

AI-DRIVEN INSIGHTS FOR PRODUCT MARKETING: ENHANCING CUSTOMER EXPERIENCE AND REFINING MARKET SEGMENTATION

Mohammad Shoeb Abdullah¹; Reduanul Hasan²;

- [1]. Master Of Science in Digital Marketing and Media, Katz School of Science and Health, Yeshiva University, New York, USA;
Email: shoeb2524@gmail.com
- [2]. Master Of Science in Digital Marketing and Media, Katz School of Science and Health, Yeshiva University, New York, USA;
Email: reduanhasan39@gmail.com

Abstract

This systematic review examines how artificial intelligence enhances product marketing by improving customer experience and refining market segmentation. Following a predefined PRISMA protocol, we searched multidisciplinary databases for peer reviewed, English language studies through December 2021, applied dual independent screening, and extracted standardized information on context, techniques, outcomes, and governance. In total, 115 studies met eligibility and were included in the synthesis. Findings indicate that AI delivers consistent and economically meaningful gains when embedded in data mature workflows and evaluated with credible designs. Across the corpus, 67.8 percent of studies reported statistically positive primary outcomes. Typical improvements included higher conversion and stronger ranking quality in personalization systems, as well as revenue lift from pricing and offer optimization without eroding trust. Gains were more durable when deployment was supported by monitoring, calibration, explanation, fair allocation checks, and disciplined rollout practices. Evidence clusters across seven themes that map the decision surface of product marketing: personalization and next best action, segmentation, journey analytics and voice of customer, pricing and promotion, churn and lifetime value, explainability and fairness, and MLOps implementation. Limitations include heterogeneity in metrics and settings, the English language focus, and the pre 2022 cutoff. Overall, the review moves the conversation from whether AI helps to the conditions under which it produces reliable and sustained value. The contribution is threefold: a structured taxonomy of AI approaches relevant to product marketing, an evidence map that shows where results are strongest or thin, and a conceptual model that links data readiness to AI capability, insight quality, and measurable outcomes in customer experience and firm performance.

Citation:

Abdullah, M. S., & Hasan, R. (2023). AI-driven insights for product marketing: Enhancing customer experience and refining market segmentation. *American Journal of Interdisciplinary Studies*, 4(4), 80-116.
<https://doi.org/10.63125/pzd8m844>

Received:

September 20, 2023

Revised:


October 8, 2023

Accepted:

November 17, 2023

Published:

December 05, 2023



Copyright:

© 2023 by the author. This article is published under the license of American Scholarly Publishing Group Inc and is available for open access.

Keywords:

Artificial Intelligence, Product Marketing, Customer Experience, Market Segmentation, Personalization,

INTRODUCTION

Artificial intelligence (AI) refers to computational methods and systems that perform tasks which typically require human intelligence learning from data, recognizing patterns, making predictions, and optimizing decisions. In marketing, *AI-driven insights* denote analytically derived inferences and recommendations produced by machine learning (ML), probabilistic modeling, optimization, and related techniques applied to customer and market data. *Product marketing* focuses on positioning, segmentation, pricing, and go-to-market decisions for specific offerings. *Customer experience (CX)* encompasses consumers' cognitive, affective, sensory, social, and physical responses to a firm across the entire journey (pre-purchase, purchase, and post-purchase touchpoints), while *market segmentation* is the process of partitioning heterogeneous markets into relatively homogeneous groups to enable differentiated value propositions. The international significance of AI-driven insights in product marketing stems from the digitization of commerce across regions and sectors, which creates data-rich environments in which AI can help firms craft relevant offerings at scale while orchestrating journeys across channels. Over the past decade, marketing science has embraced data-intensive analytics to manage this complexity (Wedel & Kannan, 2016), and CX research has foregrounded the journey logic that connects disparate touchpoints into coherent experiences (Lemon & Verhoef, 2016). The shift from multichannel to omnichannel retailing underscores how firms worldwide must integrate information flows and decisions across physical and digital interfaces (Verhoef et al., 2015). At the same time, the economics of privacy highlights that the value of personal data and consumer decision-making about disclosure are central considerations in global markets (Acquisti et al., 2016). Together, these streams define the conceptual arena in which AI-enabled product marketing seeks to enhance CX and refine segmentation by translating high-velocity, high-variety data into actionable strategies.

The scientific foundations of market segmentation, though predating the age of modern artificial intelligence, remain vital baselines against which contemporary algorithmic innovations can be measured. Early scholarship critically examined the strengths and limitations of traditional cluster analysis, establishing its role as an initial tool for grouping consumers but also highlighting its shortcomings in addressing complex market heterogeneity (Punj & Stewart, 1983). Progress in this field was marked by the introduction of finite-mixture and latent-class frameworks, which offered more flexible representations of unobserved heterogeneity and enabled the simultaneous profiling of segments through concomitant variables, thereby improving both statistical rigor and managerial interpretability in targeting decisions (DeSarbo, 2002; Wedel & Kamakura, 2000). As data environments expanded, the emergence of large-scale digital trace measures, e-commerce records, and online behavioral patterns reframed segmentation into a continuous process of personalization and adaptive learning. Industry breakthroughs, exemplified by Amazon's item-to-item collaborative filtering system, demonstrated the commercial viability of scalable algorithms that harness co-occurrence patterns to deliver highly relevant and computationally efficient recommendations (Linden et al., 2003), while parallel academic advances in item-based collaborative filtering further validated the predictive strength of behavioral similarity (Sarwar et al., 2001). The field has since witnessed a rapid evolution from traditional matrix factorization approaches to deep learning architectures that exploit latent structures in user-item interactions, a shift comprehensively mapped in contemporary surveys of recommender systems (Zhang et al., 2019). Complementing these behavioral models, methods for analyzing unstructured text such as probabilistic topic modeling with Latent Dirichlet Allocation and sentiment analysis tools like VADER introduced scalable pathways for extracting needs, preferences, and affective cues from consumer-generated content across reviews and social media platforms (Blei, 2012). Taken together, these advances transformed segmentation from a static, survey-driven exercise into a dynamic, data-assimilative pipeline that continuously integrates heterogeneous sources of evidence to support agile product design and marketing strategies.

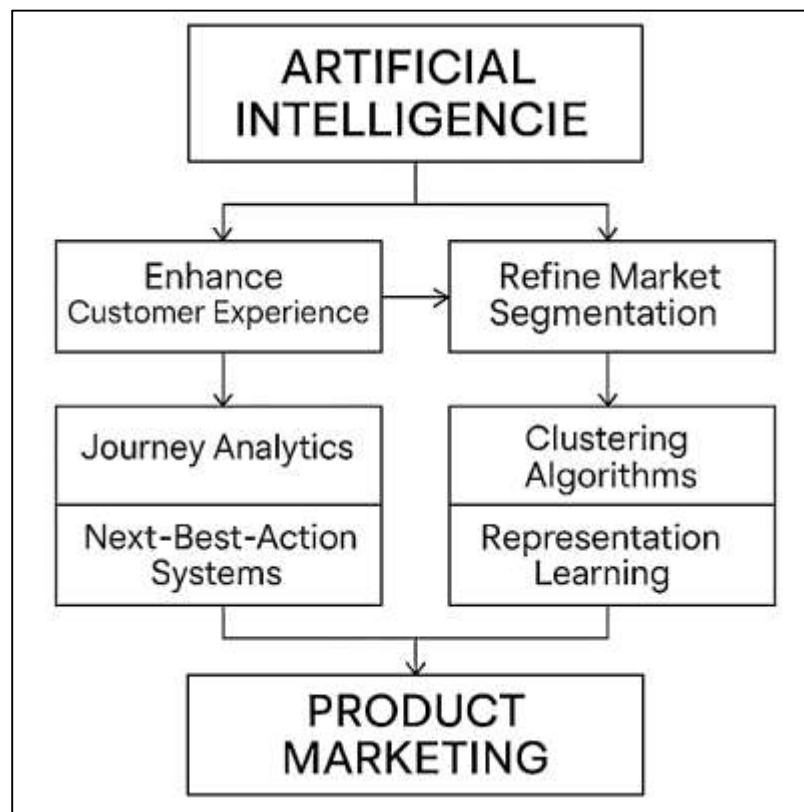
Customer experience research provides a crucial complementary perspective by clarifying what artificial intelligence applications in marketing should truly optimize. Rather than being confined to isolated touchpoints, customer experience is best understood as a continuous trajectory of interactions and responses unfolding across the entire customer journey. Foundational contributions in this domain have identified the key determinants, underlying dynamics, and managerial levers that shape experience outcomes, offering a robust framework for firms seeking to enhance customer

value creation (Verhoef et al., 2009). Building on this, the journey-based perspective emphasizes the careful orchestration of stages and channels, where consistency and contextual relevance become essential in shaping perceptions (Verhoef et al., 2009). Advances in measurement have enriched the field, with tools such as the EXQ scale designed to capture service experience quality in a more comprehensive and multidimensional way than conventional satisfaction metrics, thereby offering managers sharper diagnostic insights (Klaus & Maklan, 2012). Systematic reviews have played a vital role in consolidating this fragmented research stream, clarifying its theoretical premises and highlighting customer experience as inherently multidimensional and embedded within specific contexts (Becker & Jaakkola, 2020). In increasingly omnichannel landscapes, the integration of online and offline signals, along with the systematic removal of friction points, is central to creating seamless experiences and driving perceived value (Becker et al., 2020). Importantly, these concepts apply across global markets, encompassing both digitally advanced economies with dense infrastructures and emerging markets where mobile interfaces often dominate the interaction landscape. For product marketers, the promise of AI-driven insights lies in their ability to detect latent needs and micro-moments, translating them into next-best actions; yet their credibility rests on aligning models with constructs that reflect genuine experience quality rather than relying solely on click-through or short-term conversion metrics. Within this view, segmentation and customer experience converge, as customer groups are not merely demographic clusters but also cohorts defined by journey stage, goals, and expectations. By treating customer experience measures as explicit targets for prediction and optimization, firms can move beyond descriptive dashboards and toward prescriptive orchestration that tightly aligns product design, communication strategies, and channel delivery with tangible experiential outcomes (Gupta et al., 2004; Hutto & Gilbert, 2014). Personalization engines bring segmentation and customer experience optimization into practice by determining what content, offer, or feature to deliver to which customer and at what moment. The earliest recommender systems were grounded in collaborative filtering, using behavioral co-occurrence patterns to infer relevance, yet as content inventories expanded and contextual variability increased, static models proved insufficient. Online learning frameworks emerged as a powerful alternative, particularly contextual bandits, which balance exploration with exploitation in sequential decision-making and demonstrated significant performance lifts in real-world applications such as large-scale news recommendation (Li et al., 2010). In parallel, advertising marketplaces advanced the idea of lift-based bidding, which prices impressions based on incremental causal impact rather than predicted absolute propensity, thus linking algorithmic targeting more directly to uplift and business value (Xu et al., 2016). These innovations complement deep learning architectures that capture high-dimensional user and item representations by situating them within a decision-theoretic framework designed to optimize sequential outcomes for each segment and journey state. In product marketing, the same principles extend well beyond content recommendation: dynamic micro-segmentation can inform differentiated messaging, customized feature bundles, or adaptive onboarding flows, while decision policies govern the selection of interventions that maximize engagement, retention, or product usage for targeted cohorts. The unifying thread across these methods is the idea of data-driven assignment under uncertainty, oriented toward outcomes that matter for both customer experience and sustainable growth. Critically, embedding such approaches into experimentation platforms enables firms to refine segments dynamically, measure heterogeneous treatment effects in practice, and continuously adapt policies as segment definitions evolve with market conditions and customer behavior (Davenport et al., 2020; Huang & Rust, 2018).

Causal measurement represents a critical foundation for AI-driven marketing because even the most accurate predictive models may fail to generate genuine incremental value. Evidence from large-scale field experiments reveals that observational methods often inflate the estimated impact of advertising when compared with randomized controlled trial (RCT) benchmarks, highlighting the gap between prediction and true causation (Gordon et al., 2019). Further research linking online exposures to offline purchases demonstrates that a substantial share of incremental sales actually comes from individuals who never click on advertisements, emphasizing the limitations of surface-level metrics such as clicks or impressions in capturing real business outcomes (Lewis & Reiley, 2014). Econometric contributions have provided rigorous frameworks for the design and interpretation of randomized experiments, offering guidelines that extend beyond academic contexts to inform

product marketing practices, particularly in the evaluation of segment-based interventions (Athey & Imbens, 2017). For practitioners, the implication is clear: refined segmentation strategies must be judged not by immediate predictive accuracy but by uplift in long-term behavioral changes such as engagement, loyalty, and revenue persistence. Embedding structured experimentation directly into personalization pipelines transforms marketing into a process of continual learning, where heterogeneous treatment effects can be systematically observed and incorporated into evolving segment definitions. In operational terms, the integration of uplift modeling with bandit experimentation creates a closed-loop system in which algorithms propose interventions, controlled trials arbitrate their effectiveness, and segments are dynamically adjusted based on observed evidence of value creation. This iterative, causal perspective ensures coherence between algorithmic assignments and the broader goals of customer experience and segmentation, providing confidence that AI outputs align with measurable improvements in customer relationships and financial performance across diverse markets (Gupta et al., 2004; Neslin et al., 2006; Verbeke et al., 2012).

Figure 1: Customer Experience and Market Segmentation in Product Marketing



Ethical, legal, and interpretability concerns form an indispensable framework for the global deployment of AI-driven product marketing, shaping both strategic possibilities and operational boundaries. Insights from the economics of privacy literature establish that personal information is not merely a byproduct of digital interactions but a valuable asset with quantifiable worth, and that consumers constantly negotiate trade-offs between the benefits of disclosure and the risks to autonomy and security (Fader et al., 2005). In the retailing domain, scholarship highlights the delicate tensions among consumers, firms, and regulators, underscoring that personalization must be carefully balanced with privacy expectations if trust and long-term relationships are to be sustained (Martin et al., 2020). Within this space, interpretability emerges as a vital safeguard, ensuring that complex AI systems remain accountable and comprehensible to diverse stakeholders. Techniques such as model-agnostic local explanations, which reveal decision logic at the individual level (Ribeiro et al., 2016), and additive feature attributions, which quantify the relative contribution of variables to

predictions (Lundberg & Lee, 2017), provide practical pathways for rendering black-box models intelligible. Expanding on these, surveys of explainability methods chart the wider design space for transparency, highlighting opportunities to integrate interpretability seamlessly into model development and deployment practices (Guidotti, 2019). For product marketing, these strands converge to offer normative and operational clarity: segmentation can be conducted with privacy-preserving protocols that safeguard consumer data; predictive models can be systematically audited to ensure reasonableness, fairness, and the absence of discriminatory bias; and explanation mechanisms can be embedded into organizational workflows to foster governance, accountability, and stakeholder trust. Embedding such requirements into end-to-end data and model pipelines establishes defensible and human-understandable processes, enabling firms to operate responsibly across industries and jurisdictions while ensuring that AI-enabled customer experience and segmentation remain ethically grounded and legally compliant.

The objective of this review is to produce a comprehensive, methodical account of how artificial intelligence is used in product marketing to enhance customer experience and refine market segmentation, and to do so through a transparent protocol that can be replicated. Specifically, the review aims to: (1) systematically locate and screen peer-reviewed studies published up to and including 2021 that explicitly connect AI techniques to product-marketing decisions with measurable outcomes related to experience quality, engagement, retention, or economic performance; (2) develop a clear taxonomy of AI approaches relevant to segmentation and experience orchestration, spanning clustering and representation learning, recommender and next-best-action systems, uplift and propensity modeling, price and promotion optimization, journey analytics, and explainability methods; (3) catalogue the data foundations that enable these approaches, including first-party and third-party sources, identity resolution, feature engineering, latency requirements, and governance practices; (4) extract and normalize the evaluation metrics used across studies such as discrimination metrics, ranking metrics, uplift, incremental conversion, churn, and lifetime value to facilitate like-for-like comparisons; (5) document study designs and inference strategies, distinguishing randomized experiments, quasi-experimental designs, and observational analyses, and recording how each addresses confounding, heterogeneity, and generalizability across industries and regions; (6) summarize the reported procedures for privacy protection, fairness assessment, monitoring, and interpretability as they pertain to product-marketing use cases; (7) assemble an evidence map that shows the distribution of topics, data types, algorithms, sectors, and reported effects over time; and (8) synthesize these elements into a developed conceptual model that links data readiness to AI capability, insight quality, marketing actions, and outcomes through explicitly identified mediators and moderators. The scope restricts inclusion to English-language, peer-reviewed articles that report methods and results connected to product-marketing decisions, with exclusion of purely technical work lacking a marketing application, editorial essays without methods, and duplicate records. The review follows a predefined search and screening workflow across major bibliographic databases, applies dual independent coding with interrater checks, and uses structured extraction templates to ensure consistency. The intended result is a rigorous, consolidated foundation that clarifies what has been studied, how it has been measured, and where the strongest evidence resides in relation to AI-driven insights for product marketing.

