

Design and Evaluation of Whole-Page Experience Optimization for E-commerce Search

Pratik Lahiri
Amazon
Seattle, WA, USA
lahirip@amazon.com

Aditya Jumde
Amazon
Seattle, WA, USA
adijumde@amazon.com

Yi Liu
Amazon
Seattle, WA, USA
yiam@amazon.com

Bingqing Ge
Amazon
Seattle, WA, USA
bqge@amazon.com

Shuning Huo
Amazon
Seattle, WA, USA
shuningh@amazon.com

Mahmoud Mamlouk
Amazon
Seattle, WA, USA
mmamlk@amazon.com

Zhou Qin
Amazon
Seattle, WA, USA
qinzq@amazon.com

Lucas Scottini*
Roblox Corporation
San Francisco, CA, USA
lcostascottini@gmail.com

Wenyang Liu
Amazon
Seattle, WA, USA
lwenyang@amazon.com

Abstract

E-commerce Search Results Pages (SRPs) are evolving from linear lists to complex, non-linear layouts, rendering traditional position-biased ranking models insufficient. Moreover, existing optimization frameworks typically maximize short-term signals (e.g., clicks, same-day revenue) because long-term satisfaction metrics (e.g., expected two-week revenue) involve delayed feedback and challenging long-horizon credit attribution. To bridge these gaps, we propose a novel Whole-Page Experience Optimization Framework. Unlike traditional list-wise rankers, our approach explicitly models the interplay between item relevance, 2D positional layout, and visual elements. We use a causal framework to develop metrics for measuring long-term user satisfaction based on quasi-experimental data. We validate our approach through industry-scale A/B testing, where the model demonstrated a 1.86% improvement in brand relevance (our primary customer experience metric) while simultaneously achieving a statistically significant revenue uplift of +0.05%.

CCS Concepts

• **Information systems** → **Content ranking**; **Presentation of retrieval results**; **Retrieval effectiveness**.

Keywords

whole page optimization, search, multi-objective recommendation, causal impact

ACM Reference Format:

Pratik Lahiri, Bingqing Ge, Zhou Qin, Aditya Jumde, Shuning Huo, Lucas Scottini, Yi Liu, Mahmoud Mamlouk, and Wenyang Liu. 2026. Design and Evaluation of Whole-Page Experience Optimization for E-commerce Search.

*Work done at Amazon



This work is licensed under a Creative Commons Attribution 4.0 International License. WSDM '26, Boise, ID, USA

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2292-9/2026/02
<https://doi.org/10.1145/3773966.3779374>

In *Proceedings of the Nineteenth ACM International Conference on Web Search and Data Mining (WSDM '26)*, February 22–26, 2026, Boise, ID, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3773966.3779374>

1 Introduction

E-commerce Search Results Pages (SRPs) have evolved from simple ranked lists into dynamic, visually complex layouts [12, 20]. As shown in Figure 1, for the query "XYZ water bottle," modern SRPs interleave organic search results (displayed in a grid) with themed widgets (e.g., "Trending Now"), each possessing distinct visual styles. This shift fundamentally alters user interaction: visual design drives non-sequential attention patterns that differ significantly from traditional top-to-bottom list scanning [18]. Consequently, without changing the underlying item set, varying the placement of widgets creates distinct page layouts and thus distinct user experiences.

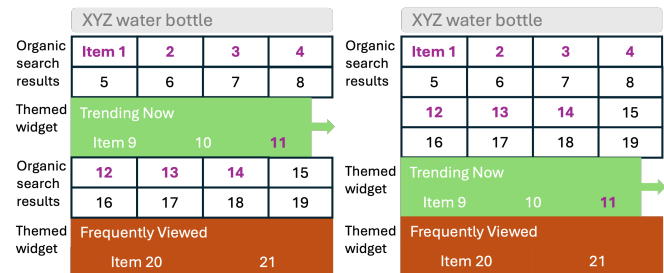


Figure 1: Different search results page layouts. Brand XYZ items are shown in bold purple.

