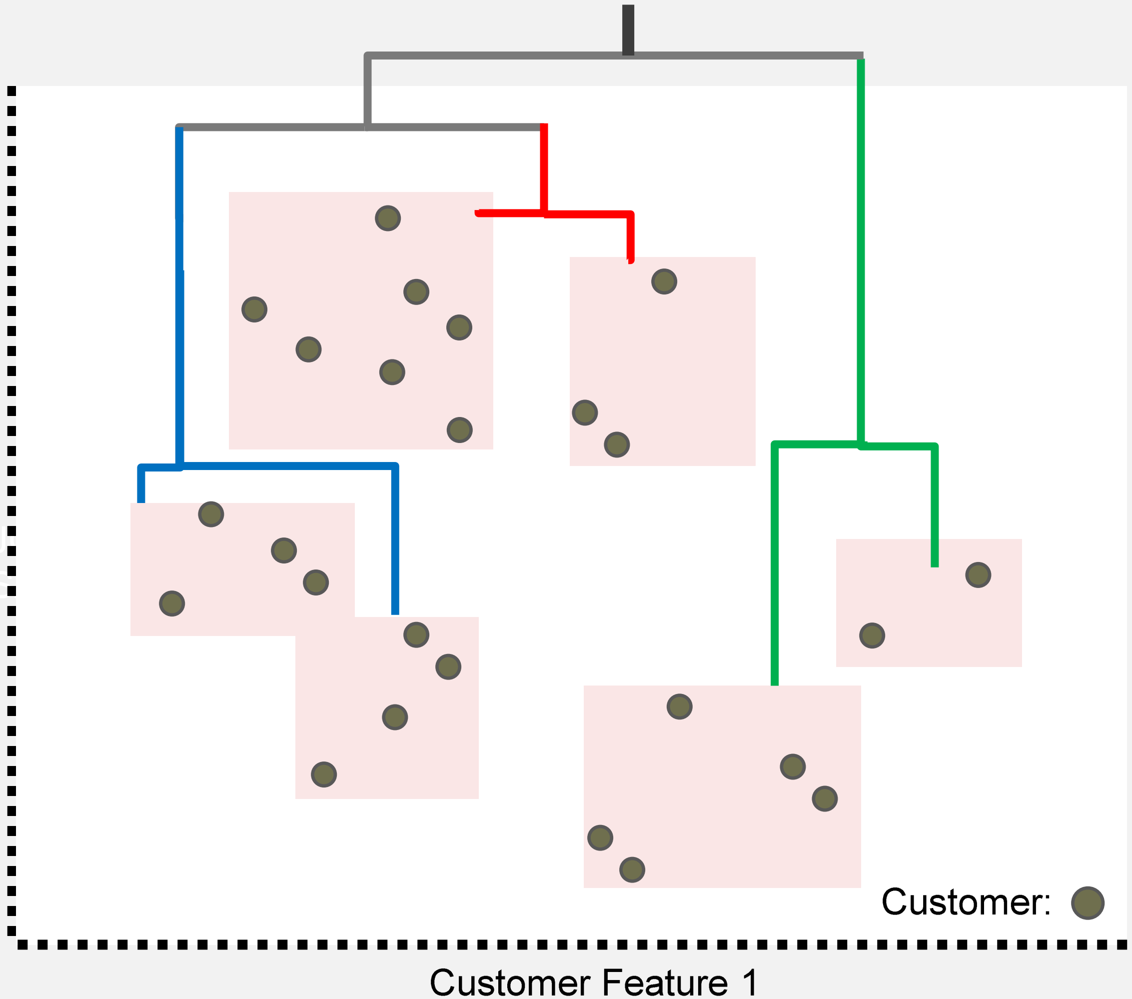

Key Considerations:
1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

Hierarchical Clustering
Given a distance level, different but nested cluster structures can be formed, resulting to a cluster hierarchy
- Computationally expensive
- Significant flexibility for business use / natural micro-segments