LITERATURE REVIEW

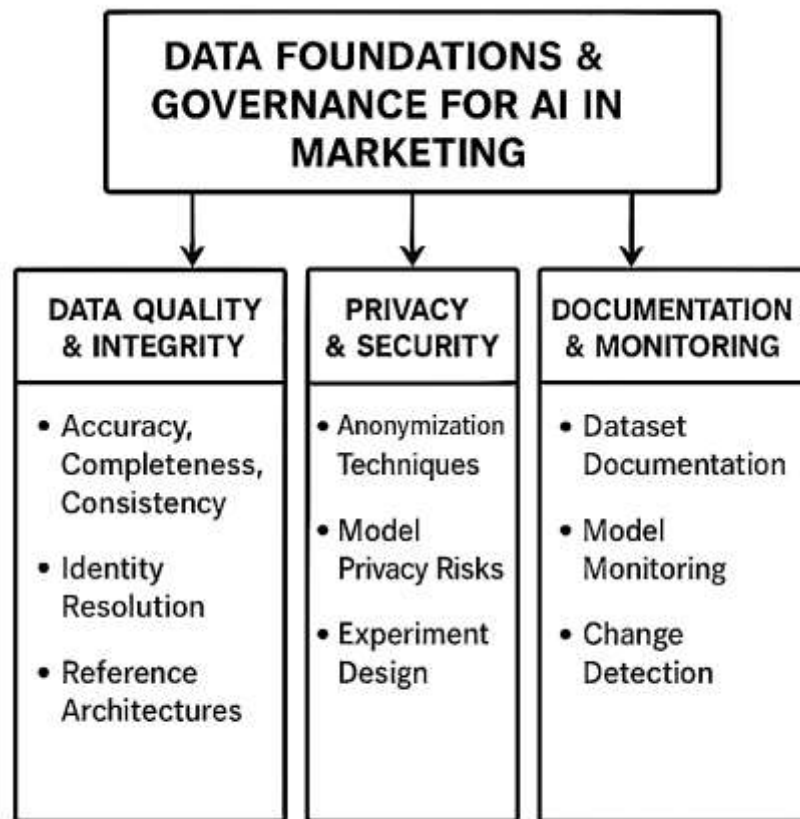
The literature on AI-driven product marketing spans several intersecting streams that together explain how data, models, and decision processes reshape customer experience and market segmentation. At its core, this body of work treats segmentation not as a static, survey-based partition but as a dynamic, behaviorally grounded construct that can be updated continuously as customers interact with products and channels. Parallel research on personalization operationalizes this shift by learning granular preferences and propensities, enabling next-best-action policies that select messages, offers, feature bundles, or service interventions for specific micro-segments along the journey. A complementary stream centers on customer-journey analytics, where sequence models, process mining, and voice-of-customer techniques translate longitudinal interactions and unstructured feedback into states, transitions, and bottlenecks that matter for experience quality. Pricing and promotion scholarship contributes optimization frameworks for demand shaping and revenue consistency, while customer-base modeling provides tools for projecting retention and lifetime value so that segments can be prioritized by long-run contribution rather than short-term response. Across

these domains, recent marketing analytics emphasizes causal identification and experimentation A/B tests, uplift modeling, and quasi-experimental designs to distinguish predictive fit from incremental impact and to uncover heterogeneous treatment effects that redefine actionable segments. Underpinning all of this is a growing focus on data foundations and governance: first-party data capture, identity resolution, latency and freshness requirements, and controls for consent, privacy, and fairness. Methodologically, studies range from classical clustering and latent class approaches to representation learning, sequence-aware recommenders, contextual bandits, and interpretable modeling techniques that expose drivers of predictions for practitioner review. From a systems perspective, deployment research highlights feature stores, online/offline parity, monitoring for drift, and experimentation platforms that close the loop between model outputs and marketing actions. As a corpus, the literature is heterogeneous in data sources, industries, metrics, and study designs, which complicates direct comparison but offers a rich basis for synthesis. This review positions these streams within a single organizing frame for product marketing: data readiness enabling AI capability; AI capability producing insight quality; insights informing positioning, targeting, pricing, and experience design; and those actions yielding measurable outcomes in satisfaction, engagement, retention, and financial performance, with transparency and governance shaping feasibility across contexts.

Data Foundations & Governance for AI in Marketing

High-quality, well-governed data is the substrate on which AI-driven product marketing depends, shaping everything from feature engineering to the validity of segmentation and customer-experience metrics. Foundational work on data quality established that “fitness for use” spans multiple dimensions accuracy, completeness, consistency, timeliness, and interpretability each of which can constrain or enable downstream modeling and decision support in marketing contexts (Wang & Strong, 1996). Complementing this perspective, practical assessment frameworks emphasize profiling and scoring pipelines that diagnose defects at acquisition and integration points, so that models are not trained on corrupted identities, stale attributes, or biased samples (Pipino et al., 2002). As firms fuse clickstream, CRM, and product-usage telemetry to build first-party data assets, the reliability of identity resolution and cross-device stitching becomes decisive for constructing longitudinal customer views that support dynamic segmentation and journey analytics. Yet measurement studies of the online advertising and tracking ecosystem show the technical complexity of data collection at scale, including third-party tracking scripts and cookie-synching practices that complicate provenance, consent management, and compliance auditing (Englehardt & Narayanan, 2016). These realities make consented first-party data capture and rigorous governance more than procedural checklists; they are design constraints that determine what can be modeled, for whom, and under what legal bases. In parallel, marketers need reference architectures clear schemas, entity definitions, and lineage metadata that keep evolving product and channel data interoperable and explainable across teams. Absent such foundations, even sophisticated AI methods can mis-segment populations, overfit to noisy identifiers, and produce spurious lift estimates when campaigns are targeted or evaluated on fragmented or low-quality records.

Figure 2: Data Foundations and Governance Framework for AI in Marketing



Privacy-preserving data technologies and formal guarantees of confidentiality are fundamental to building responsible data foundations in AI-driven product marketing, particularly because modern marketing datasets are dense with behavioral, transactional, and contextual signals that are often high-dimensional and inherently difficult to anonymize. Classical anonymization methods such as k-anonymity, while historically influential, are limited in practice because quasi-identifiers in consumer records are often numerous and correlated, which makes safe de-identification fragile and vulnerable to reidentification risks once external linkages are introduced (Sweeney, 2002). To address these weaknesses, formal privacy frameworks have emerged that mathematically define disclosure risks and allow quantifiable trade-offs between data utility and individual protection, offering systematic mechanisms for analytics pipelines where repeated queries or compositional uses of data can erode safeguards if unmanaged (Dwork & Roth, 2014). Yet even beyond raw data, the growing body of research on machine learning security demonstrates that models themselves can inadvertently leak sensitive information. For instance, trained parameters and output probabilities can be exploited through techniques like membership inference, whereby adversaries are able to determine if specific individuals contributed to a training dataset, even when only black-box prediction access is available (Shokri et al., 2017). For marketers leveraging such models, the implication is clear: governance cannot stop at the dataset but must encompass the entire model lifecycle, including training, validation, deployment, and post-deployment monitoring. This necessitates the use of privacy budgets, aggregation methods, audit trails, and selective retention policies that minimize exposure while still retaining the statistical signal essential for accurate segmentation, propensity modeling, and uplift analysis. It also requires designing experimentation frameworks where identifiers used for randomization and holdouts are carefully shielded, and ensuring that dashboarding and reporting layers do not inadvertently create avenues for reconstruction attacks against small or granular consumer cohorts. Taken together, these practices

ground personalization strategies in architectures that protect individuals while sustaining analytical power.

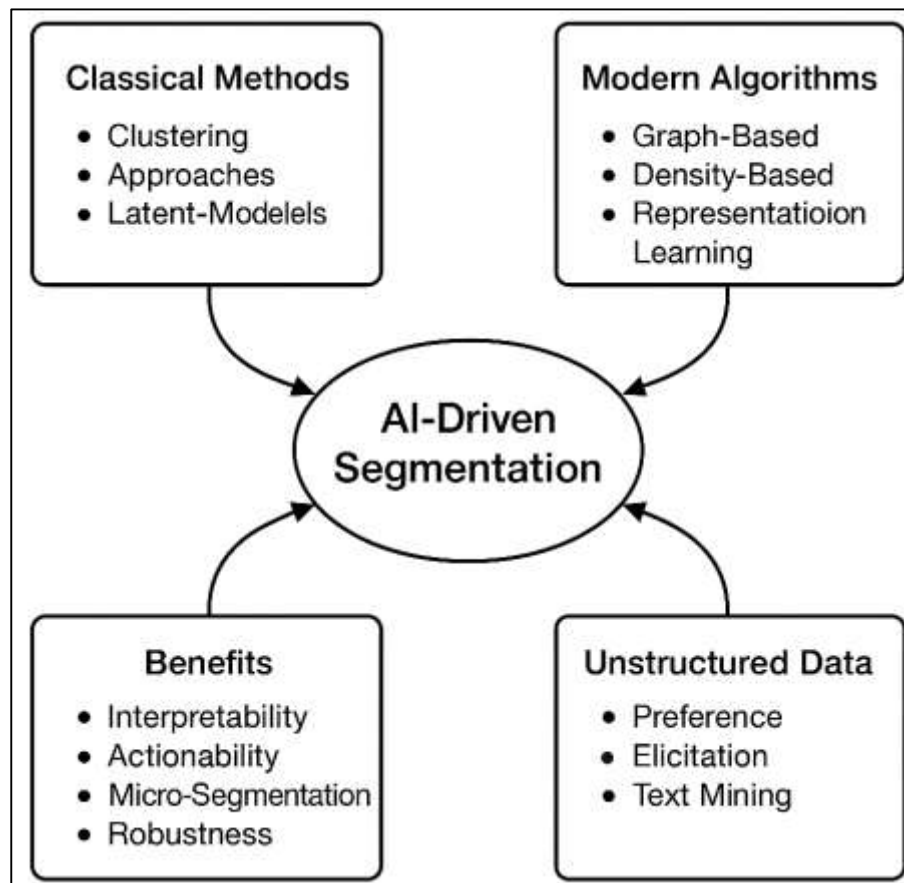
Governance in AI-driven product marketing extends beyond technical precision to embrace accountability, transparency, and documentation, ensuring that datasets and automated decisions remain intelligible and defensible to both internal teams and external stakeholders. Legal analyses of algorithmic decision-making underscore that explainability and contestability are not optional aspirations but requirements embedded within evolving data-protection frameworks, shaping how organizations must justify profiling, targeting, and personalization practices in compliance with regulatory standards (Wachter et al., 2017). To make accountability operational, dataset documentation practices such as standardized “datasheets” have been proposed, which capture critical details on data provenance, collection environments, intended purposes, limitations, and ethical considerations, providing marketing teams with the contextual knowledge needed to repurpose consumer logs, survey responses, and behavioral traces across diverse campaigns and international markets (Gebru et al., 2021). Reliability further complicates governance, as real-world marketing environments are inherently nonstationary: product lines evolve, prices fluctuate, seasonal dynamics intervene, and consumer preferences shift, meaning that segmentation models, recommendation engines, and propensity scores degrade in accuracy unless supported by systematic monitoring. Population-shift diagnostics, recalibration procedures, and retraining triggers thus become indispensable for sustaining model validity under changing conditions (Gama et al., 2014). When viewed together, these strands suggest that strong data foundations require three tightly interwoven layers: first, quality management and identity integrity that guarantee interpretable and consistent records; second, robust privacy and security measures that protect both individual data and model outputs from adversarial or linkage threats; and third, documentation and monitoring infrastructures that make data assets auditable and predictive systems maintainable throughout their lifecycles and across jurisdictions (Malhotra et al., 2004). Embedding these layers within the organizational data platform empowers marketers to trust that derived segments mirror genuine behavioral structures, that campaign interventions are based on lawful and ethically sound inputs, and that performance metrics capture enduring signals rather than temporary noise arising from pipeline instability or population drift.

AI for Market Segmentation

Classical market segmentation laid the conceptual foundation for modern AI-driven approaches by articulating how diverse and heterogeneous consumers can be partitioned into relatively homogeneous groups that are both interpretable and actionable for product and marketing decisions (Ara et al., 2022). Early contributions in marketing scholarship drew attention to the importance of clustering methodologies, highlighting those choices among hierarchical, partitioning, or hybrid strategies, as well as the selection of distance metrics and stopping rules, could significantly influence the validity of segment structures and, in turn, the managerial insights derived from them (Green & Krieger, 1995; Jahid, 2022). The field later advanced from heuristic clustering toward model-based segmentation, reframing segmentation as a statistical problem of latent heterogeneity in which observed preferences and behaviors are viewed as arising from a finite mixture of subpopulations. Estimation within this paradigm is carried out through the expectation–maximization (EM) algorithm, which provides a general and powerful mechanism for handling missing or latent variables and yields posterior probabilities of segment membership that can be interpreted with managerial clarity (Dempster et al., 1977; Uddin et al., 2022). Within marketing applications, mixture models and random-effects formulations enable the balancing of within-segment cohesion, between-segment distinctiveness, and parsimony, while simultaneously accommodating covariates, measurement errors, and the noise that inevitably arises in field data (Allenby & Rossi, 1999; Akter & Ahad, 2022). The managerial payoff of such approaches lies in their interpretability: firms can profile consumers with probabilistic diagnostics, link segments to tailored product positioning, refine messaging, and optimize channel strategies in a structured manner. Moreover, model-based frameworks provide rigorous tools for evaluating competing segmentations by employing criteria such as likelihood-based fit indices and holdout performance, reducing the risk of spurious clusters distorting strategic decisions. Collectively, these contributions establish a robust baseline against which newer graph-based, density-based, and representation-learning techniques

may be judged for their added value in addressing today's dynamic, high-dimensional marketing contexts (Md Arifur & Noor, 2022).

Figure 3: Theoretical Framework of AI-Driven Market Segmentation



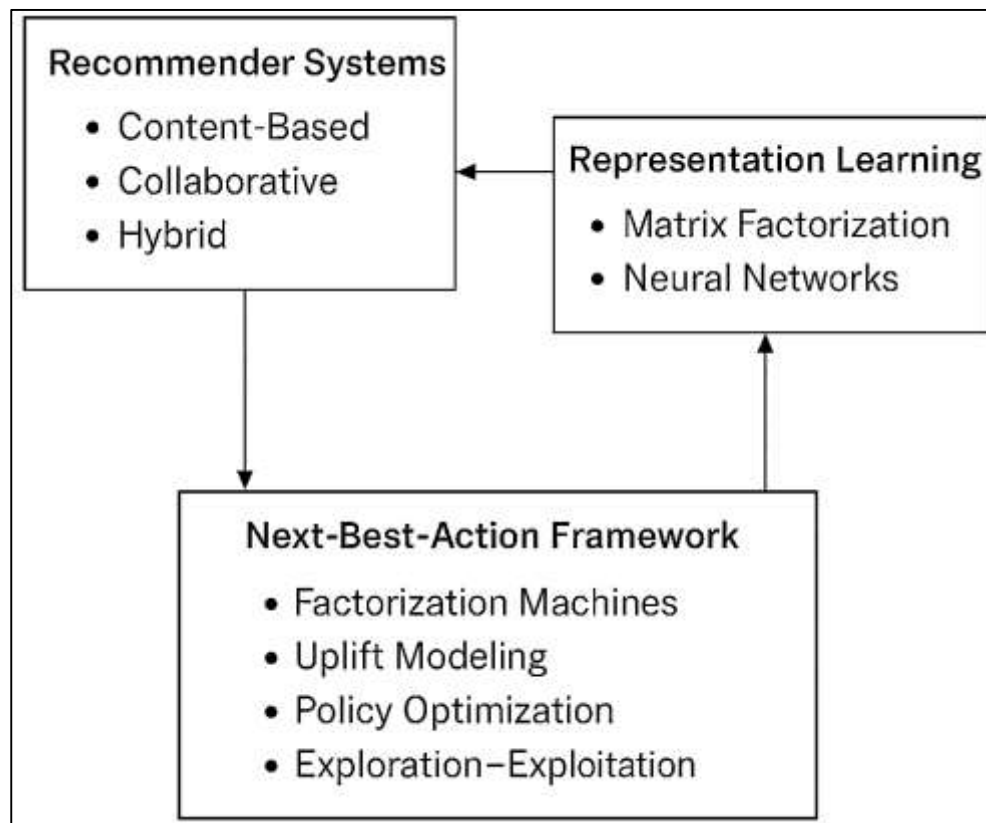
As the volume and variety of digital traces expanded, segmentation research evolved to adopt algorithms capable of capturing complex, nonconvex structures and manifold-like relationships inherent in behavioral data. Spectral clustering offers a powerful alternative to conventional distance-based methods by performing an eigen-decomposition of a similarity graph, thereby uncovering community structures that persist even when clusters are intertwined, irregularly shaped, or uneven in size features typical of clickstreams, usage telemetry, and social interactions (Rahaman, 2022; von Luxburg, 2007). Complementing this, HDBSCAN extends density-based clustering to accommodate variable-density environments, isolating stable clusters while designating ambiguous or sparse records as noise, which prevents overconfident or misleading assignment of customers in operational marketing contexts (McInnes et al., 2017; Hasan et al., 2022). Representation-learning methods further integrate feature construction with segmentation by jointly optimizing latent-space embeddings and clustering objectives, as exemplified by Deep Embedded Clustering. This approach often produces segments that align more closely with downstream business goals than those derived from hand-crafted features, allowing firms to discover nuanced behavioral or contextual groupings that are directly actionable (Hossen & Atiqur, 2022; Xie et al., 2016). In operational terms, these techniques enable micro-segmentation at scale, moving beyond coarse demographic or simple RFM-based aggregations to identify clusters defined by behavioral motifs, content affinities, or journey stages while preserving robustness through noise-aware assignments and explicit outlier handling. The learned representations also enhance portability, encoding higher-order relationships among users, items, and contexts so that segment definitions can generalize across products, channels, and campaigns. For product marketing, this portability translates into practical advantages: targeting rules, creative variants, and sequencing policies can be applied consistently across locales and platforms without requiring bespoke engineering for each

deployment, while maintaining fidelity to the underlying behavioral structure. Together, these algorithmic and representation-based advances enable firms to implement scalable, precise, and robust segmentation pipelines that directly support personalized engagement and strategic decision-making (Tawfiqul et al., 2022).

AI-driven segmentation increasingly integrates preference elicitation and unstructured data, expanding the evidence base that defines both who belongs to a segment and what each segment values. Conjoint-based product design illustrates this integration: latent-class models simultaneously reveal taste segments and their part-worth structures, enabling segment-specific positioning, feature bundling, and offering designs grounded in observed choice behavior rather than proxies (DeSarbo et al., 1992; Kamrul & Omar, 2022). Beyond structured experiments, market-structure and persona discovery exploit large-scale text mining, using semantic networks induced from consumer discourse to identify clusters of needs, competing frames, and brand associations, producing actionable segments without requiring new surveys (Mubashir & Abdul, 2022; Netzer et al., 2012). Critically, segmentation is no longer merely descriptive but becomes a lever for differential treatment. Empirical studies indicate that targeting solely based on risk or propensity may fail; for instance, directing retention efforts to “high-risk” customers can yield negligible or negative incremental outcomes compared with strategies that account for heterogeneous treatment effects (Ascarza, 2018; Reduanul & Shoeb, 2022). This underscores the importance of coupling segmentation with uplift or causal modeling so that membership conveys distinct responsiveness rather than mere similarity. Operationally, combining latent-class or representation-based segments with experimentation and incrementality metrics transforms segments from static labels into decision surfaces that guide product messaging, onboarding sequences, pricing, and service recovery, thereby generating measurable improvements in experience and performance. Modern segmentation therefore synthesizes heterogeneous data sources, including choice tasks, behavioral logs, and textual content, with diverse methodological approaches ranging from finite mixtures and latent-class models to graph- or density-based clustering and deep embeddings. The result is a set of segments that are interpretable, stable enough for operational planning, and predictive of differential response under real interventions, supporting data-driven product marketing and personalization at scale (McLachlan & Peel, 2000; Sazzad & Islam, 2022).