To manage the trade-off between exploring new layouts and exploiting known ones, bandit algorithms are widely used [10, 11, 14–17]. Yet, applying slate optimization methods [6, 7, 10, 19] to this domain reveals four gaps: (1) **Multi-Objective Nature**: The system must balance engagement, revenue, and long-term user satisfaction; and (2) **Delayed Reward Signals**: success metrics, such as shopping frequency and long-term spend, are not observable on

the same day; (3) **Complex Satisfaction Functions**: User satisfaction in 2D grids depends on the interplay of content and position, not just a simple discount factor; (4) **Heterogeneous Content**: Perception varies between organic search results and widgets.

We formulate SRP optimization as a contextual bandit problem, where the action is the presented SRP and contextual features are used to characterize content heterogeneity. To balance between multiple objectives, we propose a composite scalar reward defined as a weighted sum of page non-abandonment, short-term revenue, and a designed long-term impact based user satisfaction metric. Most prior work rewards SRP optimization using clicks [3, 16, 22] or short-term revenue [24, 25]. Some studies consider user satisfaction [2, 5], all measured via human annotation which suffers from scalability issues, annotator noise, and potential mismatch between annotation and real user preference [9, 13, 21, 23]. We hypothesize that presenting high-quality SRPs generates positive long-term effects independent of immediate clicks or purchases, such as increased return visits and shopping frequency, ultimately leading to a higher long-term spend. We therefore use an estimate of expected long-term spend as a metric to quantify user satisfaction. Our contributions are threefold:

- (1) We propose a causal framework that uses quasi-experimental variation to link whole-page quality to long-term customer satisfaction, measured 12-week spend.
- (2) We model position-dependent user sensitivity by partitioning the SRP into regions, and estimating region-specific user satisfaction, rather than assuming a continuous decay.
- (3) We validate our method through online A/B testing with three treatments: no user satisfaction in the reward, click-through rate as a proxy, and our proposed satisfaction metric, which achieves the best performance.

2 Methodology

2.1 DV-WPX: A Causal Inference Framework

We propose DV-WPX (Downstream Value of Whole Page Experience), a causal framework that addresses **delayed reward signals gap** by linking observable whole-page quality to long-term customer satisfaction via quasi-experimental variation. The model's identification strategy rests on the assumption that, conditional on historical customer features, variations in quality metrics across search events within a query (defined as keyword, search index alias, and ranking function combination), event date, and delivery locations are as good as experimental. This assumption holds exactly when variations stem from experimental assignments and approximately when variations arise from natural experiments like supply side shocks that are conditionally independent from unobserved customer characteristics driving downstream revenue (rev). The model approximates a setting where two equivalent shoppers (A and B) issue identical search queries on the same day from the same delivery locations, but experience different levels of search quality. This difference could emerge from experimental assignments or natural variations in search results. The model then estimates the difference in 12-week revenue between these shoppers after accounting for differences in short-term outcomes. Quality metrics (except pricing and duplication) are analyzed across distinct page

regions which can be adjusted (like a hyperparameter). This regional analysis strategy accounts for the varying impact of quality metrics based on their visibility and position in search results. The DV-WPX model's theoretical foundation rests on a structural equation that links downstream customer engagement to search quality and short-term behavior. For a search event e , the relationship is formalized as:

$$rev_{t(e)+\delta_L} = W(rev_{t(e)+\delta_S}(Q_e), A_{t(e)+\delta_S}(Q_e), Q_e)$$

where $rev_{t(e)+\delta_L}$ is cumulative long-term revenue (12 weeks post-search), $rev_{t(e)+\delta_S}$ is short-term revenue (2 weeks post-search), $A_{t(e)+\delta_S}$ denotes engagement metrics, and Q_e is the vector of quality metrics. The welfare function $W(\cdot)$ maps short-term signals to long-term outcomes.

