

ORCAS: 18 Million Clicked Query-Document Pairs for Analyzing Search

Nick Craswell nickcr@microsoft.com Microsoft Bellevue, USA Daniel Campos dacamp@microsoft.com Microsoft, University of Washington Bellevue, USA Bhaskar Mitra bmitra@microsoft.com Microsoft, University College London Montréal, Canada

Emine Yilmaz emine.yilmaz@ucl.ac.uk University College London London, UK

Bodo Billerbeck bodob@microsoft.com Microsoft Melbourne, Australia

ABSTRACT

Users of Web search engines reveal their information needs through queries and clicks, making click logs a useful asset for information retrieval. However, click logs have not been publicly released for academic use, because they can be too revealing of personally or commercially sensitive information. This paper describes a click data release related to the TREC Deep Learning Track document corpus. After aggregation and filtering, including a k-anonymity requirement, we find 1.4 million of the TREC DL URLs have 18 million connections to 10 million distinct queries. Our dataset of these queries and connections to TREC documents is of similar size to proprietary datasets used in previous papers on query mining and ranking. We perform some preliminary experiments using the click data to augment the TREC DL training data, offering by comparison: 28x more queries, with 49x more connections to 4.4x more URLs in the corpus. We present a description of the dataset's generation process, characteristics, use in ranking and other potential uses.

CCS CONCEPTS

• Information systems \rightarrow Web log analysis; Retrieval models and ranking; • Computing methodologies \rightarrow Neural networks.

KEYWORDS

user behavior data, web search, deep learning

ACM Reference Format:

Nick Craswell, Daniel Campos, Bhaskar Mitra, Emine Yilmaz, and Bodo Billerbeck. 2020. ORCAS: 18 Million Clicked Query-Document Pairs for Analyzing Search. In Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20), October 19–23, 2020, Virtual Event, Ireland. ACM, New York, NY, USA, 7 pages. https: //doi.org/10.1145/3340531.3412779

CIKM '20, October 19–23, 2020, Virtual Event, Ireland

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6859-9/20/10...\$15.00

https://doi.org/10.1145/3340531.3412779

1 INTRODUCTION

When people search the Web, they reveal their information needs through their queries and clicks. Such click logs become an important asset. For example, given the ambiguous query *panda*, click logs can tell us the relative popularity of different interpretations, such as animals, movies, songs, restaurants and technology products. We can ask: Which documents are clicked for the query *panda*? Which topics are most popular in panda-related queries? What topics are most popular in panda-related documents? Popularity can also vary by context, for example Australia's PANDA Helpline is more commonly clicked by Australian users, but less so by US users.

Not many click datasets are publicly available, for two reasons. One is privacy. Users enter sensitive and private information into a search engine, so a query can reveal information that should not be shared. By analyzing a stream of queries from the same user, we may see them search for their own name, address, or other details. Using these, and other hints, we can discover a lot about a person [3]. Serious problems are raised by the release of queries linked to user IDs, or even session IDs. Even without such links, it is dangerous to release a tail query that was only typed by one user.

The other barrier to sharing click data is the commercial value of the data. In a particular country or language, the search engine with the most search traffic has an advantage over smaller competitors, because it has more information about user needs. Sharing that data at scale would help competitors, including new entrants in the market, and potentially help search engine optimizers. It could also reveal information about the workings of the engine, for example if results lists were provided including rank positions of the clicked and unclicked results.

Given these barriers, one option is to provide anonymized click data, such as the Yandex search personalization challenge [22]. The 35 million search sessions in the dataset (a Kaggle competition with 194 teams) have anonymized user IDs, queries, query terms, URLs, URL domains and clicks. Providing these as numeric IDs rather than strings greatly reduces concerns about privacy and commercial value. It can't be used to build a competing search engine because it has URL IDs rather than URLs. It can't be used to identify a user who entered their name as a query, because we just have a query ID and some term IDs.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

The main disadvantage of anonymization via IDs is that we can not add new relevance judgments, or apply deep models that bring to bear text embeddings. It is also not possible to for instance discuss the meaning of the term *panda* because we do not know which termID maps to that term.

