Degree Project in Computer Science and Engineering

Second cycle, 30 credits

Evaluating Deep Topology-Preserving Models for Behavioural Customer Segmentation in Open Banking Data

A Comparative Evaluation of Autoencoder and Self-Organizing Map Hybrid Architectures

GRETA JONSSON

Evaluating Deep Topology-Preserving Models for Behavioural Customer Segmentation in Open Banking Data

A Comparative Evaluation of Autoencoder and Self-Organizing Map Hybrid Architectures

GRETA JONSSON

Master's Programme, Machine Learning, 120 credits
Date: January 31, 2026

Supervisors: Ebba Blom, Éva Székely
Examiner: Amir H. Payberah
School of Electrical Engineering and Computer Science
Host company: Kreditz AB
Swedish title: Utvärdering av djupa topologibevarande modeller för beteendebaserad kundsegmentering i Open Banking-data
Swedish subtitle: En jämförande utvärdering av Autoencoder- och Self-Organizing Map-hybridarkitekturer

Abstract

The introduction of the European Union's second Payment Services Directive (PSD2), which mandates secure, customer-consented access to bank account and transaction data, has enabled access to detailed, transaction-level financial data. This development creates new opportunities for behavioural customer segmentation based on observed financial activity rather than static demographic attributes. However, such data are high-dimensional, heterogeneous, and largely unlabeled, posing significant challenges for traditional clustering methods in terms of robustness, interpretability, and stability. This thesis evaluates the suitability of topology-preserving and deep representation learning approaches for behavioural customer segmentation in an open banking context.

Using anonymised and categorised transactional data from nearly 10,000 individuals, four unsupervised segmentation pipelines are compared: direct constrained K-means clustering, a Self-Organizing Map (SOM)-based approach, a sequential autoencoder (AE) followed by SOM, and a topology-regularised AE–SOM architecture inspired by the Deep Embedded Self-Organizing Map (DESOM) framework. All pipelines are evaluated under identical conditions using internal clustering metrics, stability analysis, and qualitative interpretability through visualisation and cluster profiling.

The results demonstrate a clear performance hierarchy, with the DESOM-based representation learning approach consistently achieving the most compact, well-separated, and stable clusters. While standalone SOMs provide strong visual interpretability in the original feature space, deep representation learning significantly improves structural cluster quality, albeit with some loss of feature-level transparency. Overall, the findings indicate that hybrid models combining autoencoders and Self-Organizing Maps are viable and effective methodologies for behavioural segmentation of PSD2-enabled transactional data, with trade-offs between interpretability and representational power that should be carefully considered in practical applications.

Keywords

Self-Organizing Maps, Clustering, Visualisation, Deep learning, Autoencoders, Open banking, Customer Segmentation

Sammanfattning

Införandet av Europeiska unionens andra betaltjänstdirektiv (PSD2), som kräver säker och kundgodkänd åtkomst till bankkonton och transaktionsdata, har möjliggjort öppen bankverksamhet och tillgång till detaljerad transaktionsdata på individnivå. Denna utveckling skapar nya möjligheter för beteendebaserad kundsegmentering baserad på observerad finansiell aktivitet snarare än statiska demografiska attribut. Samtidigt är sådan data högdimensionell, heterogen och till stor del oetiketterad, vilket medför betydande utmaningar för traditionella klustringsmetoder avseende robusthet, tolkningsbarhet och stabilitet. Denna avhandling utvärderar lämpligheten hos topologibevarande metoder och djupa representationsinlärningsansatser för beteendebaserad kundsegmentering i en öppen bankkontext.

Studien baseras på anonymiserad och kategoriserad transaktionsdata från cirka 10 000 individer och omfattar en jämförande analys av fyra oövervakade segmenteringspipelines: (i) direkt klustring med begränsad K-means, (ii) en metod baserad på Self-Organizing Maps (SOM), (iii) en sekventiell arkitektur där en autoencoder (AE) följs av en SOM, samt (iv) en gemensamt tränad AE–SOM-modell inspirerad av ramverket Deep Embedded Self-Organizing Map (DESOM). Samtliga metoder utvärderas under identiska experimentella förhållanden med hjälp av interna klustringsmått, stabilitetsanalyser samt kvalitativ bedömning av tolkningsbarhet genom visualisering och klusterprofilering.

Resultaten uppvisar en tydlig prestandahierarki, där den gemensamt tränade AE–SOM-modellen genomgående genererar de mest kompakta, väldefinierade och stabila klustren. Fristående SOM-modeller erbjuder hög visuell tolkningsbarhet i det ursprungliga attribututrymmet, medan införandet av djup representationsinlärning avsevärt förbättrar den strukturella klusterkvaliteten, dock på bekostnad av viss transparens på attributnivå. Sammantaget indikerar resultaten att hybridmodeller som kombinerar autoencoders och Self-Organizing Maps utgör ändamålsenliga och effektiva angreppssätt för beteendebaserad segmentering av PSD2-baserad transaktionsdata. De observerade avvägningarna mellan tolkningsbarhet och representationsförmåga bör emellertid beaktas vid praktisk implementering.

Nyckelord

Självorganiserande Kartor, Klustring, Visualisering, Djupinlärning, Autoencoder, Open Bankning, Kundsegmentering

Acknowledgments

I would like to thank Ebba Blom for her guidance and support throughout this thesis, and Kreditz AB for providing the opportunity to conduct this work. I would further like to thank Éva Székely and Amir H. Payberah for their valuable feedback and commitment throughout the project.

Stockholm, January 2026
Greta Jonsson

Contents

List of Figures

List of Tables

List of acronyms and abbreviations

Adam	Adaptive Moment Estimation
AE	Autoencoder
API	Application Programming Interface
BMU	Best Matching Unit
CSN	Centrala Studiestödsnämnden
DBI	Davies-Bouldin Index
DESOM	Deep Embedded Self-Organizing Map
EU	European Union
GMM	Gaussian Mixture Model
MAE	Mean Absolute Error
MI	Mutual Information
MSE	Mean Squared Error
OLS	Ordinary Least Squares
PCA	Principal Component Analysis
PSD2	Payment Services Directive 2
QE	Quantization Error
RFM	Recency-Frequency-Monetary
SCA	Strong Customer Authentication
SOM	Self-Organizing Map
TE	Topographic Error
TPP	Third-Party Provider
VaDE	Variational Deep Embedding
VAE	Variational Autoencoder

Chapter 1

Introduction

Customer segmentation is the process of grouping customers into distinct segments based on shared characteristics or behavioural patterns and constitutes a fundamental task in financial services. It enables institutions to better understand customer heterogeneity and supports data-driven decision-making in areas such as service design, risk management, and customer engagement.

Traditional segmentation approaches frequently rely on demographic variables, which risk introducing systematic bias and often fail to capture the underlying complexity of individual financial behaviours. Although financial institutions operate under regulatory frameworks that restrict the use of certain demographic attributes in credit and risk assessments, demographic information remains widely used in customer segmentation practices.

The introduction of Payment Services Directive 2 (PSD2) and Open Banking frameworks has enabled access to detailed transactional data that reflects customers' actual financial behaviour. When combined with modern data analysis tools, this development creates new opportunities for behavioural segmentation that can support more accurate, transparent, and fair decision-making. At the same time, PSD2 transactional datasets present significant methodological challenges: they are high-dimensional, heterogeneous across individuals, and largely unlabeled, limiting the applicability of supervised learning approaches.

These challenges motivate the use of unsupervised learning methods that can uncover latent behavioural structures without relying on predefined labels. In this thesis, the suitability of unsupervised learning approaches for customer segmentation based on PSD2 transactional data is investigated, with particular emphasis on representation learning and clustering performance under realistic data constraints.

1.1 Purpose

The purpose of this thesis is to evaluate the methodological suitability of combining Self-Organizing Map (SOM) and Autoencoder (AE), either jointly or sequentially, for unsupervised clustering of customers based on PSD2 transactional data. Rather than developing new algorithms or producing an operational segmentation solution, the study examines whether the evaluated pipelines can generate meaningful behavioural clusters that incorporate comprehensive transactional information, remain interpretable in terms of clearly distinguishable cluster characteristics, and are suitable for downstream analytical applications such as tracking customer transitions or analysing the size and dynamics of behavioural groups over time.

1.2 Problem Definition

This section defines the problem addressed in this thesis by specifying the research questions and the evaluation metrics used to assess the proposed methodologies.

1.2.1 Research Question

The overarching research question of this thesis is:

> To what extent can AE based and SOM based approaches be considered suitable methodologies for customer segmentation when applied to PSD2 transactional data?

This research question is examined through the following sub-questions:

1. Do sequential or jointly trained AE and SOM based approaches produce coherent and distinguishable customer segments compared to traditional SOM or K-means clustering?

2. To what extent can the resulting segments be interpreted and characterised in terms of identifiable behavioural patterns derived from PSD2 transactions?

3. How stable and robust are the segmentation results across multiple runs, and do the evaluated approaches appear feasible for use with real-world financial transaction data?

4. Do any of the approaches demonstrate properties that make them suitable for potential future applications, such as tracking customer transitions or monitoring the evolution of segment structures over time?

These pipelines are selected to enable a structured comparison between direct clustering, topology-preserving mapping, sequential representation learning, and joint learning approaches.

1.2.2 Evaluation Metrics

The approaches will be evaluated through both quantitative and qualitative analysis. Quantitative evaluation examines internal clustering metrics to assess structural properties of the segmentations, while qualitative evaluation focuses on interpretability and visual clarity.

1. **Training Behaviour Monitoring:** Monitor training progress for the AE and the SOM using reconstruction loss, Topographic Error (TE) and Quantization Error (QE) in order to ensure convergence and adequate representation quality.

2. **Cluster Quality Assessment:** Evaluate cohesion and separation of the resulting clusters using the internal validity metrics Silhouette Score and Davies-Bouldin Index (DBI), and compare these outcomes across the different clustering approaches.

3. **Stability Analysis:** Assess the robustness of the segmentation by repeating each method with multiple random initialisations and parameter configurations, and examining the consistency of the resulting cluster assignments.

4. **Interpretation and Visualisation:** Analyse visual representations including SOM grid visualisations, U-matrix plots and component planes to explore behavioural patterns. Identify and describe the key characteristics that distinguish the clusters and evaluate the interpretability of each approach.

1.3 Overview of Experimental Pipelines

To address the research questions defined in Section 1, this thesis evaluates a set of alternative unsupervised learning pipelines for customer segmentation based on PSD2 transactional data. Each pipeline represents a distinct methodological strategy for transforming high-dimensional transactional features into behavioural segments, differing in how representation learning, topology preservation, and clustering are combined.

The pipelines are designed to form a progression from simple baseline methods to more advanced hybrid architectures, allowing the incremental contribution of each modelling component to be assessed. This structure supports a systematic evaluation of how representation learning and topology-aware methods affect clustering quality, stability, and interpretability.

In total, four experimental pipelines are considered:

1. **Pipeline 0 (Baseline: K-means)**
   Direct constrained K-means clustering on the standardised transactional feature space.

2. **Pipeline 1 (SOM + K-means)**
   A SOM is trained on the original feature space to obtain a topology-preserving mapping, followed by constrained K-means clustering.

3. **Pipeline 2 (AE + SOM + K-means)**
   An AE is first used to learn a compact latent representation, which is then structured using a SOM before clustering with constrained K-means.

4. **Pipeline 3 (Deep Embedded Self-Organizing Map (DESOM) + SOM + K-means)**
   A topology-regularised AE inspired by the DESOM framework is used to learn structured embeddings, which are subsequently mapped with a SOM and clustered using constrained K-means.

These pipelines are intentionally designed to be comparable under identical data, feature, and evaluation conditions. By analysing their relative performance, the thesis examines whether increased model complexity leads to more coherent and interpretable behavioural segments, and whether such improvements justify the added methodological and computational complexity.

1.4 Goals

The specific objectives of this thesis are to operationalize the evaluation of the proposed pipelines. While the purpose defines the overarching aim of assessing methodological suitability, the goals focus on the concrete steps required to perform this assessment.

To achieve this, the thesis will:

1. Design and construct features that capture customers' financial behaviour based on PSD2 transactional data.
2. Train an AE to learn latent behavioural representations from the engineered feature space.
3. Apply SOMs and K-means clustering to the original feature space.
4. Apply SOMs and K-means clustering to the latent representations learned by the AE.
5. Evaluate the DESOM approach.
6. Compare all approaches using internal validity metrics and visual analysis, with particular emphasis on interpretability and methodological suitability for customer segmentation.

A successful outcome will provide insight into whether AE- and SOM-based methods can produce meaningful, stable, and interpretable behavioural segments when applied to PSD2 transactional data.

1.5 Research Methodology

This research adopts a quantitative, experimental approach grounded in unsupervised machine learning. The methodology consists of implementing and comparing multiple clustering frameworks, including SOM combined with K-means clustering and a jointly trained AE-SOM model, under identical data and feature conditions. Data preparation and feature engineering constitute the foundation of the study, followed by model training, parameter selection, and comparative evaluation using internal validation metrics and visualisation-based analysis. The chosen methodology is motivated by the need to balance advanced representation learning with interpretability, ensuring that the results are both scientifically rigorous and practically relevant.

1.6 Delimitations

This thesis is limited to the evaluation of existing unsupervised learning methods applied to anonymised transactional data provided by Kreditz. As a result, the findings may not generalise directly to other datasets, geographical regions, or financial systems. Demographic attributes and external contextual variables are excluded, as the analysis focuses solely on behavioural segmentation derived from financial transactions. Predictive modelling tasks, including churn prediction and credit scoring, are outside the scope of this study.

1.7 Structure of the thesis

The remainder of this thesis is structured as follows. Chapter 2 provides the necessary theoretical background, covering customer segmentation, open banking, and the machine learning concepts used in this study, including SOMs and AEs. Chapter 3 describes the data, preprocessing steps, experimental pipelines, and evaluation metrics. Chapter 4 presents the results of the experiments, analysing both quantitative metrics and the behavioural characteristics of the identified clusters. Finally, Chapter 5 summarises the main findings, discusses the limitations of the study, and suggests directions for future work.

Chapter 2

Background and Related Work

2.1 Customer Segmentation

Customer segmentation is the process of dividing a market or customer base into distinct groups of individuals who exhibit similar characteristics, needs, or behaviours. It enables organisations to identify and target specific customer segments with tailored products, services, and marketing strategies, thereby supporting data-driven decision-making in areas such as marketing, product design, and risk assessment. The concept was first introduced by Wendell R. Smith in 1956 [1], who argued that mass markets were evolving towards differentiated marketing driven by heterogeneous customer needs. Since then, segmentation has become a central concept in marketing and strategic decision-making across a wide range of industries.

Traditional segmentation approaches often rely on demographic, geographic, psychographic, or behavioural variables. However, the increasing availability of digital behavioural data, particularly from transactional sources, has enabled more dynamic and data-driven segmentation methods. Behavioural segmentation, which classifies customers based on actual usage patterns or spending behaviour, provides a more accurate representation of customer needs and preferences than demographic segmentation alone.

