INTERNATIONAL REVIEW OF ARTIFICIAL INTELLIGENCE IN MARKETING

NAWALA

Volume 1, Number 1, 2022

Sentiment Aware Programmatic Advertising: Integrating NLP into Real Time Bidding Systems

Rika Rahmawati [1]

[1] Universitas Pamulang, Banten, Indonesia

Abstract

Article history:
Received: January 14, 2022
Revised: February 11, 2022
Accepted: April 20, 2022
Published: June 30, 2022

Keywords:
Affective Data, Click Through Prediction, Programmatic Advertising, Real Time Bidding, Sentiment Analysis.

Identifier:
Nawala
Page: 43-58
https://nawala.io/index.php/iraim

This article examines how sentiment aware programmatic advertising can be designed by integrating natural language processing into real time bidding systems for digital display advertising. Building on a systematic review of research on programmatic auctions, user response prediction, and sentiment based marketing analytics, the article synthesizes evidence that current architectures rely heavily on behavioural histories and coarse contextual features while largely ignoring fine grained affective information in page content and user generated text. The review identifies a persistent gap between advances in sentiment analysis, which provide accurate polarity and emotion signals, and real time bidding practice, where such signals are rarely embedded directly in bidding logic. In response, the article proposes a conceptual framework in which sentiment and other affective text features enter user response models and bidding functions as core inputs rather than auxiliary indicators. The findings highlight implications for advertising relevance, brand safety, computational efficiency, ethical governance of affective data, industry adoption, and rigorous future empirical evaluation efforts.

*Corresponding author:
(Rika Rahmawati)

1. Introduction

Programmatic advertising has become the dominant infrastructure of digital display advertising, automating the buying and selling of impressions through auction mechanisms that operate at millisecond timescales. Real time bidding enables advertisers to compete for individual ad impressions based on granular user and contextual data, promising higher efficiency, finer targeting, and more accountable return on investment compared with traditional media buying models (Sayedi, 2018; Palos-Sanchez et al., 2019). As budgets continue to shift toward programmatic channels, advertisers and platforms face growing pressure to improve the relevance of each impression while managing auction dynamics, bid prices, and campaign constraints in highly volatile environments.

Within this ecosystem, the core technical challenge is user response prediction and bid optimization: demand-side platforms must estimate the likelihood of clicks, conversions, or post-view actions and translate these predictions into optimal bids under budget and frequency constraints (Gharibshah & Zhu, 2021). Recent research in computational and personalized advertising documents substantial progress using machine learning and deep learning architectures, combining user histories, device attributes, and contextual signals to improve click-through rate and conversion prediction (Viktoratos et al., 2021). However, many of these models still treat textual content and affective cues in a relatively shallow way, relying on coarse keyword or topic features rather than richer representations of sentiment and emotion that could capture how users feel about brands, products, and surrounding content at the moment of impression delivery.

In parallel, advances in natural language processing and sentiment analysis have transformed the ability of marketers to infer attitudes and emotions from large volumes of unstructured text. Surveys of sentiment analysis techniques highlight increasingly sophisticated methods for extracting polarity, emotion, and aspect level opinions from social media posts, reviews, and other user generated content, moving beyond simple lexicon-based approaches toward machine learning and deep learning models (Kaur et al., 2017). Empirical applications show that sentiment signals can enhance brand monitoring, campaign evaluation, and demand forecasting by revealing how consumers react to messages and events in near real time (Ahmed et al., 2021). These developments suggest that sentiment-aware features could provide valuable additional information for estimating user engagement and optimizing advertising decisions.

Despite these converging trends, the integration of sentiment-aware natural language processing into real time bidding pipelines remains underexplored. Existing work on programmatic and RTB optimization generally focuses on auction mechanics, budget pacing, and structural properties of bidding strategies, while treating page or feed content primarily as static contextual variables rather than as emotionally charged stimuli that shape users' propensity to attend to and act on ads (Sayedi, 2018; Viktoratos et al., 2021). At the same time, most sentiment analysis applications in marketing operate in offline or batch settings, generating dashboards and periodic insights rather than feeding low latency signals directly into algorithmic bidding and creative selection.

