

Original Article

Enhancing Customer Segmentation Granularity with DEA-Driven Federated Convolutional Autoencoders: A Deep Learning Approach

Pentareddy Ashalatha¹, G. Krishna Mohan²

¹Department of Computer Science & Applications, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur, Andhra Pradesh, India.

²Department of Computer Science & Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur, Andhra Pradesh, India.

¹Corresponding Author : ashalathacambam@yahoo.co.in

Received: 06 July 2025

Revised: 18 August 2025

Accepted: 03 September 2025

Published: 30 September 2025

Abstract - In the fast-changing environment of data-driven business models, successful customer segmentation is critical for targeted marketing, resource allocation, and sustainable customer retention. Conventional segmentation techniques are based on centralized designs and simple clustering algorithms, which are not capable of coping with the high-dimensional, distributed, and privacy-constrained nature of contemporary customer data. This research overcomes these limitations by proposing a new framework that combines Federated Convolutional Autoencoders (FCA) with Data Envelopment Analysis (DEA) to improve segmentation granularity and interpretability in a privacy-friendly way. The main goal of this research is to propose a decentralized, smart segmentation model that not only discovers hidden behavioral patterns in distributed datasets but also analyzes the operational efficiency of each customer segment for strategic decision-making. The suggested FCA model learns compressed, non-redundant feature representations from local nodes with unsupervised convolutional autoencoders and federates them through learning protocols without sharing raw data, and DEA is subsequently used to quantify segment efficiency based on crucial input-output parameters. The model is tested with a synthetic customer dataset that mimics multi-regional consumer behavior, and the implementation is performed using TensorFlow and PyTorch for deep learning and Python-based linear programming software for DEA. Experimental results of 95% accuracy show that the proposed method generates more sophisticated and strategically useful segments than traditional methods. In summary, the combination of FCA and DEA provides a scalable, interpretable, and privacy-preserving answer to customer segmentation, providing the foundation for smart, actionable, and secure customer analytics in distributed data settings.

Keywords - Federated Learning, Convolutional Autoencoder, Customer Segmentation, Data Envelopment Analysis, Distributed Deep Learning.

1. Introduction

In the hyper-competitive digital economy of today, companies are more dependent than ever on data-driven initiatives to know, keep, and expand their customers [1]. With the amount and intricacy of customer data growing exponentially across multiple transactional, behavioral, and demographic sources, there has been an increased need for more detailed and smarter customer segmentation techniques [2]. Conventional customer segmentation methodologies, typically based on statistical or rule-based clustering techniques such as K-means or decision trees, are insufficient to identify consumer behaviour's complex and multifaceted nature in contemporary commerce [3]. These traditional models generally demand central access to data, presume uniform customer behavior patterns, and find it difficult to

handle high-dimensional or noisy data, and therefore are less compatible with dynamic and distributed settings like global retail chains, multi-regional service providers, and e-commerce sites [4]. Recent efforts have brought in deep learning-based remedies, especially unsupervised representation learning techniques such as autoencoders, in order to learn better latent patterns in consumer data [5]. Although these approaches represent a much-needed step forward in modeling complexity and abstraction, they continue to suffer from limitations of data centralization, scalability, and explainability [6]. Furthermore, centralized learning models also present severe privacy and security threats, especially in domains that are regulated by strict data protection laws [7]. This situation has introduced FL as a potential solution, where local devices or nodes jointly train



models without exchanging raw data, thereby maintaining privacy while tapping into distributed intelligence [8]. In this regard, FCA present an innovative solution, where the representational capability of convolutional autoencoders is coupled with the privacy-preserving features of federated learning [9]. Still, even sophisticated federated models experience challenges regarding the quality and relevance of local data contributions, which may result in less-than-optimal global models due to noisy, redundant, or non-representative local data. At the same time, there is still a chronic gap in interpreting and assessing the practical value or efficiency of derived customer segments since most deep learning models are black boxes that provide little insight into the real utility of various customer groups from a business point of view. This lack of transparency presents a key bottleneck in matching machine learning outputs with actionable business strategies [10].

To overcome these weaknesses, this research suggests an integrated approach that combines the representational power of Federated Convolutional Autoencoders with the decision-analytic discipline of Data Envelopment Analysis (DEA). The fundamental driver of this hybrid design is to improve the granularity and accuracy of customer segmentation in distributed data environments, and to compare the relative effectiveness of customer segments in a way that is both understandable and actionable to stakeholders. By locally training convolutional autoencoders on every client node, such as regional shopping outlets or cloud-segregated departments, the model learns the spatial and temporal complexities of customer behavior while respecting data privacy standards. Autoencoders, when federated through periodic weight sharing and aggregation (using methods like FedAvg), enable the development of a global model that captures heterogeneous behavioral patterns without divulging sensitive information. Convolutional layers focus on detecting sequence dependencies and space hierarchies in customer purchases, while autoencoder compression and reconstruction processes make it possible to extract non-redundant, high-valued feature representations.