This paper describes a new dataset, the Open Resource for Click Analysis in Search (ORCAS). Rather than providing user-level or session-level information, which can be used for personal identification, we focus on query-document connections. We use queries that have been repeated across many users. ORCAS data has clicks for TREC documents only, and only for English-speaking users in the United States. Combined with the filtering and aggregation, the dataset is intended to provide a useful resource for ranking TREC documents and also Web search mining more generally, without releasing personally identifying or commercially valuable information.

2 RELATED WORK

There have been several previous attempts at releasing click datasets. The AOL data [21] and MSN data [28] were two of the initial attempts towards releasing click data, which came with question marks over their use.

The AOL dataset contained ~ 20 million queries together with the domains of associated clicked URLs from $\sim 659,000$ users, collected over a 3-month period from March to May 2006. It was found to allow personal identification of some individual people [3], which was a major setback for the release of any future similar datasets. As a result of this, the dataset was retracted and is no longer available.

The MSN data [28] contained ~ 15 million queries that were collected from the users of the MSN search engine from the United States. The dataset was released in 2006 and again in 2009, but the agreement required a limited-time use, usage has now timed out, and it has not been released again since 2009.

Some years later, as part of the Yandex search personalization challenge, Yandex released a dataset with \sim 35 million search sessions [22]. The dataset was created using search activity of users over 27 days as the training data, with a separate test data containing user activity over 3 days. As mentioned in the previous section, this dataset contains URL IDs as opposed to the actual URL itself, which limits the type of research that can be conducted with the dataset.

The Sogou dataset ¹, released by the major Chinese search engine company Sogou, contains ~ 25.1 million queries and ~ 43.5 million user interactions in the form of submitted queries and clicked Web page search results. Almost all queries in this dataset contain solely characters from the simplified Chinese alphabet – which has around 3,500 commonly used characters [26].

Since the aforementioned datasets have problems related to availability and use cases, a significant amount of research has been conducted on proprietary datasets instead. Proprietary datasets containing queries with associated clicked URLs have widely been used for ranking purposes [9, 13, 27], as well as for identifying related queries [1, 5, 15, 16, 25]. However, such proprietary datasets are not made available to the wider research community, which has been a major limitation for research. The TREC Deep Learning (DL) Track [8] attempted to address the need for large amounts of training data by releasing large scale training datasets that are based on human relevance assessments, derived from MS MARCO [2]. The track focuses on the document retrieval and passage retrieval tasks. In the document retrieval task, 367,013 queries together with associated relevance labels were made available, corresponding to a corpus of 3.2 million documents. The training datasets released as part of the track are sparse, with no negative labels and often only one positive label per query, analogous to some real-world training data such as click logs.

While the Deep Learning Track is a step forward towards making a large scale dataset that can be used for training information retrieval systems publicly available, the datasets released by the track are still based on human relevance assessments as opposed to real usage. The ORCAS dataset released as part of this paper can be seen as complementary to the data released as part of the Deep Learning Track since it is based on click logs.

Table 1 shows a comparison of different proprietary datasets, datasets released as part of the TREC DL Track, and the ORCAS dataset in terms of size (number of queries, URLs and query-URL pairs) and availability of the datasets, as well as the application areas for which the datasets have primarily been created for.

3 DATASET COLLECTION

The ORCAS dataset was created through log mining, aggregation and filtering.

Search engine logs can include information about the user's query, the ranking of URLs, the URLs that were clicked, as well as post-click behavior such as dwell time [10]. We focus on click events, since in aggregate these are a good positive signal for relevance. Before aggregation we apply some minimal filtering, to eliminate clicks that had very negative post-click signals, such as a short dwell, since this can indicate that the user was disappointed with the result.

We aggregate click events at the query-URL level, without considering context such as other URLs or their rank order. For the purpose of collating this dataset, we filter for users that are located in the United States and speak English. But we do not collect any other data about the users such as their search session or any other short term or long term personalizing information [6]. Such things could potentially be studied using the anonymized Yandex personalization challenge data [22], but not using our aggregation to the query-URL level. Furthermore, rather than counting how popular query-URL pairs are during aggregation, we simply note which query-URL pairs are present, to avoid revealing too much information about event popularity in our logs. We will describe in Section 4 how it is possible to recover some information about popularity, even though we have removed the per-pair popularity information.

