

Blending Advertising with Organic Content in E-commerce via Virtual Bids

Carlos Carrion¹, Zenan Wang¹, Harikesh Nair², Xianghong Luo¹, Yulin Lei¹, Peiqin Gu¹,
Xiliang Lin¹, Wenlong Chen¹, Junsheng Jin¹, Fanan Zhu¹, Changping Peng¹, Yongjun Bao¹,
Zhangang Lin¹, Weipeng Yan¹, Jingping Shao¹,

¹ JD.COM

² Stanford University

{carlos.carrion,wangzenan5,xianghong.luo,yulin.lei,peiqin.gu,xiliang.lin,
chenwenlong17,jinjunsheng1,zhufanan,pengchangping,baoyongjun,linzhangang,paul.yan,shaojingpingg}@jd.com,
harikesh.nair@stanford.edu

Abstract

It has become increasingly common that sponsored content (i.e., paid ads) and non-sponsored content are jointly displayed to users, especially on e-commerce platforms. Thus, both of these contents may interact together to influence their engagement behaviors. In general, sponsored content helps brands achieve their marketing goals and provides ad revenue to the platforms. In contrast, non-sponsored content contributes to the long-term health of the platform through increasing users' engagement. A key conundrum to platforms is learning how to blend both of these contents allowing their interactions to be considered and balancing these business objectives. This paper presents a system built for this purpose and applied to product detail pages of *JD.COM*, an e-commerce company. This system achieves three objectives: (a) Optimization of competing business objectives via Virtual Bids allowing the expressiveness of the valuation of the platform for these objectives. (b) Modeling the users' click behaviors considering explicitly the influence exerted by the sponsored and non-sponsored content displayed alongside through a deep learning approach. (c) Consideration of a Vickrey-Clarke-Groves (VCG) Auction design compatible with the allocation of ads and its induced externalities. Experiments are presented demonstrating the performance of the system. Moreover, our approach is fully deployed and serves all traffic through *JD.COM*'s mobile application.

Introduction

Presently, it is quite common that sponsored content (i.e., paid ads) and non-sponsored content (i.e., curated material selected and displayed to users free of charge to brands; hereafter referred to as organic content) are jointly displayed to users, especially in e-commerce. An example is the *recommended products* section commonly displayed in the detail pages of products. In principle, the development of a system for the allocation of ad content must consider the presence of organic content in order to avoid duplicated content, adverse competition between ads and organic content, and other possible interactions. A key conundrum for the optimal allocation of ads is satisfying multiple competing objectives (e.g. click-through rate [hereafter CTR] for ads and/or organics, ad revenue, time spent in pages of certain

advertised product categories, etc.). The multiple objectives may arise as representative proxies of various components of long-term platform health; or due to the fact that such a complex system usually involves many teams with different focuses in midsize to large companies. The multiple objectives may be competing with each other when improving one degrades another, leading to complex trade-offs.

In this research, we describe the development of a system for the personalized allocation of ad content in the presence of organic content and multiple objectives. Our system is applied to the product detail pages in *JD.COM*'s mobile application, an e-commerce company. A concrete example of this environment is the *product recommendation* section of the page presented in Figure 1 for several e-commerce vendors. In the example from *JD.COM* shown in the last panel, there are 6 positions available to display recommended ad or organic content, labeled "P1-P6". On *JD.COM*'s mobile app, the product recommendation section is located on the product page between a panel containing product summary and user reviews, and a panel containing a detailed description of the product. Our system focuses on personalized ad allocation at these positions. All users arriving at hundreds of millions of product detail pages available on the app can potentially be exposed to such ads.

Our system has three distinctive features. Firstly, it allows analysts to flexibly accommodate the multiple competing objectives in a virtual-bid formulation discussed subsequently. It presents a novel approach for the computation of these virtual bids without requiring explicit elicitation of these bids or even the specification of constraints from stakeholders. This new approach has yet to be presented in the literature of multiobjective optimization for personalized ad allocation systems (Agarwal and Chen 2016). In addition, the virtual bids have an economic interpretation; they represent the implicit valuation of the platform for each of these multiple objectives, i.e., the willingness to pay. Thus, these virtual bids allow the platform to express its desired trade-off between these objectives as its implicit valuations for each of these competing objectives of interest relative to the ad revenue, e.g., marginal rates of substitution between ad revenue and user engagement for ads. This aids interpretability and may have independent business decision value for other related questions e.g. deciding optimal ad load.

A second feature is the explicit modeling of users' click

behavior as a function of their characteristics, the individual characteristics of each sponsored content displayed, and more importantly, the influence exerted by other ad and organic content displayed alongside. Thus, our system offers a unified modeling solution that allows the platform to implement personalized ad allocation while capturing flexibly relevant *joint effects* (i.e., any substitution and/or complementarity effects within ads and between ads and organics).

The third feature of our allocation of ads mechanism is its compatibility with the Vickrey-Clarke-Groves (VCG) auction mechanism (Roughgarden 2016). In our system, the *joint effects* and the *virtual-bid approach* induce an allocation mechanism requiring a payment mechanism that properly constructs charges recognizing the *externalities* imposed by an arbitrary ad content on other ad and organic content. Generally, the widely used *Generalized Second Price* (hereafter GSP) auction mechanism does not consider such externalities in the payment mechanism (Gomes, Immorlica, and Markakis 2009; Roughgarden and Tardos 2012), whereas the VCG mechanism does consider them.