The causal effect of quality changes is captured through the total derivative.

$$\frac{\partial rev_{t(e)+\delta_L}}{\partial Q_e} = \frac{\partial W}{\partial rev_{t(e)+\delta_S}} \frac{\partial rev_{t(e)+\delta_S}}{\partial Q_e} + \frac{\partial W}{\partial A_{t(e)+\delta_S}} \frac{\partial A_{t(e)+\delta_S}}{\partial Q_e} + \frac{\partial W}{\partial Q_e}$$

This decomposition reveals three distinct channels through which search quality affects downstream revenue. The first term captures how quality changes influence downstream revenue through their immediate impact on short-term sales. The second term represents the indirect effect through short-term engagement metrics. The final term measures the direct effect of quality on long-term outcomes, independent of its short-term impacts.

To estimate these effects empirically, the model employs a linear specification.

$$Drev_{i,c,q,z}^t = \sum_{s=1}^S \beta_s^t X_{i,c,q,z}^{t,s} + \sum_{j=1}^J \theta_j^t M_{i,c}^j + \sum_{k=1}^K \gamma_k^t H_{i,c}^k + \alpha_q^t + \zeta_z^t + \epsilon_{i,c,q,z}^t$$

with corresponding surrogate equations:

$$X_{i,c,q,z}^{t,s} = \sum_{j=1}^J \phi_j^{t,s} M_{i,c}^j + \sum_{k=1}^K \tau_k^{t,s} H_{i,c}^k + \alpha_q^{t,s} + \zeta_z^{t,s} + \epsilon_{i,c,q,z}^{t,s}$$

Here, i indexes search events, c represents customers, q denotes query characteristics, and z indicates delivery ZIP codes. The β coefficients capture the causal effects of quality metrics (X), while θ and γ represent the effects of short-term metrics (M) and historical controls (H), respectively. Fixed effects α_q and ζ_z account for query-specific and geographical variations.

The estimation employs Double Machine Learning (DML) [4] to address potential confounding and ensure robust inference. The DML process consists of three carefully designed stages. First, the 0th stage implements an iterative de-averaging algorithm that removes fixed effects across queries and ZIP codes. This process requires 20 iterations to ensure convergence of residuals and downstream estimates. The data is then split into a 90 percent training fold and 10 percent testing fold to enable out-of-sample validation. The first stage of DML focuses on obtaining unbiased residuals of the de-averaged versions of both target and surrogate metrics. This is accomplished through linear ML models with two-fold cross-fitting, a technique that prevents overfitting by ensuring that the residuals for each observation are computed using models trained on different subsets of the data. The feature set includes comprehensive historical metrics and auxiliary controls. The second stage represents the core of the causal inference, implementing either

OLS or LASSO regression of the residualized target metric on residualized surrogates. For LASSO estimation, the model employs sophisticated hyperparameter tuning through a 20-point grid search with 3-fold cross-validation. This ensures improved model selection while maintaining computational feasibility. The final DV-WPX metric is computed as:

$$DVWPX_i^{t,E} = \sum_{s=1}^S \beta_s^{t,E} X_i^{t,s}$$

This formulation aggregates the estimated causal effects (β_s^t) of individual quality metrics ($X_{i,c,q,z}$) into a single, interpretable measure of search quality's impact on long-term customer engagement.

2.2 User Satisfaction Metric Derived from DV-WPX

The DV-WPX framework provides a general mechanism for translating observable whole-page quality signals into long-term user satisfaction. While agnostic to the specific quality signal, it enables the construction of satisfaction metrics whose components are weighted by their downstream impact rather than short-term interactions. Here, we instantiate DV-WPX using brand alignment on brand-sensitive queries as a concrete example to evaluate long-term satisfaction optimization in an online ranking system.

We design the Pixel and Region Weighted Whole-Page Brand Match Rate (PR-WP-BMR), which measures how well page content aligns with a branded query (e.g., “Nike shoes”). The metric aggregates brand match rates across all items on the page, weighting each item by visual prominence (i.e., pixel coverage) and coarse page position (i.e., page region).