The full set of query-URL data, aggregated based on a subsample of Bing's 26-month logs to January 2020, would still be too commercially valuable, potentially covering billions of queries. It could also potentially reveal information known only to one person. We apply several strict filters. First, keep only query-URL pairs where the URL is present in the 3.2 million document TREC DL corpus. We needed to choose some small subset of all clicked documents,

¹http://www.sogou.com/labs/

Table 1: Size comparison of query-URL pair datasets used in Web mining and ranking studies. We leave a blank if a size was
not provided, and use the filtered size in cases where filtering was applied before use. Our new dataset is larger than several
of the previous datasets, suggesting those past results could be replicated or extended using ORCAS.

Paper	Queries	URLs	Q-U pairs	Availability	Primary focus of paper
Beeferman and Berger [4]	244K	362K	1.9M	proprietary	Related Q
Wen et al. [25]	1M/week	Encarta	500K/week	proprietary	Related Q
Xue et al. [27]	862K	507K		proprietary	Ranking
Craswell and Szummer [9]	202K	505K	1.1M	proprietary	Ranking
Baeza-Yates and Tiberi [1]	7.5M	973K		proprietary	Related Q
Mei et al. [15]	637M	585M		proprietary	Related Q
Huang et al. [13]			100M	proprietary	Ranking
ORCAS	10.4M	1.4M	18.8M	open	Ranking
TREC DL [8]	367K	320K	384K	open	Ranking

Table 2: Summary of ORCAS data. Each record in the main file (orcas.tsv) indicates a click between a query (Q) and a URL (U), also listing a query ID (QID) and the corresponding TREC document ID (DID). For use in ranker training, we also provide files in TREC format for qrels, queries and runs. The run file is the top-100 using Indri query likelihood, for use as negative samples during training.

Filename	Size	Records	Data in each record
orcas.tsv orcas-doctrain-qrels.tsv orcas-doctrain-queries.tsv	1.76GB 410MB 322MB	18.8M 18.8M 10.4M	QID Q DID U QID DID QID Q
orcas-doctrain-top100	52.8GB	983M	QID DID score

so aligning with a TREC seems like a suitable choice. Second, we apply a k-anonymity filter, keeping only queries that were typed by k different users, for a high value of k. This makes it impossible for our dataset to contain a query with information that is only known to fewer than k users. Finally, we applied filters to remove potentially offensive queries, for example queries related to pornography or hate speech.

Two of our design decisions relating to ranking and position bias can now be described in more detail in relation to our *k*-anonymity requirement.

Firstly, although ORCAS includes clicked URLs, it does not provide the rankings of clicked and unclicked URLs that users saw. If released without an anonymity requirement, such ranking information could be too commercially valuable, revealing information about query popularity, ranking variation and user preferences. With a high enough k threshold, as used here, the data size would be reduced by orders of magnitude. The requirement is that the same ranking was seen by k users who all selected the same URL. Because rankings vary from user to user and also vary over time, very few cases reach the k threshold.

Secondly, although it would be possible to correct for position bias before aggregation, ORCAS does not do this. Correcting for position bias would not remove any URLs from the dataset, because results with k or more clicks tend to be relevant, independently of position. Instead, position bias correction could allow us to include lower-ranked documents in ORCAS, which had fewer clicks due to position but were equally relevant. However, for privacy reasons we can not include results clicked by fewer than k users. Therefore we focus on clicks in aggregate, uncorrected by position bias. This approach, of using uncorrected clicks as a positive signal, has proven useful in a variety of studies such as those in Table 1.

Our main dataset becomes 18.8 million records with four columns per record:

QID: 10103699 Q: why is the sky blue DID: D1968574

U: http://www.sciencemadesimple.com/sky_blue.html

The document ID (DID) is the same as the one used in the TREC corpus. To avoid revealing the unreleased query IDs in our held out test set, we assigned a disjoint set of query IDs (QIDs) to the ORCAS queries. This means the same query string can occur in TREC DL and ORCAS data, with different QIDs.

