10.48047/jocaaa.2026.35.01.65

From Static Rules to Autonomous Intelligence: How AI-Driven Bidding is Redefining Retail Advertising Economics

Ameya Gokhale
UCLA Anderson School of Management, USA

Abstract

Retail advertising has entered a structural inflection point. Rule-based bidding systems, long relied upon for their transparency and control, are increasingly unable to cope with the scale, speed, and complexity of modern retail media ecosystems. Drawing on a decade of hands-on experience designing and operating AI-driven bidding systems across global retail and advertising platforms, this article presents a practitioner framework for goal-based, machine-learning-powered bidding. It examines the technical architectures that enable real-time optimization, the economic mechanisms through which these systems create value, and the governance challenges that determine whether AI bidding delivers sustainable advantage or short-term gains. Rather than treating AI bidding as a point solution, the article argues that durable differentiation emerges from integrating machine intelligence with business judgment, proprietary data, and organizational processes. AI-enabled bidding systems redefine the rules of value creation for retail advertisers in digital marketplaces. Rule-based bidding based on conditional statements toward static digital features and indicator variables cannot operate efficiently in multidimensional feature spaces where hundreds of interacting signals across user context, product attributes, inventory availability, competitive dynamics, and temporal patterns determine auction outcomes. Industrial applications use multi-layer neural networks or gradient-increased decision trees for learning protocols to maximize business objectives while meeting severe latency requirements (measured in milliseconds). Production systems typically involve complex designs using feature engineering pipelines, distributed inference, and systems for dealing with selection bias and exploration, where solutions must balance short- and longterm returns from exploration. Value is usually generated by better temporal allocation, various types of automated micro-segmentation to capture long-tail demand, dynamic budget reallocations, as well as operational integration with inventory and margin structures. Major challenges to scale are interpretability and trust, lack of training data for cold-start performance, attribution challenges that lead to misalignment of optimization with business KPIs, technical integration into multiple business and technology systems, and evolving privacy regulations. This practitioner model introduces five interrelated layers: metaobjective design to true business value; value modeling to reconceptualize predictive business outcomes with operating constraints; policy optimization to support continuous learning; measurement validation necessary to avoid pitfalls that accompany incrementality-aware causality approaches; and governance to ensure that value is delivered consistent with business strategy and regulation. In this approach, sustainable competitive advantage comes from embedding machine intelligence into proprietary firstparty data assets and operational processes via demand orchestration infrastructure rather than having superior algorithms.

Keywords: Autonomous Bidding Systems, Retail Media Optimization, Machine Learning Infrastructure, Goal-Based Advertising, Demand Orchestration

10.48047/jocaaa.2026.35.01.65

1. Introduction

Retail media has evolved into one of the fastest-growing segments of digital advertising, driven by the convergence of first-party data, closed-loop measurement, and performance-oriented buying models. Early campaign management relied heavily on static rules: predefined bid multipliers, day-parting schedules, and heuristic adjustments encoded by human operators. While these systems offered control, they were fundamentally constrained by human cognition and operational latency.

As retail platforms scaled to millions of products, advertisers, and auctions per second, these constraints became structural rather than incidental. The emergence of AI-driven bidding systems represents not incremental optimization, but a shift in how advertising value is created, allocated, and governed. This article synthesizes technical design patterns and economic insights drawn from production-scale systems to articulate how autonomous bidding reshapes retail advertising economics—and where its limits remain. Today's AI-driven bidding platforms harness advanced neural network architectures alongside ensemble learning methods to digest massive quantities of real-time information. Everything from temporal patterns and device characteristics to demographic markers and behavioral indicators gets synthesized into actionable bidding decisions within milliseconds [2]. The necessary raw computational power will require the infrastructure to be able to handle millions of inference requests every second and still maintain the microsecond-level responsiveness that real-time auctions of an advertising company require.