Recent advances in machine learning have further expanded the analytical possibilities for customer segmentation. Neural network-based methods enable the discovery of latent behavioural patterns in large-scale financial datasets, allowing institutions to move beyond manual feature grouping towards automated and adaptive segmentation pipelines that can evolve as customer behaviour changes [2]. This development is particularly relevant in the context of open banking, where granular transaction data can reveal complex behavioural structures that traditional segmentation methods often fail to capture.

2.2 Open Banking Data

In the European Union (EU), open banking was formally introduced through the Revised PSD2 (Directive (EU) 2015/2366) [3], which came into effect in January 2018. The objective of PSD2 is to increase competition in the financial sector, enhance consumer protection, and promote innovation in payment services. By mandating that banks provide licensed Third-Party Providers (TPPs) with secure access to customer account data and payment initiation through standardised Application Programming Interfaces (APIs). PSD2 ensures that consumers retain control over their data while benefiting from new value-added financial services [4].

This regulatory framework enables greater consumer choice and encourages the development of innovative digital services across the European Single Market [3]. At the same time, PSD2 reinforces trust and security in financial transactions by requiring Strong Customer Authentication (SCA) and compliance with technical standards defined by the European Banking Authority [4].

From a data-analytic perspective, open banking data provide an opportunity to study customer behaviour based on observed financial activity rather than self-reported or proxy indicators. These datasets consist of transaction records that can be categorised and aggregated into spending profiles across domains such as groceries, housing, transport, or leisure [2]. As demographic and personally identifiable information is excluded, the analysis focuses on behavioural patterns, supporting privacy preservation while enabling robust and interpretable segmentation.

In this thesis, categorised and anonymised transactional data form the basis for evaluating clustering and representation learning methods. Such data capture customers' real financial behaviour, providing a suitable setting for investigating how unsupervised learning approaches can enhance segmentation quality and interpretability in open banking applications.

2.3 The RFM Framework

A widely used framework for behavioural feature design in finance and marketing is the Recency-Frequency-Monetary (RFM) model [5]. The RFM approach summarises customer activity using three dimensions:

- **Recency (R):** the time elapsed since the customer's most recent transaction,
- **Frequency (F):** the number of transactions made within a given time window, and
- **Monetary (M):** the total or average monetary value of transactions over the same period.

These features serve as behavioural indicators of engagement and value, allowing segmentation models to differentiate between, for instance, frequent high-value customers and infrequent low-value customers. Due to its simplicity, interpretability, and empirical effectiveness, the RFM framework has become a foundational approach for customer scoring and segmentation [5].

Recent research has extended the classical RFM framework by incorporating additional behavioural and contextual dimensions. Liao et al. [6] proposed a multi-behaviour RFM model that integrates features such as transaction diversity and enhanced recency measures, demonstrating improved segmentation stability and accuracy. Similarly, Mancisidor et al. [7] showed that representations learned by AEs can automatically extract latent analogues of RFM dimensions from raw transactional data, capturing non-linear dependencies that manual feature construction may overlook.

Despite its widespread use, the classical RFM framework is not always directly applicable to open banking data, motivating adaptations that better reflect the structure and constraints of transaction-level datasets.

2.4 Machine Learning Concepts

This chapter introduces the key machine learning concepts underpinning this study. It covers the transformation of transactional data into behavioural features, methods for unsupervised clustering, and the use of deep learning techniques to learn interpretable latent representations. Together, these concepts provide the methodological foundation for the analysis of financial customer segmentation in later chapters.

2.4.1 Feature Engineering and Selection

Feature selection is the process of identifying the most informative variables in a dataset that contribute significantly to the learning objective or clustering structure. In unsupervised learning, particularly clustering, effective feature selection is crucial, as irrelevant or redundant variables can obscure the underlying structure of the data and degrade cluster quality. The goal is to retain features that capture meaningful behavioural variation while reducing noise, dimensionality, and computational complexity.

In the context of financial customer segmentation, feature engineering and selection play a central role in constructing behavioural representations from raw transactional data. Without careful preprocessing and aggregation, such data can lead to biased or unstable clusters. Therefore, designing features that summarise customer behaviour in a compact and interpretable manner is a critical step in producing robust segmentations.

In open banking and PSD2-compliant datasets, feature engineering must also balance behavioural richness with privacy constraints. Since demographic and personally identifiable information are excluded, engineered features must be derived solely from transactional patterns. Common approaches include aggregating spending at the category level, computing average transaction sizes, or measuring variability in expenditure across time and categories [2]. These aggregated and anonymised features preserve privacy while maintaining analytical value and interpretability.

While inspired by the principles of RFM, the feature representations used in this thesis differ from the classical formulation. Recency and frequency are not explicitly modelled, as such measures can be unstable in short observation windows and sensitive to account linking gaps in open banking data, whereas aggregated monetary patterns provide a more robust basis for unsupervised clustering. This adaptation aligns with the constraints and opportunities of open banking data and provides a flexible behavioural representation suitable for unsupervised learning.

In this work, transactional data are aggregated at the customer level to construct behavioural representations suitable for unsupervised learning.

2.4.2 Clustering

Clustering is a fundamental task in unsupervised learning that aims to group similar observations based on underlying characteristics or behavioural patterns. In contrast to supervised methods, clustering does not rely on labelled data, but instead seeks to uncover latent structure within a dataset.

A wide range of clustering algorithms has been developed, each relying on different assumptions regarding data distribution, cluster shape, and similarity measures. Among the most widely used is the centroid-based K-means algorithm. Another, although less commonly applied, approach is the topology-preserving SOM. The following sections introduce both methods and discuss their relevance for customer segmentation.

K-means

The K-means algorithm is one of the most commonly used clustering techniques, due to its conceptual simplicity and computational efficiency. It partitions a dataset into k clusters by minimising the sum of squared distances between data points and their corresponding cluster centroids. The algorithm iteratively alternates between two steps. First, assigning each data point to the nearest centroid according to a distance metric, typically Euclidean distance. Secondly updating each centroid as the mean of the points assigned to it. This process continues until convergence, which typically occurs when cluster assignments stabilise or the objective function reaches a local minimum.

Despite its popularity, K-means has several limitations. It assumes that clusters are approximately spherical and of similar size, and it is sensitive to feature scaling and centroid initialisation. These assumptions can be problematic when analysing complex, high-dimensional behavioural data such as financial transactions, where cluster structures may deviate from simple geometric forms. Nevertheless, K-means remains a widely used baseline for customer segmentation due to its efficiency and its ability to provide clear centroid-based representations of cluster characteristics.

In financial applications, K-means is often applied to engineered behavioural features, such as RFM indicators, to group customers according to transaction history or engagement level. Studies such as Bartels [2] have demonstrated that even standard clustering algorithms like K-means can identify meaningful behavioural segments from open banking data, providing a strong foundation for further methodological development. However, when applied to more complex or non-linear financial datasets, K-means may struggle to capture subtle behavioural distinctions, motivating the exploration of more flexible and topology-aware clustering methods such as the SOM.

2.4.3 Self-Organizing Maps

The SOM, introduced by Kohonen in 1990 [8], is an unsupervised neural network model designed for dimensionality reduction and topology preservation. The SOM projects high-dimensional data onto a typically two-dimensional grid while maintaining neighbourhood relationships between samples. This dual capability enables both clustering and visualisation, making SOM particularly suitable for exploring complex, high-dimensional datasets such as financial transaction records.

Each neuron in the SOM represents a prototype vector in the input space. During training, the neuron whose weight vector most closely matches an input sample, referred to as the Best Matching Unit (BMU), and its neighbouring neurons are updated to become more similar to that input. Over time, this competitive and cooperative learning process organises the map such that neighbouring neurons respond to similar input patterns. Dense regions of the map correspond to clusters of similar behaviour, while sparse regions indicate boundaries between clusters.

SOMs have been applied in marketing and finance to identify patterns in customer behaviour that are not readily captured by demographic or linear models. For example, Yanık and Elmorsy [9] used SOMs to cluster credit card transactions and reveal spending-based customer segments. Similarly, Huang et al. [10] combined SOM with K-means for online retail segmentation, reporting improved cluster stability and interpretability. Liao et al. [6] further enhanced SOM adaptability by introducing dynamic learning mechanisms in a multi-behaviour RFM framework, achieving lower QEs and TEs. From a visual analytics perspective, Kovalerchuk et al. [11] demonstrated that SOM component planes and U-matrices can effectively illustrate relationships between financial attributes, thereby improving transparency in decision support systems.

In the context of open banking, Bartels [2] demonstrated that customer segmentation can be performed effectively using categorised and anonymised transaction data, highlighting the value of behavioural features in privacy-preserving analytics. While that study employed traditional clustering techniques rather than SOM, the findings underscore the potential of transaction-based segmentation and motivate further exploration of topology-preserving and representation-learning approaches. More recent work, such as Zhang et al. [12], has reaffirmed SOM's effectiveness as a front-end mapping layer for clustering high-dimensional behavioural datasets.

Overall, the SOM provides a powerful framework for interpretable and topology-aware clustering. To further explore behavioural structure, this thesis investigates the use of AEs for learning latent representations, which are then visualised and analysed using a SOM. This setup enables an examination of the topology of the latent representation space in a customer segmentation context.

2.4.4 Autoencoder

An AE is an unsupervised neural network architecture designed to learn compact, information-rich representations of input data through reconstruction. The model consists of two main components: an encoder, which compresses the input into a lower-dimensional latent space, and a decoder, which attempts to reconstruct the original data from this latent representation. By minimising the reconstruction error between the input and the output, the AE learns a lower-dimensional representation of the underlying data distribution. This enables AEs to perform non-linear dimensionality reduction, often outperforming traditional linear methods such as Principal Component Analysis (PCA) when applied to complex or high-dimensional data.

This thesis hypothesises that AEs are particularly suitable for financial and transactional data, where features are often high-dimensional, sparse, and non-linearly correlated. By learning latent representations that capture underlying behavioural patterns, AEs can facilitate downstream tasks such as clustering or anomaly detection. In the context of customer segmentation, the encoder's latent space can be used as input to clustering algorithms in order to improve the separation and compactness of behavioural groups.

Mancisidor et al. [7] demonstrated the usefulness of AEs for financial segmentation by applying a Variational Autoencoder (VAE) to bank transaction data. Their study showed that the learned latent representations captured meaningful customer behaviour patterns and outperformed manually engineered features in downstream segmentation and predictive tasks. Similarly, a study by the Saint Francis College research group [13] compared deep learning-based and traditional segmentation methods, concluding that AE-derived representations can support real-time decision-making in banking applications due to their ability to generalise complex customer profiles.

Building on this idea, Jiang et al. [14] proposed the Variational Deep Embedding (VaDE) framework, which integrates a VAE with a Gaussian Mixture Model (GMM) in the latent space. The model jointly optimises representation learning and clustering objectives, achieving improved cluster separation and stability compared to sequential pipelines. This line of research established a foundation for hybrid deep clustering approaches, where representation learning and cluster formation are performed simultaneously.

A related development is the SOM-VAE model proposed by Fortuin et al. [15], which combines a VAE with a SOM to produce interpretable and discrete latent representations. SOM-VAE balances reconstruction accuracy with topology preservation, enabling both quantitative clustering and qualitative visualisation. The model has been successfully applied to time-series and behavioural data, demonstrating its suitability for applications where interpretability is as important as representational capacity.

Forest et al. [16] further advanced this concept with the DESOM, which jointly trains an AE and a SOM through a combined loss function that enforces both reconstruction fidelity and topological structure in the latent space. This joint optimisation enables the model to learn representations that are both compact and organised according to neighbourhood relations. DESOM demonstrated superior performance in producing well-separated and interpretable clusters across several benchmark datasets, indicating strong potential for domains requiring structured yet explainable segmentation, such as financial customer analysis.

Khacef et al. [17] similarly explored the integration of unsupervised feature extraction with SOMs, showing that deep neural encoders can improve the quality of SOM mappings by providing more informative input representations.

Recent developments in the financial domain have continued to build on these principles. Jin [18] proposed a self-supervised, multi-view deep clustering framework for bank customer segmentation and risk assessment, using contrastive learning to align multiple behavioural perspectives. The study demonstrated that combining deep representation learning with clustering objectives leads to more stable and interpretable segments compared to conventional unsupervised methods. Earlier work by Liao and Hsu [19] explored adaptations of SOMs for categorical transactional data, highlighting the potential of topology-preserving approaches for modelling high-dimensional financial behaviour.

The integration of AEs and SOMs therefore represents a promising hybrid approach that balances interpretability and representational power. While previous research has demonstrated the effectiveness of AE-SOM architectures for image, sensor, and general high-dimensional numerical data, their application to financial or behavioural datasets remains limited. In particular, existing studies have not systematically explored such models on categorised transactional data, where spending behaviour is represented through aggregated expenditure categories.

This thesis contributes by extending the AE-SOM framework to a novel data domain, namely categorised and anonymised financial transaction data collected under the PSD2 open banking framework. By evaluating both sequential and jointly trained AE-SOM models against traditional clustering baselines such as SOM and K-means, the study aims to assess whether deep representation learning can enhance segmentation quality, stability, and interpretability in behavioural financial analysis. In doing so, it provides one of the first systematic examinations of topology-preserving deep clustering applied to structured transactional data, offering methodological insights relevant to both future research and practical customer analytics applications.

2.4.5 Summary

This chapter reviewed the theoretical foundations and prior research relevant to customer segmentation in financial services. It traced the evolution from traditional demographic-based segmentation towards behavioural approaches enabled by open banking and PSD2, where anonymised transaction data allow the analysis of actual spending behaviour while preserving privacy. The chapter then discussed key machine learning concepts underpinning this study, including the role of feature engineering in transforming raw financial data into behavioural representations, and the application of clustering methods such as k-means and SOMs for identifying latent customer groups. Finally, it introduced AEs and hybrid AE-SOM frameworks as contemporary approaches that combine deep representation learning with topology preservation, enabling more robust and interpretable segmentation of complex transactional data. Together, these concepts provide the theoretical and methodological foundation for the empirical analysis presented in the following chapters.

Chapter 3

Methods

3.1 Data Description and Preprocessing

The study utilises anonymised and categorised transactional data provided by Kreditz AB. As part of its standard processing pipeline, Kreditz performs transaction enrichment and categorisation, assigning each transaction to a predefined semantic category (e.g. *salary*, *rent*, *groceries*). This categorisation is completed prior to data delivery and constitutes the initial preprocessing step applied to the raw transaction data.

The dataset provided for this study consists of transaction records from 9,853 Swedish individuals. Each record contains a transaction date, a monetary amount, and an associated transaction category as assigned by Kreditz. All data are fully anonymised, and no personally identifiable or demographic information is included. The transactions were collected during the past year and cover a six-month observation period.

Prior to analysis, additional preprocessing was conducted to ensure consistency and suitability for unsupervised modelling. Transactions were aggregated at the category level to reduce granularity and to emphasise behavioural spending patterns. Feature engineering was then performed by computing the monthly monetary spending per category for each customer over the six-month observation window, resulting in a stable, fixed-length representation of individual financial behaviour.