We also provide the same data in TREC format, as qrels, queries and Indri rankings for use in negative sampling (see Section 5). These comprise the full ORCAS data release², as described in Table 2.

Dataset examples

Figure 1 shows a sample of the ORCAS data, for some documents related to various meanings of the term *panda*. Some popular URLs are about Panda Express restaurants, the Python library Pandas, Pandas Syndrome, Giant Pandas or Red Pandas. Since URLs are too

²https://microsoft.github.io/TREC-2020-Deep-Learning/ORCAS



Figure 1: Example click data relating to the string 'panda'. Query nodes are on the left and URL nodes are on the right. A query-URL connection means an edge exists in ORCAS data. The topics of the URLs are: *food* Panda Express, *py* Pandas Python library, *med* Pandas syndrome, *giant* Giant pandas, and *red* Red pandas

large to fit in the figure, we label each URL node with the type of panda it covers: food, py, med, giant and red. For example, one of the nodes labeled 'giant' is the Wikipedia page https://en.wikipedia.org/ wiki/Giant_panda. The queries were selected for being high-degree nodes, both in the global corpus and for this set of documents. A query such as *pandas syndrome* is only connected to documents of one topic. Whereas a query like *panda* is ambiguous, connected to documents of multiple topics.



Figure 2: Degree distribution of query nodes and URL nodes in the bipartite query-URL graph.

Absence of an edge can indicate that the document is less relevant, for example the Python Pandas node is not adjacent to *panda*, because it is less likely that a user who typed that query wants Python Pandas. Edge absence can also indicate that the document or query is less popular, and also indicate specific patterns of retrieval in the underlying search engine. This may explain why the Red Panda URL is not connected to *panda* or *pandas*, despite being on-topic.

The figure covers several topics, but in the ORCAS dataset its possible to identify several more panda topics, such as Kung Fu Panda movies, the debut single of rapper Desiigner and Panda Antivirus. Outside the TREC documents and US click data used here there are even more meanings of panda, such as the Australian PANDA Helpline or the PandA gateway at Kyoto University.

To illustrate the mining of related queries, in the following paragraph we give an example based on the seed query *orcas* (QID=2126294). Related queries can be found by taking a two-step walk. The first step reaches eight URLs. The second step reaches 198 queries. We can rank the queries according to the volume of paths, assuming the queries that are reachable by more distinct paths are more related to the seed query. The most similar queries (with 4–5 paths) are about whales: *orca whale*, *orca facts*, *orca*, *orca whales* and *killer whale facts*. Somewhat less related (with 2 paths) are queries about Orcas Island: *orcas island*, *orcas island washington*, *orcas island wa* and *orcas island*.

4 DATASET ANALYSIS

This section describes characteristics of the dataset and compares it to other datasets in the literature.

Query-URL datasets have been used for improving document ranking, finding related queries, and other log mining applications, as was indicated in Table 1. For reasons of commercial sensitivity, ORCAS does not contain "edge weights", which could indicate the relative popularity of URLs for a query, or vice versa. However, it is still possible to recover some popularity information from OR-CAS. Specifically, queries that are more popular tend to accumulate connections to a greater number of URLs, and URLs that are commonly clicked tend to accumulate connections to a greater number of queries. Popularity is indicated by higher degree in the bipartite query-URL graph.

Figure 2 shows the degree distributions of nodes in the ORCAS bipartite graph. Skewed degree distributions, which are better viewed



Figure 3: ORCAS queries are shorter than TREC DL queries. ORCAS queries contain a variety of query terms, from the most common term *how* is in 387K queries, and 613K different terms that each occur in one query (for example *egyptlive*).

Table 3: The most common first word in TREC DL queries is *what*, in 39.4% of the queries (145K queries). The most common first word in ORCAS queries is *how*, in 3.7% of queries (387K queries). TREC DL queries were selected in a way that favored natural language questions, so have less variety of first words, and more focus on question words. ORCAS queries were selected based on the TREC DL documents, so still have question words, but also have words in the top-10 that are more rare in TREC DL such as *www*, *the*, *best* and *free*.