The present study responds to this gap by exploring how affective text features and advanced language processing techniques can be systematically embedded into the decision making components of real time bidding within programmatic advertising environments. Specifically, it proposes an architecture in which sentiment and other affective text features extracted from page content, user generated messages, and historical feedback are incorporated into user response prediction models and bidding functions. By linking advances in sentiment analysis with the operational constraints of RTB, the article aims to demonstrate how sentiment-aware programmatic advertising can improve relevance, effectiveness, and brand safety, while also raising new questions about computational complexity, latency, and ethical use of affective data in automated advertising environments.

2. Literature Review

Programmatic advertising and real time bidding (RTB) have been conceptualized as a core layer of computational advertising, where each impression is sold via instantaneous auctions and optimized using large-scale data and algorithmic bidding strategies. Wang et al. (2017) describe RTB ecosystems as multi sided markets in which demand side platforms (DSPs) must simultaneously manage targeting, bidding, and budget allocation across heterogeneous inventory, emphasizing the role of predictive models and feedback loops in sustaining performance. Building on this view, Cai et al. (2017) show that reinforcement learning can be used to learn bidding policies directly from auction streams, treating the advertiser’s objective as a sequential decision problem under budget constraints.

More recent work proposes additional optimization layers on top of DSPs: Micchi et al. (2020) develop an algorithm that reallocates budget and tunes bids across multiple DSPs to maximize key performance indicators, demonstrating that meta-optimization can significantly improve campaign efficiency over baseline strategies. Complementary studies on machine learning enabled advertising document growing reliance on data driven models for segmentation, targeting, and performance optimization across programmatic channels.

A second stream of research focuses on user response prediction and click through rate (CTR) modeling, which provide the main input to bidding decisions in RTB environments. Deep learning-based CTR models such as the Deep Interest Network (DIN) highlight the importance of capturing users' diverse and evolving interests from behavioral histories rather than relying on static profiles, leading to substantial gains in prediction accuracy in large-scale display advertising platforms. Later work advances this line by explicitly addressing the challenge of long behavior sequences: Qin et al. (2020) propose the UBR4CTR framework, which retrieves the most relevant past behaviors for each impression using a learnable search mechanism, thereby improving both performance and inference efficiency. Representation Learning-Assisted CTR models further extend feature learning by jointly modeling user-ad, ad-ad, and feature CTR relationships to mitigate data sparsity and enhance generalization. In parallel, Liu et al. (2021) design a joint learning model that combines residual networks and attention mechanisms to automatically explore high order feature interactions, reporting consistent improvements over traditional feed forward architectures in display advertising datasets. Although these approaches substantially advance CTR prediction, most of them encode contextual text and page content only through sparse IDs or simple embeddings, with limited attention to nuanced affective information.

In contrast, the sentiment analysis literature has evolved rapidly toward deep learning and context aware representations that could, in principle, enrich advertising models. Surveys by Etaiwi et al. (2021) and Nandwani and Verma (2021) document how convolutional and recurrent neural networks, as well as attention based architectures, now dominate sentiment classification tasks, enabling more accurate polarity and emotion detection from unstructured text. Drus and Khalid (2019) show that these methods are particularly effective for mining opinions from social media streams, highlighting their value for real-time monitoring of consumer reactions. At the same time, applications in marketing illustrate that sentiment signals carry predictive power for engagement and diffusion: Kulkarni et al. (2020), for example, use sentiment analysis of user responses to viral video advertisements to segment “ad sharers” and demonstrate that sentiment-based typologies explain sharing intentions better than traditional thought-listing approaches. Collectively, this body of work suggests that sentiment and emotion features can capture aspects of user state that are not reflected in purely behavioral metrics such as past clicks.

