INTERNATIONAL REVIEW OF ARTIFICIAL INTELLIGENCE IN MARKETING

NAWALA

Volume 1, Number 2, 2022

Programmatic Advertising Optimization Using AI Bidding Strategies

Julius Gofinda Prasta[1*]

[1] Universitas Diponegoro, Semarang, Indonesia

Abstract

Article history:
Received: August 24, 2022
Revised: September 8, 2022
Accepted: October 13, 2022
Published: December 30, 2022

Keywords:
Artificial Intelligence, Bid Optimization, Programmatic Advertising, Real-Time Bidding, Return on Ad Spend.

Identifier:
Nawala
Page: 88-103
https://nawala.io/index.php/iraim

This article examines how artificial intelligence is used to optimize bidding strategies in programmatic advertising and real time bidding environments, focusing on the tension between short term performance gains and broader marketing and consumer outcomes. The study conducts a systematic literature review of peer reviewed journal articles published between 2016 and 2021, asking which AI and optimization methods are applied to bidding, which objectives and constraints they address, and what impacts they report. Across the reviewed studies, profit oriented models, reinforcement learning policies, control based pacing, and advanced click or conversion prediction generally outperform rule-based bidding on efficiency metrics such as clicks, conversions, and return on ad spend, but often under narrow objectives and proprietary data settings. The article discusses these results by grouping studies according to algorithmic approach and optimization focus, and by contrasting technical findings with evidence on media quality, privacy concerns, and brand outcomes. The main conclusion is that future work should integrate multi objective optimization and consumer centric evaluation into AI bidding research and practice.

*Corresponding author:
(Julius Gofinda Prasta)

1. Introduction

Programmatic advertising has become a central mechanism for automating the buying and selling of digital media, enabling advertisers to target individuals in real time across devices and platforms. Instead of negotiating placements in advance, impressions are evaluated and priced one by one as users load webpages or open apps. Decisions about whether to show an ad and how much to bid are made within milliseconds based on streams of user, context, and campaign data. This automation promises greater efficiency, precision, and accountability than traditional buying models, yet it also generates new concerns related to opacity, consumer perceptions, and the uneven distribution of value across the ecosystem (Samuel et al., 2021; Shehu et al., 2021). As a growing share of advertising budgets is transacted programmatically, understanding how to design and optimize these systems becomes a strategic priority for advertisers, agencies, and platforms (Wang et al., 2017; Yun et al., 2020).

At the heart of programmatic advertising is real-time bidding (RTB), where each ad impression is auctioned individually and advertisers submit bids conditional on predicted user response and campaign constraints. The advertiser's problem is to allocate a finite budget over a massive stream of heterogeneous opportunities while meeting key performance indicators such as clicks, conversions, or viewability. In practice, many campaigns still rely on heuristic rules, simple bid multipliers, and manual tuning, which are ill suited to the noisy, non-stationary nature of auction markets and can lead to under- or over-spending as conditions change. Control-theoretic work therefore reframes bidding and budget pacing as a dynamic feedback problem, treating bid levels and impression volumes as variables in a stochastic control system that must be stabilized around desired performance targets (Sang et al., 2018; Karlsson, 2020). These perspectives highlight that effective optimization requires not only accurate prediction but also robust decision policies that adapt to uncertainty in real time.

The rise of artificial intelligence has opened up new possibilities for tackling this complexity. Learning-based bidding frameworks seek to jointly model user response, market prices, and profit-oriented objective functions, allowing bids to adapt to evolving auction conditions and budget trajectories. Instead of optimizing intermediate metrics in isolation, such as click-through rate, these approaches can optimize directly for long-run profit or return on ad spend under budget constraints (Ren et al., 2017; Wang et al., 2017). AI-enabled strategies can absorb high-dimensional signals, capture non-linear relationships between context and value, and continuously update policies as new data arrive. At the same time, measurement and governance challenges in computational advertising, such as attribution bias, data sparsity, and the potential for intrusive targeting practices, raise questions about how far AI-driven optimization can proceed without jeopardizing user trust or regulatory compliance (Yun et al., 2020; Samuel et al., 2021).

