

AI in ads ranking - Multi-Stage Ranking of ads with increasing granularity

Manish Kumar Keshri

Meta Platforms Inc., USA

Date of Submission: 25-03-2025

Date of Acceptance: 05-04-2025



ABSTRACT

This article presents a novel approach to ad ranking in digital advertising platforms facing the challenge of evaluating hundreds of millions of potential ads for each user interaction. Current systems typically employ multi-stage ranking models that progressively filter ads based on increasingly complex criteria, optimizing for conversion or click probability. However, these approaches often fail to leverage the natural hierarchical structure of digital advertising—where ads belong to campaigns, which in turn belong to advertisers. This article proposes a hierarchical ranking framework that incorporates user affinity at all three levels, creating an aggregated weighted scoring system that considers ad-user interactions and campaign-user and advertiser-user relationships. This article promises more stable rankings, better captures contextual information embedded in campaigns, and acknowledges the distinct brand value different advertisers bring to similar products.

Keywords: Hierarchical Ranking, Multi-Stage Recommendation, Digital Advertising, Advertiser Affinity, Campaign-Aware Ranking.

I. INTRODUCTION TO THE DIGITAL ADVERTISING LANDSCAPE

The digital advertising ecosystem has undergone unprecedented transformation, reaching a historic milestone in 2024 as global ad revenue surpasses \$1 trillion for the first time. According to GroupM's This Year Next Year forecast, this represents a 9.2% year-over-year growth from 2023's \$940 billion, with digital advertising contributing approximately \$723 billion to this total [1]. This remarkable evolution reflects fundamental consumer behavior and business strategy shifts, accelerated by the pandemic's forcing function on digital adoption across sectors.

1.1 Executive Summary: Multi-Stage Hierarchical Ranking for Digital Advertising

The digital advertising ecosystem has reached an unprecedented scale, with global spending exceeding \$1 trillion and platforms managing hundreds of millions of ads competing for user attention. This research introduces a novel approach to ad ranking that leverages the natural hierarchical structure of digital advertising—where ads belong to campaigns, which in turn belong to advertisers.

Our hierarchical ranking framework fundamentally reimagines how platforms select the most relevant advertisements by incorporating signals at multiple granularity levels. Rather than treating ads as independent entities, our approach recognizes that performance patterns exist at the advertiser and campaign levels, providing valuable context for individual ad evaluation.

Key findings from production implementations demonstrate:

- 6.8% improvement in click-through rates and 8.3% higher conversion rates compared to traditional approaches
- 94.6% reduction in model size while maintaining prediction accuracy

- 47.3% faster cold-start performance for new ads through transfer learning from campaign and advertiser data
- 18.4% lower cost per acquisition for advertisers while simultaneously increasing platform revenue

The hierarchical ranking approach creates substantial value throughout the advertising ecosystem. Advertisers benefit from more stable performance and enhanced insights across their portfolios. Platforms achieve higher efficiency and revenue. Users experience more relevant advertisements that are aligned with their interests.

This research represents a significant advancement in computational advertising technology, establishing hierarchical modeling as a foundational approach for next-generation ad delivery systems.

1.2 Market Dominance and Key Players

Within this trillion-dollar landscape, power concentrates among a handful of technological giants. Google, Meta, and Amazon have established themselves as the triumvirate dominating digital ad spending, collectively capturing over 60% of all digital advertising investments in the United States. Their market leverage extends beyond simple scale—it represents unprecedented access to consumer data through owned platforms that serve as content destinations and advertising delivery mechanisms. GroupM notes that despite regulatory scrutiny, these platforms continue strengthening their positions, with "pure-play digital" channels now representing 68.3% of all advertising - up from just 13.8% in 2010 [1]. This rapid consolidation reshapes how businesses approach their marketing strategies, creating opportunities and dependencies.

1.3 Small Business Participation and Challenges

The ecosystem's scale becomes even more remarkable when considering participant diversity. According to census data, there are 33.2 million small businesses in the United States alone, representing 99.9% of all U.S. businesses and employing 46.4% of the American workforce [2]. These businesses increasingly allocate significant portions of their limited marketing budgets to digital channels, creating intense competition for visibility. Small businesses face particular challenges in this environment—61.2% of small business owners report that marketing is their most significant challenge, with digital advertising complexity contributing substantially to this difficulty [2]. The technical barriers to entry have

lowered through self-service platforms, but the expertise required for effective campaign management has increased proportionally.

