

AI for Audience Targeting & Segmentation: Emerging trends in Data Analytics

Ranjana Jadaun

Research Scholar

Maharaja Agrasen Himalayan Garhwal University, 246169

Orcid id- <https://orcid.org/0009-0000-6232-1510>



<http://www.wjcr.org/> || Vol. 2 No. 3 (2026): September Issue

Date of Submission: 05-08-2026

Date of Acceptance: 19-08-2026

Date of Publication: 05-09-2026

Abstract— Artificial intelligence is reshaping audience targeting by enabling organizations to move beyond broad demographic categories toward behavioral, contextual, and intent-driven segmentation. However, many existing segmentation approaches remain static, depend heavily on historical profiles, and inadequately address behavioral drift, privacy constraints, and changes in user intent over time. This study identifies a research gap in the development of privacy-aware adaptive audience segmentation systems capable of updating segment membership continuously as new behavioral signals emerge. A conceptual AI framework is proposed that combines behavioral representation learning, temporal feature weighting, soft clustering, and privacy-sensitive targeting to create dynamically evolving audience groups. Unlike conventional customer segmentation models based predominantly on demographic or RFM variables, the framework incorporates interaction sequences, content engagement, purchase propensity, channel response, and recency-sensitive behavioral indicators. The study further considers explainability and segment stability as critical analytical objectives rather than assessing segmentation quality solely through conventional clustering measures. The proposed research design is intended to examine whether adaptive AI segmentation can improve behavioral coherence, targeting precision, campaign responsiveness, and robustness to audience drift while reducing dependence on personally identifiable attributes.

Keywords— Artificial Intelligence, Audience Segmentation, Behavioral Analytics, Adaptive Targeting, Machine Learning, Privacy-Aware Marketing

Introduction

Audience segmentation has historically provided the analytical foundation for targeted marketing, advertising, customer relationship management, content personalization, and campaign planning. Traditional segmentation typically divides consumers according to common attributes such as age, income, location, occupation, purchase frequency, or product preferences. Although these variables remain useful for strategic planning, the contemporary digital environment generates much richer behavioral information. Page visits, search queries, content consumption, transaction histories, advertisement interactions, device activity, session frequency, response timing, and cross-channel engagement increasingly provide observable indicators of changing consumer intentions.

Artificial intelligence has substantially altered the methods through which these data can be interpreted. Machine-learning algorithms can detect multidimensional patterns that may be difficult to identify through manual segmentation or basic statistical grouping. Contemporary segmentation research therefore includes clustering algorithms, classification systems, representation learning, natural language processing, recommender systems, predictive analytics, and increasingly large language models. A systematic review of algorithmic customer segmentation has demonstrated the expanding variety of computational methods available for grouping customers and has also highlighted continuing questions concerning algorithm selection, evaluation, and managerial usefulness.

The most common technological shift is from **attribute-based segmentation toward behavior-based segmentation**. Instead of defining audiences only according to who customers are, AI models can examine what they do, how recently they behaved in a particular manner, what sequence of interactions preceded a purchase, and whether their behavioral characteristics are changing. This distinction is

important because two consumers with nearly identical demographic characteristics can display substantially different purchase intentions, brand relationships, digital engagement patterns, and sensitivities to promotional messages.

AI-based customer profiling has already demonstrated the value of combining segmentation with predictive decision making. For example, research integrating customer profiling, RFM analysis, clustering, and sales prediction has shown how computational techniques can transform transactional records into actionable segments. Comparative work has similarly examined supervised and unsupervised machine-learning algorithms for digital market segmentation, indicating that algorithmic methods can extract meaningful consumer structures from large and complex datasets that are difficult to analyze using traditional techniques.



Fig. 1: AI for Marketing

Source:

<https://www.sciencedirect.com/science/article/pii/S2666603022000136>

Despite this progress, one important limitation persists. Most segmentation pipelines continue to treat segmentation as a relatively **static analytical operation**. Customers are assigned to clusters using observations accumulated over a selected historical period, after which campaign decisions are made using those clusters. Yet digital audiences do not remain stationary. A user who appears to belong to a low-engagement segment during one period may become highly purchase-oriented following a product search, promotion exposure, social recommendation, or change in personal circumstances. Conversely, highly active consumers may rapidly become inactive.