	-	TREC DL			ORCAS		
Word	Count	%	[Rank]	Count	%	[Rank]	
what	144575	39.4%	[1]	362169	3.5%	[2]	
how	50113	13.7%	[2]	386611	3.7%	[1]	
where	16091	4.4%	[3]	46331	0.4%	[6]	
who	12957	3.5%	[4]	38522	0.4%	[9]	
is	10742	2.9%	[5]	41512	0.4%	[8]	
when	10002	2.7%	[6]	37885	0.4%	[10]	
which	6138	1.7%	[7]	7347	0.1%	[89]	
can	5736	1.6%	[8]	20354	0.2%	[21]	
average	5389	1.5%	[9]	9906	0.1%	[52]	
define	4620	1.3%	[10]	27469	0.3%	[14]	
www	20	0.0%	[379]	138528	1.3%	[3]	
the	2289	0.6%	[15]	95685	0.9%	[4]	
best	203	0.1%	[58]	86119	0.8%	[5]	
free	16	0.0%	[487]	41514	0.4%	[7]	

on a log-log scale such that a few nodes are very highly connected and many nodes have only one connection, are common for datasets of this sort. Very popular queries (such as *weather*) have hundreds of URLs, while very popular URLs (such as www.outlook.com) have thousands of queries.

It may be useful to leverage this popularity information in log mining and ranking studies, despite the absence of edge weights. For example, studies of query autocomplete [23] have sometimes used the AOL and MSN logs mentioned in Section 2. Such studies need not only queries but also per-query popularity information. It would be possible to use the popularity information from Figure 2 to make the ORCAS queries usable in the same way.

ORCAS queries were selected based on connection to the TREC corpus, click aggregation, anonymity filtering and other filtering criteria. They were not otherwise selected to be natural language question answering queries. By contrast, the TREC DL queries were selected in the creation of the MS MARCO question answering task [2]. This would lead us to expect that ORCAS queries and TREC DL queries have somewhat different characteristics.

Two such characteristics are query length and vocabulary. Figure 3 shows that the ORCAS queries tend to have fewer words than TREC DL queries. Table 3 compares the first word of queries. TREC DL queries are much more likely to have the word *what* at the start of a query (39.4% vs 3.5%), although the ORCAS data has more *what* queries in total (362K vs 145K). Both querysets have a healthy distribution of common and rare terms overall, we illustrate this for ORCAS data (also Figure 3), but TREC DL has a similarly healthy vocabulary distribution. The differences in query length and the likelihood of having question words at the start mainly reflect the selection criteria of TREC DL queries, which are a subset of query traffic suitable for the MS MARCO question answering dataset, favoring natural language questions.

5 RETRIEVAL EXPERIMENTS

To study the effectiveness of the ORCAS data for training deep neural models, we conduct preliminary retrieval experiments on the document ranking task in the TREC 2019 Deep Learning Track [8]. We describe these experiments and corresponding results here.

Data. The TREC deep learning benchmark for document reranking provides a large training dataset containing 367,013 queries and 384,597 positively labeled query-document pairs from the MS MARCO dataset [2]. For every query, a set of 100 documents is retrieved using Indri [24] and provided as part of the dataset. Feature-based [14] and representation learning-based [17] learning to rank models are typically trained on this data by employing optimization objectives that contrast relevant and non-relevant documents for a given query. Nonrelevant documents can be sampled either from the collection distribution or from a distribution that is more biased towards documents with at least partial matches with the query—*e.g.*, the Indri top 100 retrieved results. Previous work [18] has indicated that negative documents sampled at random from the collection.

Following a similar design to the MS MARCO training dataset, we generate a complementary dataset containing 10,405,342 queries and 18,823,602 positively labeled query-document pairs. This is approximately 28 times bigger than the MS MARCO training dataset in terms of the number of queries and approximately 49 times bigger in terms of the number of positive labels. To be consistent with the MS MARCO training data, we also provide the Indri top 100 retrieved results for each query.

Model. We adopt a public implementation³ of a Transformerbased ranking model with query term independence [20] as our

³https://github.com/bmitra-msft/TREC-Deep-Learning-Quick-Start

base architecture. Both query and document terms are first encoded using a shared term embedding model. The document term embeddings are then contextualized using stacked transformer layers. We compute the cosine similarity between every pair of query and document term embeddings and then employ windowed kernlepooling [11, 12] and multiple feedforward layers to estimate the match between the document and each query term. Finally, the scores are linearly combined across query terms.