To capture visual effects, we partition the SRP into three regions, regardless of whether items appear as standalone results or within widgets: Top (positions 1–8), Middle (positions 9–16), and Bottom (positions beyond 16). PR-WP-BMR is computed as a weighted sum of region-level, pixel-weighted brand match rates:

$$\text{PR-WP-BMR} = w_{\text{Top}} \cdot \overline{\text{BMR}}_{\text{Top}} + w_{\text{Mid}} \cdot \overline{\text{BMR}}_{\text{Mid}} + w_{\text{Bot}} \cdot \overline{\text{BMR}}_{\text{Bot}}$$

where $w_{\text{Top}} + w_{\text{Mid}} + w_{\text{Bot}} = 1$ are weights for different regions.

A key design choice is how to set the region weights $\{w_r\}$. As a baseline, we derive weights from the empirical distribution of click-through rates (CTR) across page positions, reflecting short-term engagement patterns. Alternatively, we derive region weights using DV-WPX, which estimates the marginal downstream value of an incremental brand match in each region. These effects are normalized to obtain relative weights that capture the long-term importance of brand alignment at different page locations. Empirically, DV-WPX assigns most weight to the Top and Middle regions, with negligible weight on the Bottom. Specifically, we use CTR-based weights (0.60, 0.25, 0.15) and DV-WPX-based weights (0.63, 0.37, 0) for the Top, Middle, and Bottom regions, respectively.

By combining pixel-based visual prominence with region-level weighting, PR-WP-BMR models satisfaction as a function of both layout and content, addressing **complex satisfaction functions** beyond a simple 1D position discount.

2.3 Page Template Ranker

To evaluate the DV-WPX-derived user satisfaction metric online, we integrate it into a production page template ranker.

2.3.1 Problem Formulation. In an industrial page template recommender, business constraints restrict how content can be displayed. There are C possible ways to interleave themed widgets with search results, but not all arrangements are eligible due to business rules and UI constraints. We define a page template for each eligible way to rank content—two examples are shown in Figure 1. Let $P = \{P_1, \dots, P_k\}$ denote the set of page templates, where each template P_i admits a specific subset of eligible items $C_i \subseteq C$.

The action space is $A = \{a_1, \dots, a_k\}$, where each action a_i corresponds to selecting a page template and ordering its eligible items. Given an objective function $F : A \rightarrow \mathbb{R}$ that predicts the reward of an action, the ranker selects the action that maximizes F . In practice, this corresponds to choosing, for each search request, a page template from the pool of eligible templates P .

Based on business needs, our problem is to learn a decomposition of F that optimizes for multiple objectives, including long-term delayed rewards.

2.3.2 Input Features. The ranker uses a rich set of input signals, referred to as the 3Cs: Context, Customer, and Content. Context features include, but are not limited to, marketplace, device type, inferred query specificity, and product categories at multiple granularities. Customer features capture status-related signals such as membership. Content features describe the relevance and value of each rankable, including aggregated measures of relevance, brand alignment, product-type alignment, and other value-related signals. These features are computed in a content-type-aware manner (e.g., separate aggregation and calibration for search result vs. widgets), allowing the model to capture different user responses to **heterogeneous content**.

2.3.3 Modeling. The ranker performs multi-objective optimization over business outcomes (i.e., revenue), engagement, and whole-page satisfaction. Engagement is measured by binary page non-abandonment, while whole-page satisfaction is captured through the metric defined in the previous section (i.e., PR-WP-BMR).

We train a separate predictive model for each objective. Continuous objectives use Bayesian linear regression [1], while binary objectives use Bayesian probit regression [8]. Training uses historical impression logs, with inputs consisting of the 3Cs features and the displayed template, and targets corresponding to each objective. Linear regression models minimize root mean squared error, while probit models maximize area under the ROC curve. Models are refreshed daily using incremental training, sampling 50% of data from the most recent day to adapt to evolving traffic and content.