In addition to the advantages of efficiency, such a technological breakthrough fundamentally changes the nature of competition in retail advertising. The availability of advanced optimization systems that had previously been restricted to businesses with large analytical capacity has been brought nearer to the masses. Small retailers now compete on more equal footing with industry giants. The shift from static rules to autonomous intelligence marks a complete reconceptualization of value creation and distribution across digital advertising ecosystems. This article delivers a thorough technical examination of AI-driven bidding systems—their architectural foundations, economic ramifications, and the implementation hurdles retailers encounter during the adoption of these transformative technologies.

2. Why Rule-Based Bidding Breaks at Scale

Rule-based bidding systems encode business logic through explicit conditional statements: if conversion rates rise, increase bids; if performance drops, pull back spend. This approach fails as dimensionality increases. Modern retail auctions are influenced by hundreds of interacting signals, including user context, product attributes, inventory availability, competitive dynamics, and temporal effects.

These systems are also inherently reactive. Adjustments occur only after performance changes manifest in aggregated reports, introducing lag in environments where auction conditions evolve in milliseconds. Over time, maintenance costs grow as rules proliferate, interact unpredictably, and decay under data drift. The result is a fragile system that consumes human effort while systematically underperforming. The fundamental problem with rule-based bidding systems is their inability to scale to the exponential explosion of interactions among variables that characterizes the modern retail advertising space. Rulebased systems can scale to fewer than twenty key performance indicators, after which the number of rules is unmanageable to write. Modern retail media platforms create feature spaces with potentially complex interactions among demographic characteristics, behavioral trajectories, contextual signals, competitive signals, inventory states, and temporal dimensions [1].

This is the case because maintaining a good set of rules has a super-linear cognitive load. For example, adding a new variable not only creates one new rule but creates a dependency between that rule and every rule already in place. Research on machine learning for digital ad placement optimization shows that

10.48047/jocaaa.2026.35.01.65

humans deal with only a limited number of relevant rules before their performance deteriorates. Beyond a certain threshold, the individual rules will start to conflict with one another as a rule within one area of a campaign negatively impacts others.

Another potential failure mode is temporal misalignment. That is, rules do not necessarily match the current state of the market. The current market state is constantly changing in real-time bidding domains as competitors change their bidding strategy, users change contexts, and inventory changes hands. Another issue with observing and reacting to the behavior of ads in real-time bidding exchanges is that it can take hours or days for a human to observe a change, hypothesize a new rule, enact it, and observe the effect, whereas the auctions themselves typically occur within a few microseconds. This means that any rule-based approach is forever optimized for an environment that has already changed.

The temporal problem is further complicated by data drift. Consumer behavior patterns change over time due to seasonal effects, economic changes, competitors, and other factors. The data characteristics change over time, and the decision rules extracted from the historical data become less relevant. Research has found that deep learning methods developed for predicting the sales of retail stores have rigid decision rules that require constant specialized intervention to adapt to changing consumer and market behavior [5]. This maintenance burden in turn consumes analytic capacity to update tactical rules, thereby limiting the firm's ability to explore new advertising opportunities and to respond to competitive threats.

Limitation Dimension	Operational Manifestation	System Impact
Cognitive Scalability	Manual rule specification exceeds human capacity	Incomplete coverage of variable interactions
Dimensional Complexity	Hundreds of interacting signals require exponential rules	Cascading failures and rule conflicts
Temporal Responsiveness	Hours-to-days feedback loops in microsecond auctions	Perpetual misalignment with market conditions
Pattern Recognition	Limited to simple bivariate relationships	Unexploited non-linear optimization potential
Maintenance Overhead	Continuous manual recalibration required	Resource diversion from strategic initiatives
Adaptation Capability	Rules decay under data drift	Performance degradation over time
Competitive Dynamics	Static logic is vulnerable to strategic opponents	Predictable behavior exploited by competitors
Long-tail Management	Economically infeasible to manually manage niche segments	Missed revenue opportunities

Table 1: Structural Limitations of Rule-Based Bidding Systems [3, 4]

10.48047/jocaaa.2026.35.01.65

3. Architecture of AI-Driven Bidding Systems

AI-driven bidding systems replace explicit rules with learned policies optimized against defined business objectives. At their core are predictive models—commonly gradient-boosted decision trees or deep neural networks—trained on historical impression, interaction, and conversion data. These models estimate outcome probabilities or value functions under strict latency constraints.