The system has been deployed and it serves on average tens of millions of auctions per day. In the rest of this paper, we describe first our approach; and report on a battery of experiments to evaluate the performance of our approach. Also, we review the literature to contrast our approach with related work and lastly we conclude.




Figure 1: Ads and org. recommendations on product-detail pages on apps: (L-R) Amazon, Etsy, Instacart, JD.COM.

Application Set-up

Our applied problem is to decide how to do personalized allocation of ads to users arriving at a product detail page. The allocation needs to respect possible interactions with other ads and organics on the page, so as to optimize against multiple, possibly competing, platform objectives. While our procedure can be framed and implemented more generally, we present it in the context of specific decisions made in its deployment at *JD.COM*. These decisions are about the identity and the ordering of the ads. Due to organizational reasons, both the number and the location of ad slots on the page were pre-determined and not optimized.

For a user arriving at a product’s detail page, we need to select K_{ads} to be placed in P_{ads} positions given K_{orgs} organics already placed in P_{orgs} positions. The positions are

decided a priori from positions P1-P6 on the *product recommendation* section in Figure 1. The K_{ads} selected for display must be optimal for two objectives: (1) Expected ad revenue obtained from a Click-per-Cost (hereafter CPC) billing option, where the CPC charge is determined from an auction mechanism; and (2) Expected CTR from the displayed ad content. This second objective is a metric representing user engagement with the ad content. The two objectives may not be aligned due to joint effects as explained below. There are two key challenges in optimal ad allocation:

1. *Externalities*. There are generally two types of joint effects that may influence users’ click behavior: (1) identities of the ads and organics; (2) arrangement of the ads and organics. The former refers to the possible substitution and complementarity effects within the ads and also between ads and organics. For example, displaying jointly multiple ads for pencils of different colors can induce substitution within the ads, and/or complementarity between the ads and organics, when organics are erasers. The latter refers to effects produced by displaying different permutations with the same ads and organics. In many ad contexts such as traditional search advertising, ads and organics are selected and ranked separately, neglecting such joint effects between such content, so that the externalities arising in the allocation scheme are ignored, partly because it makes the resulting auction more complex (Gomes, Immorlica, and Markakis 2009; Roughgarden and Tardos 2012).

2. *Dual Objectives*. For the platform, choosing the optimal set of ads K_{ads} given a set of organics K_{orgs} is a balancing act between maximizing ad revenue and user engagement (i.e. ad click-through-rates for our case; hereafter ad-CTRs). An approach followed in the literature (i.e. multiobjective optimization of ads in recommendation systems (Agarwal et al. 2012; Agarwal and Chen 2016; Yan et al. 2020)) is to frame the problem as maximizing an unconstrained platform objective by assigning *shadow prices* to the ad-CTRs in order to evaluate both the engagement objective and the ad revenue objective in the same units (i.e., currency). However, an open challenge is to properly choose the shadow prices. Alternatively, the problem has been formulated in this literature (Agarwal et al. 2012; Yan et al. 2020) as a constrained optimization problem where the shadow prices are traded for minimum thresholds for the user engagement goal, e.g. CTR, and then optimal ad allocation is found by maximizing ad revenue subject to this constraint. The challenge now is the minimum thresholds must now be properly set. A strategy generally followed in this literature is to use a fraction of the optimal CTR obtained by solving an unconstrained maximum CTR problem as the minimum thresholds, but this creates a new hyper-parameter to set, which is the fraction. **The principled determination of these hyperparameters, which is a major focus of this article, is barely discussed in the literature (Agarwal and Chen 2016).**

Listwise Ranker System for Ads Allocation

The system in place prior to the deployment of the new system (hereafter baseline) consists of (1) a deep learning based model (hereafter *pointwise* model) for predicting the CTR of an ad as a function of the characteristics of the user and the characteristics of only that ad considered separately from others and organics; (2) an allocation algorithm that ranks ads by computing an exponentially weighted *effective cost per mille* (eCPM) for each ad, with a hyperparameter t tuned using experimentation, e.g., the weighted eCPM for an ad is $bid_{CPC} \times pCTR^t$ where bid_{CPC} is the (CPC) bid submitted by the advertiser and $pCTR$ is the predicted CTR of the ad; and (3) a GSP auction mechanism to compute the corresponding payment given the ads allocation. It is represented as the Pointwise ranker in figure 2. The new system consists of three key components in addition to the components of the baseline system: (1) An enhanced deep learning model (hereafter *listwise* model) for predicting the CTR of ads considering the joint effects from other advertisements and organics displayed together; (2) a listwise ranking of ads using a *virtual bids* composite objective function; and (3) a listwise generator to enumerate multiple lists of size 6 of ads and organic content. It is represented as Listwise ranker in figure 2. Key difference between the baseline system and the new system is the former neglects the presence of the organic content in the allocation of ads in addition to the challenges discussed previously about *externalities* and *dual objectives*.