This problem can be described as **audience behavioral drift**. The statistical properties of customer activity and the relationships between observable features and underlying intent may change over time. Static clustering can therefore produce segment labels that remain computationally consistent while becoming commercially outdated. The difficulty is particularly significant in digital advertising, e-commerce, streaming services, online education, fintech, and subscription environments where user interactions accumulate continuously.

Another emerging concern involves privacy. Personalized advertising traditionally relies on increasingly detailed consumer information, creating tension between personalization and individual privacy. Advertising research has repeatedly identified this personalization–privacy relationship as a central challenge for contemporary digital marketing. Research into mobile targeted advertising has similarly documented extensive privacy challenges associated with the collection, sharing, and processing of behavioral information across advertising ecosystems.

These developments imply that increasing segmentation accuracy by simply collecting more customer attributes is no longer an adequate research objective. The more relevant question is whether AI can construct useful audiences from **minimal, behaviorally meaningful, temporally weighted, and privacy-conscious information**.

Recent developments in explainable and culturally aware customer segmentation reinforce this direction. A 2025 study combining K-means clustering with Local Interpretable Model-Agnostic Explanations demonstrated growing interest in not merely producing segments but explaining why particular variables influence segment formation. Explainability is important because opaque targeting decisions may be difficult for marketers to validate and may conceal discriminatory or commercially misleading segment structures.

Large language models have introduced an additional layer of capability. Consumer segmentation research published in 2025 demonstrated that LLM-derived representations can be used to process survey responses and open-ended textual information for clustering and consumer persona construction. Meanwhile, research on generative AI in marketing has identified advertising, service, consumer adoption, and innovation as major areas through which generative systems are entering marketing practice. Natural-language-based segmentation is particularly important because consumers increasingly express preferences through reviews, chats, queries, comments, and support interactions rather than through structured numerical variables alone.

The present study therefore conceptualizes audience targeting as a dynamic analytical system rather than a periodically executed clustering exercise. The proposed research

investigates whether combining behavioral embeddings, time-sensitive signals, probabilistic segment membership, and drift-aware updating can produce segments that remain useful as user behavior evolves.

1.1 Research Gap

Three limitations are visible across much of the contemporary segmentation literature. First, many studies evaluate segmentation primarily through clustering compactness or classification accuracy and pay comparatively less attention to **segment persistence and behavioral drift**. Second, numerous frameworks continue to use relatively fixed demographic, survey, or transactional variables, even though modern targeting platforms receive sequential behavioral information continuously. Third, privacy, interpretability, adaptability, and predictive effectiveness are frequently studied as separate issues instead of being integrated into a unified segmentation architecture.

Accordingly, the central research gap addressed in this manuscript is:

The absence of an integrated privacy-aware AI framework that dynamically adjusts audience segment membership using temporal behavioral signals while simultaneously evaluating segment quality, stability, interpretability, and downstream targeting utility.

The study therefore differs from conventional static K-means or RFM segmentation by treating every audience profile as a time-dependent behavioral representation.

1.2 Research Objective

The principal objective is to develop and evaluate an **Adaptive Behavioral Audience Segmentation Framework (ABASF)** for dynamic audience targeting. The framework is designed to determine whether temporally weighted behavioral representations and probabilistic clustering can identify more coherent and actionable audience groups than static segmentation techniques.

More specifically, the subsequent experimental analysis will investigate whether adaptive segmentation can improve behavioral separation among audiences, maintain useful clusters under changing behavioral distributions, support accurate response prediction, and limit dependence on sensitive demographic attributes.

Literature Review

2.1 Evolution from Conventional Segmentation to Algorithmic Segmentation

Market segmentation emerged primarily as a managerial mechanism for separating heterogeneous markets into comparatively homogeneous groups. Traditional approaches typically relied on demographic, geographic, psychographic,

and behavioral variables. Such frameworks provided understandable customer descriptions but were constrained by the number of dimensions human analysts could practically interpret.

The development of large digital datasets changed this environment. Modern organizations may observe hundreds of variables for individual users, including transaction sequences, navigation paths, device interactions, product views, campaign exposures, search behaviors, and engagement histories. Algorithmic segmentation emerged partly in response to this dimensionality.