The Machine Learning Perspective

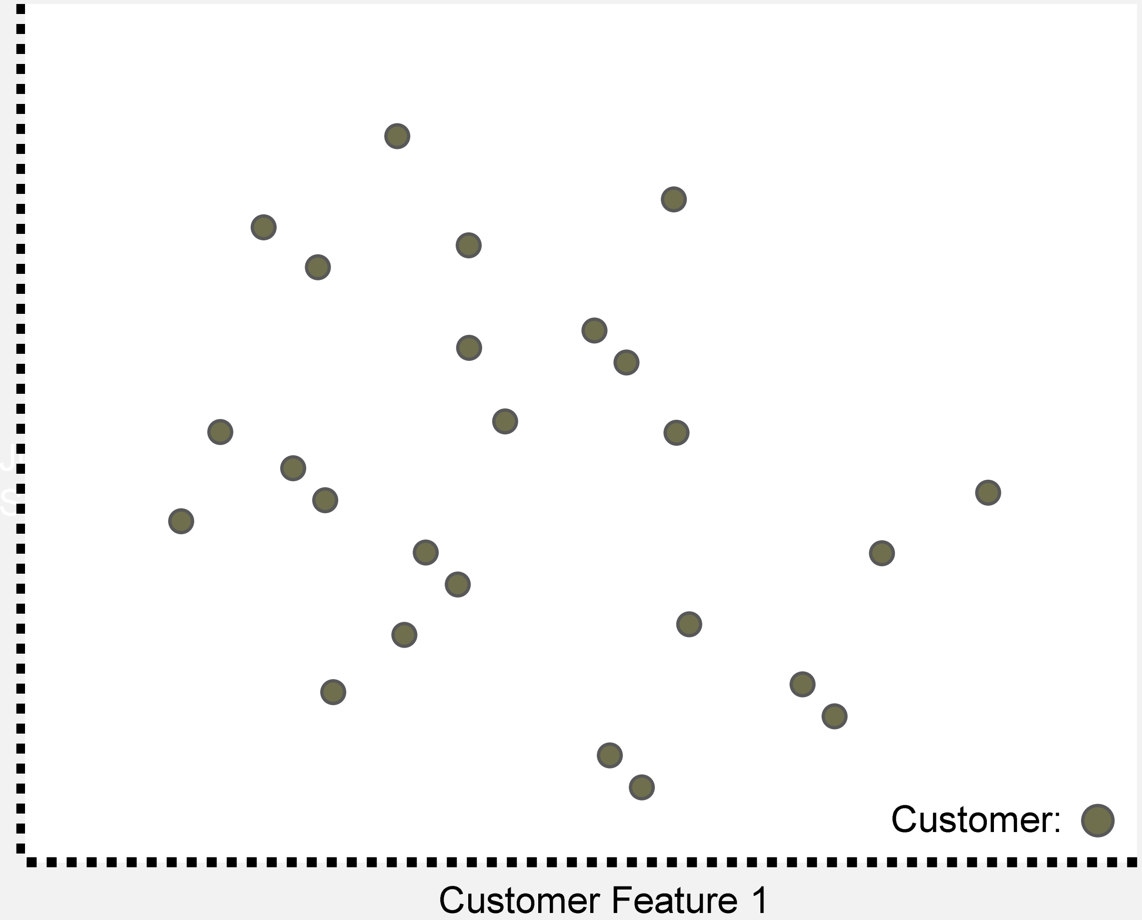

Key Considerations:

1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

- **Internal evaluation:** measures such as the silhouette and gap scores
- **External - business evaluation**: does a customer segment responds homogeneously to the same marketing actions?

The Machine Learning Perspective

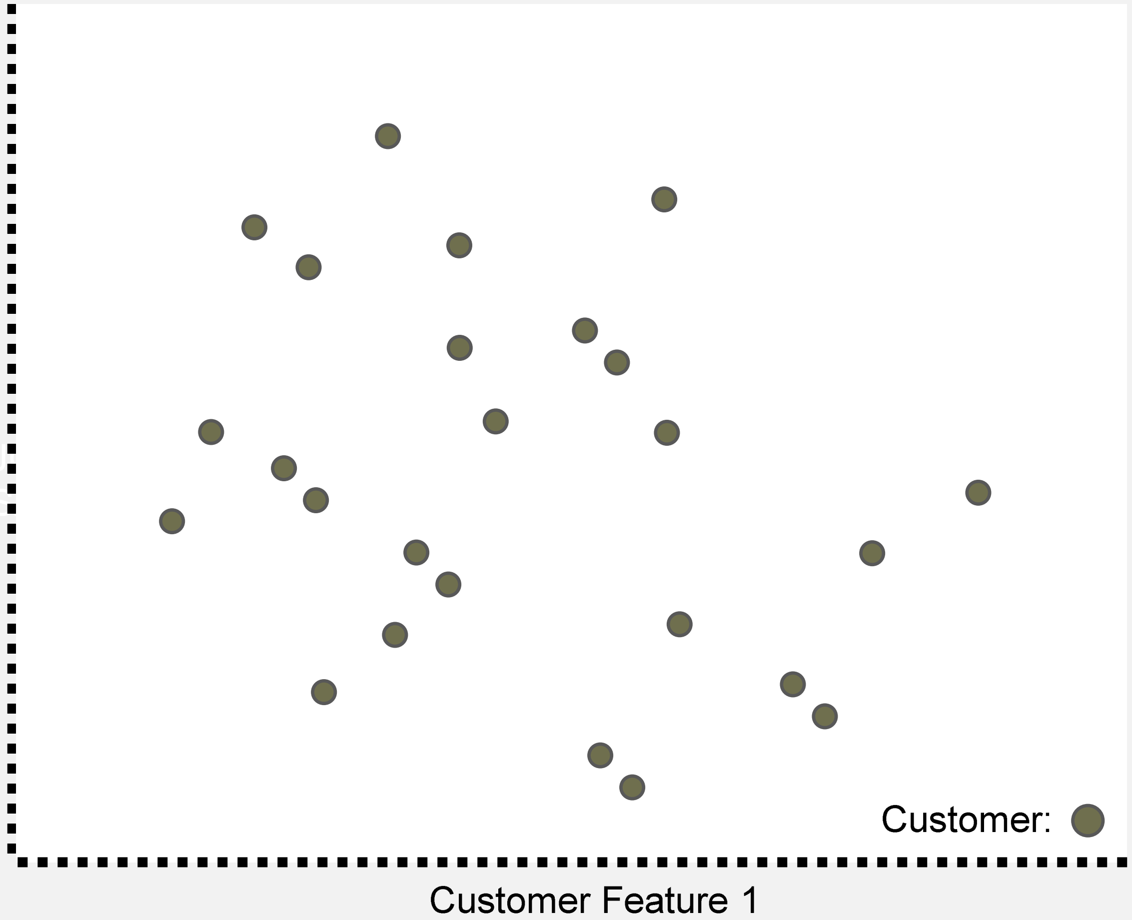

Key Considerations:
1. Knowledge Representation & Feature Engineering
2. Defining similarity
3. Algorithm Selection
4. Evaluation strategy