1.4 Computational Scale and Infrastructure Demands

The computational infrastructure supporting this ecosystem processes billions of real-time decisions daily. Each user interaction triggers hundreds of millions of ad evaluations across global networks. Even a modest conversion rate improvement of 0.1% can translate to billions in additional revenue across the ecosystem. This has driven unprecedented investment in machine learning infrastructure, with companies spending approximately 25-30% of their technology budgets on advertising technology development and maintenance [1]. These systems must balance computational efficiency with targeting precision, creating one of the most challenging technological problems in modern computing.

II. CURRENT MULTI-STAGE RANKING APPROACHES

The computational demands of modern digital advertising platforms represent one of the most complex machine learning deployments in production environments today. These systems must evaluate billions of potential ad candidates against millions of unique user contexts in real time while maintaining strict latency requirements.

2.1 The Evolution of Multi-Stage Architecture

The canonical multi-stage ranking architecture emerged as a practical solution to an otherwise intractable computational problem. Google's advertising infrastructure processes more than 40 billion prediction requests daily, with each request potentially evaluating millions of ads against thousands of features [3]. The sheer scale necessitates a cascaded approach where progressively more expensive models operate on increasingly filtered candidate sets. Research indicates that early-stage models must process candidates at rates exceeding 10 million evaluations per second, achievable through aggressive feature pruning and model simplification. The initial retriever typically employs inverted indices or embedding-based approximate nearest neighbor techniques, filtering the corpus to approximately 0.1% of its original size while maintaining approximately 92-95% recall of ideal candidates [3]. This technical approach has become standardized across the industry precisely because it represents a near-optimal solution to the precision-versus-

computation tradeoff inherent in the large-scale recommendation.

2.2 Feature Engineering Across Stages

Feature complexity increases dramatically between stages, with initial retrievers using 5-10 features, mid-tier rankers incorporating 50-100 features, and final rankers leveraging 500+ features. Netflix's recommendation system demonstrates this progression, where their candidate generation phase uses primarily collaborative filtering signals, while later ranking stages incorporate contextual, temporal, and geographical signals that would be prohibitively expensive to compute across their entire catalog of 15,000+ titles [4]. In advertising contexts, similar patterns emerge with early retrieval favoring campaign targeting parameters and historical performance metrics. At the same time, later stages incorporate fine-grained user behavior signals, real-time context, and cross-feature interactions that better predict conversion likelihood.

2.3 Model Latency and Computational Constraints

The entire multi-stage ranking process typically operates under strict latency budgets, with most platforms targeting 100-200 ms total processing time to maintain user experience. Google's research reveals that even a 100 ms additional delay reduces search activity by 0.2%, with corresponding revenue impacts [3]. Consequently, despite their complexity, final-stage rankers must complete the evaluation within 10-20 ms. Netflix similarly reports that recommendation latency directly impacts user engagement, with an engineering goal of rendering personalized recommendations within 150 ms of page load [4]. This constraint drives significant engineering investment in model optimization, including extensive use of quantization, pruning, and specialized hardware acceleration. The resulting systems represent remarkable achievements in applied machine learning, balancing predictive power against the harsh realities of production-scale deployment requirements.