2.3.4 Inference and Decision Process. At inference time, the ranker evaluates all candidate templates and applies Thompson sampling to generate predictions for each objective. For each template and context, samples are drawn from the posterior distributions of the corresponding objective models. The sampled predictions are combined into a single reward score using a weighted linear combination of objectives. The weights are determined offline based on the historical statistics (e.g., mean and variance) of each objective to

normalize their scales and ensure comparable contribution to the final score. The template with the highest aggregated score is selected for display. This scalar reward provides a practical mechanism and addresses the **multi-objective nature** of template selection.

3 Results

We evaluate the impact of incorporating user satisfaction into the page template ranker through offline and online experiments. We consider three settings: a Control without any satisfaction signal, Treatment 1 (T1) using a CTR-based PR-WP-BMR proxy, and Treatment 2 (T2) using a DV-WPX-based PR-WP-BMR to optimize long-term satisfaction. In both treatments, models also incorporate content-aware features capturing query–widget alignment in relevance, brand, and product type.

We present offline results to evaluate content-aware features improve predictions on objectives shared across models. Since PR-WP-BMR is absent from the Control, offline comparisons focus on common revenue and engagement metrics. Because T1 and T2 share the same models for these objectives and differ in the satisfaction metric, their offline results are identical and reported jointly.

3.1 Offline Data and Results

In both the Control and Treatment settings, separate models are trained to predict revenue and non-abandonment. Offline training and evaluation use one day of North America traffic, comprising approximately 1.5 million Desktop impressions and 5.6 million Mobile impressions, with a 99% / 1% train–test split.

Model performance on the held-out test set is reported in Table 1. We use root mean squared error (RMSE; lower is better) for continuous outcomes and area under the ROC curve (AUC; higher is better) for binary outcomes; positive percentage changes indicate improvement. Non-abandonment is included as an optimization objective only for Desktop in our MOO formulation, so we report it only for Desktop.

The results show that adding content-aware features improves revenue prediction on both Mobile and Desktop, and slightly improves non-abandonment prediction on Desktop.

Table 1: Offline performance comparison between baseline and treatment models.

Dev.	Metric	Eval.	Base.	Treat.	$\Delta\%$
Mob	Revenue	RMSE	47.70	44.90	+6
Desk	Non-Aband.	AUC	0.39	0.39	+1
Desk	Revenue	RMSE	8.87	8.32	+6

3.2 Online Test Results

We also evaluate the changes in a one-month worldwide A/B test on both Mobile and Desktop. Table 2 reports the rolled-up online results, comparing each treatment variant against the Control.

Incorporating a user satisfaction signal into the page template ranker improves customer experience and business outcomes regardless of how region weights are defined. Both T1 (CTR-based PR-WP-BMR) and T2 (DV-WPX-based PR-WP-BMR) deliver positive short-term revenue gains of 0.04% and 0.05%, respectively.

Notably, using DV-WPX-derived weights (T2) yields additional benefits beyond short-term performance. While T1 shows a slight decline in long-term revenue, T2 achieves a positive lift, indicating that DV-WPX better captures signals aligned with long-term customer satisfaction and value.

Although DV-WPX weights are not derived from click behavior, both T1 and T2 improve engagement, as reflected by identical gains in search CTR. This suggests that optimizing for whole-page satisfaction does not trade off short-term engagement, and that DV-WPX-based weighting can improve long-term outcomes without sacrificing immediate user interaction.

Table 2: Online performance comparison. All reported lifts are statistically significant.

Metric	Relative Lift (T1–C)	Relative Lift (T2–C)
Revenue	0.04%	0.05%
Long-term Revenue	-0.001%	0.005%
Search CTR	0.02%	0.02%

4 Discussion

This work studies how to incorporate long-term user satisfaction into whole-page optimization for e-commerce search. We introduce DV-WPX, a causal framework that maps whole-page quality to downstream value, and instantiate it with a brand satisfaction metric (PR-WP-BMR). We integrate this metric into a production page template ranker and evaluate it offline and online. Results show that satisfaction-aware optimization improves performance, with DV-WPX-based weighting providing additional long-term benefits beyond short-term engagement proxies.