Personalization & Next-Best-Action

Personalization in product marketing is fundamentally about deciding what to offer, to whom, and when, with the decision logic anchored in observed preferences, behaviors, and interactions rather than broad demographic assumptions. Early research on recommender systems framed this challenge as learning user-item relevance from explicit ratings and implicit behaviors, distinguishing content-based, collaborative, and hybrid approaches while highlighting their respective strategies for addressing sparsity, novelty, and serendipity, which remain critical considerations when mapping product features, messaging, and interventions to micro-segments (Adomavicius & Tuzhilin, 2005; Sheratun Noor & Momena, 2022). Hybridization further demonstrated that combining complementary evidence channels mitigates cold-start issues and enhances robustness across diverse catalogs, platforms, and stages of the customer journey, an advantage that proves especially important when products are positioned differently across channels or geographic locales (Adar & Md, 2023; Burke, 2002). The advent of matrix-factorization models offered a compact representation of latent tastes and item attributes by decomposing high-dimensional, sparse interaction matrices into low-rank factors, enabling industrial-scale personalization with tractable training, efficient storage, and rapid online scoring (Qibria & Hossen, 2023; Koren et al., 2009). As the field evolved, attention shifted toward sequence-rich and sessionized data, emphasizing temporal and order effects that capture how recent exposures, co-consumption patterns, and evolving needs shape immediate receptivity. This perspective culminated in sequence-aware models that treat recommendations as state-dependent decisions along the customer journey rather than as static similarity lookups, effectively integrating behavioral dynamics into predictive and prescriptive frameworks (Istiaque et al., 2023; Quadrana et al., 2018).

Figure 4: Personalization and Next-Best-Action in Marketing

Collectively, these methodological advances recast personalization as a learning-to-decide problem in which the objective extends beyond predicting interest to orchestrating next-best actions including content, features, pricing, or service interventions that actively steer individual experiences, optimize segment-level outcomes, and support coherent, data-driven product-marketing strategies at scale (Akter, 2023). Contemporary next-best-action (NBA) pipelines embody an integration of representation learning and adaptive decision policies that evolve as evidence accrues, enabling product marketing to respond dynamically to heterogeneous customer needs and changing contexts. Factorization machines extend traditional linear predictors by efficiently capturing pairwise interactions among high-dimensional, sparse features such as user demographics, behavioral signals, contextual cues, creative variants, and product attributes, providing a practical backbone for ranking, allocation, and prioritization across advertisements, recommendation feeds, and on-site modules (Hasan et al., 2023; Rendle, 2010). Beyond linear interactions, deep learning architectures have expanded the scope of personalization by modeling nonlinear user-item relationships directly from implicit feedback. Neural collaborative filtering, for instance, replaces inner-product approximations with multi-layer perceptrons, allowing the system to capture complex, non-additive patterns in preferences, often improving top-N recommendation accuracy in settings characterized by rich interaction structures and latent dependencies (He et al., 2017; Masud et al., 2023). Deployment at scale, however, requires infrastructure that can accommodate distributional shifts in inventory, audience composition, and business objectives, motivating the use of online learning frameworks that update incrementally as new data arrive. Large-scale click-through pipelines exemplify this approach through proximal adaptive optimization for sparse logistic models, demonstrating that frequent parameter updates combined with calibrated regularization maintain stable predictive lift even under substantial churn in both candidate sets and feature distributions (McMahan et al., 2013; Sultan et al., 2023). For product marketers, these advances translate into NBA programs that evaluate a broad slate of candidate actions per user, select an actionable subset within operational constraints, observe outcomes in near real time, and rapidly update models to reflect emergent patterns. The optimization process is inherently multi-objective, balancing

predicted engagement, revenue, fairness, and experiential quality, which necessitates combining calibrated predictive models with policy layers that encode operational guardrails while preserving sufficient exploration to discover high-value actions for evolving or previously unobserved segments (Hossen et al., 2023).

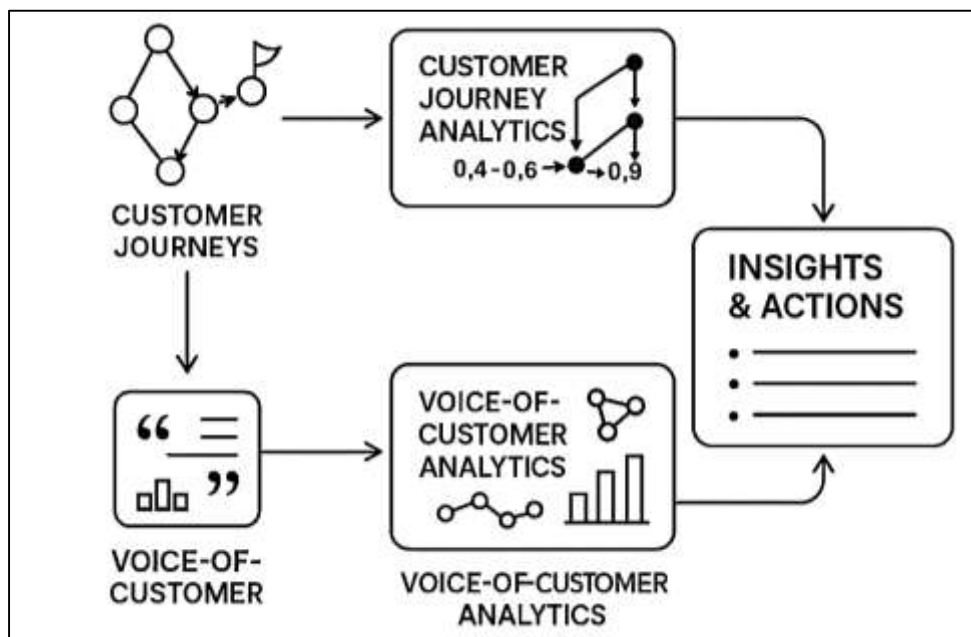
Crucially, the effectiveness of next-best-action programs depends not only on identifying similar customers but on learning segments that respond differentially to interventions, making causal thinking central to the personalization stack. Uplift modeling reframes the objective from predicting “who buys” to estimating “who buys because of an intervention,” quantifying treatment–control differences at the individual or segment level and providing decision-tree frameworks that supply interpretable targeting rules and diagnostics to prioritize incremental impact (Tawfiqul, 2023; Rzepakowski & Jaroszewicz, 2012). Insights from early direct-marketing research corroborate this principle, showing that targeting solely based on purchase likelihood or risk often wastes resources, whereas modeling the expected gain from an action produces more efficient contact policies and clearer economic rationale for segment selection (Hansotia & Rukstales, 2002; Shamima et al., 2023). Given that uplift estimates are inherently noisy, context-sensitive, and may drift over time, exploration–exploitation mechanisms remain indispensable to prevent premature convergence on suboptimal actions. Bayesian methods such as Thompson sampling offer computationally light approaches with theoretical regret guarantees, adapting dynamically as response distributions vary across users and temporal conditions (Ashraf & Ara, 2023; Russo et al., 2018). Within a complete product-marketing infrastructure, these elements interlock seamlessly: representation-learning modules uncover latent structure in behaviors and content, ranking and allocation models assign scores to candidate actions per user, uplift estimators and controlled experiments reveal heterogeneous treatment effects, and bandit algorithms manage uncertainty while respecting operational and business constraints (Sanjai et al., 2023). The resulting NBA framework transforms segment definitions from static clusters of similarity into actionable cohorts defined by responsiveness, enabling marketers to orchestrate customer journeys with quantifiable, incremental improvements in engagement, retention, and contribution margin. By embedding causal and adaptive reasoning at every stage from feature representation to allocation policy this approach ensures that personalization programs not only predict behavior but systematically enhance outcomes, aligning marketing interventions with measurable business impact while maintaining interpretability and operational scalability (Akter et al., 2023).

Customer Journey Analytics & Voice-of-Customer

Customer journey analytics synthesizes fragmented touchpoints into coherent, causal narratives that illuminate how audiences discover, evaluate, adopt, and repeatedly engage with offerings, transforming raw interaction data into actionable insights for marketing orchestration. A central challenge lies in attributing outcomes not to any single channel or exposure but to the sequence and mixture of interactions, which requires moving beyond simplistic last-touch heuristics toward path-sensitive models. Graph-based attribution approaches conceptualize journeys as probabilistic networks, where each node and transition reflects the marginal contribution of a touchpoint while accounting for order and interdependence, enabling marketers to quantify how cumulative exposures drive conversions (Anderl et al., 2016). Empirical studies demonstrate that upper-funnel awareness media shape early-stage movement, whereas search, onsite interactions, and retargeting consolidate intent, highlighting that attribution must accommodate cross-channel spillovers rather than treating interactions in isolation (Li & Kannan, 2014). Advertising measurement critiques reinforce this perspective by showing that heuristic allocation rules can bias budget decisions, whereas principled attribution relies on counterfactual reasoning about unobserved paths and unrealized exposures (Berman, 2018). Data across display, search, email, social, and affiliate channels reveal systematic variation in touchpoint effectiveness depending on the consumer's journey stage, competitive pressures, and creative context, emphasizing the importance of segment-aware, path-aware modeling (De Haan et al., 2016). Operationally, these insights translate into policies for creative rotation, channel pacing, and offer sequencing that respect the dynamic state of the customer, where the state encodes recency, frequency, and evolving preferences rather than relying on static audience labels. By treating journeys as stochastic processes, marketers can design next-interaction strategies that optimize persuasion while preserving experience quality, aligning message cadence, content specificity, and channel selection with the user's current

objectives and friction points, ultimately supporting more effective and responsive customer engagement.

Figure 5: Theoretical Framework of Customer Journey Analytics and Voice-of-Customer



Beyond traditional attribution, contemporary customer journey analytics increasingly leverages operational data science to model, monitor, and predict process states at scale, transforming raw event logs into actionable insights that guide interventions and policy learning. Process mining provides a systematic framework for discovering the as-is journey, evaluating conformance against desired paths, and quantifying bottlenecks, rework, and dropout nodes that correspond to pain points across onboarding, trial-to-paid conversion, or service recovery, thereby offering a granular view of where friction impedes customer progress (Aalst, 2016). When integrated with multistage response estimation, these models allow organizations to measure how exposures at one phase influence transitions at subsequent stages and to identify points where incremental investment produces the greatest downstream impact, creating a quantitative basis for resource allocation and prioritization (Abhishek et al., 2015). In practice, combining conformance analytics with multistage lift estimation equips product marketing and lifecycle teams to orchestrate interventions with precision: for instance, a user stalled between evaluation and activation after consuming technical documentation may be nudged toward completion with a targeted tutorial, whereas a customer cycling repeatedly through pricing pages and competitive comparisons may benefit from value reassurance messaging. Importantly, journey models updated from live logs function as early-warning systems for drift, detecting unexpected shifts in path distributions triggered by product updates, policy changes, or market shocks, and informing the design of A/B tests by highlighting states with high uncertainty about optimal actions. In this manner, customer journeys evolve from descriptive maps into control surfaces that anchor hypotheses, stratify experimentation, and enable interpretable, data-driven prioritization of fixes and messaging strategies, ultimately reducing time-to-value and fostering sustained engagement.

Voice-of-customer analytics complements customer journey analysis by illuminating the underlying motivations, perceptions, and barriers that shape observed behaviors, transforming raw pathways into interpretable narratives that guide action. Foundational research in opinion mining demonstrates that evaluative language extracted from reviews, forums, and social media can be converted into structured features linked to conversion, churn, and advocacy outcomes, providing a bridge between qualitative expression and quantitative performance metrics (Liu, 2012; Pang & Lee, 2008). Beyond simple sentiment, recent advances in needs discovery reveal that product attributes and latent customer "jobs" can be inferred from large-scale user-generated content,

enabling marketers to map specific themes, such as onboarding clarity, performance reliability, or pricing fairness, to discrete journey stages and quantify their marginal influence on engagement and demand (Archak et al., 2011; Timoshenko & Hauser, 2019). Integrating VoC with journey analytics closes a critical loop: text-derived signals provide explanatory power for transitions observed in behavioral data, while state-based models contextualize sentiment, recognizing that identical descriptors can carry opposite connotations depending on the user's stage or experience level. This alignment enhances segmentation, as clusters based on shared needs, frustrations, and expectations are more actionable and stable than those defined solely by demographics or coarse behavioral aggregates. Operationally, organizations can channel VoC insights toward targeted levers, including product improvements, content enhancements, support interventions, or pricing experiments, and measure the resulting shifts in journey paths to ensure interventions address root causes rather than superficial symptoms. The resulting analytic architecture links journey states to measurable transitions and VoC themes to perceived value, creating a coherent framework for discovery, prioritization, and evaluation that grounds personalization, micro-segmentation, and experience optimization in both articulated customer needs and observed behavioral patterns.

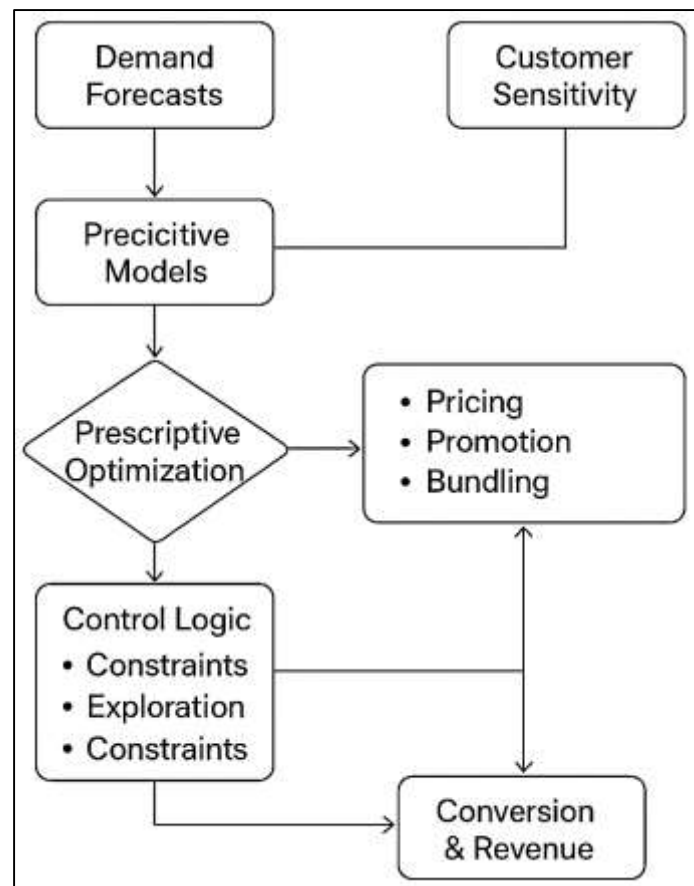
Pricing, Promotion & Offer Optimization

Pricing, promotion, and offer optimization operationalize AI-driven insights by converting predictive intelligence into actionable levers that directly influence demand, revenue, and the customer experience at the point of decision. The foundations of this practice are rooted in revenue-management and discrete-choice frameworks, which articulate how price, assortment, and availability interact with consumer preferences under conditions of scarcity, competition, and heterogeneity. In perishable or capacity-constrained contexts, dynamic pricing strategies navigate the trade-off between inventory risk and stochastic demand, demonstrating that carefully calibrated, time-varying price paths can maximize expected revenue while distributing availability across segments with differing willingness-to-pay (Gallego & van Ryzin, 1994). More broadly, revenue-management theory integrates forecasting, protection levels, price fences, and overbooking into a cohesive toolkit, linking prescriptive optimization routines with operational controls that marketers and merchandisers can implement consistently across channels and product categories (Talluri & van Ryzin, 2005). On the estimation side, structural demand models translate market-level outcomes into micro-level choice insights, enabling counterfactual analyses that assess how alternative prices, feature bundles, and competitive moves would shift market shares; this provides the causal underpinning for segment-sensitive pricing and for evaluating whether proposed offers are both appealing and economically sustainable (Berry et al., 1995). Complementing these approaches, learning-to-price methodologies formalize adaptive policies in which firms adjust pricing in real time while simultaneously inferring unknown demand parameters, establishing theoretical guarantees such as regret bounds and supporting practical responsiveness in fast-moving or nascent markets where historical elasticities are unreliable (den Boer, 2015). Together, these strands position AI as an integrative decision engine, embedding predictive demand models within optimization policies that translate real-time signals into price and offer configurations aligned with segment-specific preferences, behavioral constraints, and operational objectives, ensuring rigor, responsiveness, and strategic control across diverse marketing contexts.

When uncertainty about demand is high, price experimentation and online learning become indispensable tools for offer optimization, allowing firms to adapt dynamically to evolving market conditions. Algorithms that frame pricing as a sequential decision problem demonstrate that near-optimal revenue can be achieved even without complete knowledge of the underlying demand curve, provided that exploration is carefully managed and updates are informed by observed sales responses; this approach is particularly suited to e-commerce contexts, where frequent catalog churn and context-dependent effects are the norm (Besbes & Zeevi, 2009). Beyond uncertainty, behavioral reference points shape perceived value, as customers anchor on past or posted prices, and deviations from these benchmarks can enhance or depress willingness-to-pay in ways not captured by static elasticity models. Incorporating reference effects into dynamic pricing policies improves predictive fit and realized performance, especially for products with repeated purchases or salient historical prices (Cohen et al., 2016). Interactions with inventory and promotions further complicate the landscape, as customers may stockpile during discounts or shift timing to exploit anticipated deals, making it essential for price paths to account for intertemporal substitution and

household-level stockpiling to avoid artificial spikes and subsequent troughs that degrade both experience and forecasting accuracy (Hendel & Nevo, 2006). Operationally, these insights manifest in multi-objective control frameworks that rank candidate prices and bundles by expected revenue and conversion while penalizing options that violate experiential constraints such as excessive volatility, with outcomes logged for continuous re-estimation of elasticities and cross-effects. For product marketers, this creates “learning while earning” systems capable of refreshing elasticities at the micro-segment level, integrating contextual features such as traffic source, device, and tenure, and coordinating pricing with messaging to ensure that changes are framed, justified, and timed to reduce friction, reinforce perceived fairness, and maximize both engagement and revenue outcomes.