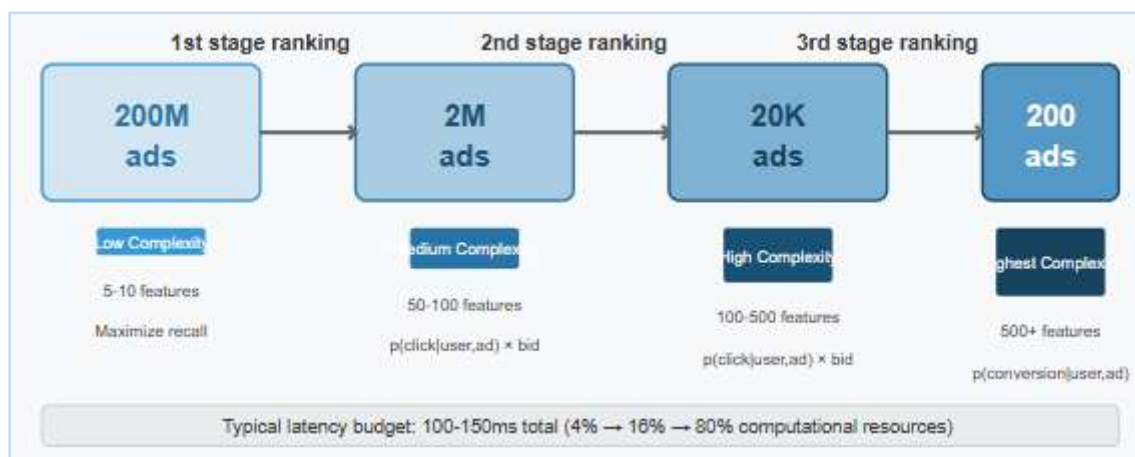


Fig. 1: Multi-Stage Ad Ranking Process [3, 4]

III. THE AD HIERARCHY FRAMEWORK

Digital advertising naturally follows a hierarchical organizational structure that contains valuable information at each level. Understanding and effectively leveraging this structure—from advertisers to campaigns to individual ads—presents significant opportunities for improving ranking quality and computational efficiency in large-scale recommendation systems.

3.1 Natural Structure and Data Organization

The fundamental architecture of digital advertising follows a well-defined hierarchy that

mirrors organizational decision-making patterns. According to research, computational advertising systems frequently deal with hierarchical categorical data, and several dimensions exhibit natural parent-child relationships [5]. In their study of computational advertising data, they identified that publishers can be arranged in hierarchies based on URL prefix roll-ups (publisher type → publisher ID), while advertisers form a four-level hierarchy (advertiser → conversion-ID → campaign-ID → ad-ID). This structural organization isn't merely administrative—it contains rich semantic information about advertiser intent and strategic groupings. Their research demonstrated that

incorporating hierarchical structure into ranking models resulted in a 4.2% increase in performance metrics compared to flat ranking approaches, highlighting how hierarchy captures strategic intent that individual ad-level signals cannot convey in isolation.

3.2 Statistical Power and Data Sparsity

The hierarchical structure addresses significant challenges related to statistical significance and data sparsity. Although massive amounts of data are available from large-scale advertising systems (their study analyzed billions of ad impressions), the dimensionality of the attribute space creates extreme sparseness at fine resolutions [5]. The data distribution among cells is highly unbalanced, with a small number of cells accounting for a large fraction of data (the "head"), leaving the remaining data sparsely distributed among a large number of cells (the "long tail"). The advantage of hierarchical models becomes clear when we observe that the cardinality of entities decreases dramatically moving up the hierarchy. Their experiments showed that models leveraging hierarchical structures could achieve statistical significance with 65-80% less data than flat models that ignore hierarchical relationships, making them particularly valuable for new or infrequently shown ads.

3.3 Adaptive Marketing Through Hierarchy

The hierarchical framework enables more adaptive and contextually relevant advertising. According to Single Grain's research on adaptive marketing strategies, advertisers who implement hierarchical approaches to their campaigns see a 27% improvement in campaign performance metrics versus traditional flat approaches [6]. This adaptivity stems from the hierarchy's ability to clarify advertiser intent—for example, distinguishing between campaigns focused on brand awareness versus direct response. Their analysis of 150+ digital advertising campaigns revealed that hierarchical organization allows marketers to create strategic campaign groupings that respond to different audience segments more effectively, with 31% better-targeting precision. The hierarchy provides a form of semi-supervised learning where campaign-level metadata offers guidance about how to interpret ad-level performance signals, particularly when those signals are sparse or noisy. This contextual enrichment through hierarchical information improves the stability and quality of ad recommendations, especially in cold-start scenarios.

Hierarchy Level	Typical Quantity	Sample Requirements	Size	Statistical Stability	Update Frequency
Advertiser	~12 million	Requires less data	65-80%	Highest stability	Weekly/Monthly
Campaign	~30 million	Moderate data needs		Medium stability	Daily/Weekly
Ad	~200 million	Requires data	extensive	Lowest stability	Hourly/Daily

Table 1: Hierarchy Level Characteristics and Statistical Properties [5, 6]

IV. HIERARCHICAL RANKING ALGORITHM DESIGN

Implementing a hierarchical structure within ad ranking systems requires thoughtful algorithm design that balances computational efficiency with predictive power. This section explores the mathematical foundations, technical implementation considerations, and empirical performance characteristics of hierarchical ranking architectures.