While the results are encouraging, several directions remain for future work. First, although we focus on brand relevance as a concrete instantiation, the DV-WPX framework can be extended to other whole-page quality dimensions such as visual diversity or cross-component coherence. Second, the current DV-WPX model relies on a 12-week observation window, which may be long for rapidly evolving retail settings; exploring shorter or adaptive horizons is an important next step. Finally, while we use fixed positional regions in this work, future extensions could consider dynamic region definitions that adapt to device form factors and user interaction patterns. Despite these opportunities, our approach demonstrates a practical and scalable solution that has been successfully deployed in a real-world e-commerce system.

5 Ethical Considerations

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- [1] Shipra Agrawal and Navin Goyal. 2014. Thompson Sampling for Contextual Bandits with Linear Payoffs. arXiv:1209.3352 [cs.LG] <https://arxiv.org/abs/1209.3352>
- [2] Peter Bailey, Nick Craswell, Ryen W. White, Liwei Chen, Ashwin Satyanarayanan, and S. M.M. Tahaghoghi. 2010. Evaluating search systems using result page context. In *Proceedings of the Third Symposium on Information Interaction in Context* (New Brunswick, New Jersey, USA) (IiIX '10). Association for Computing Machinery, New York, NY, USA, 105–114. doi:10.1145/1840784.1840801