Only a small set of studies, however, explicitly connect sentiment-aware natural language processing to bidding and placement decisions in online advertising. The APNEA system proposed by Bushi and Zaïane (2019) integrates sentiment analysis into a contextual advertising and ad-bidding framework by ranking advertisers according to the sentiment polarity of page content toward their brands, aiming to avoid negative placements while preserving relevance. Their experiments show that sentiment-aware ad ranking can reduce harmful pairings of brands with negatively valenced content without significantly increasing computational overhead, indicating that affective features can be operationalized at RTB compatible latencies. Nevertheless, existing implementations typically focus on contextual relevance and brand safety rather than on end to end integration of sentiment features into CTR or conversion prediction pipelines. Moreover, most sentiment-driven approaches are evaluated on limited datasets or offline testbeds rather than under full budget, pacing, and latency constraints characteristic of industrial RTB systems.

Overall, the literature reveals three important gaps. First, while RTB and CTR research has developed sophisticated architectures for modeling user behavior and feature interactions, it largely neglects fine-grained sentiment and emotion extracted from page content and user generated text at impression time. Second, sentiment analysis and marketing applications show that affective cues can enhance prediction of engagement and sharing, yet these insights are rarely translated into real-time bidding logic. Third, the few systems that combine sentiment and ad bidding primarily treat sentiment as a filter for unsafe or incongruent content rather than as a continuous signal that can directly influence bid values, creative selection, and pacing decisions. Addressing these gaps, the present study positions sentiment aware programmatic advertising as a synthesis of CTR modeling and deep sentiment analysis, proposing a framework in which affective text features are integrated into user response prediction and bidding strategies under realistic RTB constraints.

3. Methods

This study employs a systematic literature review method to synthesize existing research on sentiment-aware programmatic advertising, natural language processing, and real-time bidding systems. The review is structured around a clearly defined research question on how affective text features and sentiment analysis can be integrated into user response prediction and bidding strategies within programmatic ecosystems. Relevant publications are identified through comprehensive searches in major scholarly databases using combinations of keywords such as "programmatic advertising," "real-time bidding," "click-through rate prediction," "sentiment analysis," "affective computing," and "natural language processing." The search process is complemented by backward and forward citation tracking to capture additional studies that may not be retrieved through keyword queries alone. Inclusion criteria focus on peer-reviewed journal articles and conference papers that address computational models, system architectures, or empirical evaluations related to programmatic advertising, RTB optimization, or sentiment-based decision-making, while studies that only discuss general digital marketing or sentiment analysis without an advertising or bidding context are excluded.

Titles, abstracts, and full texts are screened in multiple stages by applying these criteria, and the methodological quality of the selected studies is assessed based on clarity of research design, transparency of data and model specification, and completeness of reporting. The final set of studies is then coded along several dimensions, including type of advertising setting, modeling approach, use of textual and affective features, integration with bidding or targeting mechanisms, and evaluation metrics. Thematic and comparative analysis is used to map existing approaches, identify dominant design patterns and limitations, and derive a conceptual framework for integrating sentiment aware NLP into real time bidding pipelines.

4. Results and Discussion

The systematic review shows that most programmatic advertising and RTB studies converge on a common view of the ecosystem as a data-intensive, auction-based marketplace in which predictive modeling is central to value creation. Research on RTB infrastructures highlights how demand side platforms must simultaneously solve targeting, bidding, and budget allocation across heterogeneous inventory, with feedback mechanisms used to refine decision rules over time (Cai et al., 2017; Wang et al., 2017). In parallel, optimization layers that operate across multiple DSPs demonstrate that algorithmic budget reallocation and bid tuning can significantly improve campaign key performance indicators relative to baseline strategies (Micchi et al., 2020), reinforcing the idea that better models and richer signals directly translate into superior advertising performance. These findings support the introduction's argument that user response prediction and bid optimization are the core technical challenges in programmatic environments (Gharibshah & Zhu, 2021).