Extra considerations:
1. Need for robust models
2. Big Data → high computational complexity

Personalized Campaigns

The Minimalistic Campaign Model

A campaign *unit* consists of:

Audience

- Segment (or micro-segment)
- Manually defined (threshold-based query) or automatically discovered (clustering) or both

Content

- Auto-generated or manually curated – includes an offer or not

Context

- Location or time-zone (e.g., at home, specific time-zone or event trigger)

Channel

- e.g., Mobile app, or user account or SMS

Evaluation criteria & time period

- KPIs such as R/F/M improvements at a time period (e.g. 1 week)
- Transitions from “poor” to “rich” micro-segments

A Working Example from the Betting Industry

Let's create a campaign !

This campaign will focus on the following:

Content
- A 10% discount offer for a betting amount above 50 €

Audience
- Age : any
- Game : Football or Basketball
- Medium-High Value Customers (Healthy)
- No churn customers in this example (assumption)

Context
- Customers connected on the mobile app

Channel
- Push notification in mobile app account

Duration of the campaign
- 1 Week

There will be ***two variations*** *of this campaign altering one of these characteristics and forming* ***Campaign Unit A*** *and* ***Campaign Unit B***

The multi-dimensional Customer Personas of our Campaign Example

We could possibly display our customers as previously: Arbitrary colored data points over the information-poor 2D plane, **instead…**

The multi-dimensional Customer Personas of our Campaign Example

We could possibly display our customers as previously: Arbitrary colored data points over the information-poor 2D plane, **instead…**

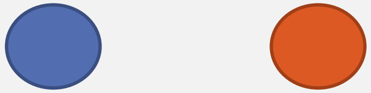

The multi-dimensional Customer Personas of our Campaign Example

We could possibly display our customers as previously: Arbitrary colored data points over the information-poor 2D plane, **instead…**

We augment each point with the game they mostly prefer

The multi-dimensional Customer Personas of our Campaign Example

We could possibly display our customers as previously: Arbitrary colored data points over the information-poor 2D plane, **instead…**

We augment each point with the game they mostly prefer

And then with their value:

- Red hat → Likely to churn
- Green hat → Healthy

The multi-dimensional Customer Personas of our Campaign Example

We could possibly display our customers as previously: Arbitrary colored data points over the information-poor 2D plane, **instead…**

We augment each point with the game they mostly prefer

And then with their value:

- Red hat → Likely to churn
- Green hat → Healthy

And then with their age:

- Black moustache → 30-45
- White beard → 45-60+

The multi-dimensional Customer Personas of our Campaign Example

We could possibly display our customers as previously: Arbitrary colored data points over the information-poor 2D plane, **instead…**

We augment each point with the game they mostly prefer

And then with their value:

- Red hat → Likely to churn
- Green hat → Healthy

And then with their age:

- Black moustache → 30-45
- White beard → 45-60+

Our campaign example focuses on these personas

A/B Testing

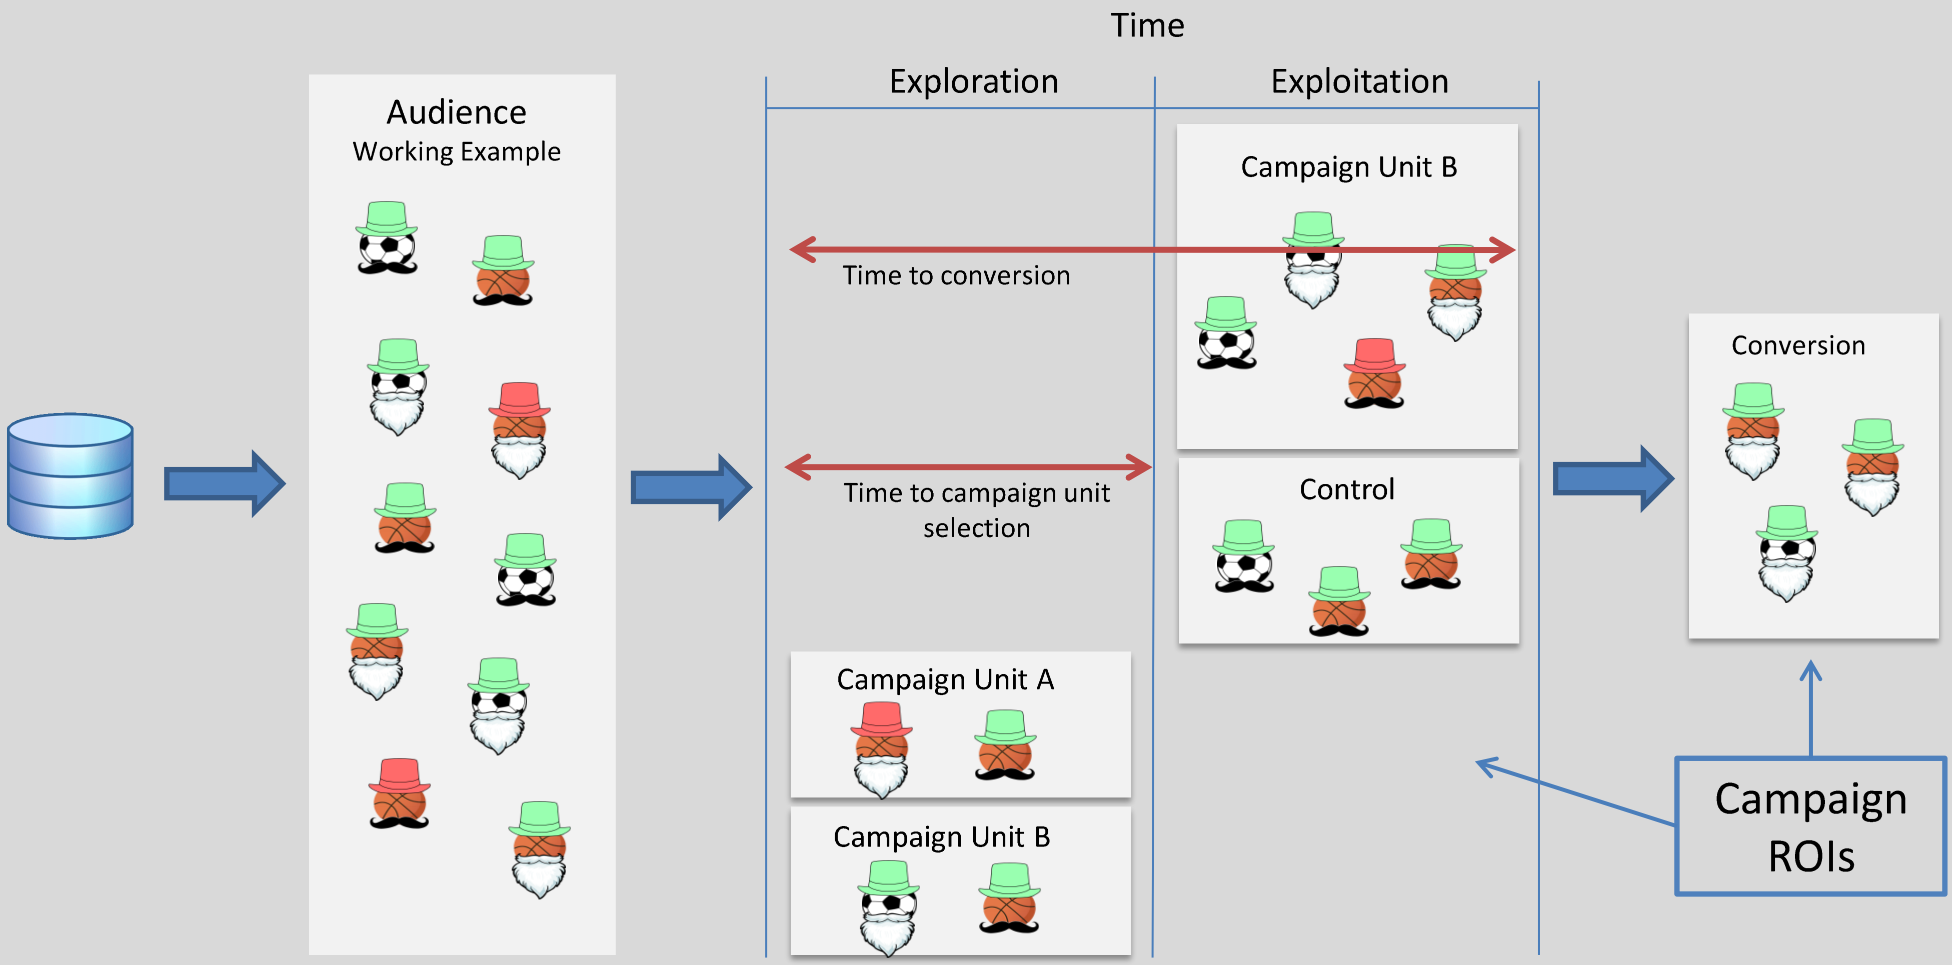

Time
Exploration
Exploitation
Audience
Working Example
Time to conversion
Time to campaign unit selection
Campaign Unit B
Control
Campaign Unit A
Campaign Unit B
Conversion
Campaign ROIs