The unit of analysis in the new system is a *mixed tuple*, denoted by ω . ω is of size 6 composed of K_{ads} ads and K_{orgs} organics (i.e. $K_{ads} + K_{orgs} = 6$) placed in 6 positions P1-P6 as shown in Figure 1. For an arriving user, the new system takes as inputs, a ranked list of candidate ads $A_{1:N_{ads}}$ of size N_{ads} , and a ranked list of organic content $O_{1:N_{orgs}}$ of size N_{orgs} . The list of candidate ads is ranked using the baseline; and the organic list is ranked based on a model outside of ads. Figure 2 presents a schematic.




Figure 2: System overview for an arbitrary product's detail page in the the company's mobile app

For an arriving user, a candidate ω is generated by choosing the top $K_{orgs} < 6$ organics of the ranked list $O_{1:N_{orgs}}$ and $K_{ads} = 6 - K_{orgs}$ ads from the ranked list $A_{1:N_{ads}}$. The set of possible ω -s is denoted Ω , whose size is $K_{ads}! \times \binom{N_{ads}}{K_{ads}}$ as the order of ad placement matters. Since there is a set

P_{orgs} of positions for organics that are fixed, and the placement of the top K_{orgs} organics from $O_{1:N_{orgs}}$ in this set are decided a priori, the organic components of all candidate ω -s considered will be the same. However, the identity and order of the K_{ads} placed in P_{ads} will vary across candidate ω -s, and are what we optimize over.

Optimization Problem We let $x(\omega)$ denote the $K_{ads} + K_{orgs}$ dimensional vector of predicted CTRs for the components of ω . The fact that $x(\cdot)$ is indexed by ω reflects the joint effect of the ads and organic components on each other. Denote the K_{orgs} organic components in ω as ω^o , and the K_{ads} ads as ω^a . Then, we denote the predicted CTR of the j th organic element ω_j^o of ω^o as $x_j^o(\omega)$; and the predicted CTR of the j th ad element ω_j^a of ω^a as $x_j^a(\omega)$. We find the best allocation by searching for the ω that solves the following *virtual bids problem*,

$$\max_{\omega \in \Omega} \sum_{\omega \in \Omega} \sum_{j=1}^{K_{ads}} v^a x_j^a(\omega) + \sum_{\omega \in \Omega} \sum_{j=1}^{K_{orgs}} b_j x_j^o(\omega) \quad (1)$$

The term to be maximized is a composite objective, reflecting the platform's interest in serving ad allocations that respect both the expected ad-CTR and the expected ad revenue. The *virtual bids problem* converts all objectives to the same units (money), so that the overall objective is additive. The first term, which we refer to as money-metric ad-CTR uses a parameter v^a to convert expected ad-CTR into money metric terms; while the second term representing ad revenue is already in monetary terms. This is a linear optimization where we search for the $\omega^* \in \Omega$ that maximizes this composite objective. We call this formulation the *virtual-bid problem*, because the virtual bid v^a can be interpreted as the platform's valuation for one click of ad content. This is analogous to the advertisers' CPC bids submitted to the platform, which are the advertisers' valuations for each click on their ads. In a VCG auction context, this formulation has the interpretation of the platform and the advertisers bidding jointly for the user impression. However, the challenge of this formulation is that it is difficult to elicit from the platform, its implicit valuations encapsulated in the virtual bids. Assume for a moment that v^a is known, and that for each $\omega \in \Omega$, $x(\omega)$ is given. Thus, a linear search algorithm (i.e. search for the tuple with the maximum composite objective value out of the set of tuples) can find the optimal ω^* for the *virtual bids problem*. In practice, the size of Ω is likely to grow very large, leading to concerns about latency in ad-serving. A solution is the heuristic: construct ω choosing only the top N'_{ads} of the ranked list of ads $A_{1:N_{ads}}$ where $N'_{ads} \ll N_{ads}$, restricting $|\Omega| = K_{ads}! \times \binom{N'_{ads}}{K_{ads}}$. This heuristic assumes that the top N'_{ads} are likely to have the largest joint effects. N'_{ads} is treated as a hyperparameter that is tuned from experimentation.

An alternative to solving the *virtual bids problem* is to set a minimum desired ad-CTR threshold, C , and find the allocation that maximizes ad revenue subject to this constraint. This formulation is followed in the literature, e.g. (Agarwal et al. 2012; Yan et al. 2020), and thus our approach may be understood as its *Lagrangian relaxation* (Bertsimas and

Weismantel 2005; Fisher 1981) and our virtual bids can be interpreted as shadow prices that penalize ad revenue for the violation of implicit ad-CTR constraints.

Obtaining Virtual Bids from Historical Data If one interprets the virtual bids as the platform’s willingness to pay for one click of ad content, in principle, one could consider eliciting v^a directly from key decision-makers. However, it is difficult for decision-makers to articulate this construct in monetary terms (an *expressivity* problem). Also, the willingness to pay may be dynamic, and change based on platform competition. Further, heuristically picking v^a can reduce ad revenue and potentially ad-CTR (due to externalities). This motivates a data-driven approach to obtaining virtual bids.