To ensure comparability across features and to mitigate the influence of differing scales, the resulting feature set was standardised using z-score normalisation. The standardisation was applied according to the following formula:

$$z = \frac{x - \mu}{\sigma} \tag{3.1}$$

where x denotes the original feature value, μ represents the mean of the feature, and σ denotes its standard deviation. This transformation centres each feature around zero and scales it by its variability, enabling fair comparison across spending categories.

3.1.1 Feature Selection

A deliberately restrictive feature selection strategy was adopted in this study. Rather than excluding transaction categories with overlapping semantics or similar behavioural meaning, correlated categories were aggregated where appropriate. This choice was motivated by the structure of the provided dataset, which already contained a predefined set of transaction categories intended to represent customers' financial activity comprehensively.

Although more aggressive feature reduction could potentially improve clustering performance for certain methods, the objective of this thesis is to evaluate segmentation approaches under the constraint that all available transactional information contributes to the customer representation. Retaining the full category set allows for a fair assessment of methods designed to learn latent structure from high-dimensional data, particularly AE-based models.

Feature aggregation decisions were informed by a combination of Pearson correlation analysis, Mutual Information (MI) analysis, and domain knowledge provided by Kreditz. This approach reduced redundancy while preserving interpretability and maintaining a comprehensive behavioural representation.

3.1.1.1 Pearson Correlation

The Pearson correlation coefficient measures the strength and direction of the linear relationship between two continuous variables. Given two random variables X and Y, their Pearson correlation coefficient r_{XY} is defined as:

$$r_{XY} = \frac{\mathrm{cov}(X, Y)}{\sigma_X \sigma_Y} \tag{3.2}$$

where $\mathrm{cov}(X, Y)$ denotes the covariance between X and Y, and σ_X and σ_Y are their standard deviations. The coefficient r_{XY} ranges from -1 to 1, where values close to 1 indicate a strong positive linear relationship, values close to -1 indicate a strong negative relationship, and values near 0 indicate weak or no linear correlation.

To identify redundancy among features, the Pearson correlation matrix $\mathbf{R}$ was computed, where each element r_{ij} represents the correlation between feature i and feature j:

$$\mathbf{R} = \begin{bmatrix} 1 & r_{12} & \cdots & r_{1p} \\ r_{21} & 1 & \cdots & r_{2p} \\ \vdots & \vdots & \ddots & \vdots \\ r_{p1} & r_{p2} & \cdots & 1 \end{bmatrix} \tag{3.3}$$

Highly correlated feature pairs (e.g., $|r_{ij}| > 0.9$) were reviewed, and features representing similar behavioural concepts were aggregated to reduce redundancy while preserving interpretability.

3.1.1.2 Mutual Information

MI is an information-theoretic measure that quantifies the amount of shared information between two random variables. It reflects how much knowledge of one variable reduces uncertainty about another. For two discrete random variables X and Y, MI is defined as:

$$I(X;Y) = \sum_{x \in X} \sum_{y \in Y} p(x,y) \log \left(\frac{p(x,y)}{p(x)p(y)} \right) \tag{3.4}$$

where $p(x,y)$ denotes the joint probability distribution of X and Y, and $p(x)$ and $p(y)$ are their marginal distributions. For continuous variables, the summation is replaced by an integral.

Higher values of $I(X;Y)$ indicate stronger dependency between variables, capturing both linear and non-linear relationships. In this study, MI scores were computed using the `mutual_info_regression` function from the `sklearn.feature_selection` module, which provides a non-parametric estimation of dependency suitable for continuous variables.

MI analysis was used as a complementary measure to Pearson correlation to assess feature dependency and support aggregation decisions. No further feature elimination was performed, as the methodological objective is to retain the full transactional information content and allow the AE to uncover latent relationships within the data.

3.2 Models

This section describes the models used to construct and evaluate the proposed customer segmentation pipelines. The focus is on methodological transparency rather than architectural novelty: all models were selected for their suitability in unsupervised, high-dimensional, and unlabeled transactional data settings. Each model is therefore presented together with its training procedure, hyperparameter selection strategy, and evaluation criteria, with particular attention paid to reproducibility and interpretability. The models are combined in different configurations across pipelines to isolate the effects of dimensionality reduction, topology preservation, and joint optimisation on clustering behaviour.

3.2.1 Self-Organizing Map

The SOM was implemented using the `MiniSom` Python library [20]. The input to the SOM consisted of the standardised (z-score normalised) feature set described in the previous section. For the sequential and joint AE-SOM approaches, the SOM was instead trained on the latent representations extracted from the bottleneck layer of the trained AE.

The SOM grid size was determined empirically based on the dimensionality and sample size of the dataset, following a commonly used heuristic in which the total number of neurons N scales proportionally with the square root of the number of samples n:

$$N \approx 5 \times \sqrt{n}. \tag{3.5}$$

This heuristic provides a balance between map resolution and computational efficiency. The grid aspect ratio was determined using the first two PCA components of the input data, such that the SOM layout reflects the dominant variance directions:

$$\text{Grid Ratio} = \frac{\lambda_1}{\lambda_2}, \tag{3.6}$$

where λ_1 and λ_2 are the eigenvalues associated with the first and second principal components, respectively. Aligning the SOM grid with the principal variance directions helps reduce distortion effects such as edge clustering and supports more faithful topology preservation [8].

For all pipelines, SOM weights were initialised using PCA-based initialisation, providing a structured starting point aligned with the principal directions of the input space. Compared to random initialisation, this approach accelerates convergence and improves topological ordering of the resulting map. Training was performed using the batch learning algorithm provided by `MiniSom`, in which neuron weights are updated simultaneously based on the aggregated influence of all training samples in each epoch. Batch learning was selected for its improved stability and reproducibility.

A grid search was conducted to identify suitable SOM hyperparameters. The explored parameter space is summarised in Table 3.1.

Table 3.1: Grid search parameters for SOM hyperparameter tuning.

Parameter	**Values Tested**
Neighborhood radius (σ)	12, 8, 6, 4, 2
Learning rate (α)	0.5, 0.3, 0.1, 0.05

During training, each input vector $\mathbf{x}$ was compared to all neuron weight vectors $\mathbf{w}_i$ to identify the BMU, defined as the neuron with the minimum Euclidean distance to the input. The BMU represents the prototype most similar to the given sample, and its neighbourhood was updated according to the chosen neighbourhood function in order to preserve topological relationships within the data.

Training duration was defined in terms of epochs, where one epoch corresponds to a complete pass through the dataset. The grid search experiments were conducted for 5 epochs, while the final SOM models were trained for 30 epochs. The total number of training iterations T is therefore given by:

$$T = E \times n, \tag{3.7}$$

where E denotes the number of epochs and n the number of samples. This training duration allowed the SOM to converge to a stable topology while avoiding excessive overfitting to local structures.

In addition to qualitative visual inspection of the trained maps, two standard SOM quality measures were computed to assess representation accuracy and topology preservation: the QE and the TE.

Quantization Error and Topographic Error

The QE and TE were computed using the built-in evaluation functions provided by the `MiniSom` library [20], ensuring consistent and methodologically aligned measurement across all SOM variants.

The QE reflects how accurately the SOM prototype vectors represent the input data. It is defined as the average Euclidean distance between each input vector $\mathbf{x}_i$ and the weight vector of its corresponding BMU:

$$\mathrm{QE} = \frac{1}{n} \sum_{i=1}^{n} \left\| \mathbf{x}_i - \mathbf{w}_{\mathrm{BMU}(i)} \right\|. \tag{3.8}$$

Lower QE values indicate that the BMU prototypes lie closer to the data distribution, signifying improved quantisation quality. In this study, QE was used to compare how well different SOM configurations captured the structure of the input or latent representation space.

The TE measures the extent to which the SOM preserves neighbourhood relationships inherent in the data. It is computed as the proportion of samples for which the first- and second-BMUs are not adjacent on the SOM grid:

$$\mathrm{TE} = \frac{1}{n} \sum_{i=1}^{n} \mathbb{I}[\mathrm{BMU}_1(i) \nsim \mathrm{BMU}_2(i)], \tag{3.9}$$

where $\mathbb{I}[\cdot]$ denotes the indicator function and $\nsim$ indicates non-adjacency on the map lattice. Lower TE values correspond to stronger topology preservation, meaning that samples that are close in the input (or latent) space are mapped to neighbouring neurons.

Together, QE and TE provide complementary measures of SOM performance by jointly capturing representation accuracy and topological consistency.

3.2.2 Autoencoder

The AE used in this study was implemented as a fully connected feed-forward neural network with a symmetric encoder-decoder architecture. The model maps an input vector $\mathbf{x} \in \mathbb{R}^d$ to a latent representation $\mathbf{z}$ through the encoder,

$$\mathbf{z} = f_\theta(\mathbf{x}),$$

and reconstructs the input via the decoder,

$$\hat{\mathbf{x}} = g_\phi(\mathbf{z}),$$

where θ and ϕ denote the learnable parameters of the encoder and decoder, respectively. The network was trained to minimise the mean squared reconstruction error,

$$\mathcal{L}_{\text{rec}} = \frac{1}{n}\sum_{i=1}^{n} \left\| \mathbf{x}_i - g_\phi(f_\theta(\mathbf{x}_i)) \right\|^2 .$$

The AE was implemented using `TensorFlow Keras` [21] in combination with `scikit-learn` [22]. The encoder consisted of progressively decreasing hidden layers that compressed the standardised input features into a lower-dimensional bottleneck representation, while the decoder mirrored this structure to reconstruct the original input. Hidden layers employed Batch Normalisation followed by LeakyReLU activation functions, with dropout applied after each hidden block for regularisation. The bottleneck and output layers used linear activation functions. Training was performed using the Adaptive Moment Estimation (Adam) optimiser together with Mean Squared Error (MSE) as the reconstruction loss, and Mean Absolute Error (MAE) was monitored as an auxiliary metric.

To ensure robust convergence and mitigate overfitting, training employed early stopping monitored on validation loss with restoration of the best-performing weights, together with a ReduceLROnPlateau learning rate scheduler that reduced the learning rate when validation loss stagnated. When a separate validation set was not provided, the data were split into training and validation sets using an 80/20 split with a fixed random seed (42) to ensure reproducibility.

An extensive grid search was conducted over a range of architectural and training-related hyperparameters to identify a robust and generalisable AE configuration. During the grid search, each candidate architecture was trained for up to 20 epochs with aggressive early stopping, while the final selected model was trained for up to 100 epochs using the same training procedure. To reduce computational cost, the grid search was performed on a randomly selected subset of up to 2,000 training samples. The explored parameter space is summarised in Table 3.2 and included variations in network depth, dropout regularisation, learning rate, L2 regularisation strength, and batch size.

Table 3.2: Grid search parameters for AE hyperparameter tuning.

Parameter	Values Tested
Hidden dimensions	[128, 64, 32], [256, 128, 64], [512, 256, 128]
Dropout rate	0.0, 0.1, 0.2
Learning rate	0.001, 0.0005, 0.0002
L2 regularisation	0.0, 0.00001, 0.0001
Batch size	64, 128

AE configurations were evaluated based on validation reconstruction loss and overall reconstruction quality, measured using MSE and MAE on validation data. The final AE architecture was selected based on a combination of low reconstruction error, stable training dynamics, and consistent performance across validation samples. The resulting latent representations were subsequently used as input for SOM-based clustering and comparative segmentation analysis.

3.2.3 Deep Embedded Self-Organizing Map

The DESOM architecture proposed by Forest et al. [16] integrates an AE with a SOM to jointly learn compact and topologically ordered latent representations. The DESOM framework combines the reconstruction objective of an AE with a topology-preserving SOM loss, enabling simultaneous optimisation of representation learning and clustering in the latent space.

In this thesis, inspiration was taken from the DESOM approach by implementing a SOM layer on top of the AE architecture described in Section 3. The AE structure itself remains identical to the configuration presented earlier; however, its loss function is extended to include a SOM-based term that enforces neighbourhood consistency in the latent space. The combined objective function is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{recon}} + \gamma\, \mathcal{L}_{\text{som}}, \tag{3.10}$$

where the reconstruction loss $\mathcal{L}_{\text{recon}}$ is the MSE between input and reconstruction, and the SOM loss $\mathcal{L}_{\text{som}}$ encourages topological ordering among the latent representations $\mathbf{z}_i$ and SOM prototypes $\mathbf{m}_k$ according to:

$$\mathcal{L}_{\text{som}} = \frac{1}{n} \sum_i \sum_k \exp\left(-\frac{d_{\text{map}}(\text{BMU}_i, k)^2}{T(t)^2} \right) \|\mathbf{z}_i - \mathbf{m}_k\|^2, \tag{3.11}$$

where $d_{\text{map}}(\text{BMU}_i, k)$ denotes the distance on the SOM grid between the BMU of sample i and neuron k, and $T(t)$ represents a temperature parameter that decreases over time, controlling the influence radius of the SOM neighbourhood function.

The parameter γ acts as a weighting factor that balances reconstruction accuracy against topological regularisation. Following the implementation guidelines from the original DESOM paper, a grid search was conducted over the hyperparameter ranges summarised in Table 3.3.

Table 3.3: Grid search parameters for DESOM hyperparameter tuning.

Parameter	**Values Tested**
Weight for SOM loss γ	0.01, 0.001, 0.0001
Initial temperature T_{max}	20, 10, 6, 2
Final temperature T_{min}	0.1

This formulation allows the AE to maintain high-fidelity reconstruction while progressively encouraging neighbouring latent representations to map to neighbouring SOM neurons. Through this joint optimisation, the DESOM aims to produce a latent space that is both compact and topologically ordered, supporting interpretable and stable clustering in the subsequent analysis.

3.2.4 K-means

K-means clustering was performed using the `KMeansConstrained` class from the `k-means-constrained` library [23], which extends the classical K-means algorithm by allowing explicit constraints on cluster sizes. Classical K-means aims to partition a dataset into k clusters by minimising the within-cluster sum of squared distances:

$$\min_{C_1,\dots,C_k} \sum_{j=1}^{k} \sum_{x_i \in C_j} \|x_i - \mu_j\|^2, \tag{3.12}$$

where μ_j denotes the centroid of cluster j, computed as the mean of the data points assigned to that cluster. The algorithm iteratively alternates between assigning each data point to its nearest centroid and updating the centroids until convergence.

To mitigate the formation of extremely small or highly imbalanced clusters, which can occur in high-dimensional spaces or when latent representations exhibit varying density, the constrained variant of the algorithm was employed. `KMeansConstrained` enforces lower and upper bounds on cluster sizes:

$$\text{size}_{\min} \leq |C_j| \leq \text{size}_{\max}, \qquad j = 1, \dots, k, \tag{3.13}$$

thereby ensuring a minimum level of cluster occupancy. In this study, the bounds were defined relative to the dataset size n as

$$\text{size}_{\min} = \frac{n}{20}, \qquad \text{size}_{\max} = \frac{n}{2} + 1, \tag{3.14}$$

which prevents clusters from becoming excessively small while still allowing substantial variation in cluster sizes. The constrained assignment step is solved via an internal optimisation routine that guarantees feasibility, after which centroid updates proceed as in standard K-means.