Production architectures balance accuracy, speed, and interpretability. Feature pipelines transform highcardinality signals into compact representations, while inference layers are optimized through quantization, caching, and distributed serving. Continuous learning loops incorporate feedback from observed outcomes while addressing selection bias through controlled exploration and off-policy evaluation techniques. Building production-scale bidding systems is substantially more complex than applying the model. Modern bidding systems are typically multi-stage pipelines that trade off predictive performance, computational and storage efficiency, interpretability, and ease of adaptation to changes. The feature engineering layer transforms the raw auction signals into representations that maximize predictive power, while being efficient enough to support real-time inference in real-world applications. Recent work on deep learning architectures for real-time bid optimizers, however, has focused on designing neural networks that are optimized for the specific characteristics of advertising data, including sparsity of user interaction histories, the high cardinality of some categorical features (e.g. product identifiers), and generalization across many different problem contexts, sometimes where little training data is available for any one context. An embedding layer is one of the first transformations that maps discrete identifiers into a continuous vector space, such that, through training, semantically similar products can be mapped to be nearby in the embedding space, which allows the model to generalize the learning of interactions into previously unseen combinations.

Architectural choices in the inference serving layer are typically made based on low and stable latency, rather than high throughput. In batch prediction systems, predictions can be made on large batches of requests to minimize the amount of computation, but in real-time bidding, the latency requirements may be in the single-digit milliseconds. In production systems tailored for low latency in real-time bidding tasks, tiered caching can be used, where frequently requested predictions are pre-computed and actioned in a first-stage, cases with intermediate complexity are inferred using a last-stage of lower complexity, while full model inference is only performed for out-of-distribution or high-value cases. This results in worst-case latency guarantees, as opposed to average-case guarantees.

Once the model has been deployed, the model serving infrastructure has to support the ability to serve new models with minimal service interruptions as the model is retrained and new training data or updated models are produced in light of performance metrics. Shadow deployments run models in parallel to production models on live requests without changing the behavior of the bidding agent. They validate any improvements, testing in a production-like deployment environment before deploying into production. Canary deployments send some traffic to the updated models to detect degraded quality or unusual behavior.

Another architecturally relevant aspect is the exploration-exploitation nature of the bidding strategy used, which has a large impact on performance. Specifically, exploitative bidding strategies that always bid according to the perceived currently optimal value tend to become stuck in a local optimum, which may systematically overlook helpful areas. Studying lightweight techniques in real-time bidding makes it clear that carefully controlled randomization, which deliberately sacrifices greedy optimization in favor of exploration, is helpful. Thompson sampling fits this requirement, as it can be interpreted as relinquishing point estimation in favor of a distribution over the model parameters, motivating exploration proportional

10.48047/jocaaa.2026.35.01.65

to uncertainty. Thus, a balance is struck between the immediate maximum benefit and informational gain such that the agent learns better policies over a longer time.