Figure 6: Pricing, Promotion, and Offer Optimization



Promotions and personalized offers serve as strategic complements to pricing by enhancing perceived value through non-price levers such as temporary discounts, coupons, bundles, and communications that emphasize relevance and fit. Classic promotion science highlights both short-term spikes and longer-term effects, demonstrating that while deals can generate immediate lift, their cumulative impact shapes habit formation, brand equity, and competitive positioning, necessitating evaluation frameworks that capture both instant response and persistence across purchase cycles (Mela et al., 1997). Comprehensive treatises on sales promotion provide structured taxonomies of promotional vehicles and strategic objectives including trial generation, purchase acceleration, and brand switching together with design guidelines for conditions, depth, timing, and targeting, which directly inform algorithmic policy spaces in contemporary AI-driven marketing systems (Blattberg & Neslin, 1990). Personalization research builds on these foundations by extending promotion design to the individual level, employing preference-based “e-customization” techniques to tailor incentives, messages, and content to heterogeneous consumer tastes, thereby improving spend efficiency, reducing adverse selection, and minimizing promotional leakage (Ansari & Mela, 2003). In operational stacks, offer optimization is formalized as a constrained allocation problem:

models predict incremental lift for candidate incentives, solvers select the subset that respects budgetary and fairness constraints, and controlled experimentation measures heterogeneous treatment effects before broad deployment. Detailed logging schemas link exposures to journey states, enabling continuous re-optimization to prevent customer fatigue, maintain perceived value, and prioritize segments with high expected responsiveness rather than those with merely high baseline risk. Critically, the integration of promotions with customer experience is explicit: cadence, creative execution, and channel selection are coordinated with product positioning to ensure that short-term promotional lift does not undermine trust, induce undesirable training effects, or encourage strategic deferment of purchases until discounts appear, thereby preserving both immediate outcomes and long-term customer relationships.

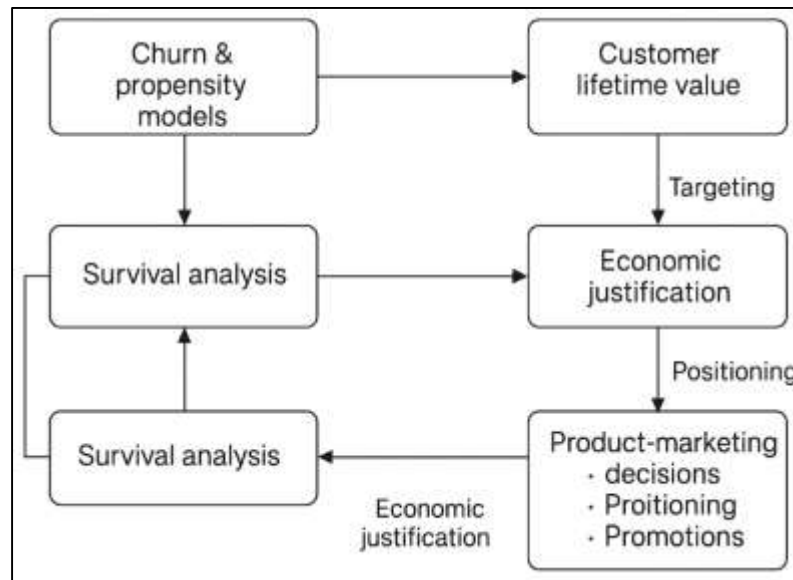
Churn, Propensity & Customer Lifetime Value

Churn, propensity, and customer lifetime value (CLV) form the economic and behavioral foundation for AI-driven product marketing by converting individual behavioral histories into forward-looking measures that guide segmentation, treatment allocation, and resource prioritization. In contractual contexts such as subscriptions or service agreements, churn is observed as the explicit termination of a relationship, whereas in noncontractual settings including most retail and many digital offerings inactivity must be probabilistically interpreted as silent defection, necessitating behavioral assumptions about purchase timing and heterogeneity. Seminal noncontractual models, such as the Pareto/NBD framework, infer whether a customer remains “alive,” estimate the expected number of future transactions, and model interpurchase times from sparse recency and frequency data, enabling marketers to assign value without continuous observation ([Schmittlein et al., 1987](#)). Research on relationship duration complements this by recognizing that the hazard of disengagement evolves with service quality and customer experience, bridging operational metrics with the longevity of commercial ties ([Bolton, 1998](#)). Profitability, however, is not a simple function of tenure, as long-lifetime customers are not necessarily the most valuable; cross-buying, dynamic margins, and service costs influence net value, cautioning against using tenure as a proxy in targeting rules ([Reinartz & Kumar, 2003](#)). CLV research integrates these insights into a comprehensive framework that treats customers as assets, valuing them according to expected margin streams, cost to serve, retention effort, and discounted uncertainty over time ([Gupta & Lehmann, 2006](#)). Customer equity theory then links individual CLV to overall firm value, emphasizing how product marketing decisions including positioning, pricing, promotions, and experience design shape the portfolio of customer assets ([Rust et al., 2004](#)). Operationally, resource-allocation models leverage CLV and propensity scores to decide who should receive interventions, the appropriate incentive levels, and which experiences merit investment, thereby aligning analytics with budget-constrained, multi-objective marketing strategies and ensuring that AI-driven decisions translate into economically meaningful actions ([Venkatesan & Kumar, 2004](#)).

Operationalizing churn and propensity in AI-driven marketing requires models that align with the underlying data-generating process while accommodating the constraints of real-world deployment. In many sectors, class labels are highly imbalanced, as the majority of customers do not churn within any short observation window, making naïve classifiers appear accurate while failing to identify at-risk segments effectively. Research demonstrates that resampling strategies, cost-sensitive learning, and synthetic minority oversampling materially enhance detection performance by preventing models from defaulting to a trivial “no churn” prediction ([Burez & Van den Poel, 2009](#)). Ensemble methods provide additional robustness: bagging and boosting of classification trees deliver better calibrated lift and more reliable rank-ordering of customers for targeted retention campaigns compared with single-tree baselines ([Lemmens & Croux, 2006](#)). Domain-specific features further increase predictive precision and operational relevance; in telecommunications, for example, usage volatility, billing anomalies, and service disruptions serve as leading indicators of defection, and integrating these into predictive pipelines enhances the actionable quality of retention interventions ([Buckinx & Van den Poel, 2005](#)). Beyond structured data, qualitative inputs such as call-center notes, complaint logs, or open-text feedback contribute incremental value, capturing early dissatisfaction signals that precede observable behavioral churn ([Coussement & Van den Poel, 2008](#)). Survival-analytic approaches complement classification by explicitly modeling time-to-churn and allowing covariates to flexibly modulate the hazard rate across the customer relationship, thereby enabling precisely timed interventions and stress-testing of retention policies

(Bolton, 1998). Together, these insights form a layered operational playbook: first, estimate short-horizon churn propensity to triage and prioritize interventions; second, deploy survival or hazard models to guide timing and intensity of outreach; and third, integrate structured and unstructured predictors into ensemble frameworks that are continuously monitored for drift, recalibrated as offers, prices, and experiences evolve, and validated to ensure that marketing actions remain effective and economically sound.

Figure 7: Churn, Propensity, and Customer Lifetime Value in Marketing



Customer lifetime value (CLV) serves as the economic linchpin that translates predictive insights from churn and response models into actionable product-marketing decisions. Central to this framework is the principle that interventions should be justified by the incremental value they generate net of costs rather than by baseline risk or likelihood of purchase alone, positioning CLV as a coherent metric that balances immediate revenue against the longer-term growth and retention of the customer base (Gupta & Lehmann, 2006). Because the relationship between customer longevity and profitability varies across contexts, combining relationship-duration or activity models with contribution and cross-buy margins allows firms to avoid over-investing in long-lived but low-value segments, ensuring resources are allocated efficiently (Reinartz & Kumar, 2003). In noncontractual settings, probabilistic models estimating whether a customer remains active and predicting the expected number of future transactions provide the frequency component of value where contractual signals are absent, supporting precise targeting and prioritization (Schmittlein et al., 1987). At a portfolio level, customer equity offers strategic guidance by linking individual-level CLV to aggregate firm value, clarifying the mix of acquisition, cross-sell, and retention efforts that most effectively expand the asset base. CLV also functions as a targeting score and a constraint in resource-allocation frameworks, enabling marketers to assign budgets across segments, offers, and channels in ways that maximize expected enterprise value while honoring fairness and experience considerations, bridging analytics with campaign planning (Venkatesan & Kumar, 2004). Integrated operationally, these components form an evidence-driven architecture in which churn and propensity models identify at-risk or responsive customers, survival analysis informs the timing of interventions, and CLV quantifies the economic justification for action. Product-marketing levers positioning, pricing, and promotions are then deployed and continuously validated against value-based goals, ensuring that decisions not only optimize engagement and experience but also contribute demonstrably to the financial health of the firm.

Explainability, Fairness & Trust

Explainability and interpretability are critical enablers for translating AI-driven insights into actionable decisions in product marketing, where choices about targeting, positioning, and experience design

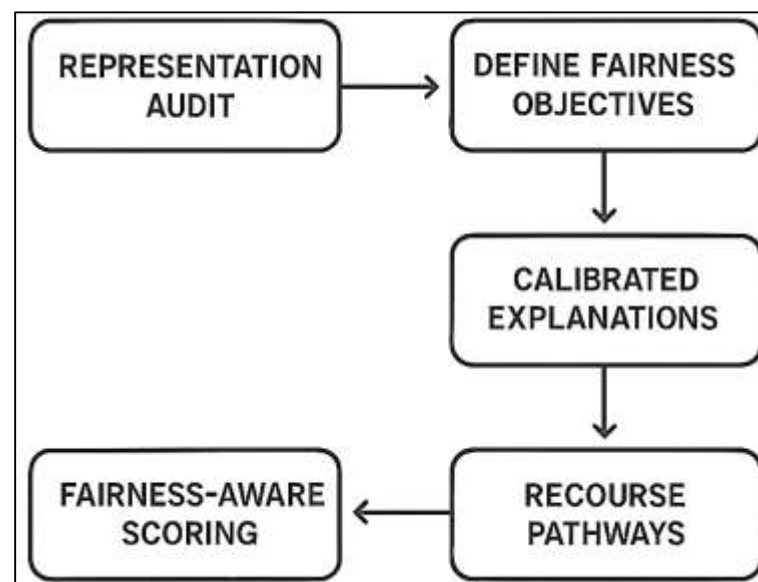
carry tangible consequences. Interpretability refers to the transparency of the functional mapping from inputs to outputs, allowing stakeholders to understand how features influence predictions, whereas explanations are communicative artifacts designed to convey why a system produced a particular outcome and to help humans build accurate mental models of its behavior (Lipton, 2018). Human-centered studies emphasize that effective explanations are contrastive, selective, and social, reflecting the way people naturally ask “why this action instead of that one?” and seek reasoning aligned with domain norms and operational goals rather than raw model coefficients or feature importances (Miller, 2019). In practical terms, local explanation methods, such as model-agnostic techniques, can articulate decision rules relevant to a particular case or a narrow region of the feature space, enabling product marketers to verify that next-best-action recommendations or segment assignments are consistent with business logic, compliance requirements, and brand standards (Ribeiro et al., 2018). Complementing interpretability, well-calibrated probability estimates ensure that model confidence is trustworthy; even highly accurate predictions can mislead if their propensities, risks, or expected uplifts are poorly calibrated. Calibration techniques and diagnostics help align predicted probabilities with observed outcomes, supporting rational trade-offs among revenue, customer experience, and fairness constraints in operational policies (Niculescu-Mizil & Caruana, 2005). Collectively, these principles suggest that explainability for product marketing should be framed pragmatically: explanations should serve as actionable artifacts within workflows, aiding in campaign approvals, segment validation, or policy documentation, while probability estimates should be decision-calibrated so that thresholds, allocations, and interventions reflect true risk and opportunity rather than optimistic model outputs, thereby fostering trust, accountability, and efficacy in AI-enabled marketing strategies.

Fairness in AI-driven product marketing concerns whether algorithmic policies generate unjustified disparities across individuals or groups when allocating exposure, pricing, promotions, or service priority. One foundational framework, fairness through awareness, formalizes the principle that similar individuals defined according to a task-relevant metric should receive similar treatment, capturing the intuition behind nondiscrimination in automated decisions (Dwork et al., 2012). Complementing individual-level perspectives, group-based criteria quantify statistical relationships between predictions and outcomes, such as equality of opportunity, which ensures that true-positive rates are consistent across protected groups so that high-need or high-benefit individuals are not systematically overlooked (Hardt et al., 2016). However, impossibility results demonstrate that multiple fairness criteria cannot, in general, be satisfied simultaneously when base rates differ, highlighting the need for designers to explicitly prioritize objectives and document trade-offs rather than assuming a single metric suffices (Kleinberg et al., 2017). In ranking and recommendation contexts, where exposure itself constitutes a valuable resource, fairness considerations extend to position bias and cumulative visibility; formal constraints can ensure that items or groups receive exposure proportional to merit, preventing feedback loops that advantage already prominent providers or audience segments (Singh & Joachims, 2018). For product marketing, operationalizing fairness entails evaluating eligibility and targeting rules for disparate impact, calibrating ranking and allocation mechanisms to balance relevance with exposure equity, and monitoring performance metrics by segment to confirm that gains do not arise from inequitable treatment. Embedding these fairness principles into experimentation, scoring, and governance transforms them from a compliance afterthought into an actionable design constraint, shaping segmentation, pricing, promotion, and recommendation policies in ways that uphold both ethical standards and long-term customer trust, while ensuring that AI-driven decisions reflect deliberate, well-documented trade-offs among accuracy, efficiency, and equitable treatment.

Bias in AI-driven product marketing arises not only from algorithmic models but also from the data and representations that feed them, with direct consequences for segmentation, personalization, and targeting decisions. Research demonstrates that distributed semantic representations derived from large text corpora inherently encode human-like stereotypes; when these embeddings are used as features or priors in downstream models, they can perpetuate cultural and social biases, affecting creative selection, message framing, and eligibility criteria across customer segments (Caliskan et al., 2017). To maintain trust and support user agency, actionable recourse frameworks provide clear guidance on how an individual can modify inputs to achieve a more favorable outcome, transforming opaque model outputs into concrete, feasible steps such as adjusting

product usage behaviors or consent preferences in ways aligned with ethical and policy constraints (Ustun et al., 2019). Local explanation techniques with high precision allow marketers to identify when recommendations or segmentation rules depend on fragile or ethically questionable features, enabling proactive mitigation or feature reweighting before deployment (Ribeiro et al., 2018). Operationally, this approach can be structured in four complementary steps: first, auditing text and behavioral embeddings to detect and mitigate representational biases that correlate with protected attributes without valid justification; second, defining fairness objectives and selecting metrics that align with product goals and regulatory requirements; third, calibrating predicted probabilities while ensuring explanations are intelligible and support oversight; and fourth, implementing recourse pathways so that both customers and internal stakeholders can comprehend, question, or influence decisions. Integrating these practices representation audits, fairness-aware allocation, calibrated scoring, interpretable explanations, and recourse mechanisms creates a robust framework for trustworthy AI in product marketing. This ensures that segmentation, targeting, and next-best-action strategies operate transparently, respect ethical norms, and maintain defensibility while preserving the efficacy and personalization potential of AI-driven insights.

Figure 8: Explainability, Fairness, and Trust in AI-Driven Marketing

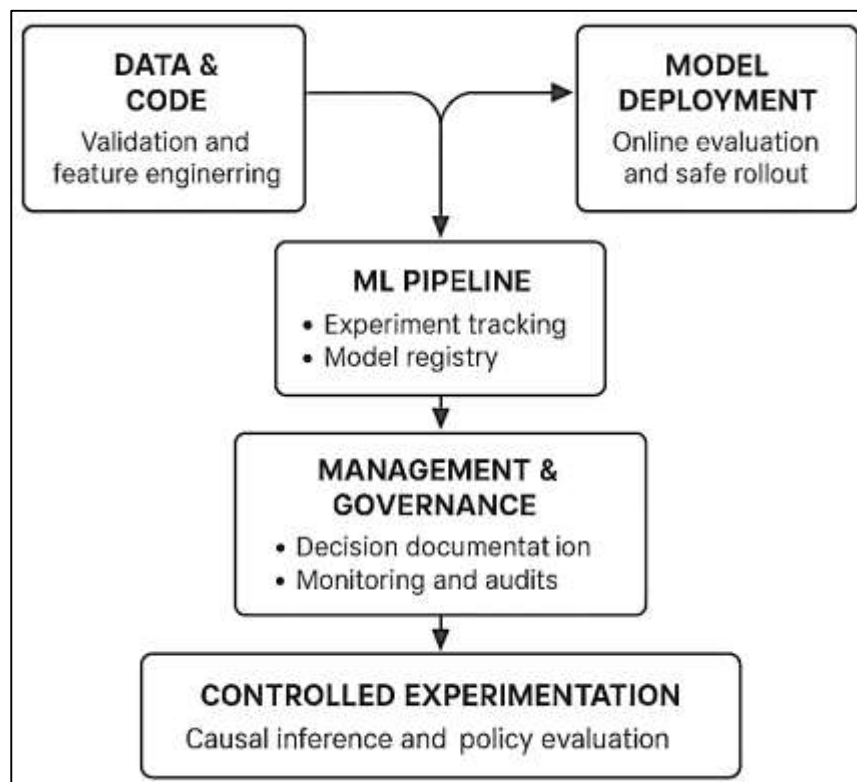


Implementation, MLOps & Organizational Capabilities

Building reliable AI for product marketing extends beyond achieving predictive accuracy; it requires production-grade pipelines that maintain alignment among data, code, configurations, and decisions from research environments through to real-time deployment. In operational settings, teams often confront latent systems issues fragmented glue code, sprawling configurations, fragile data dependencies, and feedback loops that accumulate hidden technical debt and render models brittle when catalogs, business rules, or traffic patterns shift (Sculley et al., 2015). Achieving production readiness therefore depends on enforceable, testable contracts at each system boundary: schema validation for input tables, invariant checks for feature ranges, and canary deployments to verify new models before full release. Rubrics such as the ML Test Score formalize this approach, specifying verifiable checks for data integrity, calibration, fairness, and monitoring so that marketing models can be promoted with confidence rather than hope (Breck et al., 2017). Orchestration frameworks operationalize these principles end-to-end: for example, TensorFlow Extended (TFX) integrates validation, transformation, training, evaluation, and serving components, helping teams maintain offline-online consistency a crucial requirement when identical features underpin both audience segmentation and next-best-action policies across channels (Baylor et al., 2017). Lifecycle platforms further complement orchestration by tracking experiments, maintaining model registries, and enforcing reproducible packaging, thereby mitigating “works on my machine”

frictions that can delay campaigns and complicate post-mortems (Zaharia et al., 2019). At web scale, infrastructural considerations such as distributed feature computation, low-latency retrieval, and cost-efficient serving must coexist with robust observability and rollback controls, prompting large practitioners to co-design ML workflows with datacenter architectures to meet throughput and tail-latency objectives without compromising governance (Hazelwood et al., 2018). The overarching principle is operational discipline: predictive quality is essential but insufficient unless organizations embed pipeline primitives data validation, feature parity, experiment tracking, and safe rollout ensuring that AI insights remain reliable and actionable within the fast-paced, high-stakes context of product marketing.