Reinforcement Learning: The Basics

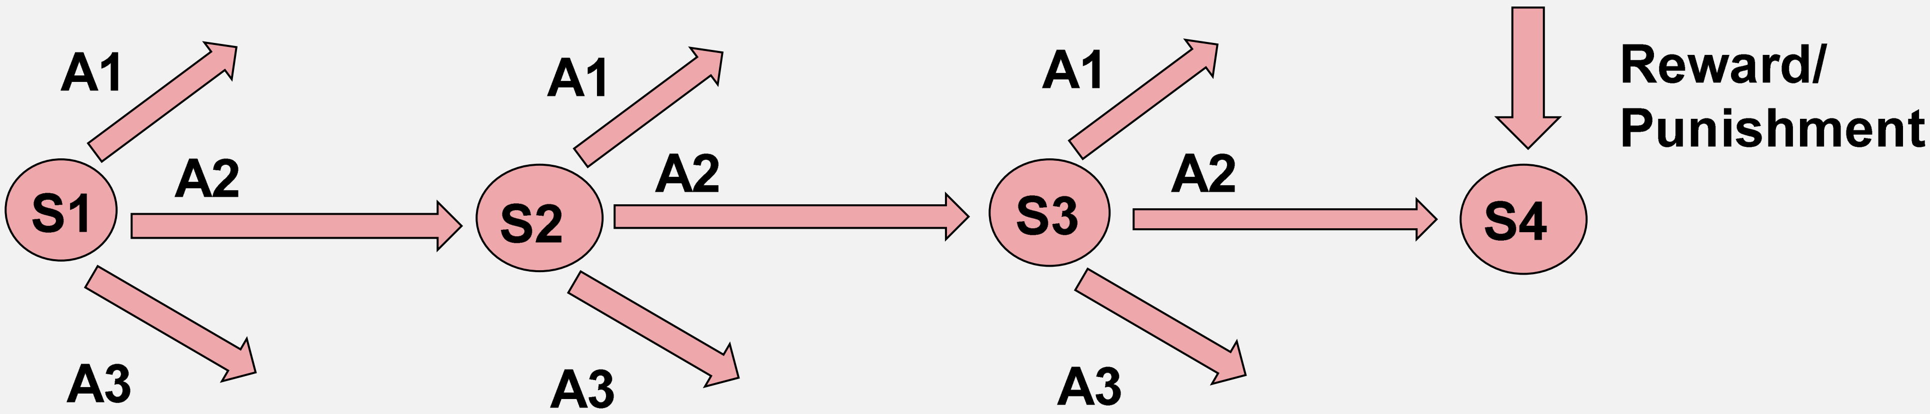

- Problems with **limited feedback** on the sequence of actions (i.e., customer interactions) we should follow to achieve a goal (no dataset of examples)
- The only feedback is a scalar evaluation of our actions (**A**) called **reward** (could be some customer behavior KPI)
- The current status of the customer is described from feature values and forms the state (***S***)
- Reinforcement learning algorithms **learn** an **action policy** that **maximizes** the long-term, **cumulative** reward.