To obtain v^a , we solve a separate sub-problem in historical data. Collect data $\mathcal{D} = \{\mathcal{X}_i(\omega) | i = 1 : N\}$ from an epoch of $i = 1, \dots, N$ past impressions with their respective set of mixer tuples \mathcal{X}_i , and for an impression i in the data, define $V_{ia}^*(v^a)$ and $V_{ir}^*(v^a)$ as the optimized values of the money-metric ad-CTR and ad revenue terms obtained by maximizing the platform’s composite objective in the *virtual bids problem* for a given virtual bid v^a . We define the *Utopia Point* of the sub-problem as the tuple $(\bar{V}_a^u(v^a), \bar{V}_r^u(0))$, where $\bar{V}_a^u(v^a) = \frac{1}{N} \sum_{i=1}^N V_{ia}^*(v^a)$, and $\bar{V}_r^u(0) = \frac{1}{N} \sum_{i=1}^N V_{ir}^*(v^a = 0)$, i.e., the best the platform could do on average across the N impressions in terms of its objective, if it cared only about each sub-objective in isolation, ignoring the other. The idea of our method is to find a v^a so that the induced money-metric ad-CTR and ad revenue come as close to the *Utopia Point* as possible on average across the N impressions. We do this by locating a point on the *possibility frontier* of ad-CTR and ad revenue that is as close as possible to the *Utopia Point*. The possibility frontier is the surface defined by all possible combinations of ad-CTR and ad revenue that can be generated by solving the *virtual bids problem* for a given vector of virtual bids that is Pareto efficient. Generally, the *Utopia Point* is not a point on this surface; if it was, the virtual bid should be 0 by definition as there is no trade-off between these two objectives. See Fig. 3 for graphical intuition.




Figure 3: Stylized Plot of Bi-dimensional Possibility Frontier and Utopia Point

Algorithm 1: Virtual Bids Search - SPSA

Input: $\mathcal{D} = \{\mathcal{X}_i(\omega) | i = 1 : N\}$ data set of mixer tuples for each impression i with their respective predicted CTRs; K maximum number of iterations; α, γ, A, a, c : hyper-parameters of SPSA, see (Spall 1998) for instructions on tuning.

Output: v^a : Optimal virtual bids

- 1: Initialize $a_k = 0; c_k = 0; \vec{g}_k = \vec{0}; \vec{v}_k^a = \vec{0}$
 - 2: Define $F_{\mathcal{D}}(\vec{v}^a)$ of size p , i.e. solve the *virtual bids problem* for each impression i in \mathcal{D}
 - 3: **for** $k \leftarrow 1$ **to** K **do**
 - 4: $a_k = \frac{\alpha}{(k+A)^\alpha}; c_k = \frac{c}{(k)^\gamma}$
 - 5: Sample $\vec{\Delta}_k$ of size p from $2(\text{Bernoulli}(0.5)) - 1$
 - 6: $\vec{v}_{plus}^a = \vec{v}_k^a + c_k \vec{\Delta}_k$
 - 7: $\vec{v}_{minus}^a = \vec{v}_k^a - c_k \vec{\Delta}_k$
 - 8: $y_{plus} = F_{\mathcal{D}}(\vec{v}_{plus}^a); y_{minus} = F_{\mathcal{D}}(\vec{v}_{minus}^a)$
 - 9: $\vec{g}_k = \frac{y_{plus} - y_{minus}}{2c_k \vec{\Delta}_k}$ (elementwise operation)
 - 10: $\vec{v}_k^a = \vec{v}_k^a - a_k \vec{g}_k$
 - 11: **end for**
 - 12: **return** \vec{v}_K^a
-

We formulate this search as a bilevel optimization problem. In the upper level, we search for a v^a that minimizes the L_2 -norm distance between the *Utopia Point* and the tuple of averages across all N impressions in the data of optimized money-metric ad-CTR and ad revenue, denoted $(\bar{V}_a^*(v^a), \bar{V}_r^*(v^a))$. $(\bar{V}_a^*(v^a), \bar{V}_r^*(v^a))$ are obtained in turn, by solving a lower-level problem for each v^a chosen by the upper level. In particular, for each chosen v^a from the upper level, the lower level solves the *virtual bids problem* for each i , and computes the averages across i of the tuple of optimized ad-CTR and ad revenue. Formally, we solve

$$\begin{aligned} \min_{v^a} F_{\mathcal{D}}(v^a) &= \left[\left(\frac{\bar{V}_a^*(v^a)}{\bar{V}_a^u} - 1 \right)^2 + \left(\frac{\bar{V}_r^*(v^a)}{\bar{V}_r^u} - 1 \right)^2 \right]^{\frac{1}{2}} \\ \text{s.t. } \bar{V}_a^*(v^a) &= \frac{1}{N} \sum_{i \in \mathcal{D}} V_{ia}^*(v^a) \\ \bar{V}_r^*(v^a) &= \frac{1}{N} \sum_{i \in \mathcal{D}} V_{ir}^*(v^a) \end{aligned} \quad (2)$$