Cluster Quality Evaluation

Cluster quality was assessed using two complementary internal validation metrics: the Silhouette Score [24] and the DBI [25]. Together, these measures capture both cluster separation and intra-cluster compactness.

The Silhouette Score quantifies the degree to which individual samples are well matched to their assigned clusters relative to other clusters. For each sample $\mathbf{x}_i$, the silhouette value is defined as

$$s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}, \tag{3.15}$$

where $a(i)$ denotes the mean intra-cluster distance between $\mathbf{x}_i$ and all other points in its assigned cluster, and $b(i)$ is the minimum mean distance from $\mathbf{x}_i$ to points in any other cluster. Silhouette values range from -1 to 1, with higher values indicating more coherent and well-separated clusters.

The overall Silhouette Score was computed as the mean of $s(i)$ across all samples. For each value of k, the average silhouette score across five independent K-means runs was used to evaluate how well the clustering captured the underlying structure of the data, both in the original feature space and in the learned SOM, DESOM, and AE-based representations.

To further assess cluster compactness and separation, the DBI was also computed. For clusters C_i and C_j, the pairwise similarity measure is defined as

$$R_{ij} = \frac{S_i + S_j}{M_{ij}}, \tag{3.16}$$

where S_i and S_j denote the average intra-cluster distances (cluster dispersions), and M_{ij} is the distance between the centroids of clusters i and j. The overall DBI is given by

$$\mathrm{DBI} = \frac{1}{k}\sum_{i=1}^{k} \max_{j \neq i} R_{ij}. \tag{3.17}$$

Lower DBI values indicate clusters that are both more compact and better separated.

3.3 Pipelines

While the conceptual motivation for these pipelines was introduced in Chapter 1, the purpose here is to describe their concrete implementation, data flow, and role in the comparative experimental design. Figure 3.1 provides a schematic overview of the four experimental pipelines and illustrates the data flow and modelling components used in each case.

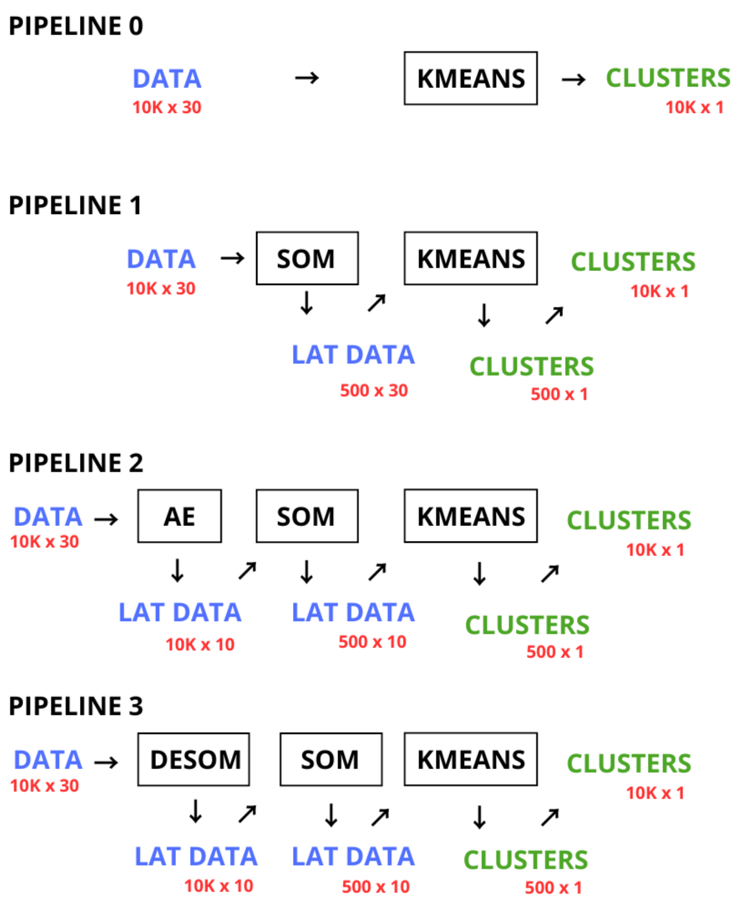

Figure 3.1: Schematic overview of the four experimental pipelines.

All pipelines are evaluated under identical conditions. The same preprocessed and standardised transactional feature set is used throughout, and clustering is performed using constrained K-means with identical cluster-size constraints, random initialisation procedures, and evaluation metrics. The pipelines differ exclusively in how intermediate representations are constructed prior to clustering.

In total, four experimental pipelines are considered. They form a progression from direct clustering in the original feature space to increasingly structured and constrained latent representations, enabling a systematic comparison of methodological complexity and segmentation quality.

1. **Pipeline 0 (Baseline: K-means)**
   Direct clustering of the standardised transactional feature vectors

using constrained K-means. This pipeline serves as a baseline to quantify the added value of topology-preserving mappings and learned representations.

2. **Pipeline 1 (SOM + K-means)**
   A topology-preserving pipeline in which a SOM is trained on the standardised feature space. Customers are represented by their corresponding SOM prototype vectors, and constrained K-means is then applied to these representations. This pipeline emphasises neighbourhood structure and feature-space interpretability through SOM-based visualisation.

3. **Pipeline 2 (AE + SOM + K-means)**
   A sequential hybrid pipeline in which an autoencoder (AE) is first trained on the standardised feature space to learn a compact latent representation. A SOM is then trained on the latent vectors, and customers are represented by the resulting SOM prototype vectors prior to constrained K-means clustering. This pipeline separates nonlinear representation learning (AE) from topology preservation (SOM).

4. **Pipeline 3 (DESOM + SOM + K-means)**
   A topology-regularised representation learning pipeline inspired by the DESOM framework, where the autoencoder is trained using a combined objective that simultaneously optimises reconstruction accuracy and SOM-based neighbourhood structure in the latent space. After training, the learned embeddings are mapped using a SOM, and constrained K-means is applied to the SOM prototype vectors. This pipeline represents the most structurally constrained approach evaluated, enforcing topological organisation during latent representation learning while retaining a consistent clustering stage across pipelines.

Chapter 4

Results and Analysis

This chapter presents the results obtained from the four clustering pipelines.

1. **Pipeline 0:** Data → K-means
2. **Pipeline 1:** Data → SOM → K-means
3. **Pipeline 2:** Data → AE → SOM → K-means
4. **Pipeline 3:** Data → DESOM → SOM → K-means

The primary objective of this chapter is to compare how each transformation step, the SOM, the AE, and the DESOM, influences cluster formation. The analysis focuses on both the robustness and interpretability of the resulting clusters.

Evaluation is performed through a combination of quantitative metrics (e.g., Silhouette Score and DBI) and qualitative assessments of the learned representations. Most importantly, the chapter investigates whether the identified clusters correspond to meaningful behavioural and financial patterns, thereby assessing whether these methods are sufficient for effective customer segmentation based on PSD2 transaction data.

4.1 Dataset

The dataset used for the main analysis consisted of 30 numerical features, of which nine represented income-related categories and 21 represented expense-related categories. All features were standardised using z-score normalisation prior to training the AE. To ensure comparable scaling across latent dimensions, the learned latent representations were additionally standardised before serving as input to the SOM.

4.2 Quantitative Cluster Evaluation

Pipeline 0 serves as the baseline model, for which constrained K-means clustering was applied directly to the standardised input features. For Pipelines 2 and 3, clustering was instead performed in the respective standardised learned representation spaces. The number of clusters was varied from $k = 2$ to $k = 15$ in order to examine a range of low- to moderately-complex segmentation solutions. Values of $k < 2$ are not meaningful for clustering, while larger values were excluded to avoid overly granular segmentations that would reduce interpretability and practical relevance in a customer segmentation context.

For each value of k, the clustering algorithm was executed five times with different random initialisations to reduce sensitivity to centroid seeding and to obtain stable performance estimates. All reported evaluation metrics therefore represent the mean across these runs.

Cluster quality was evaluated using the Silhouette Score and the DBI. The results for all four pipelines are summarised in Tables 4.1 and 4.2. To ensure comparability across pipelines, all constrained K-means models were trained using identical cluster-size constraints, requiring each cluster to contain at least five percent and no more than half of the total dataset.

4.2.1 Silhouette Scores

The Silhouette Score is well suited to this setting as it jointly reflects intra-cluster cohesion and inter-cluster separation at the observation level. This makes it particularly useful for comparing clustering performance across different values of k, where the true number of segments is not known *a priori*. High silhouette scores indicate that customers within a cluster exhibit coherent behavioural patterns while remaining clearly distinct from other clusters, a key requirement for meaningful and interpretable segmentation.

Table 4.1: Silhouette scores across all pipelines (best value per pipeline in bold)

k	Pipeline 0	Pipeline 1	Pipeline 2	Pipeline 3
2	$0.102 \pm < 0.001$	$0.130 \pm < 0.001$	0.129 ± 0.001	0.127 ± 0.002
3	$0.094 \pm < 0.001$	0.122 ± 0.001	0.161 ± 0.001	$0.194 \pm < 0.001$
4	0.087 ± 0.018	$0.139 \pm < 0.001$	$\mathbf{0.229} \pm < 0.001$	0.204 ± 0.001
5	0.116 ± 0.001	0.145 ± 0.002	$0.203 \pm < 0.001$	0.227 ± 0.001
6	0.116 ± 0.020	0.139 ± 0.025	0.206 ± 0.012	0.253 ± 0.002
7	$\mathbf{0.134} \pm \mathbf{0.013}$	0.141 ± 0.011	0.206 ± 0.001	0.271 ± 0.004
8	0.117 ± 0.017	0.148 ± 0.007	0.208 ± 0.006	0.270 ± 0.002
9	0.099 ± 0.004	0.164 ± 0.015	0.205 ± 0.009	0.250 ± 0.003
10	0.084 ± 0.008	0.178 ± 0.012	0.202 ± 0.012	0.261 ± 0.006
11	0.068 ± 0.001	0.184 ± 0.003	0.203 ± 0.006	0.268 ± 0.010
12	0.047 ± 0.004	0.190 ± 0.003	0.197 ± 0.009	0.278 ± 0.015
13	0.039 ± 0.002	$\mathbf{0.193} \pm \mathbf{0.004}$	0.198 ± 0.011	$\mathbf{0.279} \pm \mathbf{0.003}$
14	0.031 ± 0.003	0.157 ± 0.032	0.197 ± 0.007	0.274 ± 0.001
15	0.018 ± 0.003	0.136 ± 0.031	0.174 ± 0.008	0.270 ± 0.005

The results show that Pipeline 3 consistently achieves the highest silhouette scores, reaching a maximum of **0.279** at $k = 13$. Overall, the pipelines exhibit a clear performance hierarchy:

$$\text{Pipeline 0} < \text{Pipeline 1} < \text{Pipeline 2} < \text{Pipeline 3}.$$

This ordering illustrates how each additional modelling component contributes to improving the underlying cluster structure. The observed progression aligns with prior research on hybrid SOM-AE architectures, although applied here to a different dataset and problem domain.

The baseline approach (Pipeline 0) produces relatively low silhouette scores across all tested values of k, with its best performance of **0.134** at $k = 7$. Beyond this point, the scores steadily decline. This behaviour indicates that applying K-means directly to the raw feature space struggles to uncover meaningful cluster structure, likely due to noise, feature imbalance, and overlapping behavioural patterns.

Introducing a SOM layer (Pipeline 1) results in a substantial improvement. Silhouette scores are consistently higher than for the baseline, with a peak value of **0.193** at $k = 13$. This shows that the SOM effectively organises the high-dimensional input space into a structured topology, enabling K-means to form more compact and better-separated clusters.

Pipeline 2 further improves clustering performance by incorporating an AE prior to the SOM. For nearly all k, silhouette scores are higher than those of Pipeline 1. The best performance for Pipeline 2 is achieved at $k = 4$ with a silhouette score of **0.229**. These results indicate that the AE reduces redundancy and noise in the input representation, facilitating clearer and more coherent cluster formation.

Pipeline 3 performs best overall, achieving the highest silhouette scores across almost all tested values of k. Its peak score of **0.279** at $k = 13$ confirms that jointly learning the AE and SOM objectives produces a more coherent and well-differentiated latent representation for clustering. While lower values of k yield strong quantitative performance, the consistently high scores across a broad range of k suggest that the model maintains robust cluster separation even as the number of clusters increases.

Overall, the results demonstrate that increasing model complexity, first through nonlinear feature extraction with the AE and then through topology-aware structuring via the SOM-leads to more distinct and stable clusters. The standard deviations across runs remain small (± 0.000 - ± 0.03), indicating robust and reproducible clustering behaviour. Pipeline 3 consistently outperforms the other approaches, providing the most cohesive and well-separated cluster structures.

4.2.2 Davies-Bouldin Index

Table 4.2 reports the DBI values for all pipelines across $k = 2$ to 15. The DBI complements the Silhouette Score by shifting the focus from individual data points to the structure of entire clusters. Rather than averaging assignment quality across observations, the DBI evaluates how compact each cluster is relative to its separation from the most similar neighbouring cluster.

This property makes the DBI particularly valuable for identifying structural weaknesses that may be obscured by point-level averages alone. Lower DBI values indicate clusterings in which all clusters are consistently compact and well separated, reflecting a more stable and uniformly defined global cluster structure.

Table 4.2: DBI across all pipelines (lowest value per pipeline in bold)

k	Pipeline 0	Pipeline 1	Pipeline 2	Pipeline 3
2	$3.132 \pm < 0.001$	2.704 ± 0.003	2.335 ± 0.003	2.529 ± 0.014
3	2.859 ± 0.004	2.552 ± 0.001	1.855 ± 0.017	1.698 ± 0.001
4	2.627 ± 0.122	$2.265 \pm < 0.001$	1.468 ± 0.007	1.460 ± 0.005
5	$\mathbf{2.428} \pm \mathbf{0.021}$	2.237 ± 0.093	$1.588 \pm < 0.001$	1.403 ± 0.002
6	2.551 ± 0.247	2.210 ± 0.045	1.517 ± 0.047	1.281 ± 0.014
7	2.475 ± 0.204	2.040 ± 0.034	1.474 ± 0.010	1.198 ± 0.013
8	2.611 ± 0.012	1.953 ± 0.042	1.435 ± 0.010	1.209 ± 0.012
9	2.521 ± 0.047	1.930 ± 0.061	1.428 ± 0.009	1.285 ± 0.013
10	2.541 ± 0.042	1.849 ± 0.038	1.405 ± 0.030	1.229 ± 0.011
11	2.440 ± 0.011	1.850 ± 0.017	1.391 ± 0.026	1.184 ± 0.025
12	2.604 ± 0.050	1.800 ± 0.055	1.365 ± 0.037	1.156 ± 0.028
13	2.600 ± 0.046	$\mathbf{1.800} \pm \mathbf{0.027}$	1.347 ± 0.049	$\mathbf{1.139} \pm \mathbf{0.021}$
14	2.566 ± 0.065	1.809 ± 0.015	$\mathbf{1.316} \pm \mathbf{0.017}$	1.193 ± 0.016
15	2.582 ± 0.078	1.815 ± 0.044	1.385 ± 0.041	1.200 ± 0.012

The baseline approach (Pipeline 0) yields the weakest clustering performance, producing the highest DBI values across all tested values of k. Its best result of **2.428** at $k = 5$ still indicates poor cluster compactness and separation when constrained K-means is applied directly to the raw feature space. The relatively high and variable DBI values at larger k suggest that clustering in this space is highly sensitive to noise, feature correlation, and overlapping behavioural patterns. This behaviour is expected, as the dataset intentionally retains several correlated attributes, such as gambling expenditure and gambling-related income.