The Reinforcement Learning Approach

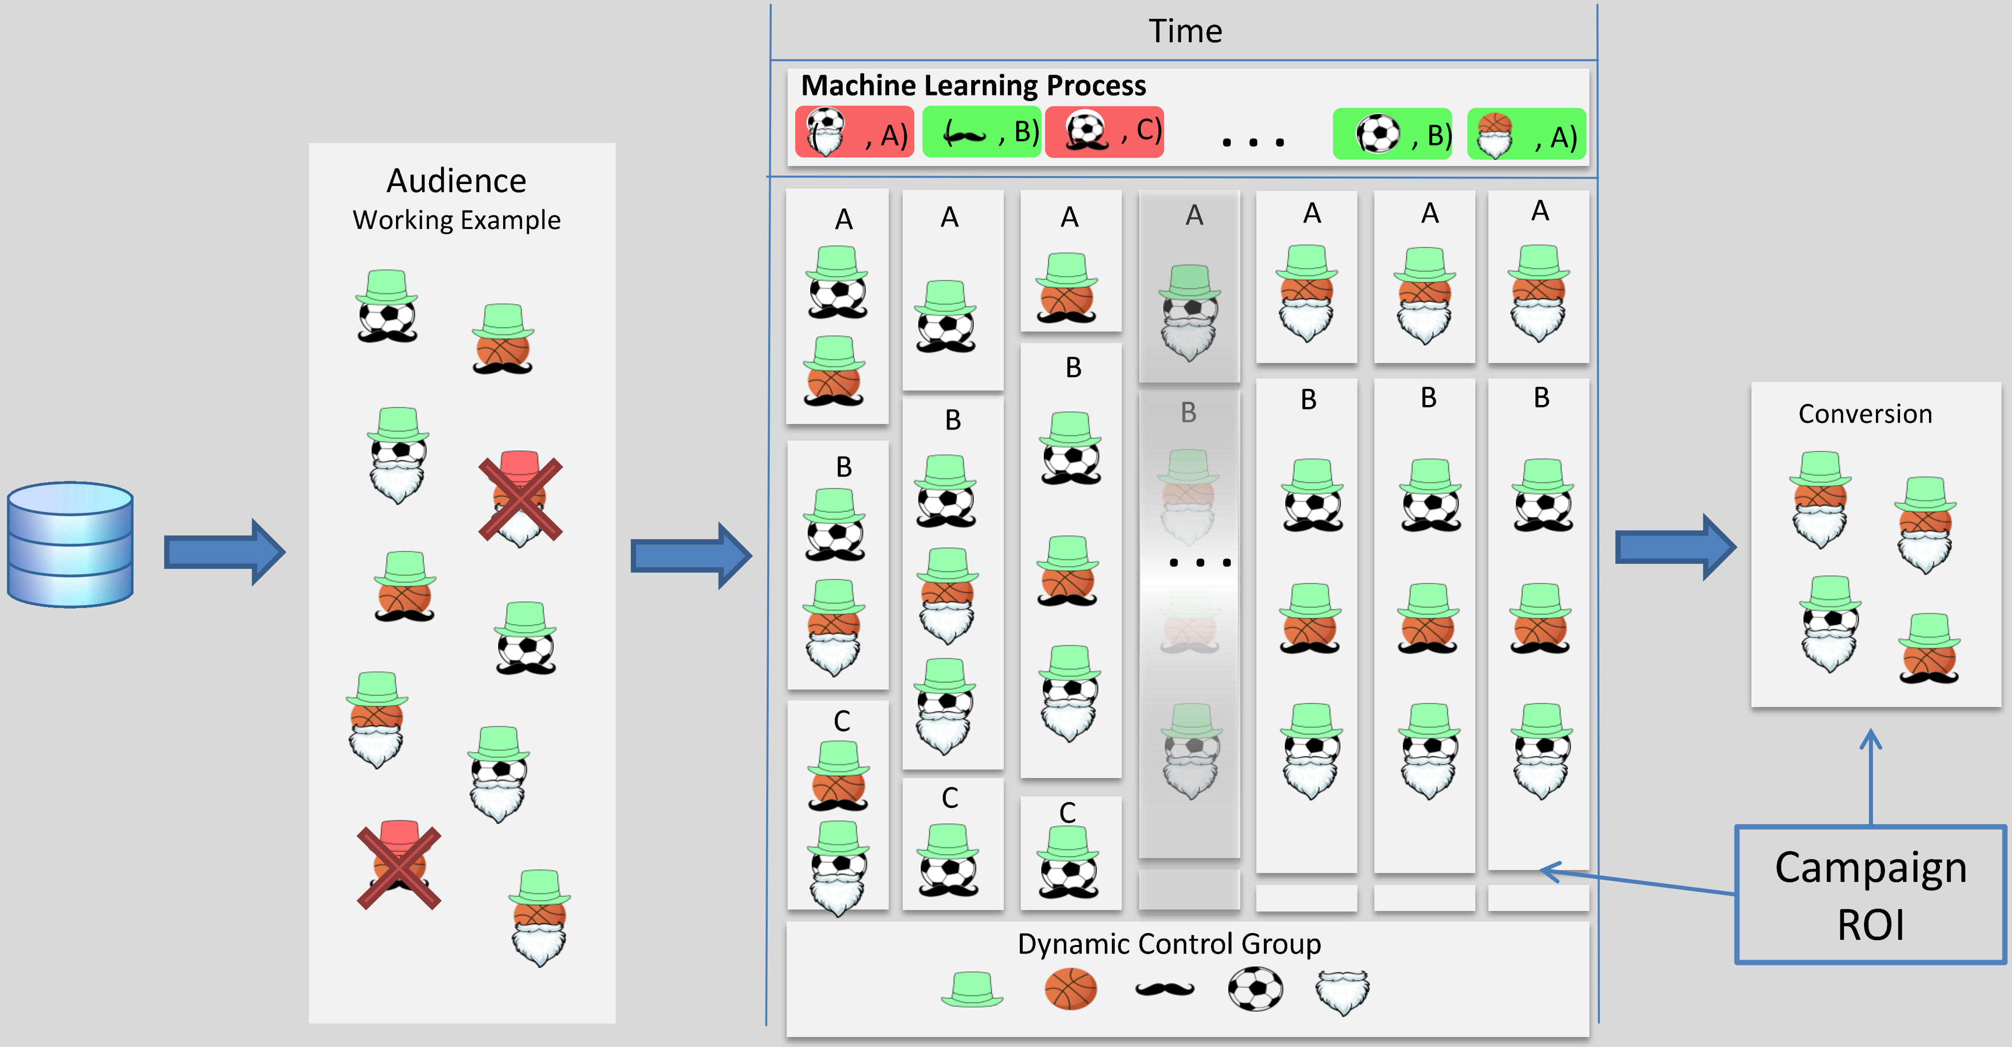

Segmento

Example of an integrated customer segmentation & targeting solution

segmento

Discover and Target Valuable Customer Segments

An integrated, data-driven customer segmentation **solution**

Automatically **discover** customer segments with common and distinct needs, preferences or value

Learn what each segment represents with detailed descriptives and unique visualizations

Observe the progress of the discovered segments over time

Design and evaluate targeted marketing **campaigns**

Solution Overview

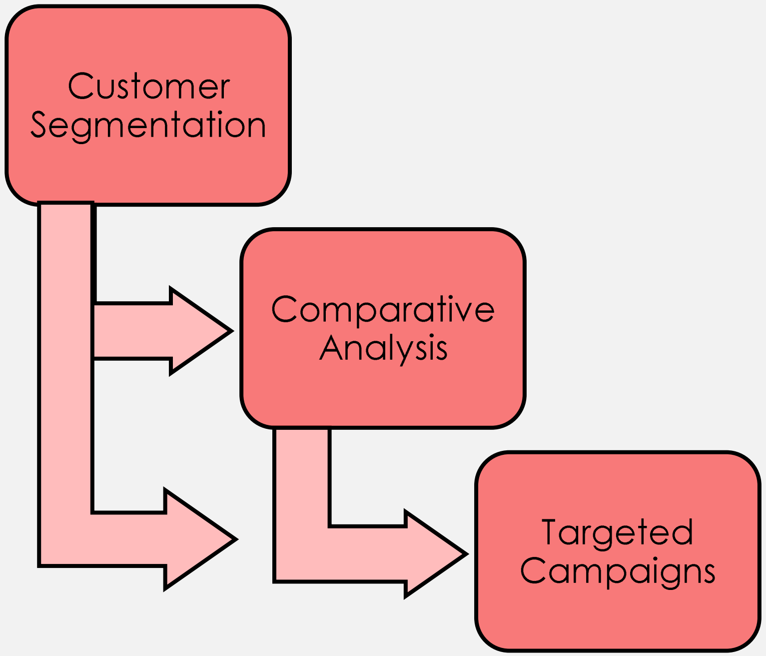

Or

Microsoft Excel Report

Customer Segmentation

Users freely select among a big variety of segmentation criteria (clustering features)

Create and save multiple customizable segmentation scenarios

State-of-the-art visualizations for each aspect of a segment's profile and the customer base in total