Architectural Layer	Technical Implementation	Performance Requirement
Feature Engineering	Embedding layers for high-cardinality variables	Compact predictive representations
Predictive Models	Deep neural networks or gradient-boosted trees	High-accuracy probability estimation
Inference Serving	Tiered caching with distributed computing	Single-digit millisecond latency
Model Updates	Shadow deployment with canary rollouts	Zero-downtime model transitions
Exploration Strategy	Thompson sampling with uncertainty quantification	Balanced exploitation and discovery
Feedback Integration	Continuous learning with bias correction	Adaptation to evolving conditions
Value Calculation	Expected value optimization frameworks	Alignment with business objectives
Serving Infrastructure	Quantization and pre-computation strategies	Consistent latency across request distribution

Table 2: Production Architecture Components for AI-Driven Bidding [5, 6]

4. A Practitioner Framework for Goal-Based Bidding

Based on observed production deployments, effective AI bidding systems follow five interdependent layers:

- Objective Design: Clearly defined, goal-based metrics aligned with true business value, not merely platform-reported proxies.
- Value Modeling: Integration of predicted outcomes with margins, inventory constraints, and fulfillment costs.
- Policy Optimization: Continuous optimization of bids as functions of predicted value rather than static thresholds.
- Measurement and Validation: Ongoing evaluation using incrementality-aware methods to detect bias and prevent metric gaming.
- Governance and Oversight: Human-in-the-loop controls, explainability surfaces, and guardrails that ensure alignment with strategic intent and regulatory requirements.

This framework shifts AI bidding from a tactical automation tool to a strategic decision system.

10.48047/jocaaa.2026.35.01.65

The practitioner framework for goal-based bidding rests on production learnings from various retail verticals and advertiser objectives. The design layer for objectives is one in which care must be exercised, both in terms of measurement and calculated alignment. Most advertisers optimize for platform-reported metrics, such as clicks and attributed conversions, without investigating whether the metrics indicate business value. A review of AI-based marketing analytics for retail strategy highlights how platform attribution models are biased by selection bias, time window, and channel aggregation bias [7].

More advanced approaches use incrementality measurement systems to estimate the causal impact of ad exposure instead of correlation. This relies on experiments, creating a randomized control group that receives less advertising or no advertising in order to compute the true advertising lift. The value modeling layer can convert incrementality estimates into dollar estimates of expected value using profit margins, customer lifetime value, inventory carrying and holding costs, or operational constraints. When retailers have both high and low margin goods, the bidding system should favor the promotion of items with positive profit margins.

They continuously adjust policy optimization mechanisms to reach a desired outcome and do so without human intervention. Previous methods for pricing, bidding, and business optimizations have been replaced by optimization problems that can accommodate predicted value, budget constraints, competing objectives, multiple time horizons, and more complex dynamics, rather than predefined rules. Artificial intelligence refers to dynamic programming and reinforcement learning for bidding policies that take sequential dependencies and the long-term effect of decisions into account [8]. If at the current time step a customer is acquired in an auction, they go into the purchase funnel and could be converted in the long term, thus leading to a higher acquisition cost.

To verify these models, it needs to take into account the selection bias in observational data, wherein models see their data for only the auctions they have won, but not lost. Off-policy evaluation methods can approximate the counterfactual behavior of hypothetical bidding policies without actually deploying them. This is important for exploration and unbiased evaluation of bidding policies. Broader applications of privacy-preserving retail analytics present this growing trend of off-policy evaluation, especially in settings where metrics can only be partially measured due to signal deprecation [10].

Finally, a governance layer in place keeps the bidding systems in alignment with changing business and regulatory requirements. Explainability surfaces provide insights into the model's decisions, building trust with business users, increasing marketing team confidence in the bidding system, and helping with diagnosing unexpected behavior. The challenge of AI interpretability cannot only be solved by technical explanations, and requires interfaces that provide a model's behavior in terms understandable by business stakeholders, and in accordance with business requirements [9].