The value of this research resides in its complete approach to resolving long-standing challenges in customer segmentation: lack of granularity, no privacy compliance, and poor interpretability. Unlike traditional models that either sacrifice data privacy or form coarse clusters that are inadequate for focused marketing strategies, the developed FCA-DEA framework offers a secure, comprehensive, and interpretable segmentation pipeline. Additionally, the federated feature promotes scalability to large networks of decentralized data sources without costly data centralization or anonymization processes. Another significant contribution of this work is its novelty in methodology - by closely integrating an unsupervised deep learning model with an established operations research tool, it forms a cross-disciplinary methodology that borrows strengths from

artificial intelligence and decision science. Not only does it enhance the pragmatism of the model outputs, but it also keeps technical innovation attuned to marketing strategists' needs, customer relationship managers' needs, and policy makers' needs.

Additionally, the system to be suggested is flexible to a diverse range of applications such as omnichannel commerce, banking services, urban IoT applications, and telemedicine services, where customer behavior is complex and sensitive. Designing the system accommodates ongoing learning and model fine-tuning via successive federated training loops, thereby ensuring that customer segments dynamically update in line with changing trend behaviors. The interpretive platform provided by DEA allows for instantaneous measurement of the performance of different segments, in turn enabling the dynamic allocation of resources, customizing the offerings, and predicting the lifetime value of prospects with higher efficacy. The built-in architecture is therefore not used alone as a technique but as a valuable asset that facilitates marketing planning, operational effectiveness, and customer strategy in a holistic and compliant context.

In summary, this work is driven by the pressing necessity of balancing deep learning innovation with practical limitations such as data privacy, distributed architectures, and deployable intelligence. Although deep learning models have exhibited superior performance in feature learning and pattern detection, they tend to be lacking in applications involving transparency, efficiency assessment, and deployment over decentralized networks. The envisioned FCA-DEA framework overcomes these limitations via a new fusion of federated deep learning and efficiency analytics, setting the stage for a new generation of smart customer segmenting systems. By overcoming the shortcomings of legacy clustering techniques, centralized architectures, and black-box models, this research hopes to establish a paradigm in customer intelligence systems that are accurate, secure, explainable, and strategically aligned with organizational objectives. The framework's scalability across various industries, its continued compliance with changing data protection standards, and the provision of valuable insights into consumer behavior make it a visionary contribution to academic and business literature. This research is delivered by leveraging the potential of sophisticated machine learning in a format that is ethical, transparent, and most beneficial to the decision-making processes in customer-focused organizations. The key contribution of the study is as follows:

- Built a federated deep learning model based on convolutional autoencoders to learn rich behavioral features from scattered customer data without sacrificing data privacy.
- Instituted the application of Data Envelopment Analysis to compare the operational efficiency of every customer group, allowing for interpretability and targeted strategy.

- Recommended a new end-to-end segmentation framework that keeps data local at client nodes, thereby meeting high data protection requirements.
- Accomplished customer groupings at fine-grained levels by taking advantage of deep, unsupervised learning and efficiency-driven validation.
- Illustrated the integration of operations research and deep learning by putting together FCA's pattern learning with DEA's performance measurement.

The order of the paper is as follows: Related works are included in Section 2. Section 3 discusses problems with the existing study. The proposed methodology is provided in Section 4. Section 5 presents the results and an overview. Conclusions and future works are covered in Section 6.

2. Literature Review

Le et al. [11] paper presents PersonalFR, a personalized federated recommender system that utilizes autoencoder-based models to learn user-item interactions without revealing raw data. The local updating of encoders and aggregation of decoders on the server side maintains data privacy. The experiments on real-world datasets show that PersonalFR offers performance close to centralized models and with lower communication overhead. The research is based mainly on the recommendation system and does not discuss its applicability in customer segmentation tasks. Furthermore, the complexity of the model might make it difficult to deploy in resource-poor environments. The non-interpretability of the autoencoder's latent features also restricts the explainability of the segmentation outcomes.

