

Aligning Ranking Objectives with E-commerce Search Intent

Dominic Seyler, Ehsan Ebrahimzadeh, Alex Cozzi, Abraham Bagherjeiran
dseyler,eebrahimzadeh,acozzi,abagherjeiran@ebay.com
eBay Search Ranking and Monetization

ABSTRACT

Learning To Rank (LTR) from implicit user feedback is the predominant approach in large scale information retrieval systems for e-commerce. The standard LTR approaches that account for search intent in training ranking models focus primarily either on enriching the feature representation of the model by estimates of the search intent, or adopting an auxiliary intent prediction task in a multi-task learning setting, or debiasing the user feedback using intent-aware propensity models. In this work, we propose intent-aware LTR schemes by stratifying the training data with respect to intent groups identified by empirical labels that correspond to distinct item desirability distributions. Specifically, through importance sampling on training queries based on the richest engagement event attributed to the engaged item on the SRP, we train rankers that align the LTR objective with the corresponding item desirability distributions for the browse and purchase intents. In order to demonstrate the efficacy of the proposed training data stratification technique in generalizing across query segments with different underlying intents, we evaluate the rankers trained with different importance sampling weights on the traffic segment identified by the Browse search experience in a major e-commerce platform. In particular, we show that the ranker trained only on queries with post click conversion signals is significantly outperformed by a ranker that relies also on non-converting queries in training. We further demonstrate a higher variance in ranking efficiency in the traffic segment identified by the browse intent, due to the more severe distribution shift.

CCS CONCEPTS

• **Information systems** → **Query intent; Information retrieval; Learning to rank; Users and interactive retrieval; Electronic commerce.**

KEYWORDS

query segmentation, learning to rank, e-commerce

ACM Reference Format:

Dominic Seyler, Ehsan Ebrahimzadeh, Alex Cozzi, Abraham Bagherjeiran. 2023. Aligning Ranking Objectives with E-commerce Search Intent. In *Proceedings of ACM SIGIR Workshop on eCommerce (SIGIR eCom'23)*. ACM, New York, NY, USA, 5 pages.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SIGIR eCom'23, July 27, 2023, Taipei, Taiwan
© 2023 Copyright held by the owner/author(s).

1 INTRODUCTION

In e-commerce, a substantial portion of customer purchases are the result of searching and exploring the item inventory of the platform. This behavior is best described as exploratory search, characterized as a combination of exploratory browsing and focused searching [2, 8]. While focused search has been a primary research focus for the information retrieval community, user studies show that roughly 40% to 65% of users' goals are informational [18], implying that searchers are seeking novel items. It is therefore essential for industry scale information retrieval systems to adjust to the user need depending on the specific search scenario [13].

In this work, we study the effect of adjusting the label distribution in the Learning To Rank (LTR) training objective to the underlying intent in exploratory and focused search scenarios in an e-commerce setting. Specifically, we evaluate the generalization performance of the rankers trained on varying label distributions across traffic segments based on the *search experience* in an e-commerce search engine. Previous work has accounted for capturing the user intent in LTR mainly by enriching the feature representations of the models. However, there are multiple advantages to dealing with user intent in a feature-agnostic setting, primarily due to the hardness of the intent prediction task in the absence of additional contextual information about the user's preferences, particularly in the early stages of the search journey.

We take an alternative intent-aware LTR approach, by stratifying the training data with respect to the label distribution that estimates the underlying item desirability distribution for different intent groups. Intent estimation policy can be a complex function of the user's query as well as the prior interactions in the search context, therefore, in our evaluations we identify the browse or focused search intent based on the type of search interaction experience. Specifically, we identify the browse intent based on the user traffic who explicitly adopt the *Browse experience* and explore the retrieved items in product category groups, in contrast to expressing a shopping intent by issuing a keyword query. We argue that due to the measurable differences in the post click conversion rates across the two search experiences, an LTR objective with a label distribution based on click engagements is more suitable for users in the exploration phase of their search journey; whereas an objective with an empirical label distribution based on purchases is more suitable for users with a focused search intent. We motivate the relationship between user engagement events (i.e., clicks and purchases) and underlying user intent (i.e., browse and focused search) by identifying distinct user behavior patterns in each segment, as surfaced in logged search data from a major e-commerce platform.