For the case of one virtual bid v^a (our application), the upper level is one dimensional; we can solve this program efficiently using the *Golden search method* e.g. (Luenberger and Ye 2010). For the case of more than one virtual bid (e.g. if organics CTR was another objective with a virtual bid v^o), we can solve the upper level using *Simultaneous Perturbation Stochastic Approximation* (SPSA) (Spall 1998). For completeness, Algo 1 presents the algorithm using SPSA. The lower level is solved impression by impression using a linear search algorithm discussed previously. While this Program is solved offline, it can be updated frequently with recently logged data so the computed virtual bids reflect relevant changes in the platform environment. The frequency of updates will be platform-specific, depending on data dis-

tribution shifts, execution time, etc. Our experiments below document that re-tuning v^a is important for effective performance (which again motivates the need for a data-driven approach to do so). Also, although in this implementation we compute one value of v^a per epoch for use in ad-serving, conceptually, it is straightforward to allow it to vary by product-category, time-of-day, and other contexts, by pooling the data and making v^a a function of these variables.

Predicting CTRs for a mixed tuple To obtain the predicted CTRs, $x(\omega)$ for each candidate ω , we develop a novel deep learning model (hereafter listwise model) here not as much in terms of its architecture and training, but its use in conjunction with the problem solutions described previously for ad allocation with externalities; hence, the description is intentionally brief (for a more in-depth discussion of such models, see for instance (Covington, Adams, and Sargin 2016; Pei et al. 2019)). The model uses as input features:(1) user; (2) ads/organics; and (3) contextual characteristics. The first and second are self-explanatory and the third category corresponds to features describing the product in the product detail page. Figure 4 presents the model architecture. Briefly, (1) Categorical features are fed into embedding layers; (2) Numerical features are fed into Fully Connected Layers along with the Embedding Layers for the categorical features; (3) Sigmoid output layers are used for the prediction. Six dense networks are used to learn product low-dimensional representations from features. The Output Layer has six Sigmoid units corresponding to the mixed tuple; the self-attention layer utilizes a multi-head self-attention mechanism to learn interactive and context information among six products. The model is trained with an entropy loss function and batch normalization is used. Training is done via an offline-online scheme, i.e., training data is logged from online experiments where traffic is randomly allocated to the two models, and also retraining is done often. The training set for the listwise model consists of the set of features and the whole observed mixed tuple with realized clicks on any of the six positions in contrast to the pointwise model which only sees the position with realized clicks. The model is large-scale as in typical ad-industry applications, with an embedding size of about 500 millions and about 4 billions parameters. In offline validation, the listwise model performs superior on key performance metrics (AUC 0.75 vs 0.65, F1 Score 0.08 vs 0.065, and Accuracy 0.65 vs. 0.55).




Figure 4: Architecture of listwise model for CTR prediction.

Experiments

This section presents experiments ran on *JD.COM*'s mobile app to assess the effectiveness of our approach on a variety of scenarios and metrics. In addition, we report experiments to demonstrate the existence of *joint effects* in our application and study how accommodating shifts in the environment is important for effectively choosing virtual bids. In all experiments, impressions on product-detail pages are randomized into treatments, wherein treatment is defined as a particular way of doing ad allocation. When reporting results, we use the lift transformation for all the metrics in order to protect proprietary information. Except in the last experiment in Section Distribution Shift and Virtual Bids, the lifts reported in all the experiments are obtained by comparing a model A against the baseline system B . Specifically, the lift transformation is $\text{lift}_A = 100 \times \frac{\text{metric}_A - \text{metric}_B}{\text{metric}_B}$. Further, the statistical significance of the results in the subsequent tables follows the convention: * for 1% significance and ** for 0.1% significance. For the figures, the confidence intervals are for 99% confidence. Also, in some figures, the x -axis or y -axis are deliberately removed to protect proprietary information. Lastly, the ads payment scheme used in these experiments remains the GSP payment scheme.

Existence of Joint Effects

There are three treatments in this experiment: Random *shuffle* of the order of optimal ads from the baseline model; Random selection of ads from the top X ($X = K_{ads} + 1$ and $X = K_{ads} + 2$) of the pre-ranked ads list and placement of these in random order on the available ad-positions. The control is the ad allocation from the baseline model. The organic content is not adjusted in any group. As for performance metrics, we consider the average ad-CTR, average ad revenue, and average organics CTR across impressions in each group. Table 1 summarizes the lift estimates for these metrics. Looking at the first row, ads shuffling is seen to only barely affect organics, but to reduce the overall ads-CTR significantly. This suggests the order of the served ads matters. Looking at rows two and three, randomly selecting and randomly placing ads is not a good solution for ads-metrics as expected. Finally, because there is no change in the selection rule of organics compared to the baseline, the evidence seen of a significant impact on organics CTR due to changes in ads identity and location indicates an externality induced by ads. These results motivate the importance of considering joint effects for ad allocation.