Framework Layer	Core Components	Strategic Purpose
Objective Design	Incrementality-based metrics aligned with business value	Avoid optimization toward invalid proxies
Value Modeling	Integration of margins, inventory, and fulfillment costs	Connect advertising to operational reality
Policy Optimization	Continuous adjustment mechanisms with sequential awareness	Enable dynamic response to conditions

10.48047/jocaaa.2026.35.01.65

Measurement Validation	Off-policy evaluation and controlled experiments	Detect bias and prevent metric gaming
Governance Oversight	Explainability surfaces with strategic guardrails	Maintain alignment and build stakeholder trust
Data Infrastructure	First-party asset development and privacysafe methods	Enable robust modeling under signal constraints
Organizational Capability	Cross-functional collaboration structures	Embed AI within demand orchestration
Performance Monitoring	Real-time anomaly detection systems	Enable rapid intervention when needed

Table 3: Practitioner Framework for Goal-Based Bidding Implementation [7, 8]

5. Economic Implications for Retail Advertising

When correctly implemented, AI-driven bidding reallocates advertising budgets toward higher marginal returns across time, audience segments, and product portfolios. Automated micro-segmentation captures value in long-tail demand, while temporal optimization concentrates spend during high-propensity periods. More importantly, integrating bidding decisions with inventory and margin data aligns marketing spend with operational reality. This coordination reduces waste, improves profitability, and changes how retailers think about advertising—from a cost center to an extension of demand orchestration. The unexpected economic consequences of AI bidding extended beyond improving efficiency and into fundamentally changing the structure of retail advertising markets. Until the rise of AI, retail advertising could be understood via limited optimization where advertisers competed on creative quality, brand strength, and budget. AI bidding shifts competition away from budget size to data quality, algorithms, and operational integration, areas wherein sustainable advantages are more likely.

Machine learning allows micro-segmentation, which enables profitable simultaneous advertising at a scale and detail that was previously impossible. Research on machine learning optimization in digital advertising finds that AI-based systems can profitably bid for thousands of long-tail keywords and microtargeted users, which are almost impossible for humans to profitably target. These tend to have low volumes and have opportunities for important revenue, and encounter less competitive pressure. The value of this is that it expands the addressable markets that were previously unavailable because the cost to manage that advertising was too high.

Temporal optimization refers to the economic benefits of shaping the tempo and intensity of advertising to match the purchase propensity of consumers. Empirical research on deep learning in retail sales forecasting reports that demand for consumer products can be characterized by complex temporal patterns that are driven by day-of-week effects, seasonality, weather, pay days, and promotions [5]. By automatically discovering and leveraging these patterns, AI systems may optimize ad spend to arbitrage temporal price discrepancies in advertising inventory, thereby optimizing for capital efficiency by avoiding wasted spend when conversions are unlikely, and increasing presence and spend through times of peak demand.

Advertising with inventory and fulfillment holds economic value when it aligns demand generation with supply chain capacity. For example, when reviewing AI-powered marketing analytics in retail strategy, it

10.48047/jocaaa.2026.35.01.65

has been observed that the best-performing applications include real-time inventory positions, fulfillment costs, and margin structures as a part of their bidding decisions [7]. Thus, the automatic promotion management solution eliminates wastage by not advertising products that run out of stock during a promotion, reduces promotion levels automatically when products approach stockout, and focuses advertising on high-margin products as long as there are enough stock levels. These economic benefits go beyond improved advertising efficiency. Advertising thus becomes a demand orchestration mechanism that spans the entire retail value chain and not just a marketing silo in its own right.

Challenge Category	Primary Risk	Mitigation Strategy
Model Interpretability	Stakeholder distrust of opaque systems	Explainability interfaces and governance frameworks
Data Sufficiency	Cold-start performance gaps	Transfer learning from related domains
Attribution Validity	Misalignment with true business value	Incrementality measurement and multitouch frameworks
Technical Integration	Cross-system coordination complexity	Modular architecture with phased deployment
Privacy Compliance	Signal deprecation under regulations	First-party data emphasis and privacysafe methods
Strategic Alignment	Short-term metric optimization versus long-term goals	Human oversight with strategic constraints
Competitive Dynamics	AI adoption is driving auction price inflation	Differentiation through proprietary data assets
Organizational Readiness	Capability gaps in technical and strategic dimensions	Deliberate capability building and change management

Table 4: Implementation Challenges and Risk Mitigation Strategies [9, 10]

6. Challenges and Failure Modes

AI bidding introduces new risks alongside its benefits. Opaque models can erode trust if explainability is not addressed. Cold-start conditions limit performance in low-data regimes. Attribution models optimized for short-term signals may misalign with long-term customer value.