In our search traffic analysis, we observe that users who engage with items in Browse Result Pages (BRP) when browsing product categories are less likely to convert on the clicked items compared to the users who issue a keyword query and engage with the items

on Search Result Pages (SRP). Therefore, we argue that browse queries based on a user-selected category drive more engagements for further exploration, and the ranking objective in this segment should be driven primarily by engagement efficiency. On the other hand, for focused queries with an underlying purchase intent, the ranking objective should be driven mostly by the conversion efficiency. Our study makes an initial step towards characterizing the fundamental sale/engagement efficiency trade-off in different traffic segments. This is achieved by training ranking policies on optimization objectives tailored to the corresponding intent class and serving requests from each traffic segment with the designated ranker for that class. Our work makes the following contributions:

- We propose training data stratification techniques for intent-aware LTR with intent strata identified by the empirical target labels that estimate distinct item desirability distributions in different search interaction scenarios.
- We identify distinct engagement patterns and measurable difference in post click conversion rates in e-commerce search interaction scenarios corresponding to browsing product category groups (i.e., exploration) versus expressing a purchase-driven intent (i.e., focused search).
- We evaluate the effectiveness of our intent-aware LTR methodology in an e-commerce ranking scenario by highlighting the variance in ranking efficiency in the intent segment identified by the browse experience, as the training label distribution of the ranking policy changes.

2 RELATED WORK

There is a remarkable body of work that focus on query intent prediction in the context of e-commerce [3, 6, 10, 14, 15, 21, 22]. While exploration has been identified as one of the major user intents in e-commerce settings [16, 17], incorporating exploratory search intent in the e-commerce ranking objective remains an open problem [19].

Once user’s browsing/exploration intent is estimated, much less work has focused on how a search system can effectively support different user behavior. For instance, Rahman et al. [11] propose a conceptual recommendation framework to support exploratory search in e-commerce, by combining ranked results from a text based search engine with a recommender system based on user interactions. Ebrahimzadeh et al. [5] propose debiased estimators with intent-aware propensity estimates based on the number of clicks in a search context. Medlar et al. [9] classify users’ search tasks in fine-grained categories, which fall within a behavioral spectrum between focused search and exploratory browsing. Rowley [12] focus on matching a browsing user to the correct product category, while the question of whether the ranking objective should vary depending on the user intent is not addressed. In our work, we proposes to train rankers for users with varying intent, where the preference for different user engagements is incorporated in the objective function.

Another line of work focuses on query understanding. For example, Wang et al. [20] parse a query into linguistic concepts that related to specific instances in a knowledge base. Dai and Callan [4] use contextual language models (such as Kenton and Toutanova [7]) in information retrieval to achieve a better understanding of

the query context. Another related area of work focuses on a better semantic understanding of queries [1]. Applying these works to E-Commerce is not trivial and needs to address many challenges, such as vocabulary gap and data sparsity [19]. Furthermore, this particular research area does not clearly distinguish users by their intent and these methods are not applicable for category-based queries without any keywords. In contrast, our browse-specific rankers are optimized to support users with browse intent and can perform on keyword-less queries.

3 METHODOLOGY

3.1 Problem Setup

We consider a generic supervised LTR setting where we observe user feedback in the form of clicks on the SRP, as well as the subsequent post click transaction event attributed to the clicked items. A search $s \in \mathcal{S}$ is characterized by the query $q \in \mathcal{Q}$ made by some user $u \in \mathcal{U}$, the SRP that is an ordered collection of items (i.e., documents) $\mathcal{D}_s \in \mathcal{D}^N$ retrieved by the engine with respect to q as well as the user’s click and post click purchase events, denoted by $c(\mathcal{D}_s) \in \{0, 1\}^N$ and $p(\mathcal{D}_s) \in \{0, 1\}^N$, respectively. It is standard to assume that ground truth relevance labels r_d are produced by some oracle according to an underlying item desirability distribution $\mathbb{P}_{r_d|s}(r_d|s)$ in the search context s . We learn a ranking policy by minimizing the statistical risk, corresponding to a search efficiency loss function, which is meant to approximate the expected number of user behavior events. Specifically, the statistical risk for a ranking function $f(d; s)$, which produces a score for each individual document $d \in \mathcal{D}_s$ given the search context s , is defined as

$$R(f) = \mathbb{E}[\mathcal{L}(f(\mathcal{D}_s), r(\mathcal{D}_s))], \quad (1)$$

where $\mathcal{L}(f(\mathcal{D}_s), r(\mathcal{D}_s)) = \ell(f(\mathcal{D}_s))^T r(\mathcal{D}_s)$ is the DCG-based ranking efficiency loss with respect to the ideal ranking defined with respect to $r(\mathcal{D}_s)$, suitably discounted by some function of the rank $\ell(f(\mathcal{D}_s)) = [-\lambda(\text{rank}_f(d))]_{d \in \mathcal{D}_s}$ attributed to the items via the scoring function f .