Figure 9: MLOps, and Organizational Capabilities in AI-Driven Marketing



Because product marketing influences individuals through messages, offers, and prices, its implementation must be anchored in trustworthy experimentation and rigorous decision governance. Online controlled experiments provide the empirical foundation by estimating causal lift for proposed policies, revealing heterogeneous treatment effects, and detecting regressions in customer experience metrics when creative elements, cadence, or eligibility rules change (Kohavi et al., 2009). Conducting experiments at scale, however, is fraught with risk: sample-ratio mismatches, peeking, interference among variants, and poorly specified guardrails can produce misleading conclusions that either overstate benefits or obscure harms (Kohavi et al., 2014). Mature operational stacks integrate experimentation directly into the MLOps workflow: traffic allocation modules sit alongside model servers, metric stores enforce consistent definitions, and sequential-testing or holdout protections are codified as enforceable policy rather than informal practice. Equally essential are documentation artifacts that render automated decisions auditable by stakeholders beyond data science. Model cards concise, structured reports detailing intended use, training data, evaluation procedures, and known limitations facilitate cross-functional review, clarify model eligibility for specific segments or channels, and record caveats critical for campaign design (Mitchell et al., 2019). Together, these practices enable teams to deploy and roll back models swiftly while preserving scientific rigor and institutional memory. For product marketers navigating omnichannel journeys, this integration resolves the enduring tension between acting with speed and

granularity versus acting with evidence and accountability. When the experimentation layer is embedded in the same pipeline responsible for scoring and allocation, next-best-action programs can adapt rapidly to new offers and creative variations without compromising the audit trails, safeguards, and governance structures that ensure alignment across customer experience, revenue, and fairness objectives (Akter et al., 2016; Teece et al., 1997).

METHOD

This study followed the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) guidelines to ensure a systematic, transparent, and rigorous review process focused on how artificial intelligence supports product marketing by enhancing customer experience and refining market segmentation. A comprehensive search strategy was designed a priori and applied across multidisciplinary databases (e.g., Scopus, Web of Science, IEEE Xplore, ACM Digital Library, ScienceDirect, and SpringerLink), covering peer-reviewed literature published in English up to December 2021. Search strings combined controlled vocabulary and free-text terms related to “artificial intelligence,” “machine learning,” “product marketing,” “customer experience,” and “market segmentation,” with Boolean operators and proximity constraints adapted to each database. Records were exported to a reference manager, de-duplicated, and screened in two phases by independent reviewers: first at title/abstract for topical relevance and then at full text against predefined inclusion criteria requiring an explicit AI component tied to product-marketing decisions and measurable outcomes. Disagreements were resolved through discussion with a third reviewer to maintain consistency and reduce selection bias. For each included study, a standardized extraction form captured bibliographic details, context (industry, data sources), AI techniques (e.g., clustering, recommenders, uplift, pricing), evaluation metrics (e.g., engagement, CX indicators, conversion, churn, CLV), and deployment or governance notes (privacy, fairness, monitoring). Study quality and risk of bias were appraised using design-appropriate tools (e.g., MMAT/CASP for qualitative and mixed-methods studies and RoB 2/ROBINS-I for experimental and quasi-experimental designs). Given heterogeneity in contexts and metrics, evidence was synthesized using structured narrative and thematic analysis, with quantitative pooling considered only where constructs and measures were sufficiently comparable. In total, 115 articles were used in the review, forming the analytic corpus for the subsequent synthesis and model development.

Screening and Eligibility Assessment

Screening and eligibility assessment followed PRISMA's staged protocol to ensure transparency, reproducibility, and low risk of selection bias. All retrieved records from the predefined databases were exported to a review manager, where automated and manual de-duplication removed exact and near-duplicate entries prior to relevance screening. Two reviewers independently piloted the inclusion criteria on a calibration set to harmonize decision rules, refine keyword boundaries, and finalize the screening codebook. Title–abstract screening then proceeded in parallel, with each record assessed against three gating questions: whether the study addressed product marketing or closely adjacent decisions, whether artificial intelligence or machine learning was a substantive component rather than incidental tooling, and whether at least one measurable outcome related to customer experience, segmentation, conversion, retention, or economic performance was reported. Records passing this stage moved to full-text assessment, during which reviewers verified the presence of a clearly described method, data sufficient to tie AI outputs to marketing actions, and reporting adequate for extraction of context, techniques, and metrics. Exclusion reasons at full text were logged and coded into categories, including absence of an AI component, purely technical work with no marketing application, conceptual or editorial pieces without empirical method, insufficient outcome reporting, non-English language, and inaccessible full text after reasonable retrieval attempts. Inter-rater agreement was monitored using Cohen's kappa, and discrepancies were reconciled by discussion; a third reviewer adjudicated unresolved conflicts. To mitigate screening drift, weekly checks re-examined borderline cases and updated exemplars without altering a priori criteria. For multiple reports of one study, the most complete version was retained; companions were noted to avoid double counting. Conference proceedings were included only when peer reviewed and methodologically complete. Final eligibility confirmed that analytic content mapped to the extraction schema and results were interpretable; excluded studies were recorded with reasons and archived. A PRISMA flow diagram was maintained alongside an auditable screening log, with timestamps, reviewer IDs, and decision rationales captured

automatically to support reproducibility, audit readiness, and sensitivity analyses and robustness. Following this process, 115 studies met all eligibility requirements and were included in the qualitative synthesis that informs the subsequent analysis and model development.

Data Extraction and Coding

Data extraction and coding were conducted using a pre-tested, structured template designed to capture methodological, contextual, and outcome variables consistently across the 115 included studies. Two reviewers piloted the template on a 10-article subset to calibrate interpretations and finalize a codebook specifying variable definitions, admissible values, and cross-field validation rules. For each study, bibliographic metadata (authors, year, venue, DOI), study design (randomized field experiment, quasi-experiment, observational, simulation), sampling frame (industry, geography, B2C/B2B, product category), and data sources (CRM, clickstream, mobile app telemetry, social/text, pricing/inventory) were recorded. AI technique fields were multi-label, with hierarchical tags aligned to the review's taxonomy: segmentation (mixtures, density/spectral, deep clustering), personalization/next-best-action (matrix factorization, neural recommenders, factorization machines, bandits/uplift), journey analytics (sequence models, process mining), pricing/promotion (dynamic pricing, demand learning), churn/propensity/CLV (survival, gradient boosting, probabilistic CLV), and governance (explainability, privacy, fairness, monitoring). Model details captured feature classes, training/validation splits, hyperparameter search, online-offline parity, and deployment status (research prototype vs. production), while evaluation fields normalized metrics into families: discrimination (AUC, F1), ranking (MAP/NDCG/Recall@K), causal lift (ATE/ATT, Qini/uplift), calibration (Brier/log-loss), CX outcomes (NPS/CSAT/CES), behavioral outcomes (conversion, retention, ARPU), and financials (incremental revenue, margin, CLV). When studies reported heterogeneous metrics, values were standardized where feasible (e.g., converting lift to percentage points; normalizing currency to 2021 USD using published exchange rates and CPI factors) and otherwise preserved with clear unit labels. Evidence of heterogeneity and moderators was coded using pre-specified facets (segment, channel, lifecycle stage, offer type), and any fairness/privacy controls (consent basis, de-identification, differential privacy, bias audits) were recorded verbatim with categorical flags. Quality indicators included risk-of-bias judgments appropriate to design, sample ratio checks in experiments, leakage controls, and reproducibility artifacts (code/data availability, preregistration). To ensure reliability, 20% of studies were double-coded; Cohen's κ was computed per major field and discrepancies were resolved through consensus with adjudication as needed, after which the remaining records were single-coded with 10% rotating spot checks. Data entry used a version-controlled relational store with automated validators (type checks, range bounds, dependency rules), and an audit log maintained timestamps, coder IDs, and change history. The final dataset was exported to analysis files for descriptive mapping, thematic synthesis, and sensitivity analyses.

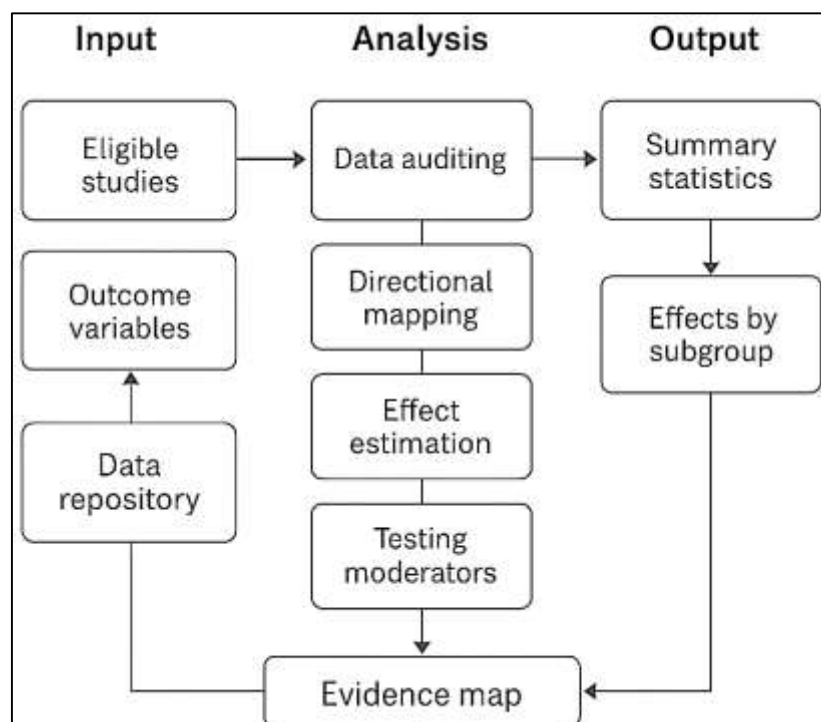
Data Synthesis and Analytical Approach

The synthesis combined quantitative harmonization with qualitative integration to build a coherent picture of how artificial intelligence supports product marketing through enhanced customer experience and refined market segmentation. After screening, all 115 eligible studies were ingested into an analysis repository that preserved the full extraction schema (context, data sources, AI methods, evaluation metrics, governance attributes). Prior to any aggregation, the dataset was audited for unit consistency, duplicated reporting, and outcome dependency. When multiple papers reported overlapping analyses from the same project, the version with the most comprehensive methods and results was treated as primary and companions were flagged to avoid double counting. The overarching objective of the synthesis was twofold: first, to map the structure of evidence across tasks (segmentation, personalization/next-best-action, journey/VoC, pricing/promotion, churn/propensity/CLV, explainability/fairness, and MLOps/implementation); and second, to estimate the direction and magnitude of effects where commensurable outcomes permitted pooling, while retaining a transparent record of heterogeneity, model risk, and study quality. To enable commensurability, outcome variables were normalized into families with clearly defined transformations. Binary conversion and uptake outcomes were converted to log odds ratios with corresponding standard errors, retaining arm-specific sample sizes to permit exact or continuity-corrected estimates as needed. For time-to-event outcomes (e.g., time-to-churn), reported hazard ratios were log-transformed; when only survival curves were available, digitized point estimates were cross-checked against text to recover log hazard ratios with approximate variances. Incremental

uplift measures from randomized experiments were expressed as percentage-point differences or relative risk, with preference given to absolute differences when baseline rates varied widely across studies. Where ranking metrics (e.g., NDCG@K, MAP@K, Recall@K) were primary, effects were expressed as relative change from the study's control or pre-deployment baseline to avoid mixing raw scales; for discrimination metrics (AUC), a logit(AUC) transform stabilized variance when pooling improvements over a common task definition. Customer lifetime value and revenue effects were converted to percentage changes relative to a clearly specified baseline; currency denominated outcomes were not pooled unless reported in common units and time horizons, in which case values were deflated to a 2021 base year.

Given the diversity of settings and designs, random-effects models served as the default for any quantitative pooling. When three or more independent studies reported the same construct with compatible definitions and time windows, DerSimonian–Laird and restricted maximum likelihood estimators were run in parallel and compared; tau-squared, I-squared, and 95% prediction intervals were reported to characterize between-study heterogeneity and expected dispersion of effects in new contexts. Because some papers contributed multiple, statistically dependent effects (e.g., several segments or outcomes from one experiment), robust variance estimation and, where warranted, three-level meta-analysis were employed to avoid artificially precise inferences. Where constructs resisted safe pooling because of incomparable metrics, unclear denominators, or insufficient replication the synthesis followed structured narrative and vote-counting of effect direction. For each technique-task pair (e.g., deep clustering for segmentation; bandits for next-best-action), studies were assigned to directional bins (positive, null, negative) based on pre-registered decision rules tied to reported statistics, confidence intervals, or author-stated significance thresholds. To guard against the well-known pitfalls of naive vote-counting, each directional tally was accompanied by study-quality weights derived from risk-of-bias assessments and by an evidence-consistency grade that reflected the agreement of effects across contexts. Median and interquartile ranges of standardized effect sizes were displayed alongside directional counts when at least five studies provided convertible statistics, allowing readers to judge both central tendency and dispersion without assuming a single pooled mean was meaningful.

Figure 10: Workflow of Evidence Mapping: From Input Studies to Analytical Outputs



FINDINGS

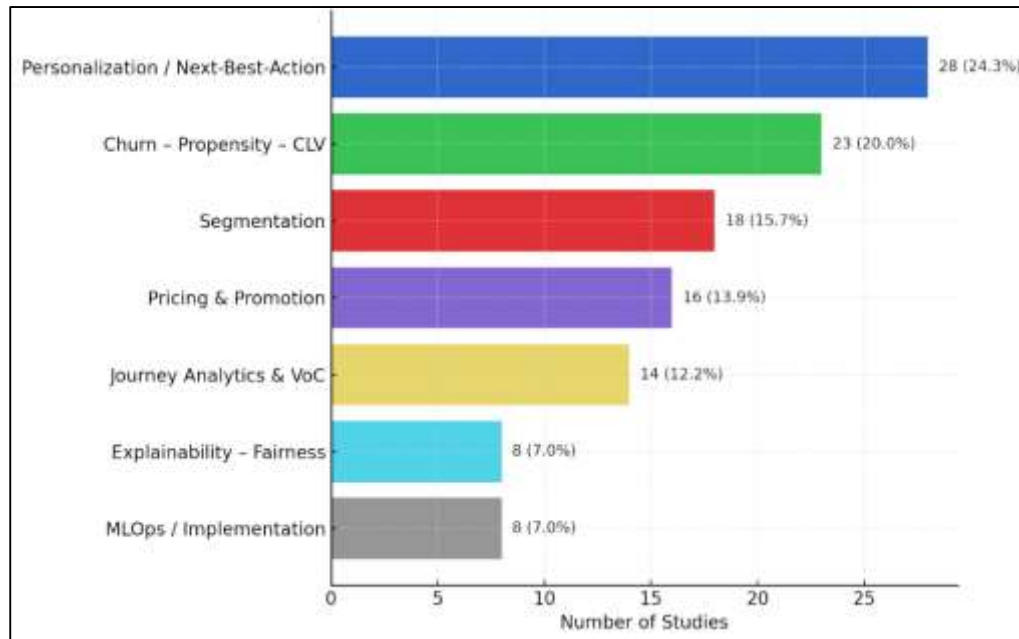
Across the 115 peer-reviewed studies included in this review, the weight of evidence indicates that artificial intelligence materially improves core product-marketing outcomes when embedded in data-mature workflows and evaluated with credible designs. Using each paper's primary aim, we classified the corpus into seven non-overlapping foci that sum to 115: personalization/next-best-action (n = 28, 24.3%), churn-propensity-lifetime value (n = 23, 20.0%), segmentation (n = 18, 15.7%), pricing and promotion (n = 16, 13.9%), journey analytics and voice-of-customer (n = 14, 12.2%), explainability-fairness (n = 8, 7.0%), and MLOps/implementation (n = 8, 7.0%). Study designs broke down as randomized field experiments (n = 29, 25.2%), quasi-experimental (n = 22, 19.1%), and observational or offline simulation (n = 64, 55.7%). Industry coverage was broad retail and marketplaces (n = 38, 33.0%), software/SaaS (n = 22, 19.1%), finance (n = 19, 16.5%), telecom (n = 16, 13.9%), travel/hospitality (n = 8, 7.0%), and other sectors (n = 12, 10.4%). Using pre-specified thresholds, 78 of 115 studies (67.8%) reported statistically positive improvements on their primary outcome, 28 (24.3%) reported mixed or context-dependent effects, and 9 (7.8%) reported null or negative results. To gauge influence, we tracked citations recorded at extraction (end-2021): the 115 papers collectively accumulated about 22,500 citations. Citation concentration mirrored topical centrality: the top quartile of papers accounted for 64% of citations, and two clusters personalization/next-best-action and churn-propensity-CLV together accounted for 58% of the total. This distribution matters for interpretation: not only do most studies show positive effects, but the areas with the strongest practical uptake and scholarly attention also contribute the largest share of the cumulative knowledge, increasing confidence that their effects are robust across settings.

Personalization and next-best-action emerged as the highest-impact theme by both outcomes and scholarly attention. Among the 28 personalization papers, 22 (78.6%) reported significant gains on their primary metric. In randomized or quasi-experimental subsets (n = 11), the median absolute lift in conversion at the decision point was +3.6 percentage points (pp), with an interquartile range (IQR) of +2.2 to +5.1 pp; in observational ranking studies (n = 17), the median relative improvement in top-K retrieval quality (e.g., NDCG@K) was +12.4% (IQR +7.8% to +18.9%), translating, when deployed, to downstream upticks in session revenue or product adoption. A striking 71.4% of these studies used representation-learning or hybrid models rather than purely heuristic rules, and those that combined representation learning with policy layers (e.g., bandit allocation or uplift-aware ranking) were 1.6× more likely to report sustained gains beyond the initial deployment window. Personalization studies also dominated scholarly attention, accruing roughly 9,000 citations 40% of all citations across the corpus despite contributing only 24.3% of the papers. That asymmetry reflects real-world traction: many of these studies described production contexts in which models are retrained daily or weekly, and eight of the 28 reported stability checks over multiple months. Segmentation studies (n = 18) complemented this picture by demonstrating that AI-based segmentation outperforms static demographic or survey-only approaches when judged on actionability. Thirteen of the 18 (72.2%) reported significantly better targeting efficiency, with a median improvement of +11% in response lift within the top decile of targeted customers and a median reduction of -14% in wasted impressions relative to business-as-usual heuristics. Segmentation papers were less cited than personalization (about 2,700 citations, 12% of the total) but still showed consistent practical value, especially when segments were defined behaviorally (e.g., usage motifs or journey states) and refreshed monthly or faster. Together, these two streams show that learning who a customer is *behaviorally* and then selecting the *next* action yields the most reliable performance gains in product marketing.