Experiments. We conduct two sets of experiments to study the usefulness of the ORCAS dataset (i) As a supplement to the MS MARCO training data, and (ii) As a document description field in addition to URL, title, and body.

For our first study, we compare the retrieval effectiveness of the Transformer-Kernel model [12] when trained on a combination of MS MARCO and ORCAS training data to training on MS MARCO data alone. When the two datasets are combined, a two-step sampling is employed for training sample selection-we first randomly select one of the two training datasets with equal probability and then randomly sample a query from the selected dataset with uniform probability. This means that our training model sees an equal proportion of samples from both datasets during training in spite of their significant difference in size. This is done intentionally to control against diverging too far from the MS MARCO query distribution, since our test queries also come from MS MARCO. We conduct two sets of experiments by sampling the negative document for training from (i) the collection distribution, and (ii) the Indri top-100 document distribution, respectively. For document representation we consider a maximum of the first 800 terms.

In our second study, we use the ORCAS dataset to generate an additional document field and compare the performance of the base model trained on MS MARCO labels with and without the ORCAS field. We restrict the document representation, across all fields, to 4000 terms and the ORCAS field, specifically, to a maximum of 2000 terms. To handle the longer document text input we employ the Conformer-Kernel [19] model as our base architecture for this study.

Across both studies, we evaluate our models on the 43 test queries from the 2019 edition of the track using the corresponding NIST labels provided as a reusable benchmark. We report MRR and NDCG for each run. The models are trained using the Adam optimizer and the RankNet objective [7] with a learning rate of 0.0001. All hyperparameters are consistent across different training runs.

Results. Table 4 summarizes the findings from the first study. Under both negative sampling settings, we find that training on the combination of the two datasets gives roughly equivalent results to training on MS MARCO data only. Although the ORCAS data has a higher number, the difference is not statistically significant on our 43 test queries.

The results from our second study is presented in Table 5. Similar to our first study we find that the inclusion of ORCAS dataset results in higher retrieval metrics but the difference is not statistically significant on our small test set.

Although these initial studies did not yield significantly better results, we posit that models with larger number of layers or learnable parameters could perform better, taking advantage of the larger size of the ORCAS dataset. Table 4: Retrieval experiment results on the document reranking task from the TREC Deep Learning benchmark to study the usefulness of the ORCAS dataset for traininig.

Training data	MRR	NDCG				
Negatives sampled from full collection						
MS MARCO only	0.798	0.505				
MS MARCO + ORCAS	0.807	0.509				
Negatives sampled from top 100 candidates						
MS MARCO only	0.909	0.574				
MS MARCO + ORCAS	0.924	0.582				

Table 5: Retrieval experiment results on the document fullranking task from the TREC Deep Learning benchmark to study the usefulness of the ORCAS dataset as an additional document field.

Document fields	MRR	NDCG
URL + Title + Body	0.902	0.616
URL + Title + ORCAS + Body	0.931	0.629

6 EVALUATING USING ORCAS LABELS

To analyze the viability of ORCAS query-URL pairs as positive relevance labels, we also used them in evaluation, to test the MRR of the 38 runs in the 2019 TREC DL document ranking task. For the 200 queries in the 2019 runs, we were able to identify 83 positive labels for 28 queries in the ORCAS dataset, whereas the official NIST labels were for 43 of the queries.

We then compared the evaluation results obtained using the subset from the ORCAS dataset with the evaluation results from the official test collection from TREC, using MRR as the evaluation metric. The comparison of the two evaluation results are shown in Figure 4. The figure reports also reports the Kendall's tau correlation between the rankings obtained using the two sets of metrics. Considering the small size of the ORCAS subset that is common with the TREC DL test collection, the correlations with the official MRR results look reasonable. These correlations would probably improve if all the queries in the ORCAS test data were to be used.