Since the underlying joint distribution of the search contexts and relevance labels is not known to the learner, the standard approach is to build an empirical risk estimate based on a sampled set \mathcal{S}^τ of search contexts with suitably defined empirical labels $\hat{r}(\mathcal{D}_s)$ to approximate the ground truth label distribution $\mathbb{P}_{r_d|s}(r_d|s)$ such that

$$\hat{R}(f) = \frac{1}{|\mathcal{S}^\tau|} \sum_{s \in \mathcal{S}^\tau} \mathcal{L}(f(\mathcal{D}_s), \hat{r}(\mathcal{D}_s)), \quad (2)$$

assuming all the training data from logged search contexts come from the same underlying distribution $\mathbb{P}_s(s)\mathbb{P}_{r_d|s}(r_d|s)$. In this work, we are primarily interested in *estimation error* in our Empirical Risk Minimization (ERM) setup, oblivious to feature representation, hypothesis class, and training optimization scheme, which have to do with optimization and approximation error in the ERM setup.

3.2 Stratified Empirical Risk Minimization

We stratify the LTR training data based on intent groups $i \in \mathcal{I}$ identified by distinct item desirability distributions

$$\mathbb{P}_{r_d|s,i}(r_d|s, i) \quad (3)$$

in different search interaction scenarios. We can therefore build stratified empirical risk estimates by adjusting the empirical labels $\hat{r}_i(\mathcal{D}_s)$ so that they approximate the underlying desirability distribution in the corresponding intent segment; that is

$$\hat{R}(f) = \sum_{i \in \mathcal{I}} \frac{\hat{\mathbb{P}}_i}{|\mathcal{S}_i^\tau|} \sum_{s \in \mathcal{S}_i^\tau} \mathcal{L}(f(\mathcal{D}_s), \hat{r}_i(\mathcal{D}_s)), \quad (4)$$

where \mathcal{S}_i^τ is the logged training data corresponding to the intent segment i and $\hat{\mathbb{P}}_i$ is some stratification distribution across the intent groups. We focus on a simple binary intent segmentation based on the underlying browse or purchase behavior of the user. In this case, our estimate $\hat{r}_i(\mathcal{D}_s)$ of the underlying intent-based item desirability distribution $\mathbb{P}_{r_d|s,i}(r_d|s,i)$ is two-fold: for users with the browse intent it is the empirical click engagement distributions $c(\mathcal{D}_s)$, whereas for users with focused search intent it is estimated the empirical distribution based on purchase signals $p(\mathcal{D}_s)$. In fact, to avoid the complexity in building our empirical estimate, let us assume that we only have a single clicked item within all the training search contexts. Thus, since we observe post-click purchases only for clicked items in converting queries, our stratified estimate can be written as

$$\hat{R}(f) = \frac{1}{|\mathcal{S}_\alpha^\tau|} \sum_{s \in \mathcal{S}_\alpha^\tau} \mathcal{L}(f(\mathcal{D}_s), \hat{c}(\mathcal{D}_s)), \quad (5)$$

where $\hat{c}(\mathcal{D}_s)$ is the suitably debiased click label (e.g. based on inverse propensity weighting) and the training search contexts \mathcal{S}_α^τ are sampled from the search context with post click purchases \mathcal{P} with an importance sampling parameter α and from the search context with no post-click purchases \mathcal{C} with importance sampling parameter $1 - \alpha$; that is $\mathcal{S}_\alpha^\tau = \{\mathcal{P}\}_\alpha \cup \{\mathcal{C}\}_{(1-\alpha)}$.

Using standard arguments, we can characterize the bias and variance of the proposed stratified empirical risk estimator with respect to the Bayes statistical risk corresponding to the underlying ground truth item desirability distribution in the segment. In fact, the generalization power of the proposed stratified empirical risk estimate can be adjusted to the corresponding query segment by an appropriate choice of the stratification parameter α .