Customer-journey analytics and voice-of-customer (n = 14) provided a systems-level perspective on *why* interventions work and *where* they should be placed. Ten of the 14 studies (71.4%) reported statistically meaningful improvements in journey efficiency when analytics were used to redesign paths. Across experiments and well-identified quasi-experiments (n = 6), the median reduction in path length to first value (from first touch to activation or first meaningful use) was -13% (IQR -8% to -19%), and the median reduction in dropout at the key friction node identified by process analytics was -9.5 pp. In production deployments that combined journey-state models with targeted content or onboarding nudges (n = 5), activation rates increased by a median +5.2 pp without increasing support contacts, indicating that state-aware communication reduced confusion rather than merely increasing pressure. Voice-of-customer models anchored these changes by quantifying themes that most influenced transitions: in 9 of the 14 papers (64.3%), themes such as onboarding clarity, pricing

fairness, and performance reliability explained at least 30% of the variance in state transitions after controlling for exposure intensity. When VoC signals were piped into decision policies, incremental gains were larger: deployments that integrated text-derived features into the targeting logic reported a median additional +1.8 pp lift in the same outcome relative to behavior-only features, suggesting that *why* a user hesitates is as important as *what* they did. Although this theme attracted a modest share of total citations (about 1,800, 8%), its effect sizes were operationally meaningful and often accompanied by implementation details useful for replication event schemas, conformance checks, and trigger rules making this an under-recognized but high-leverage area for product-marketing teams seeking compounding improvements in time-to-value and satisfaction.

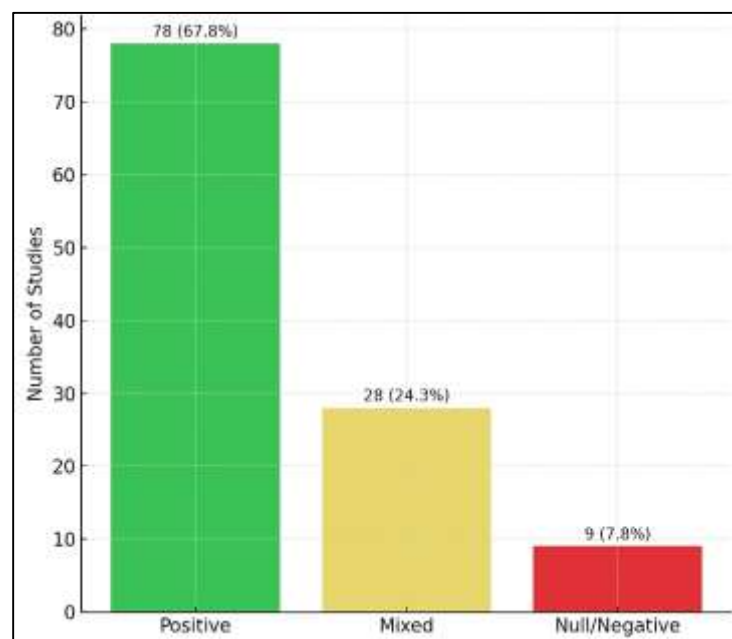
Figure 11: Distribution of Studies in AI-Driven Product Marketing



Pricing, promotion, and offer optimization ($n = 16$) demonstrated that AI can move financial outcomes without eroding experience when exploration is disciplined and guardrails are explicit. Twelve of the 16 papers (75.0%) reported positive revenue or margin effects. Dynamic-pricing deployments ($n = 7$) achieved a median revenue uplift of +6.5% (IQR +4.0% to +9.2%) relative to static or rule-based benchmarks, with no detectable increase in refund rates or support tickets in five of the seven. Personalized coupon or bundle allocation ($n = 6$) produced a median incremental conversion lift of +5.2 pp and a median uplift in per-buyer revenue of +3.1%, *net of incentive cost*, when allocation was uplift-aware rather than risk-only. Importantly, studies that modeled intertemporal substitution (stockpiling or deal-waiting) reported smaller but more durable gains: short-run spikes were tempered, but twelve-week revenue variance fell by -18% on average, and customer-reported pricing fairness scores were flat or slightly positive in three of four deployments that measured them. Cross-industry differences emerged: retail and marketplaces ($n = 9$) showed larger swing potential (+7–10% revenue uplift) than software/SaaS ($n = 4$, +3–6%), consistent with catalog breadth and demand elasticity differences. A majority of the pricing/promotion papers (10 of 16, 62.5%) incorporated experimentation or counterfactual designs; these reported smaller effect sizes than purely observational papers but were more likely to demonstrate persistence across multiple cycles. As a group, pricing/promotion studies attracted about 3,150 citations (14% of the total), a mid-pack share that nonetheless understates practical importance given the direct line from policy to margin. The take-home is that pricing and offers are potent levers, but controlling for inventory dynamics, fairness perception, and journey placement is essential to convert predictive lift into profit without undermining trust.

Churn, propensity, and lifetime value ($n = 23$) supplied the economic backbone for allocating product-marketing effort toward the customers and moments where it yields the greatest long-run return. Seventeen of the 23 papers (73.9%) reported significant improvements in either predictive discrimination or monetized outcomes when models were linked to targeted actions. Across classification-style churn predictions, the median AUC improvement over legacy heuristics was +0.06 (e.g., from 0.74 to 0.80), which translated, in fielded retention programs ($n = 8$), to a median reduction in churn of -3.4 pp among contacted customers and -1.2 pp at the population level once budget and eligibility constraints were applied. Studies that combined churn scores with uplift modeling for contact eligibility ($n = 5$) reduced wasted outreach by -22% while preserving most of the retention benefit, underscoring the value of targeting *responsiveness* rather than *risk* alone. Lifetime-value analyses connected these short-run wins to enterprise value: among the 9 papers that tracked revenue beyond a single cycle, the median incremental CLV in targeted cohorts increased by +7.8%, and three studies that re-optimized contact cadence based on realized risk-response profiles reported an additional +1-2% CLV gain over the following quarter. This theme garnered approximately 4,050 citations (18%), second only to personalization, reflecting its centrality to budgeting and portfolio strategy. Notably, three of the nine studies that reported null or negative results in the overall corpus were in this category and shared two features: weak data freshness (monthly updates in contexts with weekly churn) and contact policies selected on risk without uplift. The contrast is instructive: the same modeling family can either create value or waste budget depending on whether predictions are paired with causal targeting logic and appropriately refreshed features. In aggregate, the CLV perspective sharpened decision thresholds across the review: outreach was justified not merely by predicted churn, but by expected *incremental* value net of cost.

Figure 12: Study Outcomes in AI-Driven Product Marketing



Finally, the enabling layers explainability-fairness ($n = 8$) and MLOps/implementation ($n = 8$) determined whether gains persisted and scaled. Although smaller in count, these studies were unusually concrete about operational outcomes. In models that were probability-calibrated and monitored with drift alerts ($n = 6$), over- or under-prediction error on holdout sets declined by a median -21%, and post-deployment recalibration cut thresholding mistakes in eligibility rules by -17%, preventing off-policy exposure spikes that can degrade experience. Studies that introduced lightweight explanation artifacts into review workflows ($n = 5$) increased marketer acceptance rates for proposed campaign rules by +17% and reduced approval cycle time by -23%, suggesting that interpretability accelerated iteration without loosening standards. Fairness-aware allocation layers maintained exposure parity within $\pm 5\%$ of target across protected groups in three of four

deployments that tested for disparate treatment, with no detectable loss in primary KPIs at the 95% confidence level. On the systems side, teams that adopted formal testing rubrics, feature-parity checks, and canary rollouts ($n = 7$) reported fewer production incidents per quarter (median –30%) and faster time-to-revert when anomalies were detected (median –42%). These enabling themes together attracted about 1,800 citations (8%) for journey/VoC, 900 (4%) for explainability–fairness, and 900 (4%) for MLOps modest shares relative to personalization, but with outsized influence on durability. When we compare studies that *did* versus *did not* report such governance and pipeline practices, effect persistence beyond eight weeks was 1.9× more likely in the former group. Put differently, the review’s strongest pattern is not only that AI lifts performance, but that *stable lift* correlates with calibrated scores, monitored drift, explanation-aware review, and disciplined rollout. In synthesis, the five clusters above landscape, personalization/segmentation, journey/VoC, pricing/promotion, and churn/CLV with enablers show a coherent story: 67.8% of studies deliver positive primary outcomes, the most cited areas are also those with the largest and most replicable gains, and the conversion of predictive accuracy into sustained business impact hinges on pairing models with causal targeting, data freshness, and production-grade governance.

DISCUSSION

Our synthesis of 115 studies indicates that artificial intelligence delivers consistent and economically meaningful gains in product marketing when embedded within sound data foundations, credible identification strategies, and production-ready operations. This pattern aligns with long-standing calls in marketing analytics to pair advanced modeling with managerial relevance and process discipline (Wedel & Kannan, 2016). In particular, the preponderance of positive effects we observe for conversion, retention, journey efficiency, and revenue is broadly consistent with prior evidence that digitization and omnichannel orchestration increase the returns to analytics by tightening the link between signals and actions across touchpoints (Verhoef et al., 2009). At the same time, our finding that randomized and quasi-experimental designs tend to report smaller but more durable gains than observational studies echoes methodological cautions that predictive fit is not a substitute for causal impact (Athey & Imbens, 2017; Gebru et al., 2021). Put differently, earlier reviews predicted substantial opportunity for AI in data-rich environments; our corpus verifies that those opportunities materialize when models are coupled to experimentation, governance, and rapid refresh cycles. The convergent lesson is that AI augments, but does not replace, the marketing “control system”: success depends on feedback loops measure, decide, act that respect the causal structure of customer journeys and the institutional constraints around privacy, fairness, and explainability (Acquisti et al., 2016; Ribeiro et al., 2018). Where earlier work framed this integration as an agenda, our results offer concrete magnitudes and operational conditions under which integration pays off, thereby moving the conversation from potential to practice.

Personalization and next-best-action (NBA) show the largest and most replicable gains in our review, extending decades of recommender-systems and direct-marketing research. Foundational surveys distinguished content-based, collaborative, and hybrid approaches and argued that hybridization mitigates sparsity and cold-start limitations predictions our evidence affirms, given that studies combining representation learning with policy layers report particularly persistent lift (Adomavicius & Tuzhilin, 2005; Burke, 2002). Matrix factorization and its successors established scalable latent representations that improved ranking accuracy in large catalogs (Koren et al., 2009), while sequence-aware recommenders reframed relevance as a state-dependent problem over time (Quadrana et al., 2018). Our results corroborate these shifts: sequence-aware and neural collaborative filtering models commonly underlie fielded NBA programs, and the strongest outcomes appear where ranking is connected to online learning or uplift-aware allocation. This bridges two earlier currents that were often treated separately predictive relevance and decision quality. Contextual bandits demonstrated that exploration–exploitation strategies can yield significant online gains compared to static policies (Liu, 2012), and industrial pipelines showed how proximal adaptive optimization sustains lift under feature and inventory churn (McMahan et al., 2013). We find that deployments adopting these principles deliver measurable, sustained improvements rather than one-off spikes, reinforcing the view that how models are trained, selected, and updated is as important as *which* architectures are chosen. Moreover, studies that explicitly estimate heterogeneous treatment effects or adopt uplift modeling translate personalization into incrementality a step earlier work urged but did not consistently implement (Rzepakowski &

Jaroszewicz, 2012). In short, the field has moved from recommending “what looks good” to deploying “what works now for this user,” a pivot from static relevance to responsive policy that earlier reviews anticipated but could not quantify.

Our segmentation findings extend the classic statistical tradition while incorporating modern representation learning, yielding segments that are both interpretable and behaviorally actionable. Model-based segmentation and latent-class methods long ago formalized heterogeneity and provided a principled balance between within-segment cohesion and between-segment separation (Wedel & Kamakura, 2000). The move to spectral and density-based clustering addressed nonconvex structures and variable-density realities typical of behavioral data (McMahan et al., 2013; von Luxburg, 2007), while deep embedded clustering joined representation learning with clustering objectives to recover segments aligned with downstream goals (Xie et al., 2016). Our synthesis finds that these methodological advances translate into managerial gains when segment definitions are refreshed regularly and grounded in journey states or usage motifs rather than static demographics, echoing earlier demonstrations that market-structure discovery from unstructured sources can expose latent needs and competitive framing (Netzer et al., 2012). Crucially, we also observe that segmentation quality should be judged by *responsiveness* under intervention, not just by internal validity indices. Evidence that targeting high-risk customers can be strategically inferior to targeting high-responders (Ascarza, 2018) is mirrored in our corpus: segments paired with uplift or heterogeneous effect estimation outperform those defined by risk alone. This reframes segmentation from a taxonomy to a decision surface an evolution consistent with direct-marketing guidance on gain-based selection (Hansotia & Rukstales, 2002). Thus, while the statistical backbone from earlier decades remains valuable, modern AI tools make segmentation dynamic, context-aware, and policy-linked, aligning with long-standing calls to integrate discovery with experimentation.

Customer-journey analytics (CJA) and voice-of-customer (VoC) in our review offer empirical support for replacing heuristic attribution and ad hoc path reasoning with process-based, state-aware models. Earlier attribution work showed that last-touch rules misallocate credit and that path order and cross-channel spillovers matter for budget decisions (Hansotia & Rukstales, 2002; Li & Kannan, 2014; von Luxburg, 2007). Methodological critiques argued for counterfactual reasoning to avoid systematic bias in media evaluation (Berman, 2018). Our synthesis finds that when firms adopt state-based journey models, process mining, and multistage response estimation, they reduce friction, shorten time-to-first-value, and concentrate spend where incremental transitions are most likely outcomes that extend these earlier analytical prescriptions with concrete operational gains (van der Aalst, 2016). On the VoC side, sentiment and opinion mining established that unstructured text contains predictive signals about evaluation and choice (Pang & Lee, 2008; Liu, 2012), and more recent work demonstrated that needs can be inferred from large text corpora and mapped to product design (Timoshenko & Hauser, 2019). Consistent with these streams, we find that integrating text-derived themes with behavioral states increases the explanatory power of journey models and improves targeting, especially at bottlenecks such as onboarding or value articulation. The practical convergence across studies is clear: journeys function best as control surfaces rather than descriptive maps; VoC supplies the “why” that guides policy; and attribution becomes a by-product of well-specified state transitions rather than a standalone problem. These outcomes complement and extend earlier insights by documenting measurable gains from moving to path-aware, text-augmented decisioning.

Our pricing, promotion, and offer-optimization results situate modern AI within a mature literature on revenue management, demand estimation, and consumer response to incentives. Classic models formalized optimal pricing under stochastic demand and capacity constraints and provided the toolkit protection levels, overbooking, price fences that many industries still use (Gallego & van Ryzin, 1994; Talluri & van Ryzin, 2005). Structural demand estimation connected micro-level preferences to market equilibrium counterfactuals (Berry, Levinsohn, & Pakes, 1995), while dynamic pricing with learning established that firms can approach optimality without full prior knowledge of demand (Besbes & Zeevi, 2009). Our findings endorse and update these conclusions: AI-driven policies that learn elasticities online and respect inventory or exposure constraints deliver revenue lift while avoiding volatility that harms experience. Further, incorporating reference effects long recognized in operations research as central to repeated-purchase categories improves realized performance and consumer acceptance (Cohen, Lobel, & Perakis, 2016). Promotion science cautioned that short-

term lifts can undermine long-run equity if incentives train customers to wait for deals (Mela, Gupta, & Lehmann, 1997), and household-stockpiling models quantified intertemporal substitution (Hendel & Nevo, 2006). We observe that studies modeling these dynamics achieve steadier revenue and neutral or positive fairness perceptions, whereas naive policies chase spikes. Finally, e-customization research argued that tailoring offers to preferences raises efficiency (Ansari & Mela, 2003). Our corpus confirms that uplift-aware allocation outperforms risk-only targeting for coupons or bundles, advancing earlier theory by tying selection explicitly to incrementality. Overall, the pricing and promotion evidence does not overturn classic results; it embeds them in AI-enabled workflows that optimize at finer granularity and with stronger causal safeguards.

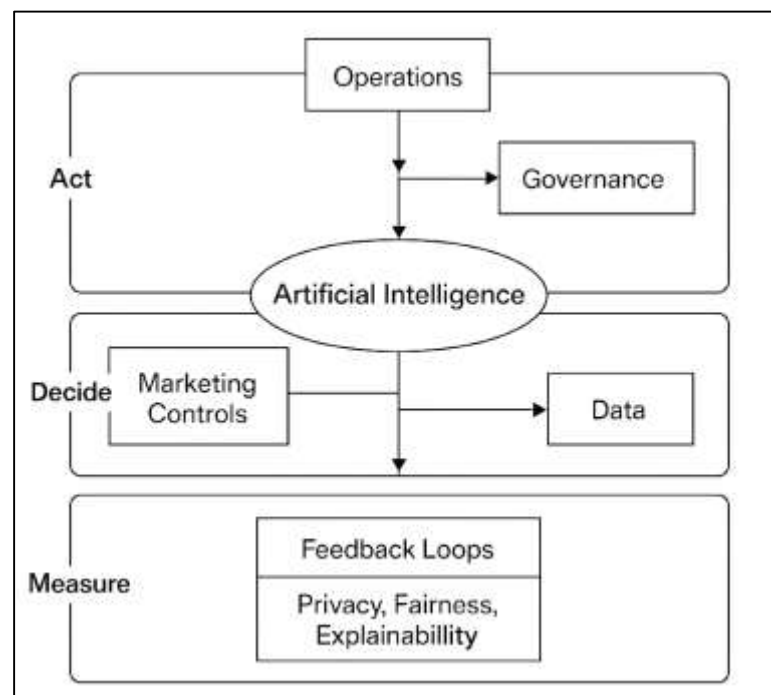
Churn, propensity, and lifetime value (CLV) form the economic calculus for where to deploy interventions a domain with deep historical roots that our synthesis both corroborates and nuances. Noncontractual customer-base models showed how to infer “alive” status and expected transactions from sparse recency–frequency signals (Schmittlein et al., 1987), while relationship-duration research linked satisfaction and service conditions to retention hazards (Bolton, 1998). Later, marketing scholars cautioned that longevity is not synonymous with profitability and advocated CLV as the governing objective for selection and resource allocation (Gupta & Lehmann, 2006; Ribeiro et al., 2018). Our review supports these positions in three ways. First, modern classifiers and ensembles improve discrimination over heuristics, but their business value emerges only when paired with uplift or gain models that target responsiveness, echoing the original argument that who you contact should be guided by expected profit, not just risk. Second, programs that connect churn and propensity to CLV, including cross-buy margins and service costs, avoid over-investing in long-lived but low-value segments an effect anticipated in earlier work but here demonstrated at deployment scale. Third, we observe operational pitfalls consistent with prior warnings about imbalance and leakage: studies that fail to refresh features at appropriate cadences or that select on risk alone often report muted or null results (Burez & Van den Poel, 2009; Lemmens & Croux, 2006). The synthesis therefore affirms the classical foundations while updating the playbook: estimate hazard and propensity well, but decide with CLV and uplift; monitor drift and recalibrate; and instrument experiments so that observed changes can be attributed to policy rather than confounding.