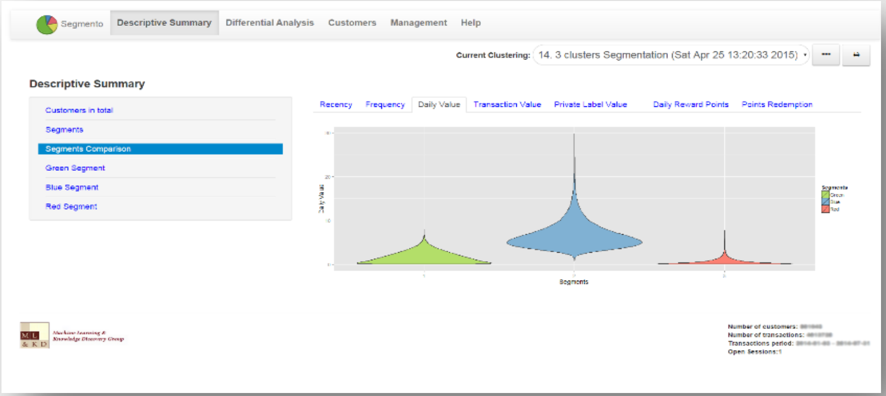

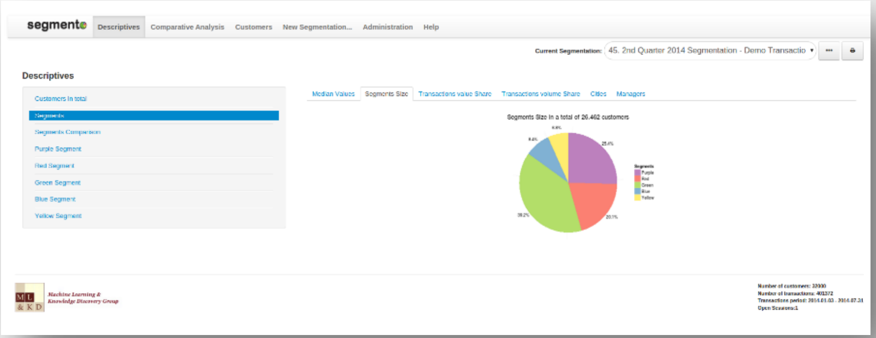

Customer Segmentation

Create segments based on quantitative data (i.e., betting preferences, type of game, teams, betting style etc.)

Example of two segments generated from 592,218 transactions and 31,675 customers in one fiscal year and four geographical locations

Segmento identified 20 customer segments

Positive and statistically significant differentiations compared to the rest of the customer base

Cluster	Category Desription	Avg. Difference	Freq. C.I. Lower	Freq. C.I. Upper
5	Μπισκότα_αλμυρά_κράκερ	+18.84%	0.54	0.61
5	Μπισκότα_γεμιστά	+15.24%	0.79	0.89
5	Ξηροί_καρποί_συσκ/νοι	+101.97%	1.39	1.68
5	Παξιμαδάκια_σνάκ	+11.71%	0.25	0.32
5	Πατατάκια	+191.06%	2.36	2.64
5	Σνακ_διάφορα	+376.11%	4.14	4.46
17	Γάλα_μακράς_διάρκειας	+129.73%	12.00	14.00
17	Μουστάρδες	+13.31%	0.33	0.33
17	Πάνα_απλή	+87.55%	1.00	1.33
17	Παιδικές_κρέμες	+55.19%	0.33	1.00
17	Σαμπουάν	+16.45%	0.67	0.67
17	Σνακ_διάφορα	+42.10%	1.00	1.33
17	Ψωμιά_τόστ	+41.77%	1.33	2.00

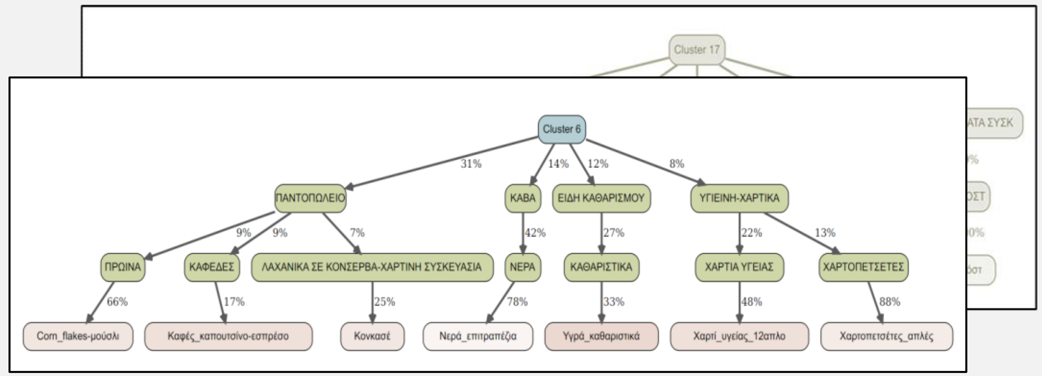

Comparative Analysis

Watch each segment progress over time, for each KPI

Information-rich plots and tables describe the effect of time over the segment's performance