4 EXPERIMENTAL EVALUATION

4.1 Evaluation Segments

While we can clearly highlight the generalization power of the proposed stratified risk estimation technique by focusing on the same strata on the evaluation dataset, we instead rely on evaluation segments that can explicitly be identified at the online serving time. Specifically, we identify the user's underlying search intent based on the the explicit search engine interaction experience that the user opts into: (A) the user expresses a browse intent by clicking on a category on the homepage, e.g. "Automotive", and explores the results on a Browse Result Page (BRP); (B) the user expresses a focused search intent by issuing a query by entering keywords in the search bar, e.g. "Audi R8 headlights", and explores the results on a Search Result Page (SRP).

To motivate the relationship between user intent, as characterized by BRPs and SRPs, with the corresponding intent strata at




Figure 1: Click Propensity Ratios of Search Result Page (SRP) and Browse Result Page (BRP) across Different Ranks.

the training time, we perform an analysis of post click conversion rates (CVR) on user search logs. In fact, we report a measurably lower CVR in BRPs compared to SRPs, with a difference $\Delta(CVR_{SRP}, CVR_{BRP}) > 50\%$. This is a strong indicator that BRP engagements are more likely to be followed by subsequent clicks via engaging in other item recommendation modules or issuing a keyword search query rather than committing to a purchase. This is a key observation as it motivates the connections between the intent classes and the corresponding engagement/conversion driven objectives. To further understand the interaction patterns across query intent segments, we perform a second study to compare click propensities at different ranks across browse and focused search query segments (Figure 1). We observe a smoother rank-decay in click propensities in browse queries and more engagement likelihood in lower ranks. This observation confirms the hypothesis that user behavior patterns vary by intent, as users with a browse intent are more likely to explore lower ranks, while users with a purchase intent are more likely to engage with top ranks.

The inherent differences in the engagement patterns between browse and focused search queries indicate that: (1) Browse queries drive engagements for further exploration, thus the ranking objective in this segment should be driven by the engagement efficiency; and (2) In contrast, focused search queries are more likely to be specific and purchase driven, thus the ranking objective should primarily be driven by efficiency of conversion events.

4.2 Experiment Setup

In this section we design our experiments to measure the empirical performance of the rankers trained on segment-specific training objectives evaluated on query segments corresponding to BRP and SRP search experiences. Given an evaluation segment, we are specifically interested in the *relative performance* of a ranker that is trained only on queries with post click purchases attributed to the engagement events versus one that can rely also on the non-converting

search contexts. We train ranking models f_α , parameterized by the importance sampling parameter α , on stratified training datasets \mathcal{S}_α^τ , where α proportion of the training queries are sampled from queries with a post click purchase event and the remaining $1 - \alpha$ proportion are sampled from the non-converting query set. By training different rankers f_α through varying the stratification parameter α , and evaluating them over different user intent segments, we establish a conversion/engagement trade-off as it relates to the underlying item desirability label distribution of the ranker.

For a logged search context s in the evaluation dataset \mathcal{S}^ϵ , suppose that the richest engagement event attributed to the engaged item on the SRP is E ; which is essentially a purchase event P if there is a post-click conversion attributed to that engaged item, otherwise it is just the click event C on the SRP. The ranking efficiency of the ranker f is measured as the average Mean Reciprocal Rank (MRR) with respect to $rank_f(E)$, the rank attributed to the engaged item by f in search contexts $s \in \mathcal{S}^\epsilon$; that is

$$MRR_E(f, \mathcal{S}^\epsilon) = \frac{1}{|\mathcal{S}^\epsilon|} \sum_{s \in \mathcal{S}^\epsilon} \frac{1}{rank_f(E)}. \quad (6)$$

4.3 Datasets

Training Data. We build distinct datasets by stratifying the queries based on the richest engagement event observed in the search context $c(\mathcal{D}_s)$ and $p(\mathcal{D}_s)$; Specifically, the proportion of queries with a purchase event in the search context is controlled by importance sampling parameter α . Note that during training, models are exposed to queries from all intent segments (both in terms of the training segments based on the underlying item desirability distribution and the evaluation segments based on the BRP/SRP search experience) to take advantage of synergies among all user experiences. To control for the size of the training dataset, for each query we sample three unengaged items from the candidate set \mathcal{D}_s at random, as negative examples. For training targets $\hat{c}(\mathcal{D}_s)$, we use debiased labels using the propensities characterized in the previous section. All training datasets contain roughly 1M queries.