Privacy regulations further constrain signal availability, increasing reliance on first-party data and robust modeling techniques. Successful systems treat these challenges as design constraints rather than afterthoughts, embedding governance and adaptability into the architecture. Challenges to deploying AIbased bidding systems in practice are by no means restricted to a technical nature but rather also organizational, calculated, and regulatory. Together, these factors drive the success of AI systems, the

10.48047/jocaaa.2026.35.01.65

deployment of these systems, and whether these deployments create enduring value or lead to other negative externalities. Hence, interpretability is an adoption and governance issue. Investigations into interpretability challenges faced by AI tools have established that marketing professionals tend to distrust systems that learn to make recommendations by learning patterns from millions of examples rather than explicit business rules [9].

This opacity leads to downstream problems: If stakeholders see the technology suddenly spending money on keywords that don't convert, or the reverse, spending far less on highly converting keywords, they're less likely to trust the spending decisions, and may push to switch back to manual control. If marketing executives don't understand the marketing rationale for the AI recommendations, they can't defend the budget to their executive stakeholders. In more regulated industries, accuracy extends to the fairness and privacy of advertising practices that utilize AI, as well as their compliance with standards of transparency and auditability.

Another closely related problem is the cold-start problem, usually seen in retail recommender systems that face new products or new markets, or at the tails of the distribution of a dataset, where the amount of available data to learn from is very small. In such cases, prediction error generally decreases logarithmically. Optimizing a retail environment is another example where transfer learning (fine-tuning models trained in other domains on relatively small amounts of data in a specific domain) can offer partially-optimal solutions but cannot eliminate the performance gap entirely, at least initially [8]. Due to the challenges of attribution, optimization objectives are misaligned with business outcomes. Most AI bidding platforms optimize for platform-defined conversion events and short attribution windows (i.e., last-click or platform-defined attribution models). In practice, however, customer purchase decisions often include multiple channels, touchpoints, and time horizons.

Conclusion

AI-driven bidding marks a structural shift in retail advertising economics. Competitive advantage no longer derives from access to automation alone, but from the ability to integrate machine intelligence with strategy, data, and organizational judgment. Retailers that approach AI bidding as core infrastructure— rather than a plug-in optimization—will define the next phase of retail media maturity. The shift from static bid management to AI-based bidding for advertising in retail search is not simply an innovation adoption cycle. It is a paradigmatic reconfiguration of value creation, value allocation, and value governance. As machine learning systems move from proofs of concept to production, the competitive dimension shifts from access to technology to data asset quality, operational depth, and organizational capacity to navigate the trade-off between calculated control and tactical flexibility. Retailers pioneering this transition have found that AI-based bidding models can be most effective as part of a demand orchestration system that connects advertising expenditure with inventory management, pricing, fulfillment, and customer lifetime value. This requires investments in model explainability, governance processes, and stakeholder education, particularly when organizations are working with advanced machine learning algorithms that are related to black boxes, and when there is a need for organizational trust in algorithmic decision-making. Third-party signal deprecation creates a need for more advanced data strategies. First-party asset development, privacy-centric measurement, better modeling, and advanced attribution techniques are expected to be critical. Attribution systems will need to move beyond platform-defined last-click attribution and toward incrementality-based measurement and analysis, allowing for true causal understanding of the effectiveness of marketing and the total long-term value of any customer. This includes the cold-start problem of predicting the performance of newly launched products and market entries, which requires transfer learning,