Introducing a SOM layer (Pipeline 1) leads to a clear and consistent improvement. By imposing a topology-preserving structure on the data, the SOM smooths the input space and reduces within-cluster dispersion, resulting in systematically lower DBI values compared to the baseline. The best performance of **1.800** at $k = 12$ indicates that organising the data into a structured latent grid enhances the robustness of the clustering process and improves cluster separation.

Combining AE-based nonlinear dimensionality reduction with SOM mapping (Pipeline 2) further strengthens cluster structure. Pipeline 2 achieves consistently lower DBI values than Pipeline 1 across nearly all values of k, reaching its best score of **1.347** at $k = 13$. This indicates that learning a low-dimensional latent representation prior to topology-aware structuring enhances both noise reduction and redundancy minimisation, yielding more compact and stable clusters.

Pipeline 3 exhibits the strongest overall clustering performance, achieving the lowest DBI values across the majority of tested values of k, with a minimum of **1.139** at $k = 13$. The joint optimisation of the AE and SOM objectives yields a richer and more coherent representation than either component alone, producing clusters that are both compact and well separated. This finding complements the silhouette score analysis, confirming that Pipeline 3 provides the most effective and stable latent space for segmentation.

Overall, the DBI results closely align with the silhouette score findings. Increasing model complexity, first through nonlinear feature extraction via the AE, and then through topology-aware structuring with the SOM, leads to progressively more distinct, stable, and well-separated clusters. While each pipeline introduces measurable improvements, Pipeline 3 remains the most robust and consistently high-performing approach across both evaluation metrics.

4.3 Component Contribution Analysis

The quantitative comparisons in Tables 4.1 and 4.2 show that Pipeline 3 achieves the strongest clustering performance across a wide range of k. However, the pipeline-level comparison does not directly quantify how much each modelling component contributes to the observed differences. A further analysis was conducted to examine the association between pipeline components and clustering quality.

For each pipeline configuration, the Silhouette Score was evaluated across multiple values of k. Since a single clustering solution must ultimately be selected, the outcome used for this analysis was the *best achieved* Silhouette Score per configuration, defined as the maximum Silhouette Score across the tested values of k.

An Ordinary Least Squares (OLS) regression model was then fitted with the best Silhouette Score as the dependent variable. The baseline was defined as constrained k-means applied directly to the input representation (Pipeline 0). Predictors included binary indicators for whether the configuration included a SOM, an AE, or a jointly trained AE+SOM (DESOM), as well as the number of input features. To examine whether associations differ with dimensionality, interaction terms between the number of features and each component indicator were included. This analysis is intended as a descriptive summary of observed trends and should be interpreted with caution given the limited number of configurations.

To account for different levels of feature granularity, the component contribution analysis was performed using multiple aggregated versions of the original feature set. The initial representation consisted of 30 transaction-based categories, including nine income-related and twenty-one expense-related features. These categories were successively aggregated into representations with 24, 18, 12, and finally 6 categories by grouping related income sources and spending types.

For each level of aggregation, the same three pipeline variants were trained and evaluated. All pipelines were trained using identical hyperparameter settings and the same grid search configurations as described in Chapter 3, ensuring consistency across experiments.

Table 4.3 reports the estimated coefficients and uncertainty intervals from the fitted OLS model.

Table 4.3: OLS regression predicting the best Silhouette Score per configuration.

Term	**Coef.**	**Std. err.**	**p-value**	**95% CI**
K-means (baseline)	0.2374	0.0390	< 0.001	[0.152, 0.323]
nr_features	-0.0037	0.0020	0.081	[-0.008, 0.001]
SOM	0.0145	0.0550	0.798	[-0.106, 0.135]
AE	0.0871	0.0550	0.142	[-0.034, 0.208]
DESOM	0.1147	0.0550	0.061	[-0.006, 0.235]
nr_features:SOM	0.0018	0.0030	0.537	[-0.004, 0.008]
nr_features:AE	-0.0031	0.0030	0.284	[-0.009, 0.003]
nr_features:DESOM	-0.0005	0.0030	0.872	[-0.007, 0.006]

The fitted model achieves $R^2 = 0.876$ (adjusted $R^2 = 0.804$) with $n = 20$ observations. While the overall fit is high, the model typically has wide confidence intervals, reflecting substantial uncertainty due to the limited sample size relative to the number of predictors. These results should be interpreted as purely descriptive of the trends observed in this specific set of experiments, rather than as statistically generalisable inferences.

The intercept term (baseline) provides the estimated Silhouette Score for the reference configuration under the coding used in the regression. The estimated coefficients for AE and DESOM are positive, with DESOM having the largest point estimate among the component indicators. The SOM coefficient is small relative to AE and DESOM. While the main effect of the number of features is negative, indicating a trend where higher dimensionality makes clustering more difficult, the interaction terms are effectively statistical noise in this sample size and do not support any conclusions about whether certain methods cope better with dimensionality than others. The primary takeaway is the consistent positive contribution of the AE-based components (AE and DESOM) to the Silhouette Score, aligning with the results from the main evaluation.

4.4 Self-Organizing Map Results

4.4.1 U-matrix Quality Evaluation

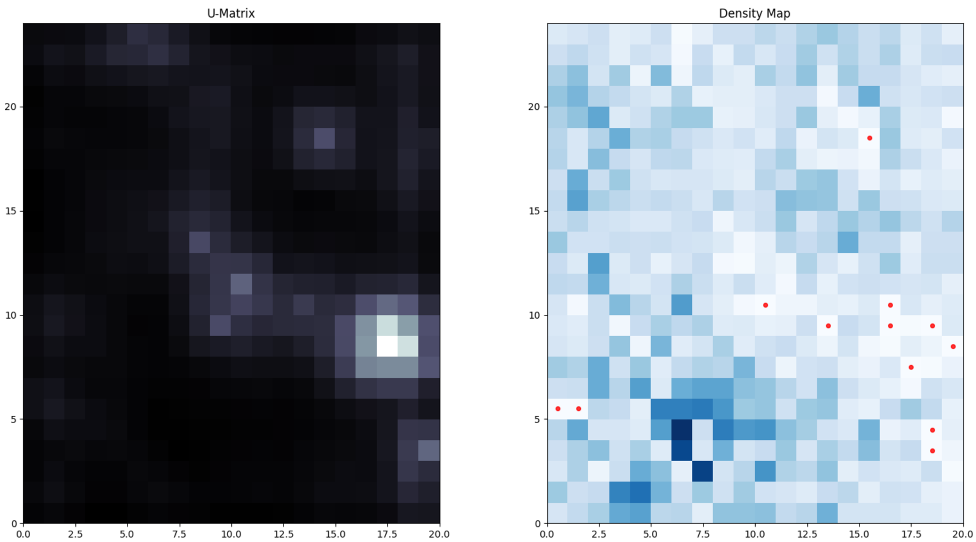

Figure 4.1: SOM visualization for Pipeline 1 showing the U-matrix and density map.

Figure 4.1 shows the U-matrix and density map for Pipeline 1. The SOM displays a coherent topological organization across the grid, indicating successful training. The U-matrix visualizes the pairwise distance between neighboring neurons: darker regions represent larger inter-neuron distances (cluster boundaries), whereas lighter regions indicate higher similarity and smoother transitions. These patterns suggest that the underlying data forms gradual, overlapping regions rather than sharply separated clusters.

The density (hit) map illustrates the distribution of BMU activations across the grid, where darker cells correspond to higher hit frequencies. The presence of only a few inactive (dead) neurons demonstrates that the data points are well distributed across the SOM, confirming efficient use of the map space.

Overall, the balanced hit distribution and coherent U-matrix indicate that the SOM for Pipeline 1 provides a faithful low-dimensional representation of the feature space, although the data structure itself does not support strong cluster separability.

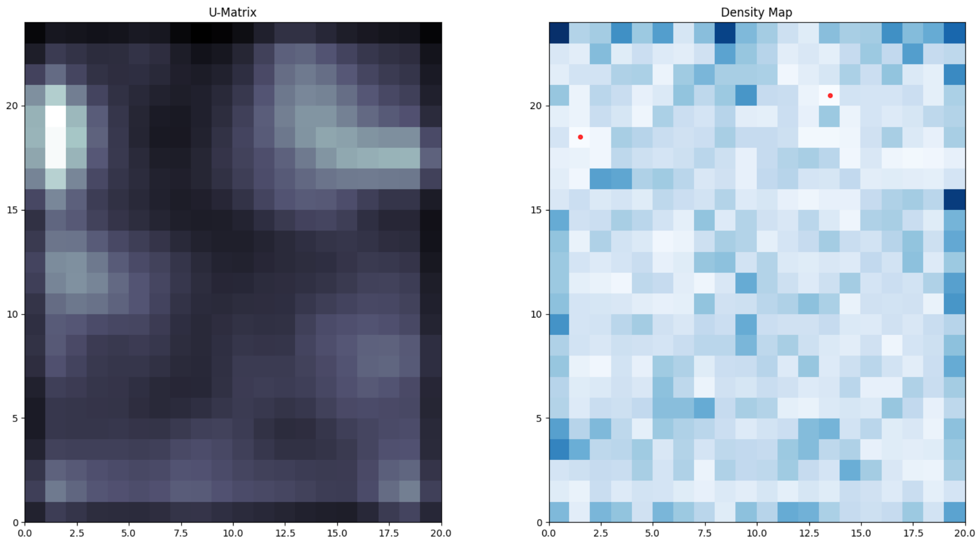

Figure 4.2: SOM visualization for Pipeline 3 showing the U-matrix and density map.

Figure 4.2 shows the U-matrix and density map for Pipeline 3. The SOM exhibits a smoother U-matrix and a more uniformly distributed density map compared to Pipeline 1. The smoother regions of the U-matrix indicate that the learned representations are more separable and topologically consistent, reflecting an improved mapping of the input space.

A visible frame pattern, with dominant neurons along the map edges, is observed in both the U-matrix and density map. This may suggest that the data distribution lies close to the boundaries of the feature space, or that the SOM grid slightly exceeds the intrinsic dimensionality of the dataset. Despite this, interior neurons remain active, indicating that the SOM has effectively captured the overall data structure.

Distinct contrast regions in the lower-right corner of the U-matrix mark clear cluster boundaries, while the smoother central regions represent continuous and homogeneous areas of the data manifold. Overall, the SOM for Pipeline 3 achieves a strong balance between cluster separability and topological fidelity, making it suitable for subsequent component-plane and feature-space analyses.

4.4.2 Quantization Error and Topographic Error

Table 4.4: SOM performance metrics across pipelines.

Pipeline	QE	TE
Pipeline 1	2.225	0.148
Pipeline 2	1.664	0.112
Pipeline 3	0.879	0.039

The quantitative metrics in Table 4.4 support the visual findings. Pipeline 3 achieves the lowest QE and TE, indicating improved mapping precision and stronger preservation of the data topology. These results confirm that the enhancements introduced in Pipeline 3 lead to a more accurate and stable representation of the input feature space.

4.4.3 Component Planes

4.4.3.1 Pipeline 1

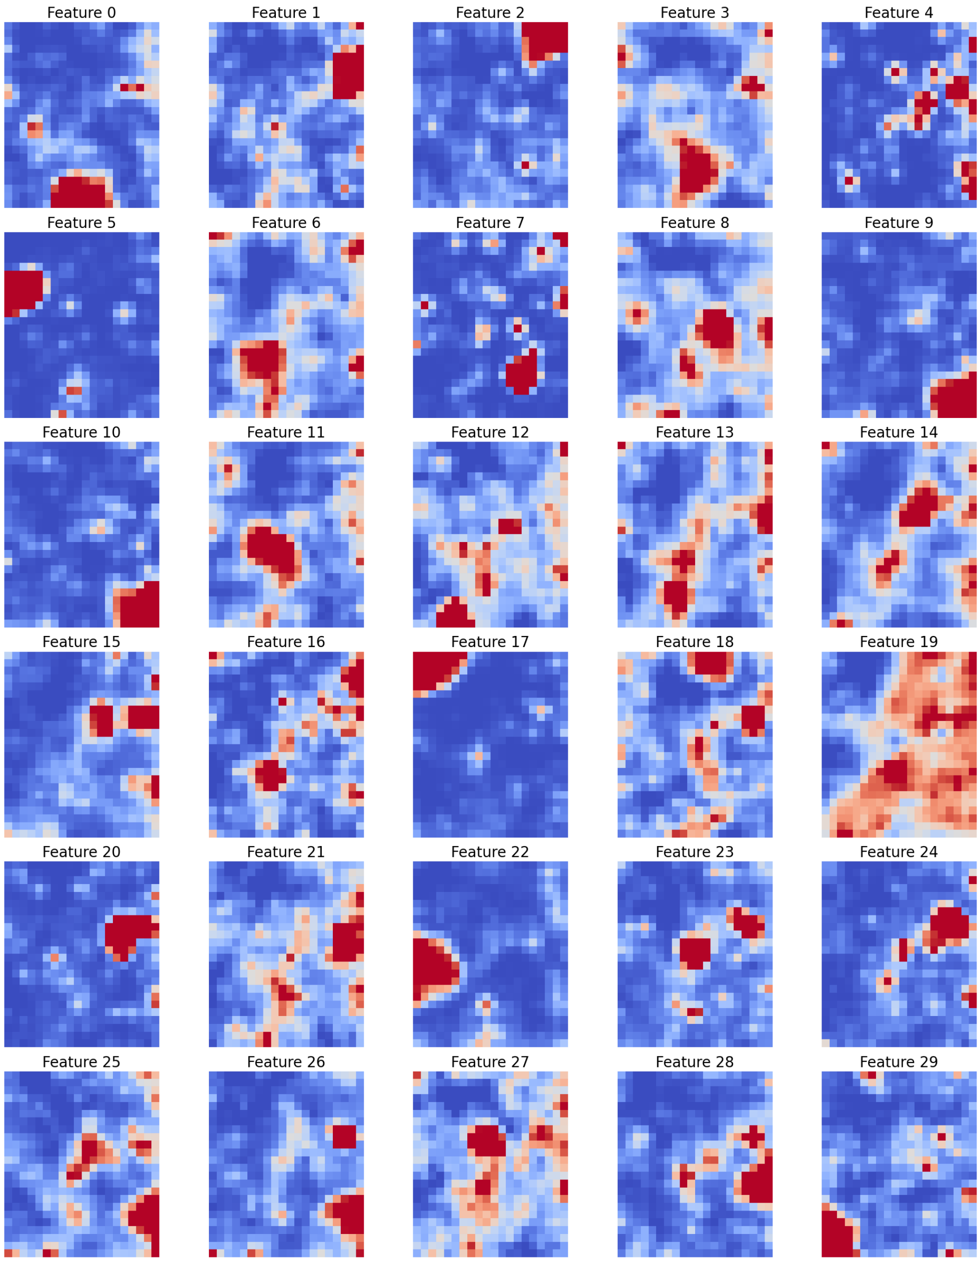

Figure 4.3: Component planes for Pipeline 1, scaled using robust scaling based on the 5th-95th percentiles.