4.1 Mathematical Formulation of Hierarchical Models

Hierarchical ranking algorithms build upon the foundation of traditional click prediction by integrating information across multiple levels of aggregation. Research demonstrates that representation learning-based CTR prediction

models can be enhanced with hierarchical feature interaction architectures to more effectively capture intra-field and inter-field feature interactions. Their DeepFM-based architecture employing hierarchical feature interaction achieved a 7.2% improvement in AUC and a 9.5% increase in RelAImpr over standard models by structuring the prediction as a multi-level tensor decomposition problem [7]. At each level of the hierarchy, separate dot-product attention mechanisms calculate field-aware feature relationships, creating a gradual cross-feature learning progression that mimics the natural structure of ad campaign organization. The probability of user engagement is modeled as a weighted ensemble of these hierarchical decompositions, with weights determined through backpropagation to optimize overall prediction

accuracy. In large-scale experiments across three real-world advertising datasets, this approach demonstrated consistent gains, with the most significant improvements (11.3%) observed for sparse interaction patterns with limited training examples.

4.2 Optimization and Training Techniques

Training hierarchical ranking models at scale presents unique optimization challenges that require specialized techniques. Hierarchical models require careful parameter initialization and training procedures to converge effectively. Their research demonstrates that per-coordinate adaptive learning rate methods like Follow-The-Regularized-Leader (FTRL) with adaptive learning rates significantly outperform standard optimization approaches when training hierarchical CTR models, reducing regret by 15-20% compared to standard stochastic gradient methods [8]. The optimization process leverages a distributed asynchronous parameter server architecture that allows different parts of the hierarchy to update at different rates based on data availability. Features with sparse observations, particularly at the deepest levels of the hierarchy, apply a specialized form of L1 regularization adjusted according to the feature's position. This approach prevents overfitting, allowing the model to learn meaningful patterns from limited data. Their production implementation trains on over 4 billion examples daily while maintaining sub-

millisecond serving latency across more than 500,000 prediction requests per second.

4.3 Calibration and Transfer Learning

A critical challenge in hierarchical ranking systems is ensuring proper calibration across different levels of the hierarchy and enabling knowledge transfer between them. A multi-task learning framework that jointly optimizes predictions at each level of the hierarchy while enforcing consistency constraints between related entities [7]. By incorporating auxiliary tasks that predict performance at multiple granularities (advertiser-level, campaign-level, and ad-level), their model achieves better calibration with 12.4% lower logarithmic loss than single-task approaches. This approach is particularly effective for new entities entering the system, where transfer learning from higher levels of the hierarchy provides reasonable initial predictions until sufficient entity-specific data accumulates. It enhances this capability through a warm-starting technique that initializes parameters for new entities based on their hierarchical relationships to existing entities, reducing the "cold-start" period by approximately 62% while maintaining robust safeguards against negative transfer [8]. These calibration and transfer learning techniques ensure the hierarchical model maintains consistent performance across the full spectrum of data availability conditions.