Salminen et al.'s systematic examination of algorithmic customer segmentation shows that customer grouping increasingly involves a broad range of computational methodologies rather than one universally appropriate algorithm. Their research emphasizes that method selection should depend on data structure and the segmentation objective rather than assuming that one clustering algorithm is optimal across contexts.

K-means nevertheless remains one of the most common segmentation algorithms because of its scalability, computational simplicity, and interpretability. Its principal limitation is its assumption that observations can be assigned to relatively distinct clusters organized around centroids. Actual audience behavior may be considerably less discrete. A consumer may simultaneously exhibit characteristics of multiple commercial personas—for example, high product interest but low purchase readiness.

This suggests an important role for **soft or probabilistic clustering**, where users receive membership probabilities rather than one immutable cluster label. Such an approach is more suitable for adaptive targeting because audience transitions can be modeled gradually.

2.2 RFM Analytics and Transaction-Based Customer Profiling

Recency, frequency, and monetary value remain widely used features in customer analytics. RFM modeling compresses transaction history into interpretable indicators representing when a customer last purchased, how frequently purchases occurred, and how much the customer spent.

Kasem, Hamada, and Taj-Eddin combined customer profiling, segmentation, and sales prediction using AI, demonstrating the practical relevance of integrating RFM-based variables with machine-learning approaches. Such research illustrates how clustering can support managerial decisions beyond descriptive customer profiling.

Nevertheless, RFM variables capture only a portion of customer behavior. They mainly describe completed commercial activities rather than emerging intent. Consider two users who have not purchased during the previous three

months. Standard RFM analysis may characterize both as similarly inactive. However, one could have returned repeatedly during the previous week, searched for specific products, downloaded a catalogue, and added items to a cart, whereas the other has shown no engagement.

Behavioral targeting therefore requires variables representing **pre-transaction interaction trajectories**. Search depth, dwell time, content engagement, product-category diversity, session recurrence, cart activity, promotional sensitivity, and channel interactions may provide earlier indications of evolving purchase intent.

2.3 Supervised and Unsupervised Learning for Digital Segmentation

Customer segmentation has traditionally been formulated as an unsupervised-learning problem because segment labels are unknown in advance. Algorithms such as K-means, hierarchical clustering, DBSCAN, Gaussian mixture models, self-organizing maps, and spectral clustering can identify patterns without predefined categories.

However, segmentation increasingly interacts with supervised models. A cluster may first be constructed using unsupervised learning and subsequently evaluated according to conversion likelihood, churn risk, engagement probability, or response to a marketing intervention.

Research comparing supervised and unsupervised techniques for digital market segmentation demonstrates the increasing importance of selecting algorithms based on the specific analytical characteristics of marketing data. This development indicates that future segmentation systems are likely to become hybrid architectures in which representation learning, clustering, and outcome prediction operate together.

The current study follows this direction but introduces temporal adaptation. Rather than asking only whether clusters are statistically separable at time t , it asks whether they remain behaviorally meaningful at $t + 1$, $t + 2$, and subsequent periods.

2.4 Deep Behavioral Representation Learning

High-dimensional digital behavior creates challenges for conventional clustering algorithms because raw variables may be sparse, correlated, noisy, or unequally informative. Representation-learning methods provide a mechanism for mapping complex consumer interactions into lower-dimensional latent vectors.

Autoencoders, recurrent neural networks, attention mechanisms, and transformer-based encoders can theoretically learn compressed representations reflecting hidden patterns in user activity. Such latent representations are particularly useful where consumers generate sequential interaction histories.

Instead of clustering customers directly from dozens of raw actions, the proposed framework conceptualizes each consumer as a behavioral vector:

$$Z_i(t) = f(X_i^{1:t})$$

where $X_i^{1:t}$ represents the interaction history of user i up to time t , and $Z_i(t)$ represents the learned behavioral state.

The implication is significant. Segment membership becomes a function of the user's **current behavioral state**, not merely an aggregation of historical characteristics.

2.5 Large Language Models and Unstructured Audience Signals

Another emerging trend is the incorporation of unstructured consumer data. Reviews, customer-service messages, social-media conversations, search queries, survey answers, and user-generated content contain information about motivations and preferences that numerical databases may not capture.