From a marketing perspective, programmatic optimization cannot be evaluated only through short-term performance metrics. Evidence suggests that low-quality environments or poorly controlled placements may undermine brand outcomes even when auction-level metrics appear favorable (Shehu et al., 2021). Furthermore, aggressive bidding strategies that chase narrowly defined outcomes can inadvertently concentrate exposure on a small set of users, exacerbate frequency issues, or neglect qualitative aspects of ad experience that matter for brand equity. Recent analytical work on near-optimal bidding in large-scale auctions underscores the potential gains from more sophisticated strategies, but also points to trade-offs between efficiency, fairness, and transparency in how algorithms participate in markets (Tunuguntla & Hoban, 2021).

Despite rapid growth in both practice and research, knowledge on AI bidding in programmatic advertising remains fragmented across disciplines and methodological traditions. Studies differ in their optimization objectives, modeling choices, evaluation metrics, and treatment of constraints, making it difficult to compare findings or translate them into robust managerial guidance (Ren et al., 2017; Tunuguntla & Hoban, 2021). Some emphasize control and stability, others profit maximization or consumer response, and few integrate these dimensions holistically. This article therefore conducts a systematic literature review of peer-reviewed work published between 2016 and 2021 that examines programmatic advertising optimization using AI-based bidding strategies. By synthesizing evidence on algorithmic approaches, performance impacts, and documented limitations, the review aims to clarify the state of the art, identify conceptual and practical gaps, and outline an agenda for future research that aligns AI bidding strategies with advertiser value, platform objectives, and consumer welfare.

2. Literature Review

Research on programmatic advertising can be grouped into several complementary streams that together frame the role of AI-enabled bidding. One stream in marketing and information systems focuses on how programmatic infrastructures reshape value creation, governance, and consumer experience. Studies highlight that algorithmic media buying offers fine-grained targeting and efficiency gains, but also intensifies concerns about opacity, privacy, and perceived intrusiveness (Yun et al., 2020; Samuel et al., 2021). Building on this, empirical work shows that programmatic campaigns can heighten users' privacy concerns over time, and that perceived usefulness and transparency are key conditions for acceptance (Palos-Sanchez et al., 2019). In parallel, evidence indicates that the quality of the media environment moderates advertising effectiveness in programmatic contexts, with low-quality sites undermining outcomes even when auction-level performance indicators appear strong (Shehu et al., 2021).

A second stream examines the optimization of display campaigns as a computational and control problem. Foundational overviews of real-time bidding describe the RTB ecosystem, the role of behavioural and contextual signals, and the centrality of bid and budget optimization to campaign performance (Wang et al., 2017). Within this stream, feedback-control approaches model RTB as a dynamic system, using controllers to stabilize key performance indicators such as effective cost per click or impression volumes under budget constraints (Sang et al., 2018; Karlsson, 2020). Complementary work develops algorithmic frameworks that jointly model user response, market prices, and profit-based objectives, including profitmaximizing bidding machines and near-optimal bid strategies in large-scale auctions (Ren et al., 2017; Tunuguntla & Hoban, 2021). More recently, meta-heuristic and machine learning methods such as genetic algorithms and additional optimization layers have been proposed to tune bidding, budget allocation, and targeting across channels and demand-side platforms, showing performance improvements over baseline rules (Miralles-Pechuán et al., 2018; Micchi et al., 2020).

A third, emerging stream bridges these perspectives by linking bid optimization to broader campaign and consumer outcomes. Empirical work on programmatic campaigns documents that optimization objectives differ widely, ranging from clicks, conversions, and cost efficiency to brand-safe reach and user experience, and that these objectives may not always align with consumer welfare or long-term brand equity (Shehu et al., 2021; Samuel et al., 2021). While AI-based bidding and optimization layers can substantially improve intermediate performance metrics, existing studies often rely on proprietary data, heterogeneous evaluation protocols, and narrow outcome definitions, which limits comparability and generalization (Miralles-Pechuán et al., 2018; Micchi et al., 2020). The current literature therefore offers rich but fragmented evidence on AI bidding strategies, with gaps in how algorithmic innovations, media quality, and consumer responses are integrated. The present systematic review responds to this fragmentation by organizing prior work from 2016 to 2021 into coherent categories of AI bidding approaches, clarifying their objectives and constraints, and assessing the extent to which they address both advertiser optimization goals and user-centric concerns.