Compare segments between different segmentations

Transitions between segments for different segmentation periods

Estimate per-day profit/loss for each customer transition

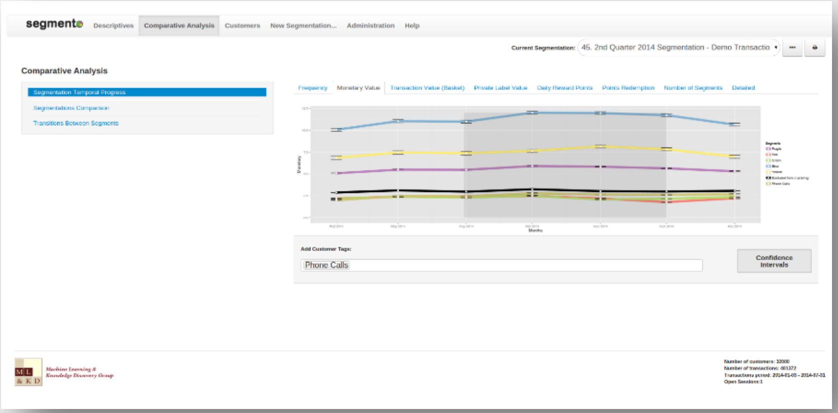

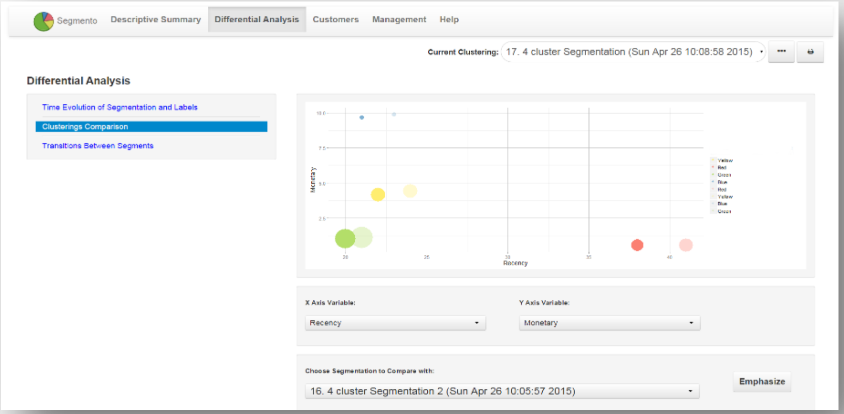

Targeted Campaigns

Identify suitable segments for a new targeted campaign

- Over 20 customer filters for finer and more focused customer selection
- Export campaign customers directly to company’s DB or even to an MS Excel file

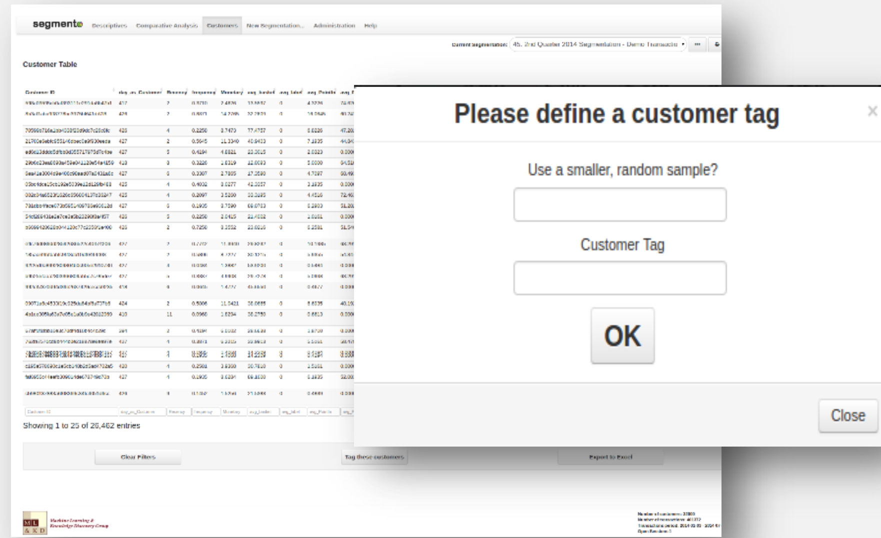

Targeted Campaigns

Evaluate campaign impact for specific customer segments

Trace customers' transitions between segments

Evaluate campaign ROI

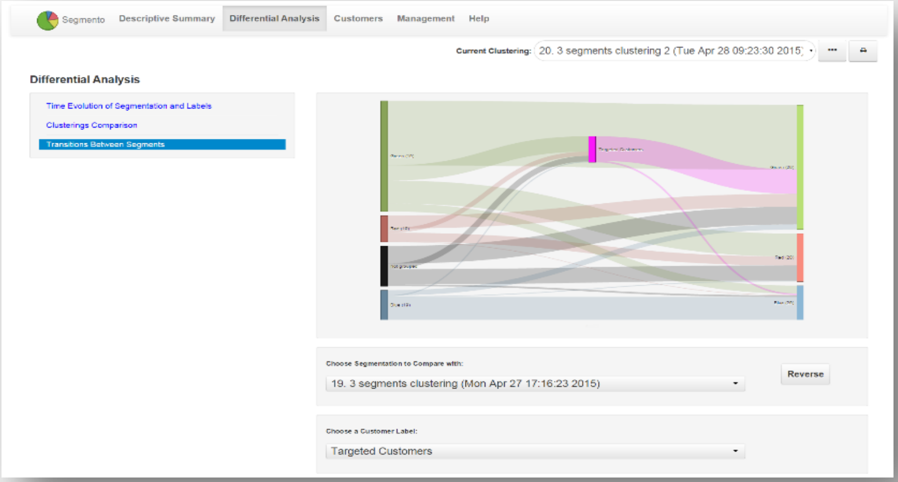

segmento
unique features

Thank You

Although this talk was not *personalized…*

…I would like to *personally* thank each one of you for attending

Intelligent Systems
Intelligent Systems Group
Computer Science Department
Aristotle University