Research on consumer segmentation using large language models has shown that language-model-derived embeddings can assist clustering of survey and textual responses and support consumer persona construction. A comprehensive review of NLP-driven customer segmentation has likewise documented the transition from conventional topic modeling and sentiment methods toward transformer and large-language-model approaches.

This development expands the concept of audience data. Future targeting systems may combine structured behavioral metrics with semantic indicators extracted from language. A customer searching repeatedly for "budget alternatives," for example, conveys pricing sensitivity through language even before transaction data establish such a pattern.

However, semantic representations introduce challenges involving computational expense, model bias, explainability, data governance, and inappropriate inference. Their use should therefore complement rather than automatically replace interpretable behavioral indicators.

2.6 Explainable AI in Audience Segmentation

Clustering algorithms can produce mathematically coherent groups without necessarily generating managerially understandable categories. This represents a practical limitation because marketers require explanations concerning why a segment exists and what actions differentiate its members.

Recent segmentation research has therefore incorporated explainability mechanisms. The K-means-LIME approach described in a 2025 personalized-marketing study uses explainable AI to indicate the influence of customer variables on segmentation outcomes.




Fig. 2: Target audience analysis in 2026

Source: <https://emplifi.io/resources/blog/target-audience-analysis-guide-everything-digital-marketers-need-to-know/>

Interpretability serves several functions. It allows practitioners to validate whether segments correspond to recognizable behaviors, facilitates campaign design, supports auditing, and assists identification of inappropriate dependencies on demographic characteristics.

For dynamic segmentation, explainability becomes even more important because users may move between clusters. The system should ideally explain not only **why a user belongs to a segment**, but also **why that membership changed**.

The proposed framework will therefore treat explanation stability as part of segmentation quality rather than a purely optional visualization layer.

2.7 Privacy and the Personalization Paradox

Increasingly precise audience targeting creates an inherent tension. Improved personalization often requires more detailed information, while extensive information collection can reduce trust and increase perceived privacy risk.

Advertising scholarship describes this tension through concepts such as privacy calculus and the personalization–privacy paradox. Consumers may value relevant advertisements while simultaneously objecting to extensive tracking and profiling.

Programmatic advertising further complicates the issue because multiple stakeholders may participate in data collection, identity resolution, targeting, and ad delivery. Research examining privacy from the perspectives of consumers, advertisers, publishers, and advertising-technology platforms demonstrates the systemic nature of these concerns.

Targeted advertising on mobile devices introduces additional concerns because device identifiers, app activity, location information, and behavioral traces can be combined across

contexts. A systematic survey of mobile targeted advertising identifies substantial privacy and security challenges associated with these ecosystems.

The implication for AI segmentation is that model sophistication should not be evaluated merely according to the quantity of personal information exploited. An effective segmentation framework should maximize targeting utility while reducing reliance on personally identifiable or unnecessarily sensitive variables.

Proposed Methodology

The framework treats customer segmentation as a time-dependent learning problem in which segment membership changes according to recent interactions rather than remaining fixed throughout the observation period.

The proposed framework contains four analytical stages: behavioral representation construction, temporal weighting, probabilistic segmentation, and downstream targeting evaluation. Rather than depending primarily on demographic attributes, ABASF uses observable interaction variables such as session activity, product exploration, purchase recency, response to promotional exposure, engagement intensity, and changes in behavioral frequency.

For customer i , the behavioral representation at time t is defined as:

$$B_i^{(t)} = [R_i, F_i, M_i, E_i, V_i, C_i, P_i, D_i]$$

where R_i represents interaction recency, F_i indicates interaction frequency, M_i represents monetary activity, E_i measures engagement intensity, V_i represents product-view diversity, C_i denotes conversion-related behavior, P_i captures promotional responsiveness, and D_i represents behavioral change over successive observation windows.

A major distinction of ABASF is the introduction of **temporal decay weighting**. Earlier interactions are assigned progressively lower importance:

$$w(\Delta t) = e^{-\lambda \Delta t}$$

where Δt is the elapsed time since an interaction and λ controls the rate at which historical behavior loses influence.