John, Shobayo, and Ogunleye [12] research compares different clustering algorithms - K-means, Gaussian Mixture Models (GMM), DBSCAN, Agglomerative Clustering, and BIRCH - on a UK-based online retail dataset with more than 500,000 records. Based on the Recency, Frequency, Monetary (RFM) framework, the study seeks to improve decision-making in the retail industry. The GMM algorithm produced the best Silhouette Score of 0.80, which reflects better clustering performance. Nonetheless, the use of centralized data processing in the study creates data privacy and scalability issues. In addition, the lack of sophisticated feature extraction methods could constrain the level of customer granularity during segmentation. The study does not provide solutions to the data heterogeneity and privacy preservation issues characteristic of distributed data settings.

Alves Gomes and Meisen's [13] review evaluates 105 publications between 2000 and 2022 for customer segmentation methodologies in e-commerce. The work recognizes a four-stage process: data collection, customer representation, segmentation, and targeting. The review indicates widespread use of RFM analysis and k-means clustering, especially their ubiquity irrespective of the size of

datasets. The integration of sophisticated approaches such as PCA and SOMs is also described in the review. Nonetheless, it observes that most research depends on hand-crafted feature selection and has no automation of feature engineering, and thus, there is a need for more advanced, automated methods in customer segmentation.

Wang [14] investigates the coupling of deep learning models with swarm intelligence algorithms for customer segmentation in online marketing. The synergy seeks to maximize clustering performance through the global search ability of swarm intelligence. The experiments reveal enhanced segmentation accuracy and marketing effectiveness. Nevertheless, the centrality of data processing creates privacy vulnerabilities. The research does not touch on the issues of data heterogeneity in distributed environments. In addition, the interpretability of the resulting customer segments is constrained, and so is integration with efficiency assessment frameworks.

Sharma, Patel, and Gupta [15] study investigates the convergence of deep learning methods with predictive analytics to enhance customer segmentation techniques. Through the use of neural networks to examine intricate customer data, the study seeks to reveal subtle behavioral patterns that may be missed by conventional methods. The method highlights the capability of deep learning to detect non-linear patterns in customer data, resulting in more precise and dynamic segmentation. Nevertheless, the research admits difficulties in model interpretability and the necessity for large computational resources, which could restrict its usage for smaller businesses or companies with limited technical infrastructure.

3. Problem Statement

Looking to provide tailored marketing, enhance customer retention, and increase decision-making effectiveness. Traditional segmentation methods, however, tend to fail to identify the subtle behavioral trends of customers, particularly when data is dispersed across different organizations or geographic locations due to privacy restrictions [16]. Conventional clustering algorithms and superficial machine learning algorithms tend to depend on centralized data processing, which poses serious data privacy issues and restricts the level of insight owing to their deficiency in learning complicated, hierarchical feature representations [17]. In addition, the inability to assess the efficacy of resultant customer groups holds back targeted strategy and resource utilization. As markets become more competitive, there is a clear need for a strong, privacy-preserving, and scalable segmentation system that can learn from distributed data and extract more profound behavioral features. The system also needs to be able to verify the quality of segments using objective metrics for efficiency to enable evidence-based decision-making. Solving these issues requires a new solution that combines federated learning with state-of-the-art deep

learning models and post-hoc efficiency analysis tools, enabling better, safer, and more actionable customer segmentation solutions for today's data-driven businesses.

4. Materials and Methods

In this research, a new customer segmentation approach, combining FCA and DEA, is proposed. This approach utilises federated learning to preserve data privacy over distributed customer data sources while learning intricate behavioral patterns. The FCA model is locally trained at different client nodes, e.g., retail outlets, where it learns customer behavior

attributes such as purchase frequency, product affinity, and expenditure patterns. After customer segments have been created, DEA is utilized as a post-processing method in order to judge the efficiency of these segments to identify the highest profit and loyalty groups. It integrates FCA's ability for deep learning together with DEA's efficiency analysis for providing a comprehensive, privacy-enabled platform for effective, actionable customer segmentation that will be applicable across business intelligence, marketing, and customer relationship management. Figure 1 represents the overall method architecture of the proposed study.




Fig. 1 Overall method architecture

4.1. Dataset Description

The dataset employed here is the Online Retail Dataset of the UCI Machine Learning Repository. The dataset is based on transaction data of a UK-based and registered non-store online retail business selling primarily distinctive gift items. Data ranges from December 1, 2010, through December 9, 2011, with a total of 541,909 records in 8 attributes.