3. Methods

This study followed a systematic literature review approach to identify and synthesize peer-reviewed research on programmatic advertising optimization using AI-based bidding strategies between 2016 and 2021. The review was guided by a set of research questions that focused on: (1) what types of AI and optimization methods have been applied to bidding in programmatic and real-time bidding environments, (2) which campaign objectives and constraints these methods address, and (3) what performance outcomes and limitations are reported. A structured search was conducted in major scholarly databases, including Scopus, Google Scholar, ScienceDirect, IEEE Xplore, and other publisher platforms that index marketing, information systems, and computer science journals. Search strings combined terms related to programmatic advertising, real-time bidding, display advertising, bid optimization, reinforcement learning, machine learning, and artificial intelligence. Only articles published in English were considered.

To align with the scope of the review, inclusion criteria were restricted to empirical or methodological studies published in peer-reviewed journals, along with at most two working papers that met academic standards of rigor. Conference proceedings, book chapters, editorials, and practitioner reports were excluded. Screening proceeded in three stages: initial removal of duplicates, title and abstract screening against the inclusion criteria, and full-text assessment for relevance to AI-enabled bidding and optimization in programmatic contexts. For each included study, a structured data extraction template captured information on study context, ad format, data sources, algorithmic approach (for example reinforcement learning, deep learning, control-based optimization, genetic algorithms), optimization objectives and constraints, evaluation metrics, and key findings. These data were then synthesized narratively and thematically, with studies grouped according to dominant algorithmic approach and optimization focus, and compared to highlight convergences, divergences, and remaining gaps in the literature.

4. Results and Discussion

The systematic review indicates that research on programmatic advertising optimization using AI bidding strategies has largely focused on algorithmic improvements to decision making in real time bidding auctions, often evaluated through simulation or proprietary campaign datasets. Across the body of work, AI based bidding systems are typically framed as dynamic optimization problems that must jointly respect budget constraints and campaign goals such as conversions or profit. Profit oriented bidding models that explicitly maximize expected surplus per impression show consistent gains over rule based or linear pacing strategies, for example through direct profit optimization and dynamic adjustment of bid prices based on auction level feedback (Ren et al., 2017). Dynamic programming and stochastic control approaches further demonstrate that auto pricing strategies can deliver higher conversion rates at lower or similar spend compared with baseline heuristics by continuously adapting bids to changing competition and remaining budget (Adikari & Dutta, 2019). Complementing these contributions, model free reinforcement learning frameworks treat bidding as a sequential decision problem and report improvements in key performance indicators when policies are trained to respond to auction level states and budget trajectories rather than fixed bid rules (Sang et al., 2018; Cheng et al., 2019; Karlsson, 2020). At the campaign orchestration level, optimization layers that sit on top of multiple demand side platforms have been proposed to allocate budgets, set target bids, and correct under delivery, with iterative algorithms such as SKOTT showing superior performance compared with single platform or myopic allocation strategies (Micchi et al., 2020).

A second stream of results emphasizes the predictive layer that underpins AI bidding, especially click through and conversion rate estimation. Many bidding strategies assume accurate response prediction and then optimize bids as a function of predicted value, which makes the quality of these models critical. Studies that explicitly focus on click through rate prediction in the context of programmatic or real time bidding environments show that more expressive machine learning architectures tend to outperform linear baselines. For instance, integrating Weighted Extreme Learning Machines with Adaboost improves area under the curve relative to conventional Extreme Learning Machines and other benchmarks on large scale RTB datasets, thereby increasing the precision with which valuable impressions are identified (Zhang et al., 2017). Similarly, work using tree-based ensembles such as extreme gradient boosting demonstrates that well-tuned models can achieve competitive ROC AUC scores while using a reduced feature set, which is attractive for real time deployment in high volume ad exchanges (Moneera et al., 2021). Other studies combine advanced response models with bidding rules inside unified frameworks, showing that when click or conversion models are tightly coupled with bid price decisions, advertisers can attain higher return on ad spend or conversion volume under fixed budgets (Miralles-Pechuán et al., 2018; Liu et al., 2020; Tunuguntla & Hoban, 2021). These findings collectively support the view that bidding optimization and response prediction are mutually reinforcing capabilities. Improvements in either layer tend to propagate into better campaign level performance, but modeling errors or bias in the predictive layer may also be amplified by aggressive bidding strategies.