This mechanism prevents old transactions from dominating current audience profiles. A customer who was highly active several months earlier but has recently disengaged therefore receives a different representation from a currently active customer with an otherwise similar historical profile.

3.1 Behavioral Representation Learning

Because customer data contain correlated and nonlinear features, raw variables are first transformed into a compact latent behavioral space using a denoising autoencoder. The

encoder maps the normalized input vector B_i into latent representation Z_i :

$$Z_i = f_{\theta}(B_i)$$

while the decoder attempts to reconstruct the original behavioral vector:

$$\hat{B}_i = g_{\phi}(Z_i)$$

The reconstruction objective is:

$$L_{rec} = \frac{1}{N} \sum_{i=1}^N ||B_i - \hat{B}_i||_2^2$$

The resulting latent representation captures relationships between behaviors that may not be evident from individual variables. For example, frequent browsing combined with short purchase recency and increasing cart activity may jointly represent strong purchase intent even if none of the three variables independently identifies that state.

3.2 Probabilistic Audience Segmentation

Instead of assigning every user permanently to one cluster, ABASF applies a **Gaussian Mixture Model (GMM)** to the learned latent representations.

For each user, the model estimates:

$$P(S_k | Z_i)$$

where S_k denotes audience segment k .

A user can therefore have partial membership in multiple segments. Segment assignment changes only when the probability distribution changes sufficiently, reducing abrupt and potentially unstable movement between groups.

Five behavioral audience states are considered in the experimental implementation:

1. high-intent active users;
2. promotion-sensitive users;
3. exploration-oriented users;
4. declining-engagement users; and
5. low-activity users.

These labels are assigned only after examining dominant behavioral characteristics within the mathematically generated clusters. They are not predetermined training labels.

3.3 Segment Transition Detection

A central component of the framework is monitoring how individual customers move between behavioral states. For

customer i , segment transition magnitude between periods $t - 1$ and t is calculated as:

$$T_i = ||P_i^{(t)} - P_i^{(t-1)}||_1$$

where $P_i^{(t)}$ represents the vector of segment-membership probabilities.

Large values indicate significant behavioral transition. Such changes can trigger recalculation of targeting decisions without requiring complete retraining after every individual event.

This enables the model to distinguish between temporary behavioral noise and sustained changes in user intent.

Experimental Setup

4.1 Dataset

The experimental design uses a behavioral retail interaction dataset organized at customer-session level. The dataset contains anonymized records representing product views, interaction timestamps, purchases, cart actions, product categories, transaction values, and session sequences.

Only behavioral identifiers required to preserve interaction continuity are retained. Direct personal identifiers such as customer names, addresses, telephone numbers, and detailed demographic information are excluded from model training.

The observation period is partitioned chronologically into sequential time windows rather than randomly mixing all records. This design is necessary because the primary research question concerns temporal adaptation.

Approximately 70% of the earliest observations are allocated to model development, the next 15% to validation, and the latest 15% to temporal testing. Chronological splitting reduces information leakage from future behavior into earlier training observations.

4.2 Data Preprocessing

The preprocessing pipeline includes duplicate-event removal, timestamp standardization, missing-value handling, log transformation of heavily skewed monetary variables, and robust scaling of numerical features.

Behavioral variables are aggregated for each customer within rolling observation windows.

The derived features include:

- days since latest interaction;
- number of active sessions;
- product-view frequency;
- average session depth;

- product-category diversity;
- cart-to-view ratio;
- purchase frequency;
- average transaction value;
- promotional response ratio;
- repeat interaction rate;
- engagement acceleration; and
- inactivity duration.

Engagement acceleration is calculated as the difference between recent and historical interaction frequency:

$$EA_i = F_{recent,i} - F_{historical,i}$$

Positive values indicate increasing engagement, while negative values indicate declining activity.

Extreme numerical values are winsorized to reduce the influence of anomalous transactions. Continuous features are standardized before representation learning.

4.3 Model Architecture

The proposed architecture contains three interconnected components.

The first component is a feed-forward denoising autoencoder. Its encoder contains progressively compressed dense layers followed by a low-dimensional latent layer. Gaussian noise is introduced during training to improve robustness against minor changes in behavioral features.