Every record indicates a single transaction and has the following fields:

- InvoiceNo: Transaction identifier.
- StockCode: Product code of the purchased item.
- Description: Short description of the item.
- Quantity: Quantity purchased per transaction.
- InvoiceDate: Time and date of the transaction.
- UnitPrice: Unit price per item (in GBP).
- CustomerID: Unique customer identifier.
- Country: Customer country of residence.

The dataset is extremely suitable for this research as it offers rich combinations of temporal, behavioral, and geographical characteristics that facilitate exhaustive analysis of the purchasing behavior of customers. The diversity in its purchase patterns, frequency of orders, and international customer base ensures a thorough assessment of segmentation

accuracy and granularity. In addition, anonymized customer identifiers enable the simulation of heterogeneous, decentralized data environments, which is imperative for evaluating segmentation performance in distributed and privacy-constrained datasets [18].

4.2. Data Preprocessing

The process of data preprocessing is very crucial when it comes to conducting proper segmentation analysis with the use of data. Data pre-processing became critical while converting raw transactional records into formal and significant patterns of customers' activities. Due to the complex and large scale of customers' information, it was necessary to make some changes to increase the quality, comparability, and usability. It allowed the determination of important trends, filtered out random variations in contrast to the true patterns, and made the extracted data relevant for model training and assessment. This preparation was of significant utility for assessing the fineness of the segmentation and for the stability of the latter learning stages.

4.2.1. Data Cleaning

Data preprocessing began by eliminating records that were missing or incomplete. More precisely, records with missing valid 'CustomerID' values were removed because customer-level examination is at the heart of segmentation. Also excluded were all non-positive 'Quantity' and

'UnitPrice' transactions, since these usually involve product returns or data entry mishaps that will skew behavioral statistics. This cleaning step guarantees that there are only significant and valid records of purchases that are used to build customer profiles, hence improving the accuracy of resulting segmentations and avoiding noise in later feature extraction and learning processes.

4.2.2. Feature Engineering

Feature engineering was applied to derive relevant behavioral patterns at the customer level. Important features were Recency, computed as the days since the last purchase of a customer, Frequency, which is the count of all transactions, and Monetary Value, the sum of all the purchase values. Other features, like Average Basket Size and Purchase Interval, were derived to represent more subtle behavior.

4.2.3. Invalid and Duplicate Records

To maintain data integrity, all redundant transactions were detected and eliminated. This meant duplicate entries that would skew the frequency or monetary calculations. On top of this, canceled transactions - usually recognizable by invoice numbers starting with the letter 'C' - were not included in the data set. Canceled entries do not represent true customer behavior and would bring inconsistencies in recency or purchase interval computations. Removing such anomalies from the data maintains the integrity of customer profiles and prevents segmentation results from being biased by administrative or transactional artifacts.

Normalization

Normalization was used to normalize all numeric features to a common range, which is important when models are sensitive to feature magnitude. This process ensures that features such as monetary value do not disproportionately affect distance-based or gradient-based learning processes. It is calculated in the Equation (1).

$$X_{scaled} = \frac{X_{max} - X_{min}}{X - X_{min}} \quad (1)$$

4.2.4. Train-Test Split

After preprocessing and feature creation, the data was split into training and test subsets with an 80-20 split at the customer level. This ensures that customers are allocated individually to the training or test set without data leakage. The training set (80%) was employed to learn customer representations and obtain segmentation structures, and the test set (20%) was kept aside to test segmentation quality and generalizability. Such separation is crucial for testing the robustness of the segmentation approach under realistic deployment conditions, especially when dealing with unseen or partially known customer behavior patterns.

Deep Learning Methods

The study discusses deep learning techniques used in customer segmentation, specifically FCA combined with

DEA. FCA facilitates privacy-preserving, distributed learning, while DEA improves segmentation granularity by analyzing efficiency scores, both constituting a solid framework for accurate and scalable customer behavior analysis.

Federated Convolutional Autoencoder (FCA)

In this study, the FCA is considered the building block of Federated Learning to derive complex patterns from spread customer data samples. Peculiarly, the FCA belongs to the federated learning family that does not keep customer data in the center; rather, different segments of data stay at the nodes. Every local node learns independently a convolutional auto-encoder on the local customer data to obtain the important customer behavior patterns including purchase frequency, product choice behavior and spending propensity. The convolutional layers are beneficial to obtain features that relate to the spatial or sequential pattern present in the customer transaction data, and the autoencoder layer is used for unsupervised learning and feature extraction.

In the downstream layers, only the model updates, as opposed to the raw data, are sent back to a central server, followed by federated averaging. The global model, in turn, is enhanced by the intelligence of all nodes with the same importance, and at the same time, it is secure in terms of data privacy. The extracted features from the FCA are generalizable, which facilitates better customer clustering and segmentation.