- [3] Lucas Bernardi, Themistoklis Mavridis, and Pablo Estevez. 2019. 150 Successful Machine Learning Models: 6 Lessons Learned at Booking.com. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Anchorage, AK, USA) (KDD '19). Association for Computing Machinery, New York, NY, USA, 1743–1751. doi:10.1145/3292500.3330744
- [4] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2024. Double/Debiased Machine Learning for Treatment and Causal Parameters. arXiv:1608.00060 [stat.ML] <https://arxiv.org/abs/1608.00060>
- [5] Aleksandr Chuklin and Maarten de Rijke. 2016. Incorporating Clicks, Attention and Satisfaction into a Search Engine Result Page Evaluation Model. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management (CIKM'16)*. ACM, 175–184. doi:10.1145/2983323.2983829
- [6] Mukund Deshpande and George Karypis. 2004. Item-based top-N recommendation algorithms. *ACM Trans. Inf. Syst.* 22, 1 (Jan. 2004), 143–177. doi:10.1145/963770.963776
- [7] Beyza Ermis, Patrick Ernst, Yannik Stein, and Giovanni Zappella. 2020. Learning to Rank in the Position Based Model with Bandit Feedback. arXiv:2004.13106 [cs.LG] <https://arxiv.org/abs/2004.13106>
- [8] Thore Graepel, Joaquin Quiñero Candela, Thomas Borchert, and Ralf Herbrich. 2010. Web-scale Bayesian click-through rate prediction for sponsored search advertising in Microsoft's bing search engine. In *Proceedings of the 27th International Conference on International Conference on Machine Learning (Haifa, Israel) (ICML'10)*. Omnipress, Madison, WI, USA, 13–20.
- [9] Marti A. Hearst. 2009. *Search User Interfaces* (1st ed.). Cambridge University Press, USA.
- [10] Daniel N. Hill, Houssam Nassif, Yi Liu, Anand Iyer, and S.V.N. Vishwanathan. 2017. An Efficient Bandit Algorithm for Realtime Multivariate Optimization. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. ACM, 1813–1821. doi:10.1145/3097983.3098184
- [11] Jaya Kawale. 2019. A Multi-armed Bandit Framework for Recommendations at Netflix. SlideShare Presentation. <https://www.slideshare.net/slideshow/a-multiarmed-bandit-framework-for-recommendations-at-netflix/102629078> Accessed: [Add your access date].
- [12] Pratik Lahiri, Zhou Qin, and Wenyang Liu. 2024. Offline multi-objective optimization (OMOO) in search page layout optimization using off-policy evaluation. (2024). <https://www.amazon.science/publications/offline-multi-objective-optimization-omoo-in-search-page-layout-optimization-using-off-policy-evaluation>
- [13] Yiqun Liu, Ye Chen, Jinhui Tang, Jiashen Sun, Min Zhang, Shaoping Ma, and Xuan Zhu. 2015. Different Users, Different Opinions: Predicting Search Satisfaction with Mouse Movement Information. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (Santiago, Chile) (SIGIR '15)*. Association for Computing Machinery, New York, NY, USA, 493–502. doi:10.1145/2766462.2767721
- [14] Yi Liu and Lihong Li. 2021. A map of bandits for e-commerce. *arXiv preprint arXiv:2107.00680* (2021).
- [15] Yizhi Mao, Miao Chen, Abhinav Wagle, Junwei Pan, Michael Natkovich, and Don Matheson. 2019. A Batched Multi-Armed Bandit Approach to News Headline Testing. arXiv:1908.06256 [cs.LG] <https://arxiv.org/abs/1908.06256>
- [16] Themis Mavridis, Soraya Hausl, Andrew Mende, and Roberto Pagano. 2020. Beyond algorithms: Ranking at scale at Booking.com.
- [17] Rishabh Mehrotra, Niannan Xue, and Mounia Lalmas. 2020. Bandit based Optimization of Multiple Objectives on a Music Streaming Platform. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (Virtual Event, CA, USA) (KDD '20)*. Association for Computing Machinery, New York, NY, USA, 3224–3233. doi:10.1145/3394486.3403374
- [18] Jakob Nielsen. 2006. F-Shaped Pattern For Reading Web Content (Original Eyetracking Research). Nielsen Norman Group. <https://www.nngroup.com/articles/text-scanning-patterns-eyetracking/> Accessed: [Add your access date].
- [19] Lijing Qin, Shouyuan Chen, and Xiaoyan Zhu. 2014. Contextual Combinatorial Bandit and its Application on Diversified Online Recommendation. In *SDM*. <https://api.semanticscholar.org/CorpusID:7362645>
- [20] Zhou Qin, Kai Yuan, Pratik Lahiri, and Wenyang Liu. 2024. Cooperative Multi-Agent Deep Reinforcement Learning in Content Ranking Optimization. arXiv:2408.04251 [cs.LG] <https://arxiv.org/abs/2408.04251>
- [21] Anne Schuth, Katja Hofmann, and Filip Radlinski. 2015. Predicting Search Satisfaction Metrics with Interleaved Comparisons. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval (Santiago, Chile) (SIGIR '15)*. Association for Computing Machinery, New York, NY, USA, 463–472. doi:10.1145/2766462.2767695
- [22] Yue Wang, Dawei Yin, Luo Jie, Pengyuan Wang, Makoto Yamada, Yi Chang, and Qiaozhu Mei. 2016. Beyond Ranking: Optimizing Whole-Page Presentation. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (San Francisco, California, USA) (WSDM '16)*. Association for Computing Machinery, New York, NY, USA, 103–112. doi:10.1145/2835776.2835824
- [23] Yue Wang, Dawei Yin, Luo Jie, Pengyuan Wang, Makoto Yamada, Yi Chang, and Qiaozhu Mei. 2018. Optimizing Whole-Page Presentation for Web Search. *ACM Trans. Web* 12, 3, Article 19 (July 2018), 25 pages. doi:10.1145/3204461
- [24] Liang Wu, Diane Hu, Liangjie Hong, and Huan Liu. 2018. Turning Clicks into Purchases: Revenue Optimization for Product Search in E-Commerce. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval (Ann Arbor, MI, USA) (SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 365–374. doi:10.1145/3209978.3209993
- [25] Weiru Zhang, Chao Wei, Xiaonan Meng, Yi Hu, and Hao Wang. 2018. The Whole-Page Optimization via Dynamic Ad Allocation. In *Companion Proceedings of the The Web Conference 2018 (Lyon, France) (WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1407–1411. doi:10.1145/3184558.3191584