Recent trends in Management and Commerce

Vol: 7(1), 2026

REST Publisher; ISBN: 978-81-936097-6-7

Website: https://restpublisher.com/book-series/rmc/

AI in Marketing and Customer Engagement

Archana. D

Arignar Anna College (Arts and Science), Krishnagiri, Tamilnadu, India.

Abstract: *Artificial intelligence (AI) is redefining marketing and customer engagement by enabling fine-grained consumer behavior analytics, scalable personalization and recommendation, efficient digital and social media execution, conversational customer relationship management (CRM) through chatbots and virtual assistants, and predictive analytics for customer retention. This paper synthesizes technical methods, data and platform requirements, and governance patterns across these five sub-themes. We highlight cross-cutting operations—MLOps/LLMOps, privacy and consent management, bias/fairness, explain ability—and propose a practical roadmap with performance metrics for trustworthy, value-accretive deployment.*

Key words: *Artificial Intelligence in Marketing, Customer Personalization,* ***Artificial*** *Intelligence in Marketing,*

1. INTRODUCTION

AI capabilities are becoming foundational to modern marketing organizations. Behavioral signals from web, mobile, and offline channels now feed machine learning (ML) pipelines that predict intent, segment audiences, personalize content, and orchestrate Omni channel journeys. Generative AI (GenAI) accelerates creative production and conversation via chatbots, while predictive models guide retention offers and next-best actions. Yet, effective adoption requires robust data foundations, production operations, and safeguards addressing consent, privacy, and fairness. This paper addresses five sub-themes central to practice: (1) AI-based consumer behavior analysis, (2) personalization and recommendation systems, (3) AI in digital and social media marketing, (4) chatbots and virtual assistants in CRM, and (5) predictive analytics for customer retention. We summarize methods, architectures, metrics, and governance, and present an implementation roadmap aligned to regulatory expectations and customer trust.

Contributions

- An integrated view of architectures and controls across the five domains;
- Concrete performance and business metrics;
- Guidance on MLOps/LLMOps and responsible AI for marketing.

AI-Based Consumer Behavior Analysis

Data Foundations Behavioral analytics rely on first-party data (web/mobile events, purchases, email interactions), zero-party data (declared preferences), and selected third- party/partner data where permitted. Unification via a customer data platform (CDP) or lake house with identity resolution supports event-level analytics and audience activation. Privacy-preserving design- granular consent, purpose limitation, and data minimization is mandatory, especially under GDPR/CCPA-like regimes.

2. METHODOLOGY

Descriptive and Unsupervised Learning: Clustering (k-means, Gaussian mixtures), topic modeling, and embedding's to derive segments and themes from behavior and text [1]. Dimensionality reduction (PCA/UMAP) and sequence embedding's enable compact representations. Sequence and Causal Inference: Markov models, sequence mining, and

recurrent/Transformer models capture path dynamics; uplift modeling and causal ML distinguish correlation from incremental effect [2]. Attribution and Media Mix: Shapley value methods and Bayesian media mix models estimate channel contributions under signal loss (e.g., cookie deprecation) [3].

A. Operationalization and Metrics Analysts deploy reusable feature pipelines (recency– frequency–monetary, session features, content affinities). Production metrics include data freshness, identity match rates, and coverage. Business metrics include segment stability, lift in engagement, and attribution confidence intervals.

Personalization and Recommendation Systems

Approaches

- Collaborative Filtering: Matrix factorization and neural CF leverage user–item interactions for recommendations.
- Content-Based and Hybrid: Text, image, and metadata embeddings (e.g., sentence and vision models) combined with interaction signals [4].
- Contextual and Sequence-Aware: Contextual bandits for exploration–exploitation; sequence models (SASRec/Transformers) predict next-item; reinforcement learning (RL) optimizes long-term rewards across sessions [5].
- Generative Personalization: LLMs generate copy and offers conditioned on profiles and guard-railed via retrieval-augmented generation (RAG) to ensure factuality and compliance.
- Ranking, Constraints, and Safety Two-stage architectures (candidate generation + ranking) balance recall and precision; business constraints (diversity, fairness, inventory, compliance) are applied via constrained ranking or re-ranking. Safety filters prevent sensitive or non-compliant content; fairness checks monitor differential performance across protected groups where applicable.
- Evaluation Offline: AUC, NDCG, MRR for ranking; calibration for propensity models. Online: A/B tests with guardrails (minimum performance, bias checks). Business metrics: CTR, CVR, average order value (AOV), revenue per session, and long-term value (LTV) impact. Explore–exploit policies should include privacy-aware exploration caps.

AI in Digital and Social Media Marketing

Creative and Copy Generation GenAI accelerates asset creation: headline/body variants, product descriptions, and imagery. Effective deployments use brand style guides, policy constraints, and human-in-the-loop review. Fine-tuning or prompt-conditioning on brand corpora improves fidelity; retrieval to approved claims prevents hallucinations.