10.48047/jocaaa.2026.35.01.65

conservative initialization, and patient capital. It also requires an architecture that trades off predictive accuracy against the latency in making inferences, incorporates continual learning methods that update in response to changing market environment conditions, and implements controlled exploration methods that avoid optimization myopia. Going forward, differentiating capabilities will include proprietary data moats from first-party customer interactions, economies of coordination across the retail value chain from deeply integrated operations, and the embedding of human judgment and calculated thinking alongside machine intelligence in the retail organization. The retailers that will shape the next chapter of retail media maturity will see AI bidding as more than a technology deployment program, but as a vehicle of transformational change that rethinks how marketing, operations, and strategy partner to drive customer demand at scale. But ultimately, success depends on whether the question of whether to implement algorithmically driven bidding becomes a much harder question of the way to construct systems whereby artificially clever and human services are combined to yield a comparative advantage not just in algorithms but in data, integration of services, and organizational capabilities.

References

[1] Suranjit Kosta and Maganti Sushanth Babu, "Machine Learning-Based Optimization of Digital Ad Placement: A Literature Review," ResearchGate, 2024. [Online]. Available: https://www.researchgate.net/publication/391008328_MACHINE_LEARNING-BASED_OPTIMIZATION_OF_DIGITAL_AD_PLACEMENT_A_LITERATURE_REVIEW

[2] Annette Nyembe, "Deep Learning for Real-Time Bid Optimization," Madgicx, 2025. [Online]. Available: https://madgicx.com/blog/deep-learning-for-real-time-bid-optimization

[3] Yunli Wang et al., "Scaling Laws for Online Advertisement Retrieval," arXiv:2411.13322v1 [cs.IR], 2024. [Online]. Available: https://arxiv.org/html/2411.13322v1

[4] Subhash Vinnakota, "Optimizing Real-Time Bidding (RTB) Latency in Ad Exchanges: A Comprehensive Analysis," Journal of Computer Science and Technology Studies, 2025. [Online]. Available: https://www.researchgate.net/publication/394748259_Optimizing_Real-Time_Bidding_RTB_Latency_in_Ad_Exchanges_A_Comprehensive_Analysis

[5] Yuta Kaneko and Katsutoshi Yada, "A Deep Learning Approach for the Prediction of Retail Store Sales," Conference: 2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW), 2016. [Online]. Available: https://www.researchgate.net/publication/313454339_A_Deep_Learning_Approach_for_the_Prediction_of_Retail_Store_Sales

[6] Ananth Kamath, "Deconstructing Real-Time Bidding: A Lightweight Approach," Medium, 2025. [Online]. Available: https://medium.com/@ananthkamath98/deconstructing-real-time-bidding-alightweight-approach-3d4f727f1d7d

[7] Tahmina Rainy, "AI-Driven Marketing Analytics for Retail Strategy: A Systematic Review of Data-Backed Campaign Optimization," ResearchGate, 2025. [Online]. Available: https://www.researchgate.net/publication/391803018_AI-DRIVEN_MARKETING_ANALYTICS_FOR_RETAIL_STRATEGY_A_SYSTEMATIC_REVIEW_OF_DATA-BACKED_CAMPAIGN_OPTIMIZATION

[8] Siyu Zhang et al., "Dynamic Marketing Resource Allocation with Two-Stage Decisions," J. Theor. Appl. Electron. Commer. Res., 2022. [Online]. Available: https://www.mdpi.com/0718-1876/17/1/17

10.48047/jocaaa.2026.35.01.65

[9] Alison Hu, Gina Schaefer, and Sanmitra Bhattacharya, “Interpretable AI," Deloitte, 2025. [Online]. Available: https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificialintelligence/articles/ai-interpretability-challenge.html

[10] Ali Ben Mrad and Brahim Hnich, "Intelligent attribution modeling for enhanced digital marketing performance," Intelligent Systems with Applications, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2667305324000139