Through identifying the behavior features within the distributed sources, the FCA plays a crucial role in improving the level of segmentation detail. When integrated with the DEA-based evaluation, this federated deep learning architecture enhances the prospects of accurate, private, and efficient customer segmentation in different and dispersed domains.

In order to perform privacy preservation and establish efficient customer segmentation, Figure 2 sketches out a model of the FCA. Each tensor node, say, a retail branch or a regional client, possesses customer data that is private and kept securely to meet the data protection regulations. These nodes perform the learning of their own FCA models here, which is a deep learning structure that can decipher specific customer behaviors such as their buying habits, preferred products, and buying rates. Essentially, instead of transmitting raw data, only the weights learned in the model of the nodes are transmitted to a centralized server. It adopts FedAvg aggregation as the key federated learning process, which amalgamates the local model updates into one global model. It allows all the nodes to interact, learning from each other while at the same time maintaining the privacy of each dataset. After these processes, the updated global model is delivered to each node so that it can be modified further in the next passing of the training phase.




Fig. 2 FCA model architecture

It goes round and round until the convergence is achieved, and then each node comes up with a consolidated FCA model that depicts the customer segments at the global level. It is especially effective when there is a need for many segments and strict privacy requirements, and the proposed framework is useful for business intelligence, marketing, and customer relationship management applications in distributed settings. Convolutional autoencoder forward pass calculated in Equation (2).

$$\hat{x}_i = D(E(x_i)), L_{recon} = \|x_i - \hat{x}_i\|^2 \quad (2)$$

Where, x_i original input data from client I, $E(.)$ Encoder network (convolutional layers), $D(.)$ Decoder network (deconvolutional layers), $\hat{x}_i$ deconstructed output, L_{recon} Mean Squared Error (MSE) loss for reconstruction. The local model update is calculated using the Equation (3).

$$w_i^{(t+1)} = w_i^{(t)} - \eta \cdot \nabla L_{recon}(w_i^{(t)}) \quad (3)$$

Where, $w_i^{(t+1)}$ Model parameters at client i at round t, η Learning rate, ∇L_{recon} Gradient of the reconstruction loss. FedAvg for global model aggregation calculated in Equation (4).

$$w^{(t+1)} = \sum_{i=1}^N \frac{n_i}{n} w_i^{t+1} \quad (4)$$

Data Envelopment Analysis

DEA can be used efficiently as a post-processing tool to analyze the relative efficiency of segmented customer groups across various client nodes. DEA is a non-parametric linear programming technique applied to measure the performance

of Decision-Making Units (DMUs) by comparing several input and output variables. In the present study, every customer cluster derived via the FCA model on local clients can be considered a DMU. Inputs can include customer-related features like transaction frequency, purchase value, or engagement levels, while outputs can include profitability, loyalty score, or conversion rates. Applying DEA to such FCA-derived segments, one is able to determine which groups of customers are operating efficiently and where strategic intervention is needed. By embedding DEA within a federated learning framework, this efficiency evaluation is carried out without breaching the confidentiality of the data of individual clients.

Figure 3 shows the framework of the DEA method. The diagram classifies three fundamental methodologies: DDEA, DBPI, and WDEA. These methodologies are used to analyze the productivity and efficiency of DMUs over a period of time, so the diagram is significantly applicable for longitudinal research. DDEA takes inter-temporal linkages and carry-over activities across periods into account, whereas DBPI facilitates measurement of changes in productivity by combining Malmquist indices or comparable constructs.

WDEA, however, constructs overlapping windows of data to measure performance trends with more stability. This systematic classification facilitates choosing the right DEA model depending on data availability and research goals. To further improve the diagram, you might include concise descriptors below each category to clarify its specific analytical purpose, and optionally add examples of applications or associated tools. The relative efficiency of DMU is given in the Equation (5).




Fig. 3 DEA workflow

$$Efficiency = \frac{\sum_{r=1}^S u_r y_{rj}}{\sum_{i=1}^m v_i x_{ij}} \quad (5)$$

5. Results and Discussion

Based on the presented evidence of the theoretical and computation-based model of combining Federated Convolutional Autoencoder (FCA) with the Data Envelopment Analysis (DEA), this paper has highlighted how advanced customer segmentation and efficiency evaluation can be efficiently set up. A central achievement of the FCA component was thus the ability to make relevant and non-redundant feature extractions from distributed customer data across multiple client nodes while preserving privacy. These extracted features could facilitate the identification of more detailed and fine-grained customer segments using the proposed model. Further, after segmentation was done, the DEA component was used to examine the efficiency of the

customer groups based on multiple input-output variables. That is why the DEA model managed to determine efficient and inefficient segments and chose effective marketing strategies. From the results presented, it is clear that the integration of deep learning with efficiency analysis provided reliable and explainable insights about the model, indicating the applicability of the model to the real-world retail dataset. In general, the model gave satisfactory results when it comes to privacy preservation and decision-making.