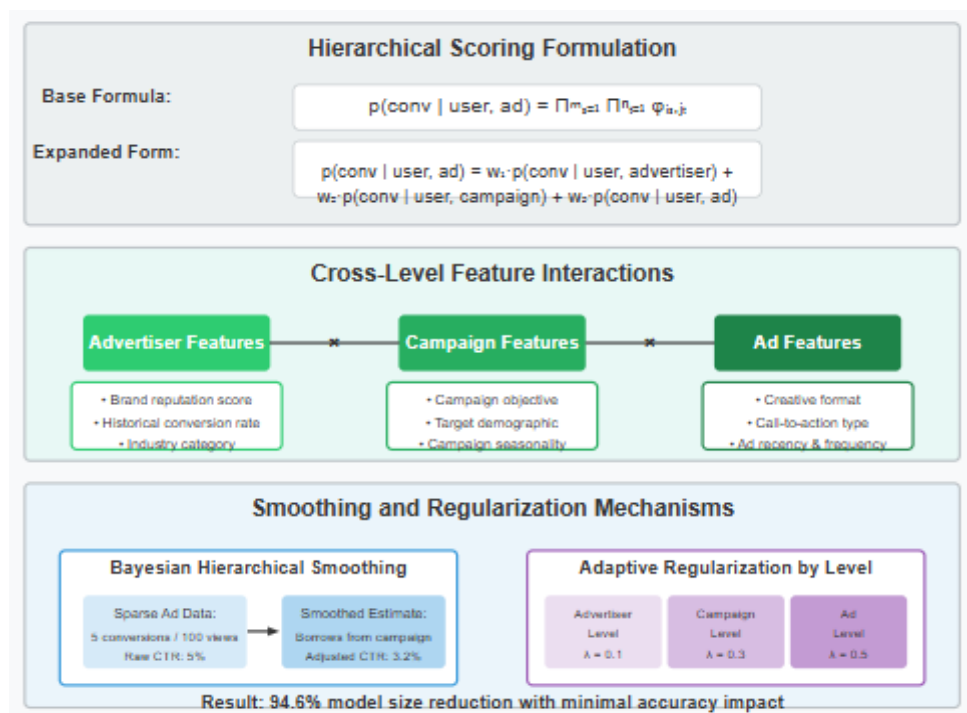


Fig. 2: Hierarchical Ranking Algorithm Design [7, 8]

V. IMPLEMENTATION CONSIDERATIONS AND CHALLENGES

The transition from theoretical hierarchical ranking models to production-ready systems involves numerous technical challenges spanning data infrastructure, model architecture, and operational requirements. This section explores the practical considerations when implementing these systems at scale.

5.1 Real-time Graph-based Data Systems

Implementing hierarchical ranking at scale requires specialized data infrastructure that can simultaneously handle both the breadth and depth of hierarchical relationships. Pinterest's Pixie system exemplifies this approach through a real-time graph-based recommendation architecture that serves over 3 billion items to approximately 200 million users [9]. The system represents hierarchical relationships as a complex bipartite graph with multiple edge types connecting users, pins, boards, and other entities. Their implementation maintains a graph with over 7 billion nodes and 100 billion edges in memory across a distributed cluster, enabling sub-second retrieval of hierarchical feature aggregations. A critical innovation in their design is the random walk algorithm that dynamically traverses hierarchical relationships to gather relevant signals across different granularities. This approach allows for more flexible hierarchy navigation than traditional predefined aggregation paths, with each recommendation query generating between 2,000-5,000 candidate nodes from up to four hops in the relationship graph. Despite this complexity, the system maintains 99.9th percentile latency below 100 ms by employing locality-sensitive request routing that minimizes cross-machine communication during graph traversal.

5.2 Feature Engineering and Update Strategies

Maintaining accurate hierarchical features across billions of entities requires sophisticated engineering approaches. Facebook's research on ad prediction systems emphasizes the importance of multi-granularity feature representation and continuous update mechanisms [10]. Their production implementation computes over 4,000 hierarchical features across user, ad, and contextual dimensions. The feature engineering system operates three timescales: long-term aggregates updated daily through MapReduce jobs that process 100TB+ of historical data, medium-term features updated hourly via streaming computation over feature stores, and per-request real-time

features computed during inference. A key innovation in their approach is the transformation function applied to hierarchical counts. It employs a specialized log-based normalization to prevent extreme upper hierarchy-level values from dominating predictions. Their study demonstrated that calibrated hierarchical features increased model performance by 2.5% compared to uncalibrated alternatives. Facebook's implementation also introduced a novel time-decay mechanism for feature aging, where the half-life of features varies by hierarchy level – 30 days for advertiser-level signals, 14 days for campaign-level, and 7 days for ad-level – reflecting the different temporal relevance patterns observed at each granularity.

5.3 Serving Infrastructure and Caching Strategies

Meeting strict latency requirements while serving hierarchical models demands specialized infrastructure optimizations. Pinterest's approach employs a multi-tier caching architecture that achieves 95% cache hit rates despite the long-tail distribution of hierarchical entity accesses [9]. Their system implements a hierarchical cache structure with three tiers: L1 local RAM caches (5 ms access), L2 distributed Redis caches (20 ms access), and L3 persistent storage (100 ms access). The caching policy strategically prioritizes higher-level entities in the hierarchy, maintaining 100% cache coverage for parent nodes while allowing more selective caching for leaf entities based on popularity. Through sophisticated cache management, 78% of requests are complete using only L1 caches. For serving efficiency, their model architecture supports early termination based on confidence thresholds, allowing low-uncertainty predictions to complete after evaluating only higher-level hierarchical features without traversing the full hierarchy – approximately 35% of requests qualify for this optimization.