Beyond algorithmic efficiency, the review highlights a growing body of marketing and consumer focused research that problematizes the outcomes of AI optimized programmatic buying. One central concern is media quality and brand safety. Evidence from large field datasets shows that programmatic buying can inadvertently shift impressions toward lower quality or less reputable websites when bids are optimized purely on cost and short-term response metrics, which in turn can depress advertising effectiveness and harm brand equity (Shehu et al., 2021). At the same time, longitudinal survey-based research with large samples of internet users suggests that while programmatic advertising can increase perceived relevance and usefulness of ads over time, it also raises concerns about privacy, tracking, and manipulation, with these attitudes evolving as consumers gain more experience with targeted communications (Palos-Sanchez et al., 2019). Qualitative and conceptual work further illustrates that consumers experience programmatic systems as both helpful and unsettling. The pursuit of more granular personalization can heighten feelings of surveillance, and automation can increase the risk of inappropriate placements that damage trust and perceived legitimacy of advertising (Yun et al., 2020; Samuel et al., 2021). Related contributions in media and advertising journals extend these insights to emerging channels such as programmatic television. Conceptual analyses describe how the programmatic model, when transferred to television, promises finer audience segmentation and cross device measurement but also introduces complexities in data governance and transparency between buyers and sellers (Malthouse et al., 2018). Taken together, these findings suggest that although AI bidding strategies can meaningfully improve short term performance, they may also increase the strategic and ethical risk profile of programmatic campaigns if they are not embedded within broader policies that govern where and how ads are delivered.

Synthesizing across these literatures, the review reveals several important gaps and tensions that guide future research. First, there is a clear methodological divide between technical studies that foreground algorithmic performance and marketing or consumer behavior studies that foreground attitudes, brand outcomes, and media quality. Very few articles combine sophisticated AI bidding models with field or experimental designs that directly measure consumer level responses, brand effects, or cross channel interactions. Second, most optimization work still adopts single objective formulations that prioritize immediate indicators such as clicks, conversions, or profit, while website quality, user experience, and brand safety are treated, if at all, as exogenous constraints rather than integrated objectives. Third, transparency and interpretability remain underexplored, even though consumer work shows that perceptions of opacity and loss of control are central to concerns about programmatic advertising. From a managerial standpoint, the results suggest that advertisers and platforms should move from purely performance driven AI bidding to multi objective frameworks that jointly optimize economic outcomes, media quality, and consumer trust. From an academic standpoint, there is a need for interdisciplinary research that connects reinforcement learning and advanced predictive modeling with rich measurements of user experience and brand equity, so that AI enabled programmatic systems can be designed not only to win auctions efficiently but also to sustain long term, trust-based relationships between brands, publishers, and consumers.

5. Conclusion

This review shows that AI enabled bidding strategies can substantially improve the efficiency and effectiveness of programmatic advertising, but also that these gains come with important trade-offs. Across the studies examined, profit oriented bidding models, reinforcement learning policies, control-based pacing mechanisms, and advanced response prediction consistently outperform rule based or manually tuned strategies on short term indicators such as clicks, conversions, and return on ad spend. At the same time, marketing and consumer focused work highlights that such performance improvements are often achieved under narrow objectives, with limited consideration of media quality, brand safety, transparency, and user welfare. The overall picture is therefore one of technically sophisticated, but often single objective, optimization in an ecosystem where stakeholders increasingly care about long term trust, equity, and governance.

The review also reveals several limitations in the existing body of research and in this article's synthesis that constrain the strength and generalizability of conclusions. Many technical studies rely on proprietary datasets, opaque feature spaces, and heterogeneous evaluation protocols, which makes it difficult to compare results or judge how well proposed methods would transfer across markets, verticals, and regulatory environments. Consumer and brand oriented research, in turn, typically treats the algorithmic layer as a black box, which limits its ability to specify exactly how particular bidding or targeting choices shape attitudes, effectiveness, and perceived intrusiveness. Methodologically, this review is restricted to peer reviewed journal articles and a small number of working papers in English, published between 2016 and 2021, and it excludes conference proceedings and practitioner reports that might contain relevant technical innovations or field evidence. These boundaries may bias the sample toward more established methods and limit the visibility of very recent or highly applied work.