5.1. Experimental Outcome

The experimental result verifies the model's better performance in customer segmentation tasks. With high accuracy, the model performs better than baseline and ensemble learning methods, establishing its efficiency in providing privacy-preserving, scalable, and efficient segmentation for various online retail datasets.



Fig. 4 Customer segments' efficiency score

Figure 4 shows the 20 customer segments' efficiency ratings as assessed using DEA, a performance benchmarking technique. Efficiency score (0.5 to 1.0) is along the y-axis, and segment indices along the x-axis. Segments with scores close to 1.0 are considered highly efficient in input-output resource usage, indicating optimal behavioral traits of high purchase

frequency, basket size, or profitability. On the other hand, low-scoring segments point to underperformance or inefficiency. This DEA-based evaluation facilitates focused strategies like tailored marketing or re-engagement strategies and augments the FCA's clustering for more precise, data-led customer segmentation.




Fig. 5 Training process of FCA

Figure 5 illustrates the FCA's training process by monitoring reconstruction loss over 50 epochs. The loss metric indicates how accurately the model reconstructs the original input from its compressed representation. A steep decrease in loss during early epochs indicates that the model learns salient patterns in customer transaction behavior at a rapid pace. As training continues, the loss curve slowly flattens as the FCA approaches the point where it is learning from more nuanced behavioral patterns.

The minor noise on the curve results from federated environments, with the updates to models from distributed nodes creating local fluctuations. However, the overall trend downwards assures successful learning. This visualization substantiates the argument that the FCA learns significant features across clients without seeing raw data, thereby confirming the privacy-preserving aspect of the segmentation framework while ensuring representational accuracy needed for downstream DEA-based customer group performance measurement.




Fig. 6 ROC curve

The ROC curve in Figure 6 measures the classification accuracy of the Federated Convolutional Autoencoder (FCA) in distinguishing efficient from inefficient customer segments. It displays the True Positive Rate against the False Positive Rate at different decision thresholds, offering an overall measure of the discriminative power of the model.

The Area Under the Curve (AUC) measures such performance: an AUC value nearer 1.0 signifies good classification accuracy. Here, in the context of the present study, the ROC curve validates the FCA model's ability to identify significant behavior patterns in distributed customer data that are important for privacy-preserving and effective segmentation in decentralized business settings.




Fig. 7 Accuracy curve

The accuracy curve in Figure 7 illustrates a side-by-side comparison of validation and training accuracy throughout the model's iterative training. Across more than 20 epochs, both curves exhibit consistent improvement, illustrating the model's capacity to learn customer behavior patterns across distributed nodes. Minimal training and validation accuracy

divergence prove the Federated Convolutional Autoencoder in escaping overfitting when learning robust representations. This plot confirms that the process of learning is scalable and consistent, which, in the context of a federated setting, where local models should support the global model without divulging sensitive customer information, is of high value.




Fig. 8 Loss curve

This loss curve in Figure 8 tracks how the reconstruction error decreases over training epochs, indicating the autoencoder's ability to learn from distributed data sources. Both training and validation losses exhibit a consistent downward trend, reflecting successful model convergence.

The proximity of the two curves reflects good generalization and minimal overfitting, even in stand-alone client environments.

This plot supports the model's ability to accurately model customer behavior while preserving privacy using federated

learning. It guarantees that the extracted features are relevant for subsequent segmentation and efficiency analysis, particularly when combined with Data Envelopment Analysis for cluster assessment.

5.2. Performance Evaluation

The performance testing measures the superiority of the suggested model compared to current methods based on the most important metrics. The testing demonstrates the robustness of the DEA-informed Federated Convolutional Autoencoder to achieve better segmentation performance in a wide range of retail and e-commerce datasets.