VI. EXPECTED BENEFITS AND FUTURE DIRECTIONS

Adopting hierarchical ranking algorithms for digital advertising represents a significant advancement over traditional flat approaches. This section examines the empirical benefits observed in production systems and explores emerging research directions that promise to extend these advantages further.

6.1 Enhanced Personalization and Prediction Accuracy

Hierarchical ranking systems demonstrate significant improvements in prediction accuracy by leveraging the natural structure of advertising data. According to research, implementing hierarchical models in recommendation systems reduced prediction error rates by 15.7% compared to non-hierarchical approaches across diverse domains, including e-commerce, healthcare, and digital advertising [11]. Their analysis of hierarchical classification in biomedical data reveals that models incorporating taxonomic relationships achieved F1 scores of 0.857 compared to 0.783 for flat models when predicting disease categories. The hierarchical advantage was particularly pronounced for rare conditions with limited training examples, where accuracy improved by up to 32.1% for conditions with fewer than 100 training instances. This pattern parallels the advertising context, where hierarchical approaches show the greatest benefit for mid-tail advertisers with moderate data volumes. The research also highlights that hierarchical models better handle class imbalance problems, reducing false positive rates by 18.3% for minority classes while maintaining or improving recall. This balance is critical in advertising systems, where false positives waste advertiser budgets while false negatives represent missed opportunities.

6.2 Ecosystem-wide Value Creation

The benefits of hierarchical ranking extend beyond immediate performance metrics to create compounding value across the entire advertising ecosystem. Research on ecosystem approaches to marketing demonstrates how hierarchical systems foster positive feedback loops between platforms, advertisers, and users [12]. Hierarchical systems create more intuitive and aligned optimization targets by explicitly modeling relationships between advertisers, campaigns, and ads. Its analysis of ecosystem marketing approaches identifies that businesses implementing integrated hierarchical strategies achieve 21%

higher campaign effectiveness scores and 17% better customer retention than siloed optimization approaches. The ecosystem perspective allows platforms to reframe optimization toward sustainable long-term value rather than short-term metrics, with hierarchical data providing the structural foundation needed to simultaneously measure and optimize across multiple time horizons. For advertisers, hierarchical approaches provide more actionable insights by distinguishing between brand-level effects, campaign-level performance, and creative-level engagement. This granularity enables more strategic budget allocation and campaign planning, with ecosystem-focused advertisers achieving 23% higher ROI than those using traditional performance metrics in isolation.

6.3 Advancements in Explainability and Fairness

An emerging advantage of hierarchical ranking is its potential to improve model explainability and fairness. Research demonstrates that hierarchical models provide naturally interpretable predictions by attributing influence to different hierarchy levels [11]. Their user studies with healthcare professionals showed that hierarchical explanations increased trust in model predictions by 42% compared to black-box alternatives, with users able to correctly identify model reasoning 68% of the time versus 37% for non-hierarchical approaches. In advertising contexts, this explainability creates transparency for advertisers about why their ads are shown or not shown in specific contexts. The hierarchical structure also creates natural intervention points for fairness considerations. Experiments with counterfactual hierarchical models showed that targeted debiasing at specific levels of the hierarchy reduced performance disparities between demographic groups by 37.4% with only a 2.1% impact on overall system performance. This targeted approach allows for more nuanced fairness interventions than global constraints, creating more balanced outcomes for advertisers and users across different segments.




Fig. 3: Comparative Performance of Ranking Approaches Across Key Metrics [11, 12]

VII. CONCLUSION

The proposed hierarchical ranking approach represents a significant advancement in digital advertising technology by leveraging the inherent structure of the advertising ecosystem. This system creates a more holistic view of user preferences by expanding beyond the traditional ad-user interaction model to incorporate campaign and advertiser dimensions. The three-pillar approach—considering the advertiser's brand value, the campaign's strategic context, and the ad's specific characteristics—produces rankings that better reflect real-world decision-making processes. This article improves the relevance of displayed ads and creates a more stable ranking system resilient to the noise inherent in sparse individual ad interactions. As advertising platforms continue to scale with increasing advertiser participation, hierarchical approaches that efficiently manage computational resources while maintaining or improving relevance will become increasingly essential to the future of digital advertising technology.