The component planes from the SOM trained in Pipeline 1 (Figure 4.3) show clear separations for most features. Features 10 and 11, which describe behavioral spending such as spendings and winnings in online casino activities, appear in similar regions of the map, reflecting their expected correlation. Feature 20, representing salary, exhibits a more diffuse and heterogeneous pattern. Although income could in principle form the basis for separating customers into “high earners” and “low earners,” the component plane suggests that this feature varies more within clusters than across them.

In contrast, several behavioral features display much stronger and more coherent separation in the SOM. This indicates that the dominant structure captured by the SOM is driven more by behavioral differences than by income level, implying that the most meaningful distinctions between customers arise from their financial activity patterns rather than their earnings alone. Further insights from the component planes indicate that individuals receiving financial aid (features 6, 23, and 18) consistently appear in regions of the SOM associated with low or absent salary values. This pattern suggests that customers who rely on various forms of financial support generally do not exhibit regular income streams, reinforcing the interpretability of the SOM representation.

These patterns indicate that the SOM has been trained effectively and is capturing meaningful relationships in the data. Coherent spatial groupings, such as correlated behavioral features appearing in the same regions, income-related features showing expected variability, and financial-aid indicators aligning with low-salary areas, demonstrate that the map reflects underlying data structure rather than noise. This supports the validity of the SOM representation and its suitability for subsequent clustering and interpretation.

4.4.3.2 Pipeline 2

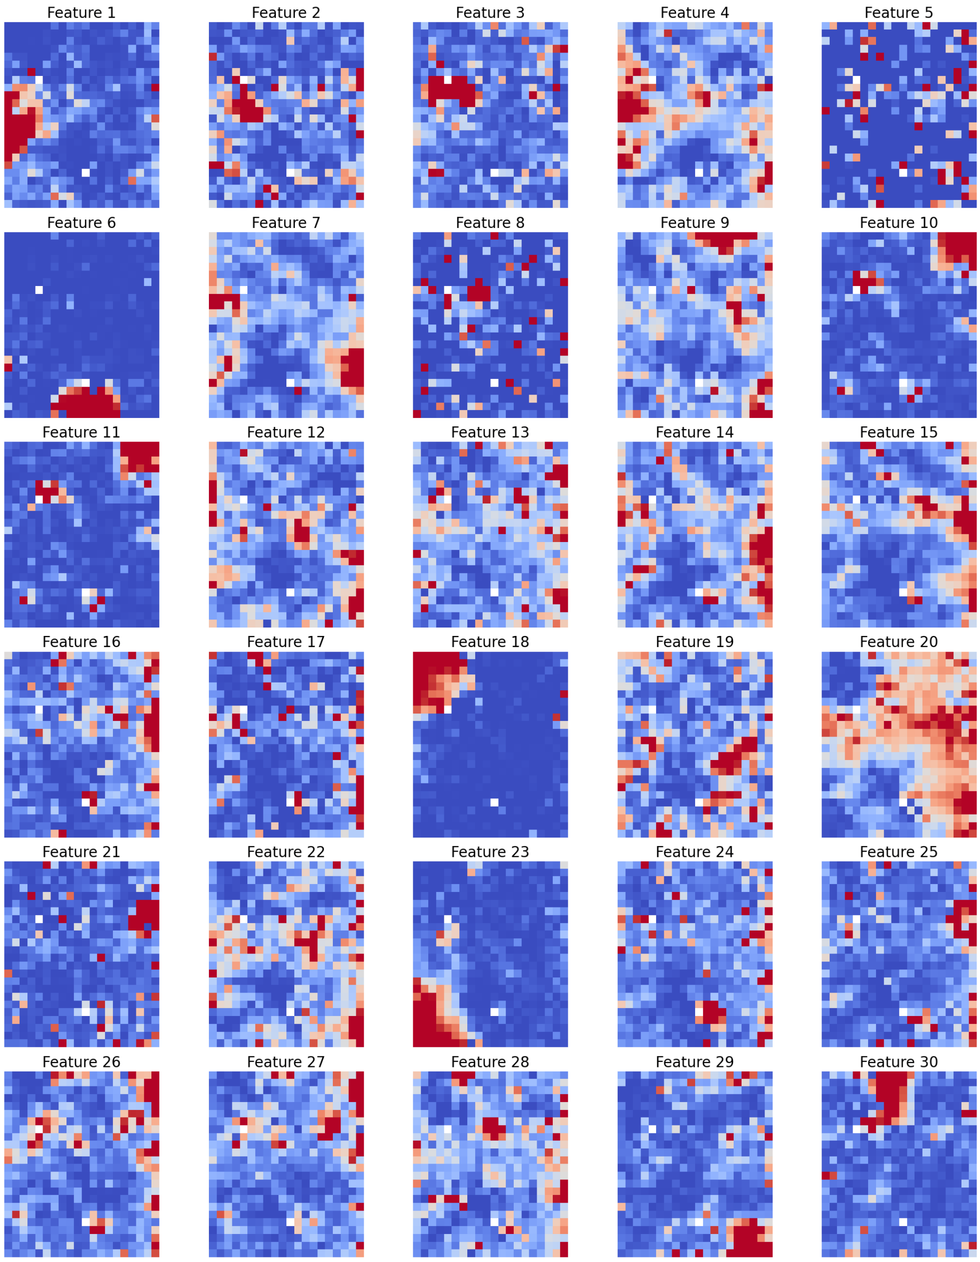

Figure 4.4: Reconstructed component planes for Pipeline 2, scaled using robust scaling based on the 5th-95th percentiles. The values represent the original features reconstructed from the SOM trained on the AE's latent space.

Figure 4.4 displays the reconstructed component planes. The SOM is trained on the latent representations and therefore only contains component planes for the ten latent features of the AE. However, to gain insight and interpretability, and to partially mitigate the black-box nature of the AE, the values for all original features corresponding to the BMU of each neuron in the SOM are displayed here.

For Pipeline 2, we can see that the SOM exhibits partial ordering for a small subset of features, specifically features 6, 10, 11, 18, 20, and 23, but overall it fails to produce strong or well-structured component planes. The strong correlation between features 10 and 11 is still preserved; however, the remaining patterns are comparatively weak and less informative than anticipated. Although some meaningful structures do emerge, they are noticeably underdeveloped. A central advantage of SOMs lies in their ability to offer intuitive and interpretable visualizations; however, in this case, the visual organization is limited.

This limitation is expected, given the high dimensionality and complexity of the dataset, which make it challenging for the SOM to enforce a coherent topological structure, particularly when it is trained only on the latent representations produced by the AE. Part of the experimental objective was to assess whether the bottleneck representation would sufficiently capture behavioral spending across all features. The resulting maps, however, suggest that this compressed representation is lacking for that purpose.

4.4.3.3 Pipeline 3

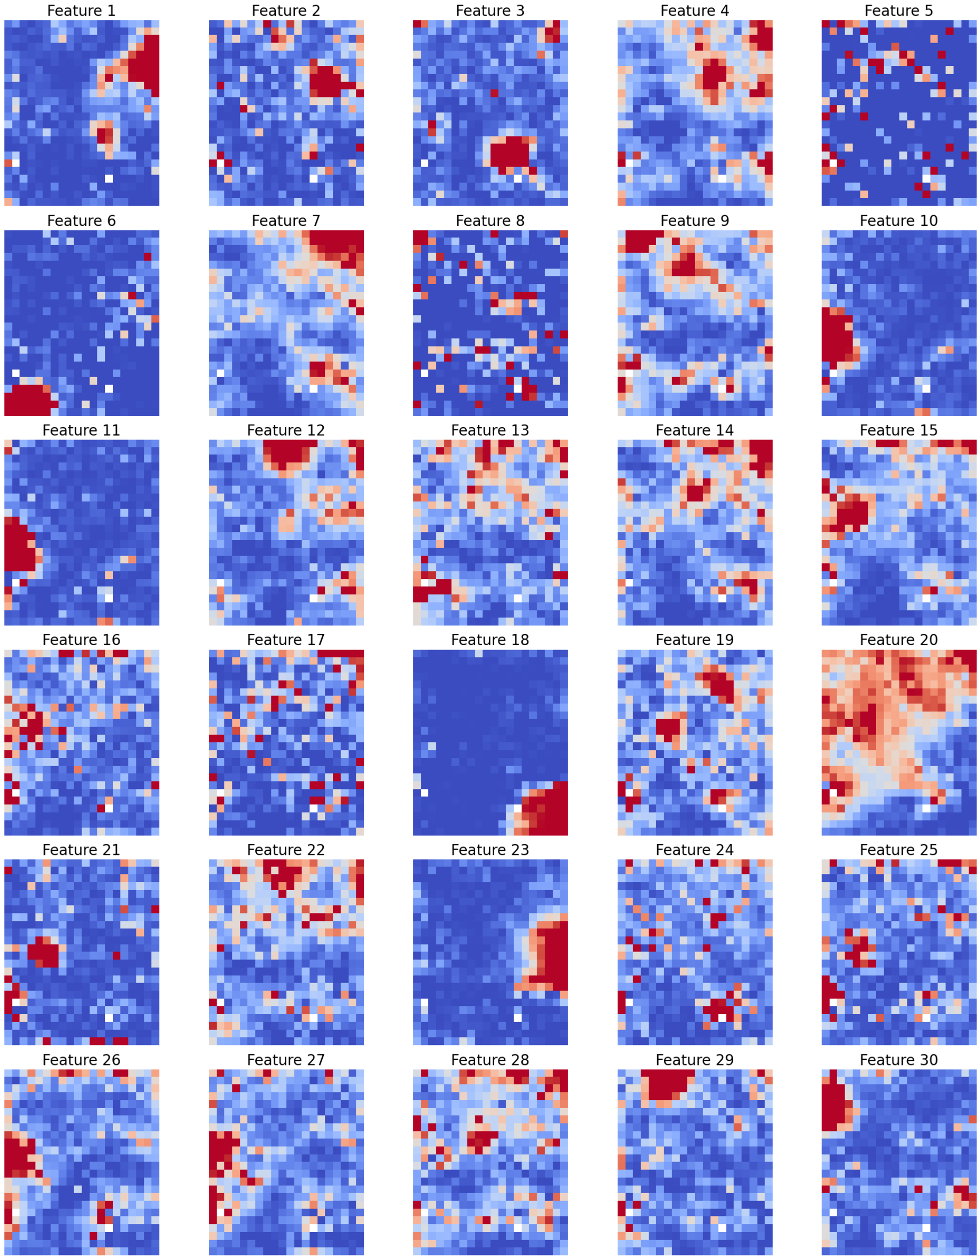

Figure 4.5: Reconstructed component planes for Pipeline 3, scaled using robust scaling based on the 5th-95th percentiles. The values represent the original features reconstructed from the SOM trained on the AE's latent space.

Lastly, Figure 4.5 shows that the component planes for Pipeline 3 are similar to those in Pipeline 2 but exhibit clearer separations for a larger number of features. This aligns with expectations, given that both pipelines use the same AE architecture, but Pipeline 3 is trained jointly with the SOM layer optimized for topology. As in the previous case, the component maps have been reconstructed to display the values of all original features corresponding to the BMU of each neuron in the SOM, thereby enhancing interpretability and mitigating the black-box nature of the model.

We know from the silhouette score that Pipeline 3 achieves the highest cluster quality among all pipelines. However, these results illustrate that Pipeline 3 prioritizes latent structural over feature-aligned visual separability: while the SOM topology has improved, the AE primarily learns nonlinear relationships between features, and these are embedded in the latent representations rather than preserved explicitly in each original feature dimension. As a result, such nonlinear interactions can be difficult to visualize in the component planes, which are inherently feature-wise projections.

This highlights a fundamental trade-off, deep learning approaches significantly improved cluster compactness and separation (as measured by internal metrics), but this came at the cost of visual interpretability in the original feature space. While Pipeline 1 offered high visual transparency, Pipeline 3 effectively traded this off for enhanced structural quality of the clusters.

4.5 Autoencoders Results

4.5.1 Autoencoders Hyperparameters

Table 4.5: AE hyperparameters used in Pipelines 2 and 3.

Parameter	Pipeline 2	Pipeline 3
Input dimensions	30	30
Encoding dimension	10	10
Hidden dimensions	[512, 256, 128]	[512, 256, 128]
Dropout rate	0.0	0.0
L2 regularisation	0.0	0.0
Learning rate	0.001	0.001
Optimizer	Adam	Adam
Batch size	128	128
Epochs	20	20
γ (DESOM)	-	0.001
t_{max} (DESOM)	-	2
t_{min} (DESOM)	-	0.1

4.5.2 Reconstruction Performance

Table 4.6 summarises the reconstruction quality achieved by both AE models on the held-out test set.

Table 4.6: Reconstruction performance of AE models in Pipelines 2 and 3.

Pipeline	MSE	MAE	Mean feature corr.
Pipeline 2 (AE → SOM)	0.1150	0.1878	0.8628
Pipeline 3 (DESOM)	0.1018	0.1729	0.8165

Because all input features were z-scored prior to training, the reported error metrics quantify reconstruction error in units of standard deviations. Both models achieve consistently low reconstruction error on the held-out test set, indicating that the AE captures the essential structure of the transactional feature space without substantial information loss. In this configuration, Pipeline 3 achieves slightly lower MSE/MAE than Pipeline 2, while Pipeline 2 shows a higher mean feature-wise correlation between original and reconstructed features. This is consistent with the models optimising slightly different objectives, where the DESOM constraints in Pipeline 3 can influence the balance between pointwise error minimisation and representation structure.

4.6 Cluster Analysis

The clustering solution was further examined by analysing the feature-wise z-scores for each cluster. This analysis enables a detailed understanding of how the groups differ in their financial behaviour. The resulting deviations provide an empirical basis for assessing whether the segmentation aligns with domain expectations and for identifying additional, non-obvious behavioural patterns. Overall, the clusters identified in Pipeline 1 can be clearly interpreted and characterised through distinct behavioural signatures, demonstrating that the algorithm has uncovered meaningful and coherent user groups rather than arbitrary subdivisions of the population.

4.6.1 Pipeline 1

For Pipeline 1, the optimal number of clusters identified using k-means was $k = 13$.