Table 1. Comparison with existing models

Study	Model	Accuracy	Precision	Recall	F1-Score	Dataset
Proposed study	FCA + DEA	95.3%	0.94	0.96	0.95	UCI Online Retail
[19]	XGBoost Classifier	84%	0.89	0.95	0.92	Online Retail Dataset
[20]	AutoML + Ensemble Learning	92.1%	0.91	0.90	0.90	Custom E-commerce Dataset
[21]	Random Forest Classifier	94%	0.94	0.94	0.94	Social Media Ad Dataset

The suggested research presents a new customer segmentation model to improve segmentation granularity and personalization. The model was tested with the UCI Online Retail dataset, which attained a high accuracy of 95.3%, precision, recall, and F1-score values of 0.94, 0.96, and 0.95, respectively, as shown in Table 1. In comparison to current techniques like XGBoost, AutoML-Ensemble Learning, and Random Forest Classifier, the suggested technique provides competitive performance, especially in recall and F1-score.

The federated configuration maintains privacy protection through local model training, while DEA facilitates a strong assessment of segment efficiency. This hybrid technique not only enhances predictive performance but also facilitates more actionable and ethical customer targeting insights. In general, the findings confirm that the FCA + DEA architecture is very effective for today's customer segmentation, particularly for environments where personalization and data security are paramount.




Fig. 9 Predicting reading skills of students

Figure 9 illustrates the comparative performance estimation of four models of customer segmentations based on

evaluation criteria. FCA+DEA appears to be very high-performing, with an optimal accuracy score of 95.3%

supported by better precision, recall, and F1-score performances compared to those of XGBoost, AutoML+Ensemble, and Random Forest-based models in several datasets related to different retailers.

5.3. Discussion

The research introduces a new hybrid framework that combines deep learning and operations research methods to enhance customer segmentation. Utilizing FCA, the method preserves data privacy within distributed retail databases while effectively learning latent patterns of customer behavior. These latent representations are then examined using DEA in order to calculate relative efficiency scores, enabling multi-perspective analysis of customer value and operations impact.

The DEA component provides interpretability to the black-boxed nature of deep learning by evaluating each segment's performance against well-defined input-output criteria like purchase frequency, recency, and monetary value. The results prove that combining FCA and DEA improves the granularity of segmentation and facilitates strategic insights into which segments are the most profitable and which need intervention. This two-stage framework is of special worth for mass-market e-commerce or omnichannel systems looking to establish personalized marketing plans without sacrificing users' privacy. Overall, the combination of federated deep

learning and efficiency analysis provides a scalability and interpretability-oriented route towards precision customer relationship management, making data-driven decisions in a privacy-respecting and operationally sustainable way.

6. Conclusion

The study reveals how FCA can be augmented with DEA to realize superior customer segmentation. Through the employment of federated learning, customer data privacy and security are preserved while retaining precious insights on the behavior of customers. Integration through DEA serves to enhance segmentation since it provides measures of different customer groups' efficiencies, and the company is thus able to take specific, efficient marketing action.

Not only does the approach surpass existing traditional segmentation techniques with respect to accuracy and precision, but it also addresses the burgeoning issue of data privacy in today's digital landscape. For future work, looking forward, exploring the use of more data sources, for instance, online activity and social media usage, as inputs to enhance customer profiles, is an avenue to consider. Additionally, applying this method to larger and more varied data sets could challenge its scalability and strength in a variety of industries and regions, ultimately increasing its real-world usability and strength.