7 CONCLUSION

ORCAS is a click dataset with more than 10 million queries, connected to documents in the TREC Deep Learning Track document corpus. The queries can be released without privacy violation because they are not connected to any user ID or session ID, and only queries that were entered by many users are included. The dataset was also filtered, for example to exclude adult and offensive queries.

Compared to datasets used in previous log mining and ranking studies based on query-URL pairs, ORCAS is of comparable size to many of them, but is also an open dataset rather than a proprietary one. This could allow the replication and extension of a variety of past log mining studies and ranking studies.

We train a ranker using a mixture of ORCAS and TREC DL data, with our initial experiments showing some sign of promising



Figure 4: MRR values obtained when the runs submitted to the document ranking task of the TREC Deep Learning Track are evaluated with ORCAS labels vs. actual Qrel from the track.

results, but no statistically significant gains. Correctly weighting and adjusting the data to achieve a significant gain on a test set of 43 queries is left as future work, but may be possible given that the ORCAS training set is 49x larger than the TREC DL set. We also show that the ORCAS clicks can be used to calculate an MRR metric, evaluating TREC runs with a Tau correlation of 0.622 compared to MRR using TREC labels. This is further evidence that the ORCAS data can be used to augment labeled data, in training and evaluation.

The dataset could have several other uses. After mining to find related queries, the query pairs could futher be used for synonym mining, for use in query rewriting. Instead, the related queries could be issued separately, then a final result list constructed by blending the results. It would also be possible to mine the ORCAS query-document pairs to identify synonyms, viewing it as a translation problem, translating from query vocabulary to document vocabulary. Query-document connections can also be viewed as a kind of relevance feedback, for query expansion that relies on documents rather than synonym mining. The queries themselves could be used whenever a query histogram is useful, for simulating query autocompletion, or sampling queries for a future test collection. Since query-document connections are at the heart of information retrieval, we hope ORCAS proves a useful dataset for these and other future applications.

REFERENCES

- Ricardo Baeza-Yates and Alessandro Tiberi. 2007. Extracting Semantic Relations from Query Logs. In Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (San Jose, California, USA) (KDD '07). Association for Computing Machinery, New York, NY, USA, 76–85.
- [2] Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. 2016. MS MARCO: A Human Generated MAchine Reading COmprehension Dataset. arXiv preprint arXiv:1611.09268 (2016).
- [3] Michael Barbaro, Tom Zeller, and Saul Hansell. 2006. A face is exposed for AOL searcher no. 4417749. New York Times 9, 2008 (2006), 8.
- [4] Doug Beeferman and Adam Berger. 2000. Agglomerative Clustering of a Search Engine Query Log. In Proc. SIGKDD. Association for Computing Machinery, 407–416.