Cluster: Travel and Dining Intensive Users (3.6%)

Table 4.7: Pipeline 1: Travel and Dining Intensive Users, Key Characteristics

Feature	**z-Score**
Travel	3.14
Food and Drinks	3.12

Users in this group display distinctly elevated spending on travel and dining, reflecting an experience-oriented lifestyle where financial behaviour prioritises leisure, consumption, and other discretionary activities.

Cluster: Low-Consumption Mainstream Users (45.7%)

Table 4.8: Pipeline 1: Low-Consumption Mainstream Users, Key Characteristics

Feature	***z*-Score**
Feature 4	-0.88

This cluster shows no strong deviations across any feature category and thus represents a neutral, low-activity baseline group. As expected, a large share of the population falls into this segment, encompassing users who exhibit predictable and non-extreme financial behaviour.

Cluster: Pension-Dominant, Low-Activity Users (8.7%)

Table 4.9: Pipeline 1: Pension-Dominant, Low-Activity Users, Key Characteristics

Feature	***z*-Score**
Pension	3.31
Salary	-1.91
Shopping	-0.82
Hobbies	-0.87
Food and Drinks	-0.93

The very high pension income, combined with the strongly negative salary component, indicates users who are retired or primarily reliant on pension payments. These individuals also exhibit low engagement across discretionary financial categories, reinforcing the interpretation of a financially stable but low-activity, retirement-oriented segment.

Cluster: Student and Centrala Studiestödsnämnden (CSN)-Funded Users (5.2%)

Table 4.10: Pipeline 1: Student and CSN-Funded Users, Key Characteristics

Feature	***z*-Score**
CSN Income	3.32
Salary	-1.23

This cluster is characterised by highly elevated CSN income combined with below-average salary levels, strongly suggesting a student demographic or young adults with limited employment income.

In addition to the clusters presented above, Pipeline 1 also identified several other behaviourally meaningful groups. These include gambling-oriented users, characterised by very high gambling and speculation-related transactions often accompanied by substantial transfer activity, as well as individuals in financial distress, distinguished by elevated payments to enforcement authorities and collection agencies. Overall, Pipeline 1 produced clusters that align well with existing domain knowledge while also revealing more nuanced behavioural patterns, demonstrating both interpretability and practical relevance.

4.6.2 Pipeline 3

For Pipeline 3, the optimal number of clusters identified using k-means was also $k = 13$.

Cluster: High Discretionary and Lifestyle Spending Users (5.2%)

Table 4.11: Pipeline 3: High Discretionary and Lifestyle Spending Users, Key Characteristics

Feature	**z-Score**
Travel	3.21
Feature 12	2.74
Shopping	2.72
Food and Drinks	1.86
Salary	1.06
Loan	0.95
Hobbies	0.89

This cluster captures high-spending households with elevated discretionary consumption across multiple categories. Compared to Pipeline 1, the AE has sharpened the definition of this group by compressing moderate earners toward the mean and isolating a more distinct upper-spending segment.

Cluster: Pension-Dominant, Low-Activity Users (7.0%)

Table 4.12: Pipeline 3: Pension-Dominant, Low-Activity Users, Key Characteristics

Feature	***z*-Score**
Pension	2.76
Hobbies	-0.92
Feature 28	-0.94
Shopping	-1.07
Food and Drinks	-1.44
Salary	-1.63

Strong pension inflows combined with weak wage income and low discretionary spending identify retirees or older individuals with stable but limited financial activity. This cluster is consistently identified in both Pipeline 1 and Pipeline 3, with similar size and feature profiles.

Cluster: Gambling, Transfer, and High Financial Turnover Users (4.5%)

Table 4.13: Pipeline 3: Gambling, Transfer, and High Financial Turnover Users, Key Characteristics

Feature	***z*-Score**
Feature 9	3.31
Feature 10	3.30
Feature 15	3.20
Feature 25	3.16
Feature 27	3.15
Feature 26	3.07
Condominium Fee	2.41
Rent	2.26
Mortgage Loans	1.90
Loan	1.16
Saving	1.13
Salary	0.82
Hobbies	0.81

This cluster functions as an outlier segment, grouping individuals whose financial profiles deviate strongly from the mean. The AE appears to aggregate heterogeneous, atypical behaviours, previously distributed across several smaller clusters, into a single, high-variance group. This improves the separation between dominant behavioural archetypes and anomalous profiles.

Cluster: CSN-Funded Leisure-Oriented Users (6.0%)

Table 4.14: Pipeline 3: CSN-Funded Leisure-Oriented Users, Key Characteristics

Feature	***z*-Score**
CSN Income	3.33
Hobbies	2.64
Condominium Fee	-0.89
Mortgage Loans	-0.99
Loan	-1.22

This cluster exhibits strong CSN income combined with low borrowing, consistent with student or education-funded financial profiles. Compared to Pipeline 1, the AE has reduced within-cluster noise, yielding a more homogeneous and clearly defined student segment.

Overall, the AE in Pipeline 3 reinforces mean-centred behavioural patterns while compressing variability among typical users. This produces clearer, more cohesive mainstream clusters alongside a single, distinct outlier group. Unlike Pipeline 1, where nearly half of the population was concentrated in one baseline cluster, Pipeline 3 distributes typical users across several interpretable segments, revealing a richer and more structured representation of financial behaviour.

4.7 Summary

This chapter presented the empirical results of the four experimental pipelines. The quantitative evaluation demonstrated a clear performance hierarchy, with the jointly trained DESOM (Pipeline 3) consistently achieving the highest Silhouette Scores and lowest DBI values. The component contribution analysis indicated that while the SOM contributes to topological ordering, the inclusion of an AE, and in particular joint optimisation, is the primary driver of structural cluster quality.

Table 4.15 provides a condensed overview of the comparative performance and key characteristics for each approach.

Table 4.15: Summary of comparative results and component contribution coefficients across pipelines.

Pipeline	**Component Contribution Coef.**	**Best Sil.**	**Best DBI**
Pipeline 0	Ref.	0.134	2.43
Pipeline 1	+0.015	0.193	1.80
Pipeline 2	+0.087	0.229	1.32
Pipeline 3	+0.115	**0.279**	**1.14**

Qualitatively, the analysis highlighted a trade-off between structural coherence and visual interpretability. While Pipeline 1 provided the clearest feature-wise component planes, Pipeline 3 produced the most distinct and behaviorally homogeneous customer segments. Notably, Pipeline 3 successfully resolved the large, heterogeneous “baseline” cluster found in simpler approaches, redistributing users into meaningful groups such as high-income households, students, and consolidated outlier segments.

These findings provide the empirical basis for the methodological assessments and conclusions discussed in the final chapter.

Chapter 5

Conclusions and Future work

5.1 Conclusions

This thesis set out to evaluate whether AE-based and SOM-based approaches, implemented either sequentially or jointly, can be considered suitable methodologies for customer segmentation when applied to PSD2 transactional data. Rather than aiming to identify an optimal or universally best segmentation method, the objective was to investigate whether the evaluated pipelines are capable of generating coherent, interpretable, and reasonably stable behavioural segments under realistic constraints on data availability and modelling assumptions.

Four pipelines were considered, ranging from direct constrained K-means clustering on the standardised feature space (Pipeline 0) to a jointly trained DESOM architecture (Pipeline 3). When evaluated using internal clustering validity metrics, a consistent performance hierarchy emerged:

$$\text{Pipeline } 0 < \text{Pipeline } 1 < \text{Pipeline } 2 < \text{Pipeline } 3.$$

This ordering was further supported by the component contribution analysis presented in Chapter 4, which examined associations between key modelling components and clustering quality. The regression analysis suggested that while the SOM contributes valuable topological organisation, the inclusion of an AE is associated with substantial improvements in structural cluster quality. The strongest association was observed for joint DESOM training, indicating that simultaneously optimising reconstruction fidelity and topology preservation yields a latent representation that is particularly well suited for clustering under the evaluated metrics.

Pipeline 0 demonstrated that clustering directly in the original feature space is, in practice, insufficient for discovering well-separated behavioural groups in high-dimensional PSD2 data. Silhouette scores remained low for larger values of k, and DBI values were consistently high. This aligns with the intuition presented in the introduction: the raw feature space is noisy and redundant, limiting the effectiveness of straightforward K-means clustering even when constraints are imposed on cluster sizes.

Introducing a SOM layer (Pipeline 1) substantially improved both internal validation metrics and qualitative interpretability. The SOM component planes revealed meaningful structure in the original feature space, including correlated gambling-related features, relationships between salary and financial aid, and regions associated with distinct behavioural profiles. The resulting clusters, selected at $k = 13$, could be characterised using intuitive behavioural labels (for example, pension-dominated users, CSN-funded students, travel- and dining-oriented lifestyles, and groups exhibiting signs of financial distress). These results directly address the first two research sub-questions, indicating that a SOM-based approach can produce coherent, distinguishable segments while supporting feature-level interpretation.

The sequential AE → SOM → K-means pipeline (Pipeline 2) further improved the structural properties of the clustering. This is consistent with the component contribution analysis, which associated AE-based dimensionality reduction with higher silhouette scores. However, the SOM component planes in the latent space became less visually informative, highlighting an important trade-off: while the AE strengthens the statistical structure of the representation, it shifts part of the interpretable structure away from individual features and into higher-order nonlinear relationships.

Pipeline 3, which jointly optimises the AE and SOM objectives (DESOM), exhibited the strongest overall performance under the evaluated internal metrics. It achieved the highest silhouette scores across most values of k and the lowest DBI values, indicating improved geometric separation and compactness in the learned representation space. The component contribution analysis supports this observation, with the DESOM configuration associated with the largest positive coefficient. Clustering in the resulting latent space produced segments that were qualitatively comparable in interpretability to those obtained in Pipeline 1 and, in several cases, exhibited more behaviourally coherent aggregate profiles. For instance, the previously dominant "low-activity baseline" cluster observed in Pipeline 1 was subdivided into several more homogeneous groups in Pipeline 3, while heterogeneous outliers were consolidated into a distinct anomaly-type cluster. These findings suggest that joint AE-SOM training can improve structural clusterability while sharpening behavioural distinctions, albeit with a shift in interpretability from the original feature space toward the learned latent representation.

It is important to emphasise that optimal SOM performance in terms of low QE and TE does not necessarily imply that the resulting clustering is best aligned with the underlying analytical objective. In the context of this study, the goal is not necessarily to obtain equally sized or well-balanced clusters. Instead, a more relevant objective may be to identify one dominant cluster representing typical behaviour, alongside smaller clusters capturing deviant or anomalous patterns. This perspective aligns with application domains such as behavioural analysis in gambling or fraud detection, where rare but meaningful deviations are often of greater interest than uniform segmentation.

From this standpoint, the embeddings used in Pipeline 3, while highly optimised for SOM training, may bias the map toward compact and homogeneous clusters, potentially reducing sensitivity to outliers. Conversely, pipelines exhibiting higher error metrics may still preserve structural characteristics that facilitate the separation of atypical behaviours. In the absence of ground-truth cluster labels, it is not possible to rigorously assess which representation best captures behaviourally meaningful deviations. This limitation highlights the need for further investigation, potentially incorporating labelled data or downstream validation tasks such as anomaly detection.

Nevertheless, all evaluated approaches provide informative and viable representations of the data. The results suggest that SOMs trained on different feature representations can yield complementary insights, and that the choice of pipeline should ultimately be guided by the specific analytical objective rather than clustering quality metrics alone.

Regarding the research questions, the main conclusions can be summarised as follows:

Coherence and distinguishability (RQ1): Both the SOM-based and AE+SOM-based approaches produce coherent and distinguishable segments that consistently outperform direct K-means clustering under the evaluated metrics. Among the considered pipelines, the DESOM configuration (Pipeline 3) exhibits the strongest structural separation.

Interpretability (RQ2): Pipeline 1 provides the most direct feature-level interpretability through SOM component planes, while Pipelines 2 and 3 trade some visual clarity for improved cluster structure. Despite this trade-off, the clusters derived from Pipeline 3 can still be characterised using intuitive behavioural narratives (e.g. students, high-income households, and financially distressed users).

Stability and robustness (RQ3): Repeated runs with different random initialisations produced relatively small variations in internal metrics, and the relative ranking of pipelines remained stable. This indicates that the evaluated methods are sufficiently robust for exploratory segmentation of PSD2 data when trained with appropriate regularisation and hyperparameter choices.

Suitability for future applications (RQ4): The pipelines appear suitable as methodological building blocks for future applications. However, operational deployment would require additional validation against external criteria, such as business outcomes or regulatory constraints, as well as mechanisms for monitoring model drift and segment stability over time.

Overall, this thesis has largely met its objectives. It demonstrates that combining SOM and AE, either sequentially or jointly, constitutes a viable methodological strategy for behavioural segmentation in the PSD2 context, while also highlighting the trade-offs associated with each design choice.

5.2 Limitations

While the results demonstrate the potential of AE-SOM pipelines for behavioural customer segmentation, several limitations should be considered when interpreting the findings.

First, the absence of ground-truth labels is an inherent challenge in unsupervised segmentation. Unlike supervised learning tasks, where performance can be evaluated against a known target such as churn or default, this study relies on internal clustering metrics, including Silhouette Score and DBI, complemented by qualitative visual inspection. Although these measures provide useful indications of cluster cohesion, separation, and stability, they do not guarantee that the resulting segments are optimal or directly actionable for specific downstream business applications.

Second, customer behaviour in this study is represented as a static aggregation over a fixed six-month observation window. While this design choice supports interpretability and comparability across models, it necessarily smooths out short-term fluctuations and sequential dependencies. Financial behaviour is often dynamic and may exhibit seasonal effects or abrupt changes. As a result, the static representations used here may fail to capture important temporal patterns or rapid shifts in individual financial situations.

Third, the feature engineering process involved manual aggregation of transaction categories informed by domain knowledge. This was necessary to manage high-dimensional and sparse transactional data and to maintain interpretability. However, this aggregation may also discard finer-grained information that more expressive representation learning models could potentially exploit. In addition, the use of z-score normalisation implicitly assumes approximately Gaussian feature distributions, an assumption that does not always hold for financial transaction data, which often exhibits skewed or heavy-tailed behaviour.

Finally, while grid searches were conducted for key hyperparameters, such as the SOM neighbourhood radius and the AE architecture, the exploration of the model design space was not exhaustive. More complex architectures or alternative regularisation strategies may yield further improvements, although such extensions fall outside the scope of the present study, which prioritises methodological evaluation and interpretability over maximum performance.

5.3 Future work

This thesis was intentionally framed as a methodological investigation rather than an attempt to develop a final operational segmentation system. As a result, the findings naturally point towards several directions for future work that build directly on the design choices made in this study and on limitations identified in both the data and the existing research literature.