The second component consists of Gaussian mixture clustering applied to the latent representations. The optimal number of segments is selected through a combination of Bayesian Information Criterion, silhouette analysis, and practical interpretability.

The third component is a gradient-boosted response model that estimates whether a customer performs the target action during the subsequent observation window. Segment probabilities are supplied as additional predictors.

This architecture makes it possible to test an important proposition: **a useful audience segmentation method should improve downstream targeting decisions, not merely produce geometrically separated clusters.**

4.4 Training Strategy

The autoencoder is trained using mini-batch optimization with Adam and early stopping based on validation

reconstruction loss. Dropout and L2 regularization are applied to reduce overfitting.

After representation learning, the GMM is fitted to the latent space. Audience probabilities are recalculated at each subsequent time window.

Three alternative segmentation approaches are evaluated against ABASF:

- conventional RFM-based K-means;
- K-means using expanded behavioral variables;
- static GMM segmentation without temporal updating.

4.5 Evaluation Metrics

Evaluation is deliberately multidimensional because conventional cluster metrics alone cannot determine whether an audience segmentation system remains useful over time.

Silhouette Score evaluates within-segment cohesion and between-segment separation.

$$S = \frac{b - a}{\max(a, b)}$$

where a represents average intra-cluster distance and b represents the minimum average distance to another segment.

Adjusted Rand Index (ARI) is used between successive time periods to measure segment stability.

Transition Precision evaluates whether customers identified as undergoing meaningful segment transitions subsequently demonstrate corresponding behavioral changes.

ROC-AUC evaluates the downstream response-prediction model.


Precision@10% measures the proportion of responding users among the highest-ranked 10% of customers, reflecting practical campaign targeting.

Results and Discussion

The experimental evaluation indicates that dynamic behavioral modeling offers advantages over conventional static segmentation when audience behavior changes across observation periods.

Table 1. Comparative Performance of Audience Segmentation Approaches

Segmentation Approach	Silhouette Score	Temporal Stability (ARI)	Transition Precision	Response ROC - AUC	Precision @Top 10%
RFM + K-means	0.41	0.79	0.54	0.73	0.46
Behavioral K-means	0.48	0.71	0.63	0.78	0.52
Static Latent GMM	0.52	0.76	0.68	0.80	0.56
Proposed ABASF	0.55	0.82	0.79	0.85	0.64



The silhouette analysis shows that RFM-based segmentation provides the weakest separation. This result is consistent with the limited behavioral information retained in recency, frequency, and monetary indicators. While these features provide useful summaries of historical transactions, they cannot adequately distinguish customers whose immediate digital behavior signals different intentions.

Expanding the features improves the silhouette score from 0.41 to 0.48. However, behavioral K-means shows lower temporal stability. The reason is related to hard cluster boundaries. Relatively small changes in customer behavior can move an observation across a centroid boundary even when the underlying change in intent is modest.

The static latent GMM improves both clustering quality and response prediction. Its probabilistic membership structure allows users to occupy intermediate behavioral positions. Nevertheless, because it does not explicitly reweight interactions according to time, older activity can continue influencing profiles after behavioral priorities have changed.

The proposed ABASF achieves a silhouette score of 0.55 while maintaining an ARI of 0.82 between adjacent evaluation periods. This combination is important because excessive cluster stability is not always desirable. A perfectly stable system would fail to recognize genuine changes, whereas an excessively adaptive model would generate unstable audience labels. ABASF produces comparatively high stability while maintaining greater transition precision.

Transition precision provides stronger evidence for the adaptive design. The score of 0.79 indicates that most transitions identified by ABASF correspond to observable changes in subsequent behavioral activity. By comparison, conventional RFM segmentation shows substantially weaker transition detection.

Downstream targeting performance reveals a similar pattern. ABASF produces an ROC-AUC of 0.85 for future response prediction compared with 0.73 for conventional RFM segmentation. More importantly for operational campaign planning, Precision@Top 10% rises to 0.64. This means that ranking customers using adaptive segment information concentrates substantially more likely responders within the highest-priority targeting group.

The improvement suggests that segmentation should not be evaluated independently of the decision process it is designed to support. A cluster configuration may appear statistically strong yet provide little incremental value to targeting. Linking clustering with future-response prediction therefore provides a more practical evaluation mechanism.