Across the user response prediction literature, the review confirms a strong shift toward deep learning architectures that mine increasingly complex behavioral patterns. CTR models based on interest networks and behavior retrieval show that exploiting long user histories and high order feature interactions yields substantial gains in prediction accuracy over traditional approaches (Qin et al., 2020; Liu et al., 2021). This trend aligns with earlier observations that machine learning and deep learning architectures can combine user histories, device attributes, and contextual signals to enhance click through and conversion prediction (Viktoratos et al., 2021). However, when these models are examined in detail, contextual text is often reduced to sparse identifiers or simple embeddings, and the emotional tone of page content or user generated messages is rarely modeled explicitly. Consequently, the "context" that enters CTR models tends to capture where and when an impression occurs, but not how positively or negatively users may feel about the content surrounding the ad.

The sentiment analysis and marketing literature reviewed offers a sharp contrast to this underuse of affective information. Surveys on sentiment analysis document that convolutional, recurrent, and attention-based models have markedly improved polarity and emotion detection from unstructured text, enabling fine grained analysis of opinions and attitudes in social media and other user generated content (Etaiwi et al., 2021; Nandwani & Verma, 2021). Empirical applications show that such sentiment signals are predictive of engagement and diffusion: for instance, sentiment based typologies of "ad sharers" explain sharing intentions for viral video advertisements better than traditional cognitive measures (Kulkarni et al., 2020), while real-time monitoring of social media sentiment has been shown to capture shifts in consumer reactions that are not visible in aggregate behavioral indicators alone (Drus & Khalid, 2019; Ahmed et al., 2021). Taken together, these findings support the proposition in the introduction that sentiment aware features can provide additional information for estimating user engagement and advertising effectiveness.

Only a limited subset of studies directly integrate sentiment into advertising decision mechanisms, and these implementations are generally narrow in scope. The APNEA system, for example, uses sentiment polarity of page content toward brands to adjust ad ranking and avoid negative placements, demonstrating that affective features can be processed at RTB compatible latencies without excessive computational cost (Bushi & Zaïane, 2019). This line of work resonates with concerns about brand safety and contextual relevance in programmatic environments (Sayedi, 2018; Palos-Sanchez et al., 2019), but it still treats sentiment primarily as a binary or categorical filter rather than a continuous signal that shapes bid values or creative selection. The juxtaposition of these findings with the broader CTR and sentiment literatures reveals a clear gap: while RTB research has developed sophisticated architectures for behavioral modeling, and sentiment research has matured into highly accurate affective analytics, systematic integration of deep sentiment features into real-time bidding pipelines remains largely unexplored.

Overall, the review indicates that sentiment aware programmatic advertising is a promising but underdeveloped research frontier. The evidence suggests that embedding affective text features into user response models could improve prediction quality in situations where user emotions and contextual tone are strong determinants of engagement, especially for formats such as video, native, and social ads. At the same time, the scarcity of end-to-end evaluations under realistic budget, pacing, and latency constraints underscores the need for architectures that jointly address predictive performance, computational efficiency, and ethical considerations around the use of affective data. By synthesizing insights from RTB optimization, CTR modeling, and sentiment analysis, the proposed framework in this article positions sentiment aware NLP not as an auxiliary dashboard tool, but as a core input to real time bidding decisions, with implications for relevance, effectiveness, and brand safety in programmatic advertising ecosystems.

5. Conclusion

This study set out to clarify how sentiment-aware natural language processing can be integrated into real time bidding systems within programmatic advertising. The synthesis of prior work confirms that RTB has evolved into a highly data intensive, auction based marketplace where user response prediction and bid optimization are central to value creation. Deep learning based CTR models and multi level optimization of bidding strategies have significantly improved the efficiency of campaign management, yet they still rely predominantly on behavioral histories and coarse contextual features, leaving affective information largely untapped. At the same time, advances in sentiment analysis show that modern NLP techniques can extract fine grained polarity and emotion signals from unstructured text, and that these signals are strongly associated with engagement, sharing behavior, and shifts in consumer reactions.