Test Data. The test datasets are based on a random sample of queries for which there is at least one click event observed in the search context, suitably separated from training sets to avoid data leakage. The two test datasets are built via stratification based on BRP and SRP intent segments and each contain around 10K to guarantee meaningfully narrow confidence intervals. For the test data, we keep all the candidate items to be re-ranked by the candidate ranker.

Note that datasets are a random sample of e-commerce search traffic in the months of April to June 2022.

4.4 Empirical Results

Table 1 offers a fundamental observation on the generalization power of the proposed stratified empirical risk estimate. Essentially, we can achieve a bias/variance trade-off by an appropriate choice of the stratification parameter α , which can in fact be adjusted based on the evaluation segment that ranker is served on. As such, the unbiased estimator of the purchase desirability intent $f_{\alpha=1}$, fails to generalize even on the traffic population with purchase intent, primarily due to the selection bias in training search context qualification. Essentially, $f_{\alpha=1}$ relies only on converting queries to

Table 1: Performance difference between purchase-driven model $f_{\alpha=1}$ and the baseline model with optimal f_{α^*} . All reported lifts are statistically significant and exceed the threshold for a meaningful effect size

Lift w.r.t Evaluation Data	$\mathcal{S}_{BRP}^\epsilon$	$\mathcal{S}_{SRP}^\epsilon$
$\Delta (MRR_C(f_{\alpha=1}, \cdot), MRR_C(f_{\alpha^*}, \cdot))$	-8.3%	-5.7%
$\Delta (MRR_P(f_{\alpha=1}, \cdot), MRR_P(f_{\alpha^*}, \cdot))$	-2.1%	-0.7%

estimate the more diverse traffic distribution. This is particularly more pronounced in the BRP segment, because as pointed out in the search traffic analysis, most users do not commit to a purchase upon engaging with an item in a BRP. In fact, by relying also on non-converting queries in training the rankers via stratified empirical risk estimates, we can improve the estimation accuracy of the target distribution for engagement efficiency.

On the other hand, if we do not over-sample converting queries and only rely on a randomized sample from the traffic population \mathcal{S}^C , both the conversion and engagement rates will be meaningfully impacted, even in population segments with primarily a browse intent, because search contexts with post click purchase \mathcal{S}^P usually offer richer information for training.

Next, by comparing the effect size of the relative performance of the rankers across query segments, we highlight a meaningfully *higher variability in engagement efficiency* in the browse segment compared to the focused query segment, highlighting strong evidence towards the choice of a click engagement driven objective for the browse segment. In fact, we report a significantly higher variance in the search efficiency metric MRR_C in the BRP segment compared to the SRP segments as the ranking model f_α varies; that is

$$\Delta \left(\text{Var}_\alpha [MRR_C(f_\alpha, \mathcal{S}_{BRP}^\epsilon)], \text{Var}_\alpha [MRR_C(f_\alpha, \mathcal{S}_{SRP}^\epsilon)] \right) > 10\% \quad (7)$$

The large variance difference is due to a significant performance drop for model f_α when α is close to 1, as they fail to generalize to the population distribution with an underlying browse intent. In fact, the best performing setting for α can be tuned optimally for the specific segments that ranking policy serves. We leave the exploration of other query segmentation strategies, beyond BRP and SRPs, as well as the best practice for fine tuning the stratification parameter to future work.

5 CONCLUSION

We presented a training data stratification technique based on intent groups characterized by distinct underlying item desirability distributions. To verify the generalization performance of the proposed scheme, we performed empirical evaluations on search traffic from a major e-commerce platform. We showed the poor generalization of the ranker trained only on queries with rich post click engagement events, and demonstrated that as we change the training target distribution of the ranker, the variance in search efficiency is measurably higher in the browse segment compared to the traffic segment identified by keyword queries. These observations confirm the hypothesis on the fundamental benefits of aligning the ranking objective with the user's search intent.