Treat group	Lift over baseline model			Sample size
	Advertisements CTR	Revenue	Organics CTR	
Shuffle	-1.44%**	-0.24%	0.03%	15.6M
$K_{ads}+1$	-1.90%**	-0.41%	1.32%**	15.6M
$K_{ads}+2$	-4.07%**	-3.01%**	1.41%**	15.7M

Sample size of control group: 31.3M

Table 1: Randomly allocating ads – Dec, 2020

Treat group	Lift over baseline model			Sample size
	Advertisements CTR	Revenue	Organics CTR	
v^a	1.06%**	6.60%**	-0.17%	10.2M
$v^a - 1$	-3.14%**	11.45%**	0.95%**	10.2M
$v^a + 1$	3.57%**	3.89%**	-0.12%	10.2M

Sample size of control group: 20.5M

Table 2: Grid around the optimal virtual bid – Dec, 2020

Virtual Bids

This experiment demonstrates the efficacy of our approach compared to the baseline approach. Also, we demonstrate that our approach for choosing a virtual bid is able to find a tipping point for ad-CTR and ad revenue (i.e., improving both objectives to a point where improving one may worsen another). We consider three treatments corresponding to ad allocations that solve the *virtual bids problem* for three different values of v^a : (1) v^a = the virtual bid chosen as discussed in Section Obtaining Virtual Bids from Historical Data; (2) $v^a - 1$; (3) $v^a + 1$. The control is the ad allocation from the baseline model. Looking at the first row in Table 2, we see the new system is able to increase ad-CTR by 1.06% and ad revenue by 6.60% relative to the baseline system. In addition, the organics CTR is not statistically significantly different compared to the baseline for the treatment v^a . Rows two and three show that v^a is a tipping point and as we decrease or increase this virtual bid one of the objectives increases or decreases. Figure 5 explores heterogeneity in ad revenue increases across product categories. Substantial heterogeneity is seen. The fact that the impact on ad revenue is different across categories (and is negative in one), suggests that fine-tuning of v^a to specific categories could improve overall revenue even further. Figure 6 compares v^a vs. the CPC/bid distributions of advertisers across 10 selected categories. There are big disparities in the bid distribution across categories. This shows that even with full information of advertisers’ bids/valuation, manual tuning of the virtual bids can be a difficult challenge for the platform.




Figure 5: Distribution of ad revenue for optimal VB vs. baseline




Figure 6: VB and Advertisers’ Bid Distributions

Impact on Diversity

This experiment has one treatment and the control: the ad allocation with v^a chosen as discussed in Section Obtaining Virtual Bids from Historical Data and the baseline. The goal of this experiment is to better understand possible sources of improvement under the new system. We demonstrate empirically that considering the joint effects explicitly induces the selection of mixed tuples with greater diversity in terms of products. For the metrics, we consider: (1) More than 1 Subcat. in the mixed tuple meaning that at least one of the ads is of a product from a different subcategory than the product on whose page the ads are served; (2) Number of Subcat. in the mixed tuple meaning the number of different unique subcategories of products tied to the ads shown; and (3) Herfindahl-Index for the unique product subcategories of the ads shown (the index is a widely used measure of diversity, see (Rhoades 1993)). Looking at Table 3, we see our approach is able to present more diverse mixed tuples than the baseline as measured in the three metrics. For the Herfindahl Index, a smaller value reflects more diversity. Notably, there is no change in diversity among organics as expected because the selection rule for organics is the same as in the baseline. Finally, Figure 7 shows this effect occurs not just at the mean: the distribution of the number of subcategories shown to users under our approach at the optimal virtual bid (termed “VB”) is shifted to the right of the baseline.

Type	Lift over baseline model		
	More than 1 Subcat.	Number of Subcat.	Herfindahl
ads	6.06%**	1.34%**	-0.80%**
organic	-0.16%	-0.02%	0.01%
overall	2.82%**	0.94%**	-0.40%**

Sample sizes - treatment and control: 21.8M & 21.8M

Table 3: Diversity for new system – Jan, 2021

Constrained Optimization

This experiment features two treatments: ad allocations based on (1) v^a chosen as discussed in Section Obtaining




Figure 7: Distribution Changes of Number of Subcat. in the mixed tuple

Virtual Bids from Historical Data; (2) Constrained optimization with manually chosen C (i.e. ads CTR) following previous methods (Agarwal and Chen 2016). Ad allocation under the baseline remains as the control. The purpose of this experiment is to compare the data-driven approach for choosing virtual bids vs. a constrained optimization using a C an analyst may consider. By analogy to the virtual bids version of the problem, another way to interpret the constrained version is it represents an ad-allocation under a virtual bids problem where the virtual bid is not picked by principled tuning. We consider the same metrics discussed in previous experiments. Table 4 indicates that the manually selected threshold C (as explained using methods reviewed in section Related Works) increases ad revenue but at the expense of ad-CTR relative to the baseline. In contrast, the data-driven selected virtual bid v^a is able to find a balanced tipping point increasing ad-CTR by 2.05% and ad revenue by 9.25% relative to the baseline. It’s worth noting this experiment was done separately from the experiment in Section Virtual Bids, and serves as a replication of the value of the new system. Being able to find a balanced tipping point both times showcases the reliability of our approach.