One important direction concerns generalisability. The behavioural representations analysed in this thesis are derived from transaction categorisation provided by Kreditz, which enables a structured and semantically meaningful feature space well suited for AE-, SOM-, and DESOM-based models. At the same time, this reliance on a specific categorisation scheme means that the learned representations and resulting clusters are partly shaped by that design choice. Applying the same pipelines to datasets constructed using alternative categorisation taxonomies, aggregation strategies, or data providers would help clarify which aspects of the observed cluster structures reflect underlying behavioural patterns and which are primarily induced by the feature construction process. Despite the widespread use of categorised transactional data in practice, such comparative analyses remain limited in the current research literature.

A second natural extension relates to the temporal nature of financial behaviour. Rather than treating customer representations as static snapshots, future work could analyse how customers transition between clusters over time by applying the trained models to rolling or sequential observation windows. Studying cluster transitions would enable the analysis of behavioural stability, gradual changes in spending patterns, and the emergence of atypical trajectories, thereby extending the framework from static segmentation towards longitudinal behavioural modelling.

Another relevant direction is to broaden the range of models considered. While this study focused on deterministic AE-based pipelines and the DESOM framework, related approaches such as the SOM-VAE model proposed by Fortuin et al. [15] integrate topology preservation with probabilistic latent representations. Such models may be better suited for capturing uncertainty and multi-modal behavioural distributions commonly observed in transactional data. A systematic comparison between deterministic and variational representation learning approaches, conducted under identical data and evaluation conditions, would therefore provide valuable insight into the trade-offs between interpretability, robustness, and representational flexibility.

Finally, future research could strengthen evaluation by incorporating limited supervised or weakly supervised signals, such as known customer outcomes or externally defined behavioural events. While the unsupervised setting adopted in this thesis is well motivated given the characteristics of PSD2 data, additional validation signals would make it possible to assess more directly whether improvements in internal clustering metrics translate into meaningful differences for downstream analytical tasks.

5.4 Reflections

This thesis sits at the intersection of machine learning and financial data analysis and therefore raises broader ethical, social, and sustainability considerations.

A central motivation of this work is the potential for behavioural segmentation to offer a fairer and more inclusive alternative to traditional demographic-based approaches. By relying exclusively on transactional behaviour rather than attributes such as age, gender, or location, the methods explored in this thesis reduce reliance on variables that are often associated with structural bias. Behaviour-based segmentation shifts the focus from assumptions about individuals to patterns derived from their actual financial activity, supporting a more transparent and evidence-based understanding of customers.

At the same time, it is important to acknowledge that certain behavioural features may still act as proxies for demographic characteristics. For example, income from student aid or pension payments is closely linked to specific life stages. While this does not eliminate the need for careful interpretation, there remains a meaningful distinction between being grouped based on observed actions (such as receiving a particular type of income) and being categorised based on inferred demographic traits. This distinction is particularly relevant in regulated financial contexts, where transparency and explainability are essential.

The clusters identified in this thesis largely align with intuitive behavioural groups, such as students, pensioners, or high-spending households, but they also reveal less obvious patterns, including segments characterised by lifestyle-driven spending such as travel and dining. These findings illustrate both the interpretability and the flexibility of the proposed framework. Different feature constructions or modelling choices could emphasise alternative aspects of behaviour, suggesting that the framework can be adapted to a range of analytical objectives.

From an economic perspective, more nuanced behavioural segmentation has the potential to support better-informed decision-making in financial services, even though such downstream impacts were not directly evaluated in this thesis. From a sustainability standpoint, the results also suggest that relatively lightweight models, such as standalone SOM-based pipelines, can achieve a favourable balance between interpretability and performance. This indicates that effective behavioural segmentation does not necessarily require large, energy-intensive deep learning models, aligning with broader principles of Green AI and responsible model development.

References

[1] W. R. Smith, "Product Differentiation and Market Segmentation as Alternative Marketing Strategies," *Journal of Marketing*, vol. 21, no. 1, pp. 3–8, Jul. 1956. doi: 10.1177/002224295602100102. [Online]. Available: https://journals.sagepub.com/doi/10.1177/002224295602100102

[2] C. Bartels, "Cluster Analysis for Customer Segmentation with Open Banking Data," in *2022 3rd Asia Service Sciences and Software Engineering Conference*. Macau Macao: ACM, Feb. 2022. doi: 10.1145/3523181.3523194. ISBN 9781450387453 pp. 87–94. [Online]. Available: https://dl.acm.org/doi/10.1145/3523181.3523194

[3] European Commission, "Directive (EU) 2015/2366 on Payment Services in the Internal Market (PSD2)," Official Journal of the European Union, Tech. Rep., 2015.

[4] European Banking Authority, "Final Guidelines on the Security Measures for Operational and Security Risks of Payment Services under PSD2," European Banking Authority, Tech. Rep., 2017.

[5] J.-T. Wei, S.-Y. Lin, and H.-H. Wu, "A review of the application of RFM model," *African Journal of Business Management December Special Review*, vol. 4, pp. 4199–4206, Jan. 2010.

[6] J. Liao, A. Jantan, Y. Ruan, and C. Zhou, "Multi-Behavior RFM Model Based on Improved SOM Neural Network Algorithm for Customer Segmentation," *IEEE Access*, vol. 10, pp. 122 501–122 512, 2022. doi: 10.1109/ACCESS.2022.3223361. [Online]. Available: https://ieeexplore.ieee.org/document/9955546/

[7] R. A. Mancisidor, M. Kampffmeyer, K. Aas, and R. Jenssen, "Learning latent representations of bank customers with the Variational

Autoencoder," *Expert Systems with Applications*, vol. 164, p. 114020, Feb. 2021. doi: 10.1016/j.eswa.2020.114020. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S0957417420307910

[8] T. Kohonen, "The self-organizing map," *Proceedings of the IEEE*, vol. 78, no. 9, pp. 1464–1480, Sep. 1990. doi: 10.1109/5.58325. [Online]. Available: http://ieeexplore.ieee.org/document/58325/

[9] S. Yanık and A. Elmorsy, "SOM approach for clustering customers using credit card transactions," *International Journal of Intelligent Computing and Cybernetics*, vol. 12, no. 3, pp. 372–388, Aug. 2019. doi: 10.1108/IJICC-11-2018-0157. [Online]. Available: http://www.emerald.com/ijicc/article/12/3/372-388/119962

[10] R. Vohra, J. Pahareeya, A. Hussain, F. Ghali, and A. Lui, "Using Self Organizing Maps and K Means Clustering Based on RFM Model for Customer Segmentation in the Online Retail Business," in *Intelligent Computing Methodologies*, D.-S. Huang and P. Premaratne, Eds. Cham: Springer International Publishing, 2020, vol. 12465, pp. 484–497. ISBN 9783030607951 9783030607968. [Online]. Available: http://link.springer.com/10.1007/978-3-030-60796-8_42

[11] C. Maçãs, E. Polisciuc, and P. Machado, "Visualization and Self-Organising Maps for the Characterisation of Bank Clients," in *Integrating Artificial Intelligence and Visualization for Visual Knowledge Discovery*, B. Kovalerchuk, K. Nazemi, R. Andonie, N. Datia, and E. Banissi, Eds. Cham: Springer International Publishing, 2022, vol. 1014, pp. 255–287. ISBN 9783030931186 9783030931193. [Online]. Available: https://link.springer.com/10.1007/978-3-030-93119-3_10

[12] W. Zhang, X. Liu, Y. Wen, S. Zhou, R. Ge, and X. Xie, "User Classification Based on SOM+K-Means Cluster Analysis," in *2023 International Conference on Applied Physics and Computing (ICAPC)*. Ottawa, ON, Canada: IEEE, Dec. 2023. doi: 10.1109/ICAPC61546.2023.00079. ISBN 9798350359244 pp. 391–394. [Online]. Available: https://ieeexplore.ieee.org/document/10497916/

[13] Master of Science in Project Management, Saint Francis College (SFC), Brooklyn, New York, USA, M. S. Khan, A. K. Gharami, Master of science in computer science, Westcliff university, USA,

F. N. Nitu, MS in Management Science & Supply Chain Management, Wichita State University, USA, M. N. Uddin, Masters of Business Administration, Major in Data Analytics, Westcliff University, USA, M. Ahmed, Master's in Public Administration, University of Dhaka, Dhaka, Bangladesh., M. K. Roy, Ms in Digital Marketing & Information Technology Management, St. Francis College, USA, S. Yezdani, and Master's in computer science, Saint Leo University, Tampa, Florida., "Deep Learning-Driven Customer Segmentation in Banking: A Comparative Analysis for Real-Time Decision Support," *International Interdisciplinary Business Economics Advancement Journal*, vol. 6, no. 08, pp. 9–22, Aug. 2025. doi: 10.55640/business/volume06issue08-02. [Online]. Available: https://www.iibajournal.org/index.php/iibeaj/article/view/71/71

[14] Z. Jiang, Y. Zheng, H. Tan, B. Tang, and H. Zhou, "Variational Deep Embedding: An Unsupervised and Generative Approach to Clustering," in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*. Melbourne, Australia: International Joint Conferences on Artificial Intelligence Organization, Aug. 2017. doi: 10.24963/ijcai.2017/273. ISBN 9780999241103 pp. 1965–1972. [Online]. Available: https://www.ijcai.org/proceedings/2017/273

[15] V. Fortuin, M. Hüser, F. Locatello, H. Strathmann, and G. Rätsch, "SOM-VAE: Interpretable Discrete Representation Learning on Time Series," 2018. [Online]. Available: https://arxiv.org/abs/1806.02199

[16] F. Forest, M. Lebbah, H. Azzag, and J. Lacaille, "Deep embedded self-organizing maps for joint representation learning and topology-preserving clustering," *Neural Computing and Applications*, vol. 33, no. 24, pp. 17 439–17 469, Dec. 2021. doi: 10.1007/s00521-021-06331-w. [Online]. Available: https://link.springer.com/10.1007/s00521-021-06331-w

[17] L. Khacef, L. Rodriguez, and B. Miramond, "Improving Self-Organizing Maps with Unsupervised Feature Extraction," in *Neural Information Processing*, H. Yang, K. Pasupa, A. C.-S. Leung, J. T. Kwok, J. H. Chan, and I. King, Eds. Cham: Springer International Publishing, 2020, vol. 12533, pp. 474–486. ISBN 9783030638320 9783030638337. [Online]. Available: https://link.springer.com/10.1007/978-3-030-63833-7_40

[18] Z. Jin, "Multi-View Clustering Model Based on Self-Supervised Comparative Learning for Bank Customer Segmentation and Risk Assessment," in *2025 6th International Conference on Computer Engineering and Application (ICCEA)*. Hangzhou, China: IEEE, Apr. 2025. doi: 10.1109/ICCEA65460.2025.11103076. ISBN 9798331543303 pp. 1335–1344. [Online]. Available: https://ieeexplore.ieee.org/document/11103076/

[19] W.-C. Liao and C.-C. Hsu, "A self-organizing map for transactional data and the related categorical domain," *Applied Soft Computing*, vol. 12, no. 10, pp. 3141–3157, Oct. 2012. doi: 10.1016/j.asoc.2012.06.014. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1568494612002906

[20] G. Vettigli, "MiniSom: minimalistic and NumPy-based implementation of the Self Organizing Map," 2018. [Online]. Available: https://github.com/JustGlowing/minisom

[21] M. Abadi and others, "TensorFlow: A system for large-scale machine learning," in *OSDI*, 2016.

[22] F. Pedregosa and others, "Scikit-learn: Machine Learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[23] J. Levy-Kramer, "k-means-constrained," 2018. [Online]. Available: https://github.com/joshlk/k-means-constrained

[24] P. J. Rousseeuw, "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis," *Journal of Computational and Applied Mathematics*, 1987.

[25] D. L. Davies and D. W. Bouldin, "A cluster separation measure," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-1, no. 2, pp. 224–227, 1979.

TRITA – XXX-EX 2025:0000
Stockholm, Sweden 2026

www.kth.se

€€€€ For DIVA €€€€

{
”Author1”: { ”Last name”: ”Jonsson”,
”First name”: ”Greta”,
”Local User Id”: ”u1XXXXXX”,
”E-mail”: ”gretaj@kth.se”,
”organisation”: {”L1”: ”School of Electrical Engineering and Computer Science”,
}
},
”Cycle”: ”2”,
”Course code”: ”DA233X”,
”Credits”: ”30”,
”Degree1”: {”Educational program”: ”Master’s Programme, Machine Learning, 120 credits”
,”programcode”: ”TMAIM”
,”Degree”: ”Degree of Master of Science in Engineering”
,”subjectArea”: ”Computer Science and Engineering”
},
”Title”: {
”Main title”: ”Evaluating Deep Topology-Preserving Models for Behavioural Customer Segmentation in Open Banking Data”,
”Subtitle”: ”A Comparative Evaluation of Autoencoder and Self-Organizing Map Hybrid Architectures”,
”Language”: ”eng” },
”Alternative title”: {
”Main title”: ”Utvärdering av djupa topologibevarande modeller för beteendebaserad kundsegmentering i Open Banking-data”,
”Subtitle”: ”En jämförande utvärdering av Autoencoder- och Self-Organizing Map-hybridarkitekturer”,
”Language”: ”swe”
},
”Supervisor1”: { ”Last name”: ”Blom”,
”First name”: ”Ebba”,
”E-mail”: ”ebba@kreditz.com”,
”Other organisation”: ”Kreditz AB”
},
”Supervisor2”: { ”Last name”: ”Székely”,
”First name”: ”Éva”,
”Local User Id”: ”u1XXXXXX”,
”E-mail”: ”szekely@kth.se”,
”organisation”: {”L1”: ”School of Electrical Engineering and Computer Science”,
”L2”: ”XXX” }
},
”Examiner1”: { ”Last name”: ”Payberah”,
”First name”: ”Amir H.”,
”Local User Id”: ”u1XXXXXX”,
”E-mail”: ”payberah@kth.se”,
”organisation”: {”L1”: ”School of Electrical Engineering and Computer Science”,
”L2”: ”XXX” }
},
”Cooperation”: { ”Partner_name”: ”Kreditz AB”},
"National Subject Categories": "10210",
"SDGs": "5, 9, 10",
”Other information”: {”Year”: ”2026”, ”Number of pages”: ”xiv,66”},
”Copyrightleft”: ”copyright”,
”Series”: { ”Title of series”: ”TRITA – XXX-EX” , ”No. in series”: ”2025:0000” },
”Opponents”: { ”Name”: ”S. Ribaric”},
”Presentation”: { ”Date”: ”2026-01-29 13:00”
,”Language”:”XXX”
,”Room”: ”via Zoom https://kth-se.zoom.us/my/payberah”
,”City”: ”Stockholm” },
”Number of lang instances”: ”2”,
”Abstract[eng]”: €€€€
1abstracts ”Abstract[swe]”: €€€€
2abstracts €€€€,
”Keywords[eng]”: €€€€
Self-Organizing Maps, Clustering, Visualisation, Deep learning, Autoencoders, Open banking, Customer Segmentation €€€€,
€€€€,
”Keywords[swe]”: €€€€
Självorganiserande Kartor, Klustring, Visualisering, Djupinlärning, Autoencoder, Open Bankning, Kundsegmentering €€€€,
}