Targeting and Bidding

- Predictive Targeting: Look-alike modeling and interest inference on first-party data with ID-graph constraints.
- Bidding Optimization: Bayesian optimization and RL adjust bids/budgets across channels respecting pacing and ROI targets; privacy-preserving conversion modeling addresses signal loss [3].
- Social Listening and Crisis Detection: NLP for sentiment, emotion, and topic shifts; anomaly detection flags brand risk. Multilingual models enable global monitoring.
- Measurement and Experimentation Causal lift tests (geo-holdouts, intent-to-treat) complement platform-reported metrics. MMM provides long-horizon guidance; instrumentality experiments validate audience and creative strategies. Report transparency includes uncertainty ranges and assumptions.

3. CHATBOTS AND VIRTUAL ASSISTANTS IN CRM

A.Use Cases

- Pre-Sales: Product Q&A, guided discovery, and lead qualification.
- Post-Sales Support: Order status, returns, troubleshooting, and knowledge base retrieval.
- Account Management: Billing queries, plan changes, personalized offers as next-best actions (NBA).

B.Architectures

- Task-Oriented Bots: Intent classification and slot filling with dialog managers; deterministic policies for high-risk workflows.

- LLM-Powered Assistants: Retrieval-augmented generation grounded in curated knowledge; tool use for API calls (order lookup, ticket creation); safety layers to prevent data leakage and off-policy responses [6].
- Human-in-the-Loop: Confidence thresholds trigger escalation; transcripts and actions are logged for quality assurance and training.

Quality, Safety, and Compliance Metrics include containment rate, first contact resolution, CSAT, average handle time, and deflection. Safety involves jailbreak resistance, PII redaction, and role-based access. Compliance requires consent capture, auditable reasoning for automated decisions, and clear handoff paths.

Predictive Analytics for Customer Retention

- Churn and LTV Modeling
- Churn: Gradient boosting, calibrated classifiers, survival analysis (Cox, AFT), and sequence models identify at-risk customers and time-to-event.
- LTV: CLV models (Pareto/NBD, Gamma-Gamma) and ML-based LTV with uncertainty; combine with marginal profitability to prioritize interventions [7].
- Uplift and Next-Best Action Uplift modeling (T-learner, U-learner, meta-learners) estimates treatment effect to avoid waste and negative reactions; constrained optimization chooses offers subject to budget, fairness, and policy limits. Multi-armed bandits or RL adjust contact policies over time.
- Experimentation and Measurement Run randomized controlled trials for retention programs; measure incremental retention, revenue, and payback period. Guardrails include fatigue caps, contact frequency constraints, and customer preference adherence.

Cross-Cutting Foundations for Trustworthy AI in Marketing

- Data, Identity, and Consent a governed data platform (CDP/lake house), feature store, and real-time events infrastructure underpin all use cases. Identity resolution requires deterministic and probabilistic matches with transparency. Consent and preference management must propagate to all models and activations; implement deletion and opt-out handling.
- MLOps and LLMOps CI/CD for models, feature/version registries, champion–challenger, canary releases, shadow deployments, and automated monitoring for performance, drift, and bias. For LLMs: prompt/version management, grounding policies, safety filters, hallucination tests, and red-teaming.
- Explain ability, Fairness, and Governance Local and global explain ability for targeting and retention decisions; fairness evaluations where applicable; clear decision rights and escalation; documentation for marketing governance and regulators. Implement data minimization and privacy-preserving techniques (differential privacy, federated learning, and synthetic data) where appropriate.

Implementation Roadmap

- Phase 1—Foundations and Quick Wins: Stand up the feature store and event pipeline; launch personalization on a high-traffic surface; deploy an FAQ chatbot with retrieval grounding; run a baseline churn model. Define success metrics: CTR/CVR lift, CSAT, deflection, and
- Incremental retention.

- Phase 2—Scale and Governance: Expand to sequence-aware recommenders and uplift-driven retention; introduce RL for bidding and send-time optimization; formalize MLOps/LLMOps,
- Consent propagation, and bias monitoring. Integrate MMM with experimentation for budget allocation.
- Phase 3—Advanced Optimization and Automation: Multi-objective optimization across channels (revenue, LTV, fairness); argentic assistants with safe tool use for CRM actions; creative optimization with closed-loop feedback; privacy-preserving collaboration with partners.

4. CONCLUSION

AI now permeates the marketing value chain—from behavior analytics and personalization to social execution, conversational CRM, and predictive retention. Durable value emerges when organizations combine solid data and operational foundations with measurable experimentation and responsible AI practices. The methods and controls outlined here offer a practical path to improved engagement, higher conversion and retention, and sustained customer

trust.

REFERENCES

[1]. I. Guyon, U. von Luxburg, and R. Garnett, Eds., Machine Learning: The Art and Science of Algorithms that Make Sense of Data, 2015.
[2]. S. R. Brodersen et al., “Inferring causal impact using Bayesian structural time-series
[3]. Models,” The Annals of Applied Statistics, 2015.
[4]. D. Chan, “Viewability and attribution in digital advertising,” Applied Marketing Science, 2018.
[5]. J. Leskovec, A. Rajaraman, and J. D. Ullman, Mining of Massive Datasets, 3rd ed., 2020.
[6]. X. Chen et al., “Deep reinforcement learning in recommendation systems: A survey,”
[7]. ACM Computing Surveys, 2021.
[8]. T. B. Brown et al., “Language models are few-shot learners,” NeurIPS, 2020.
[9]. P. Fader and B. Hardie, Customer-Base Audit: The First Step on the Journey to Customer Centricity, 2019.