Treat group	Lift over baseline model			Sample size
	Advertisements CTR	Revenue	Organics CTR	
v^a	2.05%**	9.25%**	0.32%	10.1M
C	-4.79%**	13.57%**	1.19%**	10.1M

Sample size of control group: 20.2M

Table 4: New system with Const. Opt. – Dec, 2020

Distribution Shift and Virtual Bids

The purpose of these experiments is to demonstrate the impact of distribution shifts in the environment, and how it may have an adverse impact on the metrics of interest if the virtual bids are not adjusted with proper frequency, or fixed manually to a static value. To implement this experiment,

we repeat the two-treatment experiment described in the section above, two weeks later. For ease of exposition, we call the first the T_1 -experiment and the second, implemented two weeks later, the T_2 -experiment. The T_2 -experiment is exactly the same as the T_1 -experiment except that for its first treatment, we use the optimal value of v^a obtained from the T_1 -experiment. For its second treatment, we continue to use constrained optimization with manually chosen C . Then, for each treatment, we report in Table 5 the lift in various metrics between T_2 and T_1 . Looking at the table, we see that using the same virtual bid in T_2 as was optimal in T_1 leads to worse performance; also, holding fixed the manually tuned C is bad for performance. Figure 8 shows this occurs not just at the mean, documenting that the entire distribution of revenue is shifting between the two periods. Clearly, the virtual bids require adjustment in order to function properly.

Treat group	Lift over respective T1 metrics of the same treatment		
	Advertisements CTR	Revenue	Organics CTR
$v_{T_2}^a = v_{T_1}^a$	-0.33%	-8.55%**	-0.40%
C	-1.61%**	-9.33%**	-0.64%*

Table 5: Virtual Bids on two periods Dec, 2020 & Jan, 2021




Figure 8: Distribution Shift in Revenue

Related Work

In the contextual ad-recommendation literature, the main thrust is on developing systems for serving recommendations considering the additional information available in the environment, e.g., a user interacting with a related item, which provides the context for recommendations. (Barbieri, Manco, and Ritacco 2014; Agarwal and Chen 2016; Zhang et al. 2019) provide recent treatises on this literature. In this stream, our work fits into feature-based deep supervised learning systems for item recommendations. A novelty of our approach is considering joint effects in the CTR prediction; we relax the assumption of conditional independence of the CTRs of the ads and organics displayed jointly given the context of the product featured.

Within the multiobjective optimization problem of recommendation systems, our work is closest to the stream of *constrained optimization* approaches proposed in (Agarwal et al. 2011, 2012; Yan et al. 2020). (Agarwal et al. 2011) proposed a linear constrained optimization problem to optimize the engagement time spent of users on an article’s landing page subject to CTR constraints. The minimum desired CTRs are a fraction of the optimal CTRs obtained by maximizing an unconstrained maximum CTR problem. Thus, this fraction is a hyperparameter that requires tuning. (Agarwal et al. 2012) improves the constrained optimization problem of (Agarwal et al. 2011) to allow personalization, and uses Lagrangian Duality to construct a scalable solution. (Yan et al. 2020) also formulate a constrained optimization and propose a scalable solution where the Lagrangian Duals are taken as given. For optimal performance, these duals have to be tuned. While there are many suggestions, developing a structured approach to tuning such hyperparameters is an open question. A novel contribution of this study is to present a principled approach to tuning these virtual bids. Our approach uses a bilevel optimization problem for tuning the hyperparameters of our problem (the virtual bids), which to the best of our knowledge has not been proposed in this literature. We show that tuning the virtual bids appropriately is critical for optimal performance. Furthermore, our work is easily extendable to the setting of (Yan et al. 2020) where both ads and organics are jointly optimized. Also, our formulation uses theoretical results of obtaining Pareto optimal solutions from the literature on compromise optimization (Miettinen 1999). Lastly, our work forms the allocation scheme of an Auction Design problem (Nisan 2007; Roughgarden 2016). The other part is the payments scheme. Our formulation links our allocation scheme to a VCG payment scheme which is attractive for contextual advertising with externalities, see (Lahaie et al. 2007; Varian and Harris 2014). In computing VCG payments, the objective in the *virtual bids problem* also serves as the VCG–auction payoff, thus linking the allocation and payment parts in an internally-consistent way. While our allocation scheme may be used with GSP auctions, the *externalities* present would not be properly priced (see (Gomes, Immorlica, and Markakis 2009; Roughgarden and Tardos 2012)). In the deployed version, the payment scheme is the GSP scheme as moving to a VCG system involves complex business and engineering decisions requiring further evaluation.

Deployment

Our approach was developed for the product recommendation section located for all the product detail pages accessed through *JD.COM*’s mobile app. This recommendation section is located between a panel containing a short summary of the product and user reviews, and a panel containing a detailed description of the product. The recommendation section corresponds to 6 positions available to display recommended ad or organic content (see the last panel in Figure 1). This product section serves on average tens of millions of auctions per day. The development of the approach started in May 2020. Mainly, the development of the new components needed for the Listwise ranker presented in Figure 2. This