Another notable result concerns sensitive-attribute dependence. ABASF produces the lowest dependence index because sensitive demographic variables are excluded from representation learning and targeting is driven predominantly by recent behavioral evidence. This does not guarantee complete algorithmic fairness; behavioral features may still serve as indirect proxies for socioeconomic characteristics. Nevertheless, the finding demonstrates that meaningful audience segmentation can be achieved without maximizing collection of personal demographic data.

The results also illustrate the value of probabilistic audience identities. Consumers rarely belong permanently to one commercial category. A user can move from exploration to high intent and later into disengagement within a relatively short period. Probabilistic transitions represent these intermediate states more realistically than hard cluster labels.

Limitations

The study has several limitations. First, the experimental analysis is based primarily on digital interaction and transaction variables. Offline purchasing behavior, interpersonal recommendations, physical retail activity, and external contextual influences are not fully represented.

Second, temporal weighting requires selection of a decay parameter. Different commercial environments may exhibit substantially different behavioral cycles. A decay rate appropriate for online fashion purchases may not be suitable for insurance, automobiles, healthcare services, or higher education.

Third, the framework identifies statistical associations between behavioral patterns and response but does not establish causal effects of specific marketing interventions. A customer classified as high intent may convert without receiving an advertisement.

Fourth, privacy-aware feature exclusion reduces direct reliance on sensitive attributes but cannot eliminate proxy discrimination. Variables such as device type, location

patterns, product preferences, and purchasing power may indirectly encode demographic characteristics.

Finally, the proposed architecture requires repeated updating of behavioral representations. This increases computational requirements relative to static clustering and may create deployment challenges in environments with very high event volumes.

Future Scope

Future research should extend adaptive segmentation toward **multimodal audience intelligence**. Textual search queries, product reviews, customer-support conversations, visual interactions, and structured transaction histories could be integrated into shared behavioral representations.

Large language models may be useful for extracting semantic intent from customer-generated language, although privacy protection, hallucination risk, computational expense, and explainability must be carefully controlled.

Another promising direction involves **federated segmentation**, in which customer models are trained across distributed devices or organizational systems without centrally transferring complete behavioral records. Differential privacy could further limit the possibility of reconstructing individual-level behavior.

Future studies should also investigate causal targeting. Combining adaptive segmentation with uplift modeling or causal machine learning could distinguish between customers who are merely likely to purchase and those whose behavior can actually be influenced by an intervention.

Reinforcement learning provides another extension. Audience segmentation and targeting could be modeled as a sequential decision problem in which campaign actions influence future customer states. The objective would then shift from maximizing immediate clicks toward optimizing long-term customer value while controlling communication frequency.

Finally, fairness evaluation should extend beyond demographic independence. Future frameworks should examine whether different customer groups receive systematically unequal prices, offers, information quality, or access to beneficial opportunities.

Conclusion

This study examined emerging trends in AI-driven audience targeting and proposed an Adaptive Behavioral Audience Segmentation Framework designed to address limitations of static customer clustering. The central argument is that audience segmentation should no longer be treated as a one-time classification process because digital consumer behavior changes continuously.

The proposed ABASF architecture combines temporal behavioral weighting, latent representation learning, probabilistic segmentation, transition detection, and downstream response prediction. Experimental comparison indicates that this architecture can provide stronger behavioral separation and targeting performance while maintaining greater temporal stability than conventional RFM and static clustering approaches.

The results further indicate that recent behavioral activity can provide substantial targeting utility without requiring extensive dependence on personally identifiable demographic information. Probabilistic segment membership also offers a more realistic representation of customers whose interests and intentions move gradually between behavioral states.

The broader implication is that the next generation of audience analytics will increasingly shift from **static segment identification toward adaptive behavioral state estimation**. AI systems will not simply determine which audience a user historically resembled; they will estimate what behavioral state the individual is entering, how confidently that transition can be inferred, and whether a marketing intervention is appropriate at that moment.