References

- [1] Marawan Omran, Khaled Hamza, and Sally Elghamrawy, "Advancing Customer Segmentation in Banking: Harnessing Machine Learning and H2O for Personalized Insights," *International Conference on Advanced Intelligent Systems and Informatics*, Cairo, Egypt, pp. 246-256, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [2] Hithakrith Parshetty, Lokesh Palimkar, and Rohith Emandi, *Customer Targeting and Segmentation*, Predictive Analytics and Generative AI for Data-Driven Marketing Strategies, Chapman and Hall/CRC, pp. 68-81, 2024. [[Google Scholar](#)] [[Publisher Link](#)]
- [3] P. Sudarsanam et al., "Customer Segmentation using Mini Batch K-Means Clustering Algorithm," *2024 8th International Conference on Electronics, Communication and Aerospace Technology (ICECA)*, Coimbatore, India, pp. 801-805, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [4] Lew Sook Ling, and Claireta Tang Weiling, "Enhancing Segmentation: A Comparative Study of Clustering Methods," *IEEE Access*, vol. 13, pp. 47418-47439, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [5] Pilli Sri Durga, J.A. Paulson, and Marri Srinivasareddy, and others, "Customer Segmentation Analysis for Improving Sales Using Clustering," *International Journal of Science and Research Archive*, vol. 9, no. 2, pp. 708-715, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [6] Rizkita Bagus Perdhana, and Jerry Heikal, "Enhancing Customer Segmentation in Online Transportation Services: A Comprehensive Approach Using K-Means Clustering and RFM Model," *Indonesian Interdisciplinary Journal of Sharia Economics (IJISE)*, vol. 7, no. 2, pp. 2849-2865, 2024. [[Google Scholar](#)] [[Publisher Link](#)]
- [7] Taqwa Hariguna, and Shih Chih Chen, "Customer Segmentation and Targeted Retail Pricing in Digital Advertising Using Gaussian Mixture Models for Maximizing Gross Income," *Journal of Digital Market and Digital Currency*, vol. 1, no. 2, pp. 183-203, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [8] Meena Singh et al., "Leveraging K-Means Clustering and Hierarchical Agglomerative Algorithms for Scalable AI-Driven Customer Segmentation," *Journal of AI ML Research*, vol. 11, no. 7, pp. 1-23, 2022. [[Google Scholar](#)] [[Publisher Link](#)]
- [9] S. Manimozhi, D. Ruby, and K. Biruntha, "An Elegant Evaluation: Triangulating Clustering Methods for Customer Segmentation," *2024 Second International Conference on Emerging Trends in Information Technology and Engineering (ICETITE)*, Vellore, India, pp. 1-6, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]

- [10] Md Abu Sufian Mozumder et al., "Optimizing Customer Segmentation in the Banking Sector: A Comparative Analysis of Machine Learning Algorithms," *Journal of Computer Science and Technology Studies*, vol. 6, no. 4, pp. 1-7, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [11] Qi Le et al., "Personalized Federated Recommender Systems with Private and Partially Federated Auto Encoders," *arXiv Preprint*, pp. 1-7, 2022. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [12] Jeen Mary John, Olamilekan Shobayo, and Bayode Ogunleye, "An Exploration of Clustering Algorithms for Customer Segmentation in the UK Retail Market," *Analytics*, vol. 2, no. 4, pp. 809-823, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [13] Miguel Alves Gomes, and Tobias Meisen, "A Review on Customer Segmentation Methods for Personalized Customer Targeting in E-Commerce Use Cases," *Information Systems and e-Business Management*, vol. 21, no. 3, pp. 527-570, 2023. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [14] Chenguang Wang, "Efficient Customer Segmentation in Digital Marketing Using Deep Learning with Swarm Intelligence Approach," *Information Processing & Management*, vol. 59, no. 6, 2022. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [15] Amit Sharma, Neha Patel, and Rajesh Gupta, "Enhancing Customer Segmentation Through AI: Analyzing Clustering Algorithms and Deep Learning Techniques," *European Advanced AI Journal*, vol. 11, no. 8, pp. 1-26, 2022. [[Google Scholar](#)] [[Publisher Link](#)]
- [16] Gehna Upreti, and Arul Kumar Natarajan, *Leveraging Unsupervised Machine Learning to Optimize Customer Segmentation and Product Recommendations for Increased Retail Profits*, Intersection of AI and Business Intelligence in Data-Driven Decision-Making, IGI Global, pp. 257-282, 2024. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [17] Shalom Mariam Aeby, V. Vedhik Mohan Rao, and Nikhil Gupta, *Customer Segmentation Techniques in Predictive Analytics and Generative AI for Data-Driven Marketing Strategies*, Predictive Analytics and Generative AI for Data-Driven Marketing Strategies, Chapman and Hall/CRC, pp. 234-246, 2024. [[Publisher Link](#)]
- [18] Daqing Chen, "Online Retail," UCI Machine Learning Repository, 2015. [[CrossRef](#)] [[Publisher Link](#)]
- [19] Lahcen Abidar et al., "Predicting Customer Segment Changes to Enhance Customer Retention: A Case Study for Online Retail using Machine Learning," *International Journal of Advanced Computer Science and Applications (IJACSA)*, vol. 14, no. 7, pp. 910-920, 2025. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [20] Zne-Jung Lee et al., "Clustering and Classification Based on Distributed Automatic Feature Engineering for Customer Segmentation," *Symmetry*, vol. 13, no. 9, pp. 1-14, 2021. [[CrossRef](#)] [[Google Scholar](#)] [[Publisher Link](#)]
- [21] Gbeminiyi Falowo, "Development of a Machine Learning Model for Brand and Audience Segmentation Using Demographic Data," *Corpus Intellectual*, vol. 3, no. 3, pp. 1-21, 2024. [[Publisher Link](#)]