Listwise ranker is build on top of the baseline (i.e. Pointwise ranker), and thus as a fallback the baseline ranking may be served in case of any issues. These new components are: (1) Mixed tuple candidate generation; (2) CTR Prediction for the Mixed tuple candidates; (3) Optimization using the Mixed tuple candidates; (4) Batch service for Virtual Bid Optimization. Component 1-3 are used in real-time and meet the latency requirements per impression of less than 30 ms. Component 4 is run periodically offline to tune the virtual bid due to data distribution shifts. The small traffic experiments started in July 2020 and the fraction of the traffic allocated to the tests gradually increased in multiple phases until January 2021. The test results in the Experiments section showed significant improvements in the business metrics of ad revenue and user engagement (i.e. ads CTR for our case) without impacting negatively other key metrics (e.g. organics CTR) over the baseline system. For example, table 4 reports a lift of 9.25% revenue for ads for our approach relative to the baseline. Thus, suppose the baseline brings on average \$1M USD in revenue per day then our approach brings on average \$1M + 92,500 USD per day. These numbers are for illustrative purposes only. Table 4 also reports lifts for Ads CTR and Organics CTR. We do not report the precise values of the virtual bid and other parameters of the deployed model due to these being proprietary information. Nevertheless, the impacts of the deployed model on revenue and engagement are both positive, and consistent with the test results. Thus, the new system was cleared for the official launch and completed in January 2021. Additionally, testing with full traffic continued until February 2021 to verify the stability of the results. Furthermore, a summary of the key milestones of the project from development to deployment is as follows: (1) Identifying the existence of joint effects through experimentation; (2) Revision of the Deep Learning architecture to model explicitly these joint effects; (3) Proposing a heuristic to solve the *virtual bids problem* defined in equation 1 with a latency constraint of 30 ms per auction served; (4) Proposing a system to update the virtual bids with minimal intervention that maintains the balance between the business objectives; (5) Measuring the benefits of the approach on key dimensions, e.g. diversity of the mixed tuples and the interpretation of the virtual bids within distinct product categories. Milestone 1 served as the motivation for this work, while Milestones 2-4 tackled the challenges of developing the approach. Milestone 5 studied the benefits of the approach.

Conclusion

We developed a system for ads allocation in the presence of organic content on product detail pages in e-commerce and report on the system’s deployment on *JD.COM*’s mobile app. We demonstrate empirically the presence of motivating joint effects in our application. We show that considering these effects explicitly has benefits such as increased diversity in the ads presented. Also, we demonstrate that our data-driven approach for choosing virtual bids is able to find tipping points for the objectives. Lastly, we emphasize the importance of properly updating these bids for good performance to deal with distribution shifts in the data.

References

- Agarwal, D.; Chen, B.-C.; Elango, P.; and Wang, X. 2011. Click shaping to optimize multiple objectives. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 132–140.
- Agarwal, D.; Chen, B.-C.; Elango, P.; and Wang, X. 2012. Personalized click shaping through lagrangian duality for online recommendation. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, 485–494.
- Agarwal, D. K.; and Chen, B.-C. 2016. *Statistical methods for recommender systems*. Cambridge University Press.
- Barbieri, N.; Manco, G.; and Ritacco, E. 2014. Probabilistic approaches to recommendations. *Synthesis Lectures on Data Mining and Knowledge Discovery*, 5(2): 1–197.
- Bertsimas, D.; and Weismantel, R. 2005. *Optimization over integers*. Dynamic Ideas.
- Covington, P.; Adams, J.; and Sargin, E. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*, 191–198.
- Fisher, M. L. 1981. The Lagrangian relaxation method for solving integer programming problems. *Management science*, 27(1): 1–18.
- Gomes, R.; Immorlica, N.; and Markakis, E. 2009. Externalities in keyword auctions: An empirical and theoretical assessment. In *International Workshop on Internet and Network Economics*, 172–183. Springer.
- Lahaie, S.; Pennock, D.; Saberi, A.; and Vohra, R. 2007. Sponsored Search Auctions. In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V., eds., *Algorithmic game theory*, chapter 28, 699–716. Cambridge University Press.
- Luenberger, D. G.; and Ye, Y. 2010. *Linear and Nonlinear Programming*. Springer, 3 edition.
- Miettinen, K. 1999. *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers.
- Nisan, N. 2007. Introduction to mechanism design (for computer scientists). In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V., eds., *Algorithmic game theory*, chapter 9, 209–242. Cambridge University Press.
- Pei, C.; Zhang, Y.; Zhang, Y.; Sun, F.; Lin, X.; Sun, H.; Wu, J.; Jiang, P.; Ge, J.; Ou, W.; et al. 2019. Personalized re-ranking for recommendation. In *Proceedings of the 13th ACM Conference on Recommender Systems*, 3–11.
- Rhoades, S. A. 1993. The herfindahl-hirschman index. *Fed. Res. Bull.*, 79: 188.
- Roughgarden, T. 2016. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press.
- Roughgarden, T.; and Tardos, É. 2012. Do externalities degrade GSPs efficiency? In *The Eighth Ad Auctions Workshop*.
- Spall, J. C. 1998. Implementation of the simultaneous perturbation algorithm for stochastic optimization. *IEEE Transactions on aerospace and electronic systems*, 34(3): 817–823.
- Varian, H. R.; and Harris, C. 2014. The VCG auction in theory and practice. *American Economic Review*, 104(5): 442–45.
- Yan, J.; Xu, Z.; Tiwana, B.; and Chatterjee, S. 2020. Ads Allocation in Feed via Constrained Optimization. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3386–3394.
- Zhang, S.; Yao, L.; Sun, A.; and Tay, Y. 2019. Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys (CSUR)*, 52(1): 1–38.