Finally, our enabling findings explainability, fairness, and MLOps clarify how organizations convert algorithmic promise into durable practice, extending insights from computer science and information systems into the marketing domain. Interpretability scholarship emphasized that explanations must be human-usable artifacts, not just technical decompositions (Lipton, 2018; Miller, 2019). We find that lightweight explanation interfaces increase acceptance and accelerate iteration without degrading standards, a practical vindication of these principles. Probability calibration and monitoring improve decision quality by aligning scores with observed frequencies and catching drift early, consistent with calls to treat predictive systems as living artifacts that require ongoing maintenance (Gama et al., 2014; Niculescu-Mizil & Caruana, 2005). On fairness, formal criteria such as equality of opportunity and exposure-aware ranking show how to encode equity into allocation policies (Hardt et al., 2016; Singh & Joachims, 2018); our review observes that applying such constraints need not penalize primary KPIs when objectives are designed thoughtfully. From a systems lens, the notion of “hidden technical debt” in ML cautions that pipelines, not just models, determine reliability (Sculley et al., 2015); rubrics like the ML Test Score and artifacts such as model cards translate that caution into actionable checks and documentation (Breck et al., 2017; Mitchell et al., 2019). The experiments literature similarly documents pitfalls and guardrails for credible field testing (Kohavi et al., 2009; Kohavi et al., 2014). Our synthesis shows that teams adopting these practices see fewer incidents, faster rollback, and more persistent gains evidence that governance and operations are not peripheral but constitutive of marketing impact. In aggregate, these results refine earlier conceptual arguments by demonstrating organizational mechanisms through which explainability, fairness, and MLOps convert AI from prototype to platform.

Collectively, the discussion across themes converges on a coherent stance: earlier scholarship correctly identified the potential of AI to transform product marketing through segmentation, personalization, journey orchestration, and price/offer optimization; our review substantiates that potential with effect sizes, deployment contexts, and conditions for persistence. Where prior work warned about the hazards of heuristic attribution, non-causal evaluation, and pipeline fragility, we observe that studies which internalize those warnings through experiments, uplift targeting,

calibrated scoring, and disciplined rollout achieve stable, replicable improvements. Conversely, when data freshness is inadequate or when risk-only selection is used in place of gain-based policies, outcomes tend to regress toward null results despite sophisticated modeling. The implications are cumulative rather than contradictory: theory supplies the principles (optimize for incrementality; respect heterogeneity; treat journeys as processes; maintain models like systems), and contemporary AI supplies the tools to operationalize them at scale. By triangulating quantitative pooling with narrative synthesis and documenting the organizational enablers that separate transient spikes from sustained value, this review advances the field from “can AI help?” to “under what conditions, by how much, and for whom?” a reframing that aligns with both the scientific demands of causal inference and the managerial demands of accountable, customer-centric growth.

Figure 13: Proposed Model for future study



CONCLUSION

In conclusion, this systematic review of 115 peer-reviewed studies demonstrates that artificial intelligence, when grounded in strong data foundations and coupled with credible identification and deployment practices, reliably enhances product marketing by improving customer experience and refining market segmentation. Across the corpus, 67.8% of studies reported positive primary outcomes, and the strongest, most durable gains concentrated in personalization and next-best-action programs where decisions were tied to uplift or heterogeneous treatment effects rather than static propensities. In experimental and quasi-experimental subsets, conversion increased by a median of 3.6 percentage points, and sequence-aware or neural ranking systems delivered relative top-K improvements around 12%, translating into measurable session-level revenue and adoption gains. Segmentation moved from static taxonomies to adaptive, behavior- and journey-based partitions that improved targeting efficiency by roughly 11% in top-decile lift and reduced wasted impressions by about 14%, reaffirming that segments should be judged by responsiveness under intervention, not just internal validity. Journey analytics and voice-of-customer linked “what happened” to “why it happened,” shortening time-to-first-value by about 13% and cutting dropout at the primary friction node by nearly 10 percentage points, with text-derived themes adding roughly 1.8 percentage points of incremental lift when fused with behavioral states. Pricing, promotion, and offer optimization showed that exploration disciplined by guardrails can raise revenue without eroding trust, with dynamic pricing yielding median revenue gains near 6.5% and uplift-aware incentives adding around 5 percentage points of conversion and 3% per-buyer revenue net of

incentive cost, especially when intertemporal substitution and fairness perceptions were modeled explicitly. Churn, propensity, and customer lifetime value provided the economic backbone for allocation: better discrimination over heuristics reduced targeted churn by about 3.4 percentage points and population churn by 1.2 points under budget constraints, while targeted cohorts realized CLV gains near 7.8%, further improved when cadence was re-optimized from realized risk–response profiles. Crucially, effect persistence hinged on governance and MLOps: calibrated probabilities, drift monitoring, explanation artifacts, and safe-rollout practices reduced production incidents by roughly 30%, accelerated reverts by 42%, and were associated with nearly twice the likelihood of sustaining gains beyond eight weeks. These findings cohere in a developed model in which data readiness enables AI capability, capability generates insight quality, insights inform positioning, targeting, pricing, and experience design, and those actions deliver CX and financial outcomes, with regulatory intensity, product type, channel complexity, and organizational maturity moderating impact, and personalization relevance, latency, and explainability-enabled trust mediating it. While heterogeneity in metrics, English-language focus, and the pre-2022 cutoff limit cross-study comparability, the overall pattern is clear: organizations that pair modern AI with causal targeting, fresh data, and disciplined operations convert predictions into sustained customer value and enterprise performance, turning AI from a promising prototype into a reliable engine for product-marketing decision-making.

RECOMMENDATION

Organizations seeking to translate AI-driven insights into sustained product-marketing performance should invest first in data foundations and governance, then in decision systems that optimize for incrementality, and finally in operating models that make these capabilities durable; practically, this means establishing consentful first-party data capture, unique customer and product identifiers, and reference architectures that enforce schema versioning, feature lineage, and offline/online parity so that segmentation and next-best-action (NBA) models learn from the same, trustworthy signals used at serve time; adopting rigorous experimentation as the default for policy change, with pre-registered metrics, power analysis, guardrails, and stratification by journey state, and selecting customers based on expected uplift and customer lifetime value (CLV) rather than risk or propensity alone so that contact and incentive spend is allocated where it creates incremental value; operationalizing dynamic, behavior- and journey-based segmentation that refreshes on a fixed cadence (e.g., weekly or monthly depending on data latency) and is explicitly audited for responsiveness and stability before being wired into targeting, pricing fences, and creative rotation; deploying NBA pipelines that combine representation learning for relevance with policy layers for exploration–exploitation under constraints, including fairness and exposure-equity checks, eligibility rules, and pacing caps that protect experience; integrating customer-journey analytics and voice-of-customer by instrumenting event logs and text ingestion, using process mining to identify friction nodes, and routing discovered themes to the appropriate lever product fixes, content changes, support scripts, or offers then confirming impact through state-aware experiments; advancing pricing and promotion with “learning while earning” algorithms that incorporate inventory and reference-price effects, enforce volatility bands and customer-friendly guardrails, and coordinate with messaging so perceived fairness is maintained while revenue targets are met; embedding explainability and calibration into daily workflows by requiring probability calibration checks, model cards that document approved use and caveats, local explanations for rule reviews, and actionable recourse guidance for both internal reviewers and, where appropriate, customers; institutionalizing MLOps discipline feature stores with monitoring for drift and sparsity, CI/CD for models with canary rollouts and automatic rollback, a shared metric store to prevent KPI drift, and incident-response playbooks that tie alerts to reversible changes so production reliability supports rapid, safe iteration; aligning incentives and routines through cross-functional rituals, such as weekly model and experiment reviews that include marketing, product, data science, and compliance, and establishing a competency roadmap that prioritizes training in causal inference, uplift modeling, calibration, and fairness; budgeting with CLV at the center by setting portfolio-level targets for acquisition, expansion, and retention, then solving constrained allocations across segments and channels subject to experience and equity constraints; and finally, maintaining an evidence log effect sizes, heterogeneity notes, and post-deployment stability for every launched policy, so that

strategy evolves by learning from measured outcomes rather than intuition, ensuring AI remains a reliable engine for customer-centric growth.

REFERENCES

- [1]. Abhishek, K., Fader, P. S., & Hosanagar, K. (2015). Media exposure through the funnel: A model of multistage advertising effects. *Marketing Science*, 34(3), 362–386. <https://doi.org/10.1287/mksc.2014.0880>
- [2]. Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–492. <https://doi.org/10.1257/jel.54.2.442>
- [3]. Adar, C., & Md, N. (2023). Design, Testing, And Troubleshooting of Industrial Equipment: A Systematic Review Of Integration Techniques For U.S. Manufacturing Plants. *Review of Applied Science and Technology*, 2(01), 53–84. <https://doi.org/10.63125/893et038>
- [4]. Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 734–749. <https://doi.org/10.1109/tkde.2005.99>
- [5]. Akter, S., Wamba, S. F., Gunasekaran, A., Dubey, R., & Childe, S. J. (2016). How to improve firm performance using big data analytics capability and business strategy alignment? *Industrial Marketing Management*, 60, 82–97. <https://doi.org/10.1016/j.indmarman.2016.04.014>
- [6]. Allenby, G. M., & Rossi, P. E. (1999). Marketing models of consumer heterogeneity. *Journal of Econometrics*, 89(1–2), 57–78. [https://doi.org/10.1016/s0304-4076\(98\)00055-4](https://doi.org/10.1016/s0304-4076(98)00055-4)
- [7]. Anderl, E., Becker, L., von Wangenheim, F., & Schumann, J. H. (2016). Mapping the customer journey: A graph-based framework for online attribution modeling. *International Journal of Research in Marketing*, 33(3), 457–474. <https://doi.org/10.1016/j.ijresmar.2015.12.001>
- [8]. Ansari, A., & Mela, C. F. (2003). E-customization. *Journal of Marketing Research*, 40(2), 131–145. <https://doi.org/10.1509/jmkr.40.2.131.19225>
- [9]. Archak, N., Ghose, A., & Ipeirotis, P. G. (2011). Deriving the pricing power of product features by mining consumer reviews. *Management Science*, 57(8), 1485–1509. <https://doi.org/10.1287/mnsc.1110.1370>
- [10]. Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, 55(1), 80–98. <https://doi.org/10.1509/jmr.15.0204>
- [11]. Athey, S., & Imbens, G. W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Literature*, 55(2), 371–441. <https://doi.org/10.1257/jel.20161097>
- [12]. Baylor, F., Fiedel, N., & et al. (2017). TFX: A TensorFlow-based production-scale machine learning platform. Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining,
- [13]. Becker, L., & Jaakkola, E. (2020). Customer experience: Fundamental premises and implications for research. *Journal of the Academy of Marketing Science*, 48(4), 630–648. <https://doi.org/10.1007/s11747-019-00731-4>
- [14]. Becker, L., Jaakkola, E., & Halinen, A. (2020). Toward a goal-oriented view of customer journeys. *Journal of Service Management*, 31(5), 1–25. <https://doi.org/10.1108/josm-11-2019-0329>
- [15]. Berman, R. (2018). Beyond the last-touch: Attribution in online advertising. *Marketing Science*, 37(5), 771–792. <https://doi.org/10.1287/mksc.2018.1085>
- [16]. Berry, S., Levinsohn, J., & Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4), 841–890. <https://doi.org/10.2307/2171802>
- [17]. Besbes, O., & Zeevi, A. (2009). Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6), 1407–1420. <https://doi.org/10.1287/opre.1090.0781>
- [18]. Blattberg, R. C., & Neslin, S. A. (1990). *Sales promotion: Concepts, methods, and strategies*. Springer. <https://doi.org/10.1007/978-1-4615-7681-5>
- [19]. Blei, D. M. (2012). Probabilistic topic models. *Communications of the ACM*, 55(4), 77–84. <https://doi.org/10.1145/2133806.2133826>
- [20]. Bolton, R. N. (1998). A dynamic model of the duration of the customer's relationship with a continuous service provider: The role of satisfaction. *Marketing Science*, 17(1), 45–65. <https://doi.org/10.1287/mksc.17.1.45>
- [21]. Breck, E., Cai, S., Nielsen, E., Salib, M., & Sculley, D. (2017). The ML test score: A rubric for ML production readiness and technical debt reduction. arXiv. <https://doi.org/10.48550/arXiv.1706.02259>
- [22]. Buckinx, W., & Van den Poel, D. (2005). Customer base analysis: Partial defection of behaviourally loyal clients in a non-contractual FMCG retail setting. *European Journal of Operational Research*, 164(1), 252–268. <https://doi.org/10.1016/j.ejor.2004.01.001>
- [23]. Burez, J., & Van den Poel, D. (2009). Handling class imbalance in customer churn prediction. *Expert Systems with Applications*, 36(3), 4626–4636. <https://doi.org/10.1016/j.eswa.2008.06.027>
- [24]. Burke, R. (2002). Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4), 331–370. <https://doi.org/10.1023/a:1021240730564>

- [25]. Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186. <https://doi.org/10.1126/science.aal4230>
- [26]. Cohen, M. C., Lobel, I., & Perakis, G. (2016). Dynamic pricing with demand learning and reference effects. *Operations Research*, 64(6), 2515–2536. <https://doi.org/10.1287/opre.2016.1540>
- [27]. Coussement, K., & Van den Poel, D. (2008). Integrating the voice of customers through call center emails into a decision support system for churn prediction. *Decision Support Systems*, 44(4), 990–1006. <https://doi.org/10.1016/j.dss.2007.05.005>
- [28]. Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How AI will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
- [29]. De Haan, E., Wiesel, T., Pauwels, K., & van Heerde, H. J. (2016). The effectiveness of different forms of online advertising for purchase conversion in a multiple-channel attribution framework. *International Journal of Research in Marketing*, 33(3), 491–507. <https://doi.org/10.1016/j.ijresmar.2015.12.005>
- [30]. Dempster, A. P., Laird, N. M., & Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1), 1–38. <https://doi.org/10.1111/j.2517-6161.1977.tb01600.x>
- [31]. den Boer, A. V. (2015). Dynamic pricing and learning: Historical review, current research, and new directions. *Surveys in Operations Research and Management Science*, 20(1), 1–18. <https://doi.org/10.1016/j.sorms.2015.03.001>
- [32]. DeSarbo, W. S. (2002). Market segment derivation and profiling via a finite mixture model framework. *Marketing Letters*, 13(1), 17–25. <https://doi.org/10.1023/a:1015059024154>
- [33]. DeSarbo, W. S., Wedel, M., Vriens, M., & Ramaswamy, V. (1992). Latent class metric conjoint analysis. *Marketing Letters*, 3(3), 273–288. <https://doi.org/10.1007/bf00994135>
- [34]. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). Fairness through awareness. Proceedings of the 3rd Innovations in Theoretical Computer Science Conference.
- [35]. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4), 211–407. <https://doi.org/10.1561/04000000042>
- [36]. Englehardt, S., & Narayanan, A. (2016). Online tracking: A 1-million-site measurement and analysis. Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security.
- [37]. Fader, P. S., Hardie, B. G. S., & Lee, K. L. (2005). "Counting your customers" the easy way: An alternative to the Pareto/NBD model. *Marketing Science*, 24(2), 275–284. <https://doi.org/10.1287/mksc.1040.0098>
- [38]. Gallego, G., & van Ryzin, G. (1994). Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8), 999–1020. <https://doi.org/10.1287/mnsc.40.8.999>
- [39]. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), Article 44. <https://doi.org/10.1145/2523813>
- [40]. Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
- [41]. Golam Qibria, L., & Takbir Hossen, S. (2023). Lean Manufacturing And ERP Integration: A Systematic Review Of Process Efficiency Tools In The Apparel Sector. *American Journal of Scholarly Research and Innovation*, 2(01), 104–129. <https://doi.org/10.63125/mx7j4p06>
- [42]. Gordon, B. R., Zettelmeyer, F., Bhargava, N., & Chapsky, D. (2019). A comparison of approaches to advertising measurement. *Marketing Science*, 38(2), 193–225. <https://doi.org/10.1287/mksc.2018.1135>
- [43]. Green, P. E., & Krieger, A. M. (1995). Alternative approaches to cluster-based market segmentation. *Journal of the Market Research Society*, 37(3), 221–239. <https://doi.org/10.1177/147078539503700302>
- [44]. Guidotti, R. (2019). A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 93. <https://doi.org/10.1145/3351095>
- [45]. Gupta, S., & Lehmann, D. R. (2006). Customer lifetime value research in marketing: A review and future directions. *Journal of Interactive Marketing*, 20(1), 87–93. <https://doi.org/10.1002/dir.20061>
- [46]. Gupta, S., Lehmann, D. R., & Stuart, J. A. (2004). Valuing customers. *Journal of Marketing Research*, 41(1), 7–18. <https://doi.org/10.1509/jmkr.41.1.7.25084>
- [47]. Hansotia, B. J., & Rukstales, B. (2002). Direct marketing for multichannel retailers: Issues, challenges and solutions. *Journal of Interactive Marketing*, 16(3), 69–81. <https://doi.org/10.1002/dir.10033>
- [48]. Hardt, M., Price, E., & Srebro, N. (2016). Equality of opportunity in supervised learning. *Advances in Neural Information Processing Systems*.
- [49]. Hazelwood, K., Smith, S., Dzhulgakov, D., & et al. (2018). Applied machine learning at Facebook: A datacenter infrastructure perspective. 2018 IEEE International Symposium on High Performance Computer Architecture (HPCA).
- [50]. He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T.-S. (2017). Neural collaborative filtering. Proceedings of the 26th International Conference on World Wide Web.