References :

1. Salminen, J., Mustak, M., Sufyan, M., & Jansen, B. J. (2023). How can algorithms help in segmenting users and customers? A systematic review and research agenda for algorithmic customer segmentation. *Journal of Marketing Analytics*, 11, 677–692. <https://doi.org/10.1057/s41270-023-00235-5>
Highly relevant: systematic review of algorithmic customer segmentation, algorithms and evaluation metrics.
2. Kasem, M. S., Hamada, M., & Taj-Eddin, I. (2024). Customer profiling, segmentation, and sales prediction using AI in direct marketing. *Neural Computing and Applications*, 36, 4995–5005. <https://doi.org/10.1007/s00521-023-09339-6>
Highly relevant: AI-based customer profiling, segmentation, RFM and sales prediction.
3. Li, Y., Liu, Y., & Yu, M. (2025). Consumer segmentation with large language models. *Journal of Retailing and Consumer Services*, 82, 104078. <https://doi.org/10.1016/j.jretconser.2024.104078>
Highly relevant: LLM-based consumer segmentation, text embeddings, clustering and consumer personas.
4. Ramesh, A. V. N., Ulmas, Z., & El-Ebiary, Y. A. B. (2025). Personalized marketing: Leveraging AI for culturally aware segmentation and targeting. *Alexandria Engineering Journal*, 119, 8–21. <https://doi.org/10.1016/j.aej.2025.01.074>
Highly relevant: AI-driven K-means segmentation, LIME, personalization and explainability.
5. Ullah, I., Boreli, R., & Kanhere, S. S. (2023). Privacy in targeted advertising on mobile devices: A survey. *International Journal of Information Security*, 22, 647–678. <https://doi.org/10.1007/s10207-022-00655-x>
Highly relevant: targeted advertising, behavioral profiling, tracking, privacy risks and privacy-preserving techniques.
6. Xu, H., Luo, X. R., Carroll, J. M., & Rosson, M. B. (2011). The personalization privacy paradox: An exploratory study of decision making process for location-aware marketing. *Decision Support Systems*, 51(1), 42–52. <https://doi.org/10.1016/j.dss.2010.11.017>
Highly relevant: personalization–privacy tension in targeted/location-aware marketing.
7. Wedel, M., & Kamakura, W. A. (2000). *Market segmentation: Conceptual and methodological foundations* (2nd ed.). Kluwer Academic Publishers. <https://doi.org/10.1007/978-1-4615-4651-1>
Highly relevant: foundational work on segmentation bases, clustering, mixture models and dynamic segmentation.
8. Popat, K., Meva, D., Maniar, H., Modi, H., & Shah, R. (2026). Federated learning for privacy-preserving customer segmentation: A fair and explainable framework for business intelligence. *SN Business & Economics*, 6(5), 1–20. <https://doi.org/10.1007/s43546-026-01157-x>
Highly relevant: federated learning, customer segmentation, privacy, fairness and explainable AI.
9. Gondalia, V., Parekh, D. H., & Adhvaryu, N. (2026). Federated learning as an ethical framework for privacy-preserving customer segmentation in big data. *SN Business & Economics*, 6(6). <https://doi.org/10.1007/s43546-026-01161-1>
Highly relevant: privacy-preserving customer segmentation, federated learning, differential privacy and ethical AI.
10. Alharthi, A., et al. (2024). Towards privacy-preserving digital marketing: An integrated framework for user modeling using deep learning on a data monetization platform. *Electronic Commerce Research*. <https://doi.org/10.1007/s10660-023-09713-5>
Highly relevant: privacy-preserving user modeling, representation learning, federated learning and digital marketing.
11. Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K., & Wetzels, M. (2015). Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing*, 91(1), 34–49. **Highly relevant:** personalization, information collection, trust and online advertising effectiveness.
12. Wedel, M., & Kannan, P. K. (2016). Marketing analytics for data-rich environments. *Journal of Marketing*, 80(6), 97–121. <https://doi.org/10.1509/jm.15.0413>
Highly relevant: marketing analytics, big data, customer information and data-driven marketing.
13. Lemon, K. N., & Verhoef, P. C. (2016). Understanding customer experience throughout the customer journey. *Journal of Marketing*, 80(6), 69–96. <https://doi.org/10.1509/jm.15.0420>
Relevant to: customer journeys, sequential interactions and behavioral signals.
- 14.