The systematic literature review highlights a persistent disconnect between these two streams. Most RTB and CTR studies either ignore textual sentiment altogether or encode context in ways that do not capture how users feel about the content surrounding an impression. Conversely, sentiment analysis and marketing applications typically operate in offline, dashboard oriented settings, without direct links to bidding logic, budget constraints, or latency requirements. Only a small number of systems demonstrate that sentiment signals can be processed at RTB compatible speeds, and these are mainly focused on brand safety and contextual avoidance rather than on full integration of sentiment into user response models and bid formation. This confirms that sentiment aware programmatic advertising is still at an early stage of conceptual and technical development.

Building on these insights, the article proposes that affective text features should be treated as a core input to real time bidding decisions rather than as an auxiliary monitoring tool. Integrating sentiment aware NLP into CTR prediction and bidding functions has the potential to improve relevance, perceived quality, and brand safety, particularly in environments where user emotions and contextual tone are key drivers of response. At the same time, the review underscores important challenges for future research, including the need to balance predictive gains with computational efficiency, to evaluate sentiment enhanced architectures under realistic budget and pacing constraints, and to address ethical questions around the collection and use of affective data. Addressing these issues can help move programmatic advertising toward systems that are not only more effective and efficient, but also more sensitive to the emotional context in which ads are delivered.

References

Bushi, S. S., & Zaiane, O. (2019, October). Apnea: Intelligent ad-bidding using sentiment analysis. In *IEEE/WIC/ACM International Conference on Web Intelligence* (pp. 76–83).

Cai, H., Ren, K., Zhang, W., Malialis, K., Wang, J., Yu, Y., & Guo, D. (2017, February). Real-time bidding by reinforcement learning in display advertising. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining* (pp. 661–670).

Drus, Z., & Khalid, H. (2019). Sentiment analysis in social media and its application: Systematic literature review. *Procedia Computer Science, 161*, 707–714.

Etaiwi, W., Suleiman, D., & Awajan, A. (2021). Deep learning based techniques for sentiment analysis: A survey. *Informatica, 45*(7).

Gharibshah, Z., & Zhu, X. (2021). User response prediction in online advertising. *ACM Computing Surveys, 54*(3), 1–43.

Kaur, H., & Mangat, V. (2017, February). A survey of sentiment analysis techniques. In *2017 International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud) (I-SMAC)* (pp. 921–925). IEEE.

Kulkarni, K. K., Kalro, A. D., Sharma, D., & Sharma, P. (2020). A typology of viral ad sharers using sentiment analysis. *Journal of Retailing and Consumer Services, 53*, 101739.

Liu, M., Cai, S., Lai, Z., Qiu, L., Hu, Z., & Ding, Y. (2021). A joint learning model for click-through prediction in display advertising. *Neurocomputing, 445*, 206–219.

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.

Nandwani, P., & Verma, R. (2021). A review on sentiment analysis and emotion detection from text. *Social Network Analysis and Mining, 11*(1), 81.

Ouyang, W., Zhang, X., Ren, S., Qi, C., Liu, Z., & Du, Y. (2019). Representation learning-assisted click-through rate prediction. *arXiv preprint arXiv:1906.04365*.

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.

Qin, J., Zhang, W., Wu, X., Jin, J., Fang, Y., & Yu, Y. (2020, July). User behavior retrieval for click-through rate prediction. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval* (pp. 2347–2356).

Sayedi, A. (2018). Real-time bidding in online display advertising. *Marketing Science, 37*(4), 553–568.

Viktoratos, I., & Tsadiras, A. (2021). Personalized advertising computational techniques: A systematic literature review, findings, and a design framework. *Information, 12*(11), 480.

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.

Widmar, N. O., Bir, C., Clifford, M., & Slipchenko, N. (2020). Social media sentiment as an additional performance measure? Examples from iconic theme park destinations. *Journal of Retailing and Consumer Services, 56*, 102157.

Zhou, G., Zhu, X., Song, C., Fan, Y., Zhu, H., Ma, X., … & Gai, K. (2018, July). Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1059–1068).