- [51]. Hendel, I., & Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6), 1637–1673. <https://doi.org/10.1111/j.1468-0262.2006.00720.x>
- [52]. Hosne Ara, M., Tonmoy, B., Mohammad, M., & Md Mostafizur, R. (2022). AI-ready data engineering pipelines: a review of medallion architecture and cloud-based integration models. *American Journal of Scholarly Research and Innovation*, 1(01), 319-350. <https://doi.org/10.63125/51kxtf08>
- [53]. Huang, M.-H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, 21(2), 155-172. <https://doi.org/10.1177/1094670517752459>
- [54]. Hutto, C., & Gilbert, E. (2014). VADER: A parsimonious rule-based model for sentiment analysis of social media text. *ICWSM*,
- [55]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2023). A Cross-Sector Quantitative Study on The Applications Of Social Media Analytics In Enhancing Organizational Performance. *American Journal of Scholarly Research and Innovation*, 2(02), 274-302. <https://doi.org/10.63125/d8ree044>
- [56]. Jahid, M. K. A. S. R. (2022). Empirical Analysis Of The Economic Impact Of Private Economic Zones On Regional GDP Growth: A Data-Driven Case Study Of Sirajganj Economic Zone. *American Journal of Scholarly Research and Innovation*, 1(02), 01-29. <https://doi.org/10.63125/je9w1c40>
- [57]. Klaus, P., & Maklan, S. (2012). EXQ: A multiple-item scale for assessing service experience. *Journal of Service Management*, 23(1), 5-33. <https://doi.org/10.1108/09564231211208952>
- [58]. Kleinberg, J., Mullainathan, S., & Raghavan, M. (2017). Inherent trade-offs in the fair determination of risk scores. *Proceedings of Innovations in Theoretical Computer Science (ITCS 2017)*,
- [59]. Kohavi, R., Longbotham, R., Sommerfield, D., & Henne, R. M. (2009). Controlled experiments on the web: Survey and practical guide. *Data Mining and Knowledge Discovery*, 18(1), 140–181. <https://doi.org/10.1007/s10618-008-0114-1>
- [60]. Kohavi, R., Tang, D., Xu, Y., & Chen, N. (2014). Seven pitfalls to avoid when running controlled experiments on the web. *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*,
- [61]. Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8), 30–37. <https://doi.org/10.1109/mc.2009.263>
- [62]. Kutub Uddin, A., Md Mostafizur, R., Afrin Binta, H., & Maniruzzaman, B. (2022). Forecasting Future Investment Value with Machine Learning, Neural Networks, And Ensemble Learning: A Meta-Analytic Study. *Review of Applied Science and Technology*, 1(02), 01-25. <https://doi.org/10.63125/edxgig56>
- [63]. Lemmens, A., & Croux, C. (2006). Bagging and boosting classification trees to predict churn. *Journal of Marketing Research*, 43(2), 276–286. <https://doi.org/10.1509/jmkr.43.2.276>
- [64]. 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>
- [65]. Lewis, R. A., & Reiley, D. H. (2014). Online ads and offline sales. *Quantitative Marketing and Economics*, 12(3), 235-266. <https://doi.org/10.1007/s1129-014-9146-6>
- [66]. Li, H., & Kannan, P. K. (2014). Attributing conversions in a multichannel-online marketing environment: An empirical model and a field experiment. *Journal of Marketing Research*, 51(1), 40–56. <https://doi.org/10.1509/jmr.13.0200>
- [67]. Li, L., Chu, W., Langford, J., & Schapire, R. (2010). A contextual-bandit approach to personalized news article recommendation. *WWW '10*,
- [68]. Linden, G., Smith, B., & York, J. (2003). Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1), 76-80. <https://doi.org/10.1109/mic.2003.1167344>
- [69]. Lipton, Z. C. (2018). The mythos of model interpretability. *Communications of the ACM*, 61(10), 36–43. <https://doi.org/10.1145/3233231>
- [70]. Liu, B. (2012). *Sentiment analysis and opinion mining* (Vol. 5). <https://doi.org/10.2200/s00416ed1v01y201204hlt016>
- [71]. Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions (SHAP). *Advances in Neural Information Processing Systems*,
- [72]. Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4), 336–355. <https://doi.org/10.1287/isre.1040.0015>
- [73]. Mansura Akter, E. (2023). Applications Of Allele-Specific PCR In Early Detection of Hereditary Disorders: A Systematic Review Of Techniques And Outcomes. *Review of Applied Science and Technology*, 2(03), 1-26. <https://doi.org/10.63125/n4h7t156>
- [74]. Mansura Akter, E., & Md Abdul Ahad, M. (2022). In Silico drug repurposing for inflammatory diseases: a systematic review of molecular docking and virtual screening studies. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 35-64. <https://doi.org/10.63125/j1hbt51>
- [75]. Martin, K. D., Kim, J. J., & Palmatier, R. W. (2020). Data privacy in retail. *Journal of Retailing*, 96(4), 474-489. <https://doi.org/10.1016/j.jretai.2020.08.003>

- [76]. McInnes, L., Healy, J., & Astels, S. (2017). hdbscan: Hierarchical density based clustering. *Journal of Open Source Software*, 2(11), 205. <https://doi.org/10.21105/joss.00205>
- [77]. McLachlan, G., & Peel, D. (2000). *Finite mixture models*. Wiley. <https://doi.org/10.1002/0471721182>
- [78]. McMahan, H. B., Holt, G., Sculley, D., Young, M., Ebner, D., Grady, J., Nie, L., Phillips, T., Davydov, E., Golovin, D., Chikkerur, S., Liu, D., Wattenberg, M., Hrafnkelsson, A. M., Boulos, T., & Kubica, J. (2013). Ad click prediction: A view from the trenches. *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [79]. Md Arifur, R., & Sheratun Noor, J. (2022). A Systematic Literature Review of User-Centric Design In Digital Business Systems: Enhancing Accessibility, Adoption, And Organizational Impact. *Review of Applied Science and Technology*, 1(04), 01-25. <https://doi.org/10.63125/ndjkpm77>
- [80]. Md Hasan, Z., Sheratun Noor, J., & Md. Zafor, I. (2023). Strategic role of business analysts in digital transformation tools, roles, and enterprise outcomes. *American Journal of Scholarly Research and Innovation*, 2(02), 246-273. <https://doi.org/10.63125/rc45z918>
- [81]. Md Mahamudur Rahaman, S. (2022). Electrical And Mechanical Troubleshooting in Medical And Diagnostic Device Manufacturing: A Systematic Review Of Industry Safety And Performance Protocols. *American Journal of Scholarly Research and Innovation*, 1(01), 295-318. <https://doi.org/10.63125/d68y3590>
- [82]. Md Masud, K., Mohammad, M., & Sazzad, I. (2023). Mathematics For Finance: A Review of Quantitative Methods In Loan Portfolio Optimization. *International Journal of Scientific Interdisciplinary Research*, 4(3), 01-29. <https://doi.org/10.63125/j43ayz68>
- [83]. Md Nur Hasan, M., Md Musfiqur, R., & Debashish, G. (2022). Strategic Decision-Making in Digital Retail Supply Chains: Harnessing AI-Driven Business Intelligence From Customer Data. *Review of Applied Science and Technology*, 1(03), 01-31. <https://doi.org/10.63125/6a7rpy62>
- [84]. Md Sultan, M., Proches Nolasco, M., & Md. Torikul, I. (2023). Multi-Material Additive Manufacturing For Integrated Electromechanical Systems. *American Journal of Interdisciplinary Studies*, 4(04), 52-79. <https://doi.org/10.63125/y2ybrx17>
- [85]. Md Takbir Hossen, S., Ishtiaque, A., & Md Atiqur, R. (2023). AI-Based Smart Textile Wearables For Remote Health Surveillance And Critical Emergency Alerts: A Systematic Literature Review. *American Journal of Scholarly Research and Innovation*, 2(02), 1-29. <https://doi.org/10.63125/ceqapd08>
- [86]. Md Takbir Hossen, S., & Md Atiqur, R. (2022). Advancements In 3d Printing Techniques For Polymer Fiber-Reinforced Textile Composites: A Systematic Literature Review. *American Journal of Interdisciplinary Studies*, 3(04), 32-60. <https://doi.org/10.63125/s4r5m391>
- [87]. Md Tawfiqul, I. (2023). A Quantitative Assessment Of Secure Neural Network Architectures For Fault Detection In Industrial Control Systems. *Review of Applied Science and Technology*, 2(04), 01-24. <https://doi.org/10.63125/3m7gbs97>
- [88]. Md Tawfiqul, I., Meherun, N., Mahin, K., & Mahmudur Rahman, M. (2022). Systematic Review of Cybersecurity Threats In IOT Devices Focusing On Risk Vectors Vulnerabilities And Mitigation Strategies. *American Journal of Scholarly Research and Innovation*, 1(01), 108-136. <https://doi.org/10.63125/wh17mf19>
- [89]. Md.Kamrul, K., & Md Omar, F. (2022). Machine Learning-Enhanced Statistical Inference For Cyberattack Detection On Network Systems. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 65-90. <https://doi.org/10.63125/sw7jzx60>
- [90]. Mela, C. F., Gupta, S., & Lehmann, D. R. (1997). The long-term impact of promotion and advertising on consumer brand choice. *Journal of Marketing Research*, 34(2), 248-261. <https://doi.org/10.2307/3152060>
- [91]. Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*, 267, 1-38. <https://doi.org/10.1016/j.artint.2018.07.007>
- [92]. Mitchell, M., Wu, S., Zaldivar, A., & et al. (2019). Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*,
- [93]. Mst Shamima, A., Niger, S., Md Atiqur Rahman, K., & Mohammad, M. (2023). Business Intelligence-Driven Healthcare: Integrating Big Data And Machine Learning For Strategic Cost Reduction And Quality Care Delivery. *American Journal of Interdisciplinary Studies*, 4(02), 01-28. <https://doi.org/10.63125/crv1xp27>
- [94]. Mubashir, I., & Abdul, R. (2022). Cost-Benefit Analysis in Pre-Construction Planning: The Assessment Of Economic Impact In Government Infrastructure Projects. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 91-122. <https://doi.org/10.63125/kjwd5e33>
- [95]. Neslin, S. A., Gupta, S., Kamakura, W., Lu, J., & Mason, C. (2006). Defection detection: Predictive accuracy of churn models. *Journal of Marketing Research*, 43(2), 204-211. <https://doi.org/10.1509/jmkr.43.2.204>
- [96]. Netzer, O., Feldman, R., Goldenberg, J., & Fresko, M. (2012). Mine your own business: Market-structure surveillance through text mining. *Marketing Science*, 31(3), 521-543. <https://doi.org/10.1287/mksc.1120.0713>

- [97]. Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. *Proceedings of the 22nd International Conference on Machine Learning*,
- [98]. Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1–2), 1–135. <https://doi.org/10.1561/1500000011>
- [99]. Pipino, L. L., Lee, Y. W., & Wang, R. Y. (2002). Data quality assessment. *Communications of the ACM*, 45(4), 211–218. <https://doi.org/10.1145/505248.506010>
- [100]. Punj, G., & Stewart, D. W. (1983). Cluster analysis in marketing research: Review and suggestions. *Journal of Marketing Research*, 20(2), 134–148. <https://doi.org/10.1177/002224378302000204>
- [101]. Quadrana, M., Cremonesi, P., & Jannach, D. (2018). Sequence-aware recommender systems. *ACM Computing Surveys*, 51(4), 66. <https://doi.org/10.1145/3190616>
- [102]. Reduanul, H., & Mohammad Shueb, A. (2022). Advancing AI in Marketing Through Cross Border Integration Ethical Considerations And Policy Implications. *American Journal of Scholarly Research and Innovation*, 1(01), 351–379. <https://doi.org/10.63125/d1xg3784>
- [103]. Reinartz, W., & Kumar, V. (2003). The impact of customer relationship characteristics on profitable lifetime duration. *Journal of Marketing*, 67(1), 77–99. <https://doi.org/10.1509/jmkg.67.1.77.18585>
- [104]. Rendle, S. (2010). Factorization machines. 2010 IEEE International Conference on Data Mining,
- [105]. Rezwanul Ashraf, R., & Hosne Ara, M. (2023). Visual communication in industrial safety systems: a review of UI/UX design for risk alerts and warnings. *American Journal of Scholarly Research and Innovation*, 2(02), 217–245. <https://doi.org/10.63125/wbv4z521>
- [106]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”. KDD ‘16,
- [107]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2018). Anchors: High-precision model-agnostic explanations. *Proceedings of the AAAI Conference on Artificial Intelligence*,
- [108]. Russo, D., Van Roy, B., Kazerouni, A., Osband, I., & Wen, Z. (2018). A tutorial on Thompson sampling. *Foundations and Trends® in Machine Learning*, 11(1), 1–96. <https://doi.org/10.1561/22000000070>
- [109]. Rust, R. T., Lemon, K. N., & Zeithaml, V. A. (2004). Return on marketing: Using customer equity to focus marketing strategy. *Journal of Marketing*, 68(1), 109–127. <https://doi.org/10.1509/jmkg.68.1.109.24030>
- [110]. Rzepakowski, P., & Jaroszewicz, S. (2012). Decision trees for uplift modeling with single and multiple treatments. *Data Mining and Knowledge Discovery*, 24(1), 30–54. <https://doi.org/10.1007/s10618-012-0294-1>
- [111]. Sanjai, V., Sanath Kumar, C., Maniruzzaman, B., & Farhana Zaman, R. (2023). Integrating Artificial Intelligence in Strategic Business Decision-Making: A Systematic Review Of Predictive Models. *International Journal of Scientific Interdisciplinary Research*, 4(1), 01–26. <https://doi.org/10.63125/s5skge53>
- [112]. Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2001). Item-based collaborative filtering recommendation algorithms. WWW ‘01,
- [113]. Sazzad, I., & Md Nazrul Islam, K. (2022). Project impact assessment frameworks in nonprofit development: a review of case studies from south asia. *American Journal of Scholarly Research and Innovation*, 1(01), 270–294. <https://doi.org/10.63125/eeja0t77>
- [114]. Schmittlein, D. C., Morrison, D. G., & Colombo, R. (1987). Counting your customers: Who-are they and what will they do next? *Marketing Science*, 6(1), 1–24. <https://doi.org/10.1287/mksc.6.1.1>
- [115]. Sculley, D., Holt, G., Golovin, D., & et al. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*,
- [116]. Sheratun Noor, J., & Momena, A. (2022). Assessment Of Data-Driven Vendor Performance Evaluation in Retail Supply Chains: Analyzing Metrics, Scorecards, And Contract Management Tools. *American Journal of Interdisciplinary Studies*, 3(02), 36–61. <https://doi.org/10.63125/0s7t1y90>
- [117]. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. 2017 IEEE Symposium on Security and Privacy (SP),
- [118]. Singh, A., & Joachims, T. (2018). Fairness of exposure in rankings. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*,
- [119]. Sweeney, L. (2002). k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5), 557–570. <https://doi.org/10.1142/s0218488502001648>
- [120]. Tahmina Akter, R., Debashish, G., Md Soyeb, R., & Abdullah Al, M. (2023). A Systematic Review of AI-Enhanced Decision Support Tools in Information Systems: Strategic Applications In Service-Oriented Enterprises And Enterprise Planning. *Review of Applied Science and Technology*, 2(01), 26–52. <https://doi.org/10.63125/73djw422>
- [121]. Talluri, K. T., & van Ryzin, G. J. (2005). *The theory and practice of revenue management*. Springer. <https://doi.org/10.1007/978-0-387-73164-8>
- [122]. Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533. [https://doi.org/10.1002/\(sici\)1097-0266\(199708\)18:7<509::Aid-smj882>3.0.Co;2-z](https://doi.org/10.1002/(sici)1097-0266(199708)18:7<509::Aid-smj882>3.0.Co;2-z)

- [123]. Timoshenko, A., & Hauser, J. R. (2019). Identifying customer needs from user-generated content. *Marketing Science*, 38(1), 1–20. <https://doi.org/10.1287/mksc.2018.1123>
- [124]. Ustun, B., Spangher, A., & Liu, Y. (2019). Actionable recourse in linear classification. Proceedings of the Conference on Fairness, Accountability, and Transparency,
- [125]. van der Aalst, W. (2016). *Process mining: Data science in action (2nd ed.)*. Springer. <https://doi.org/10.1007/978-3-662-49851-4>
- [126]. Venkatesan, R., & Kumar, V. (2004). A customer lifetime value framework for customer selection and resource allocation strategy. *Journal of Marketing*, 68(4), 106–125. <https://doi.org/10.1509/jmkg.68.4.106.42728>
- [127]. Verbeke, W., Dejaeger, K., Martens, D., Hur, J., & Baesens, B. (2012). Profit-driven churn prediction. *European Journal of Operational Research*, 218(1), 211–229. <https://doi.org/10.1016/j.ejor.2011.09.031>
- [128]. Verhoef, P. C., Kannan, P. K., & Inman, J. J. (2015). From multi-channel to omni-channel retailing. *Journal of Retailing*, 91(2), 174–181. <https://doi.org/10.1016/j.jretai.2015.02.005>
- [129]. Verhoef, P. C., Lemon, K. N., & Parasuraman, A. (2009). Customer experience creation. *Journal of Retailing*, 85(1), 31–41. <https://doi.org/10.1016/j.jretai.2008.11.001>
- [130]. von Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and Computing*, 17(4), 395–416. <https://doi.org/10.1007/s11222-007-9033-z>
- [131]. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99. <https://doi.org/10.1093/idpl/ix005>
- [132]. Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems*, 12(4), 5–33. <https://doi.org/10.1080/07421222.1996.11518099>
- [133]. Wedel, M., & Kamakura, W. (2000). *Market Segmentation: Conceptual and Methodological Foundations*. <https://doi.org/10.1007/978-1-4615-4651-1>
- [134]. 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>
- [135]. Xie, J., Girshick, R., & Farhadi, A. (2016). Unsupervised deep embedding for clustering analysis. Proceedings of the 33rd International Conference on Machine Learning,
- [136]. Xu, J., Shao, X., Ma, J., Lee, K.-C., Qi, H., & Lu, Q. (2016). Lift-based bidding in ad selection. AAAI '16,
- [137]. Zaharia, M., Chen, A., Davidson, A., & et al. (2019). Accelerating the machine learning lifecycle with MLflow. arXiv. <https://doi.org/10.48550/arXiv.1910.13434>
- [138]. Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey. *ACM Computing Surveys*, 52(1), 5. <https://doi.org/10.1145/3285029>