REFERENCES

- [1] Ricardo Baeza-Yates. 2017. Semantic query understanding. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1357–1357.
- [2] Marcia J Bates. 2007. What is browsing-really? A model drawing from behavioural science research.
- [3] Douglas Cirqueira, Markus Hofer, Dietmar Nedbal, Markus Helfert, and Marija Bezbradica. 2020. Customer purchase behavior prediction in e-commerce: A conceptual framework and research agenda. In *New Frontiers in Mining Complex Patterns: 8th International Workshop, NFMCP 2019, Held in Conjunction with ECML-PKDD 2019, Würzburg, Germany, September 16, 2019, Revised Selected Papers*. Springer, 119–136.
- [4] Zhuyun Dai and Jamie Callan. 2019. Deeper text understanding for IR with contextual neural language modeling. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 985–988.
- [5] Ehsan Ebrahimzadeh, Alex Cozzi, and Abraham Bagherjeiran. 2023. Intent-Aware Propensity Estimation via Click Pattern Stratification. In *Companion Proceedings of the ACM Web Conference 2023*. 751–755.
- [6] Long Guo, Lifeng Hua, Rongfei Jia, Binqiang Zhao, Xiaobo Wang, and Bin Cui. 2019. Buying or browsing?: Predicting real-time purchasing intent using attention-based deep network with multiple behavior. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1984–1992.
- [7] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*. 4171–4186.
- [8] Gary Marchionini. 2006. Exploratory search: from finding to understanding. *Commun. ACM* 49, 4 (2006), 41–46.
- [9] Alan Medlar, Joel Pyykkö, and Dorota Glowacka. 2017. Towards fine-grained adaptation of exploration/exploitation in information retrieval. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. 623–627.
- [10] Peng Mingjing, Liu Xinglin, Luo Ximing, Zhu Mingliang, Zhang Xianyong, Deng Xiangming, and Wu Mingfen. 2017. Recognizing intentions of E-commerce consumers based on ant colony optimization simulation. *Journal of Intelligent & Fuzzy Systems* 33, 5 (2017), 2687–2697.
- [11] Mohammed Mahmudur Rahman, Zulkifly Mohd Zaki, Najwa Hayaati Binti Mohd Alwi, and Md Monirul Islam. 2020. A conceptual model for the E-commerce application recommendation framework using exploratory search. *Solid State Phenomena* (2020).
- [12] Jennifer Rowley. 2000. Product search in e-shopping: a review and research propositions. *Journal of consumer marketing* (2000).
- [13] Tuukka Ruotsalo, Jaakko Peltonen, Manuel JA Eugster, Dorota Glowacka, Patrik Floréen, Petri Myllymäki, Giulio Jacucci, and Samuel Kaski. 2018. Interactive intent modeling for exploratory search. *ACM Transactions on Information Systems (TOIS)* 36, 4 (2018), 1–46.
- [14] Daniel Schellong, Jan Kemper, and Malte Brettel. 2016. Clickstream Data as a Source to Uncover con-Sumer Shopping Types in a Large-Scale Online Setting. In *24th European Conference on Information Systems, ECIS 2016, Istanbul, Turkey, June 12-15, 2016*. Research Paper 1. http://aisel.aisnet.org/ecis2016_rp/1
- [15] Daniel Schellong, Jan Kemper, and Malte Brettel. 2017. Generating Consumer Insights from Big Data Clickstream Information and the Link with Transaction-Related Shopping Behavior. (2017), 24. http://aisel.aisnet.org/ecis2017_rp/24
- [16] Parikshit Sondhi, Mohit Sharma, Pranam Kolari, and ChengXiang Zhai. 2018. A taxonomy of queries for e-commerce search. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 1245–1248.
- [17] Ning Su, Jiyin He, Yiqun Liu, Min Zhang, and Shaoping Ma. 2018. User intent, behaviour, and perceived satisfaction in product search. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*. 547–555.
- [18] Jaime Teevan, Christine Alvarado, Mark S Ackerman, and David R Karger. 2004. The perfect search engine is not enough: a study of orienteering behavior in directed search. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 415–422.
- [19] Manos Tsagkias, Tracy Holloway King, Surya Kallumadi, Vanessa Murdock, and Maarten de Rijke. 2021. Challenges and research opportunities in ecommerce search and recommendations. In *ACM SIGIR Forum*, Vol. 54. ACM New York, NY, USA, 1–23.
- [20] Zhongyuan Wang, Kejun Zhao, Haixun Wang, Xiaofeng Meng, and Ji-Rong Wen. 2015. Query understanding through knowledge-based conceptualization. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- [21] Melina Zavali, Ewelina Lacka, and Johannes De Smedt. 2021. Shopping hard or hardly shopping: Revealing consumer segments using clickstream data. *IEEE Transactions on Engineering Management* (2021).
- [22] Bichen Zheng and Bingwei Liu. 2018. A scalable purchase intention prediction system using extreme gradient boosting machines with browsing content entropy. In *2018 IEEE International Conference on Consumer Electronics (ICCE)*. IEEE, 1–4.