- [5] Doug Beeferman and Adam Berger. 2000. Agglomerative Clustering of a Search Engine Query Log. In Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Boston, Massachusetts, USA) (KDD '00). Association for Computing Machinery, New York, NY, USA, 407–416.
- [6] Paul Bennett, Ryen W. White, Wei Chu, Susan Dumais, Peter Bailey, Fedor Borisyuk, and Xiaoyuan Cui. 2012. Modeling the Impact of Short- and Long-Term Behavior on Search Personalization. In Proceedings of the 35th Annual ACM SIGIR Conference (SIGIR 2012) (proceedings of the 35th annual acm sigir conference (sigir 2012) ed.). ACM, New York, NY, USA.
- [7] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. 2005. Learning to rank using gradient descent. In Proceedings of the 22nd International Conference on Machine Learning. ACM, New York, NY, USA, 89–96.
- [8] Nick Craswell, Bhaskar Mitra, Emine Yilmaz, and Daniel Campos. 2020. Overview of the TREC 2019 deep learning track. In TREC.
- [9] Nick Craswell and Martin Szummer. 2007. Random Walks on the Click Graph. In Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (Amsterdam, The Netherlands) (SIGIR '07). Association for Computing Machinery, New York, NY, USA, 239–246.
- [10] Katja Hofmann, Lihong Li, and Filip Radlinski. 2016. Online evaluation for information retrieval. Foundations and trends in information retrieval 10, 1 (2016), 1–117.
- [11] Sebastian Hofstätter, Hamed Zamani, Bhaskar Mitra, Nick Craswell, and Allan Hanbury. 2020. Local Self-Attention over Long Text for Efficient Document Retrieval. In Proc. SIGIR. ACM.
- [12] Sebastian Hofstätter, Markus Zlabinger, and Allan Hanbury. 2019. TU Wien@ TREC Deep Learning'19–Simple Contextualization for Re-ranking. arXiv preprint arXiv:1912.01385 (2019).
- [13] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In Proc. CIKM. ACM, 2333–2338.
- [14] Tie-Yan Liu. 2009. Learning to Rank for Information Retrieval. Foundation and Trends in Information Retrieval 3, 3 (March 2009), 225–331.
- [15] Qiaozhu Mei, Dengyong Zhou, and Kenneth Church. 2008. Query Suggestion Using Hitting Time. In Proceedings of the 17th ACM Conference on Information and Knowledge Management (Napa Valley, California, USA) (CIKM '08). Association for Computing Machinery, New York, NY, USA, 469–478.
- [16] Bhaskar Mitra. 2015. Exploring Session Context using Distributed Representations of Queries and Reformulations. In *Proc. SIGIR*. ACM, 3–12.
- [17] Bhaskar Mitra and Nick Craswell. 2018. An introduction to neural information retrieval. Foundations and Trends® in Information Retrieval (2018).
- [18] Bhaskar Mitra, Fernando Diaz, and Nick Craswell. 2017. Learning to Match Using Local and Distributed Representations of Text for Web Search. In Proc. WWW. 1291–1299.
- [19] Bhaskar Mitra, Sebastian Hofstatter, Hamed Zamani, and Nick Craswell. 2020. Conformer-Kernel with Query Term Independence for Document Retrieval. arXiv preprint arXiv:2007.10434 (2020).
- [20] Bhaskar Mitra, Corby Rosset, David Hawking, Nick Craswell, Fernando Diaz, and Emine Yilmaz. 2019. Incorporating Query Term Independence Assumption for Efficient Retrieval and Ranking using Deep Neural Networks. arXiv preprint arXiv:1907.03693 (2019).
- [21] Greg Pass, Abdur Chowdhury, and Cayley Torgeson. 2006. A Picture of Search. In Proceedings of the 1st International Conference on Scalable Information Systems (Hong Kong) (InfoScale '06). Association for Computing Machinery, New York, NY, USA, 1–es.
- [22] Pavel Serdyukov, Georges Dupret, and Nick Craswell. 2014. Log-based personalization: The 4th web search click data (WSCD) workshop. In Proceedings of the 7th ACM international conference on Web search and data mining. 685–686.
- [23] Milad Shokouhi. 2013. Learning to personalize query auto-completion. In Proc. SIGIR. 103–112.
- [24] Trevor Strohman, Donald Metzler, Howard Turtle, and W Bruce Croft. 2005. Indri: A language model-based search engine for complex queries. In Proceedings of the International Conference on Intelligent Analysis, Vol. 2. Citeseer, 2–6.
- [25] Ji-Rong Wen, Jian-Yun Nie, and Hong-Jiang Zhang. 2001. Clustering User Queries of a Search Engine. In Proceedings of the 10th International Conference on World Wide Web (Hong Kong, Hong Kong) (WWW '01). Association for Computing Machinery, New York, NY, USA, 162–168.
- [26] Stewart Whiting and Joemon M. Jose. 2014. Recent and Robust Query Auto-Completion. In Proceedings of the 23rd International Conference on World Wide Web (Seoul, Korea) (WWW '14). Association for Computing Machinery, New York, NY, USA, 971–982.
- [27] Gui-Rong Xue, Hua-Jun Zeng, Zheng Chen, Yong Yu, Wei-Ying Ma, WenSi Xi, and WeiGuo Fan. 2004. Optimizing Web Search Using Web Click-through Data. In Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management (Washington, D.C., USA) (CIKM '04). Association for Computing Machinery, New York, NY, USA, 118–126.
- [28] Yuye Zhang and Alistair Moffat. 2006. Some Observations on User Search Behaviour. Austr. J. Intelligent Information Processing Systems 9, 2 (2006), 1–8.