Taken together, the evidence suggests that future research should move beyond isolated technical or attitudinal perspectives toward genuinely interdisciplinary designs. Algorithmic studies would benefit from integrating user level and brand level outcomes into their objective functions and evaluation frameworks, for example by treating media quality, frequency control, and perceived intrusiveness as explicit constraints or co objectives rather than afterthoughts. Marketing and consumer research could, in parallel, work with more transparent descriptions or open implementations of AI bidding systems in order to test how specific design choices affect attitudes, trust, and long-term brand equity. For practitioners, the main implication is that AI based bidding should not be adopted as a purely performance driven tool, but as part of a broader governance framework that specifies where campaigns may run, how data are used, and how trade-offs between efficiency and user welfare are resolved. In closing, the promise of AI in programmatic advertising lies not only in winning auctions more effectively, but in enabling advertising systems that are economically robust, socially responsible, and sustainable over time.

References

Adikari, S., & Dutta, K. (2019). A new approach to real-time bidding in online advertisements: Auto pricing strategy. *INFORMS Journal on Computing, 31*(1), 66-82.

Cheng, Y., Zou, L., Zhuang, Z., Liu, J., Xu, B., & Zhang, W. (2019). An extensible approach for real time bidding with model free reinforcement learning. *Neurocomputing, 360*, 97–106.

Karlsson, N. (2020). Feedback control in programmatic advertising: The frontier of optimization in real-time bidding. *IEEE Control Systems Magazine, 40*(5), 40-77.

Liu, M., Yue, W., Qiu, L., & Li, J. (2020). An effective budget management framework for real-time bidding in online advertising. *IEEE access, 8*, 131107-131118.

Malthouse, E. C., Maslowska, E., & Franks, J. U. (2018). Understanding programmatic TV advertising. *International Journal of Advertising, 37*(5), 769-784.

Micchi, G., Soheily Khah, S., & Turner, J. (2020). A new optimization layer for real-time bidding advertising campaigns. *Intelligent Data Analysis, 24*(1), 199-224.

Miralles-Pechuán, L., Ponce, H., & Martínez-Villaseñor, L. (2018). A novel methodology for optimizing display advertising campaigns using genetic algorithms. *Electronic Commerce Research and Applications, 27*, 39-51.

Moneera, A., Maram, A., Azizah, A., AlOnizan, T., Alboqaytah, D., Aslam, N., & Khan, I. U. (2021). Click through rate effectiveness prediction on mobile ads using extreme gradient boosting. *Computers, Materials & Continua, 66*(2), 1681–1696.

Palos-Sanchez, P., Saura, J. R., & Martin-Velicia, F. (2019). A study of the effects of programmatic advertising on users' concerns about privacy overtime. *Journal of Business Research, 96*, 61-72.

Ren, K., Zhang, W., Chang, K., Rong, Y., Yu, Y., & Wang, J. (2017). Bidding machine: Learning to bid for directly optimizing profits in display advertising. *IEEE Transactions on Knowledge and Data Engineering, 30*(4), 645-659.

Samuel, A., White, G. R., Thomas, R., & Jones, P. (2021). Programmatic advertising: An exegesis of consumer concerns. *Computers in Human Behavior, 116*, 106657.

Sang, Q., Karlsson, N., & Guo, J. (2018). Feedback control of event rate in online advertising campaigns. *Control Engineering Practice, 75*, 126-136.

Shehu, E., Abou Nabout, N., & Clement, M. (2021). The risk of programmatic advertising: Effects of website quality on advertising effectiveness. *International Journal of Research in Marketing, 38*(3), 663-677.

Tunuguntla, S., & Hoban, P. R. (2021). A near-optimal bidding strategy for real-time display advertising auctions. *Journal of Marketing Research, 58*(1), 1-21.

Yun, J. T., Segijn, C. M., Pearson, S., Malthouse, E. C., Konstan, J. A., & Shankar, V. (2020). Challenges and future directions of computational advertising measurement systems. *Journal of Advertising, 49*(4), 446–458.

Wang, J., Zhang, W., & Yuan, S. (2017). Display advertising with real-time bidding (RTB) and behavioural targeting. *Foundations and Trends in Information Retrieval, 11*(4–5), 297–435.

Zhang, S., Fu, Q., & Xiao, W. (2017). Advertisement Click-Through Rate Prediction Based on the Weighted-ELM and Adaboost Algorithm. *Scientific Programming, 2017*, 1–8.