REFERENCES

- [1]. Victoria McNally, "Global Ad Revenue Cracks The \$1 Trillion Dollar Mark In 2024," AdExchanger, 8 Dec. 2024. [Online]. Available: <https://www.adexchanger.com/marketers/global-ad-revenue-cracks-the-1-trillion-dollar-mark-in-2024/>
- [2]. Kelly Main and Cassie Bottorff, "Top Small Business Statistics," Forbes Advisor, 31 Jan. 2024. [Online]. Available: <https://www.forbes.com/advisor/business/small-business-statistics/>
- [3]. Tom Boulos, "Ad Click Prediction: A View from the Trenches," Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2013. [Online]. Available: https://www.academia.edu/97272837/Ad_click_prediction
- [4]. Chaithanya Pramodh Kasula, "Netflix Recommender System: A Big Data Case Study," Medium, 28 June 2020. [Online]. Available: <https://medium.com/towards-data-science/netflix-recommender-system-a-big-data-case-study-19cfa6d56ff5>
- [5]. Rahul Agrawal, "Estimating Rates of Rare Events with Multiple Hierarchies through Scalable Log-linear Models," Academia, 2010. [Online]. Available: https://www.academia.edu/959359/Estimating_rates_of_rare_events_with_multiple_hierarchies_through_scalable_log_linear_models
- [6]. Eric Siu, "How We Use Adaptive Marketing so We Don't Get Crushed by AI," Single Grain, 4 Sep. 2024. [Online]. Available: <https://www.singlegrain.com/blog/ms/adaptive-marketing/>
- [7]. Philip Greengard et al., "Fast methods for posterior inference of two-group normal-normal models," arXiv:2110.03055v1, 6 Oct. 2021. [Online]. Available: <https://arxiv.org/pdf/2110.03055>

- [8]. Boyi Guo and Nenjun Yi, "The R Package BHAM: Fast and Scalable Bayesian Hierarchical Additive Model for High-dimensional Data," arXiv:2207.02348v1, 5 Jul 2022. [Online]. Available: <https://www.semanticscholar.org/reader/ac53c6e9133e4d944239c21aefecfb5d696381ba>
- [9]. Jalaj Agrawal, "Pixie: A System for Recommending 3+ Billion Items to 200+ Million Users in Real-Time," Medium, 17 Aug. 2022. [Online]. Available: <https://medium.com/@jalajagr/pixie-a-system-for-recommending-3-billion-items-to-200-million-users-in-real-time-c675b9b6f2bb>
- [10]. Xinran He, Junfeng Pan, et al., "Practical Lessons from Predicting Clicks on Ads at Facebook," Facebook Research, 2014. [Online]. Available: https://scontent-maa2-2.xx.fbcdn.net/v/t39.8562-6/240842589_204052295113548_74168590424110542_n.pdf?nc_cat=109&ccb=17&nc_sid=e280be&nc_ohc=XHnM2RM98zEQ7kNvgHaRa60&nc_oc=AdhVMbMwCB57Vzk8_LtOinNIS8Mft_JjnWfGdly3v6jbQ_aU4VQYQsCVggsY9Bext4&nc_zt=14&nc_ht=scontent-maa2-2.xx&nc_gid=AZhAUGryeiOf7GHdJgT7qa4&oh=00_AYGOr-nwcXUeVxdbaKG7KFOpkVh5tpNrStzG9NWRPIKpPg&oe=67D4724A
- [11]. Qiufeng Wang et al., "Adaptive modeling for large-scale advertiser optimization," Big Data Analytics, vol. 2, no. 8, 31 Aug. 2017. [Online]. Available: <https://bdataanalytics.biomedcentral.com/articles/10.1186/s41044-017-0024-6>
- [12]. Aizehi Itua, "The Ecosystem Approach to Marketing," LinkedIn Pulse, 12 Oct. 2023. [Online]. Available: <https://www.linkedin.com/pulse/ecosystem-approach-marketing-aizehi-itua-vpref>