

Mining Implicit Relevance Feedback from User Behavior for Web Question Answering

Linjun Shou[†]

lisho@microsoft.com
STCA NLP Group, Microsoft, Beijing

Shining Bo

shiningbo@bupt.edu.cn
School of Information and
Communication Engineering, BUPT

Feixiang Cheng

fecheng@microsoft.com
STCA NLP Group, Microsoft, Beijing

Ming Gong

migon@microsoft.com
STCA NLP Group, Microsoft, Beijing

Jian Pei[‡]

jpei@cs.sfu.ca
School of Computing Science, Simon
Fraser University

Daxin Jiang

djiang@microsoft.com
STCA NLP Group, Microsoft, Beijing

ABSTRACT

Training and refreshing a web-scale Question Answering (QA) system for a multi-lingual commercial search engine often requires a huge amount of training examples. One principled idea is to mine implicit relevance feedback from user behavior recorded in search engine logs. All previous works on mining implicit relevance feedback target at relevance of web documents rather than passages. Due to several unique characteristics of QA tasks, the existing user behavior models for web documents cannot be applied to infer passage relevance. In this paper, we make the first study to explore the correlation between user behavior and passage relevance, and propose a novel approach for mining training data for Web QA. We conduct extensive experiments on four test datasets and the results show our approach significantly improves the accuracy of passage ranking without extra human labeled data. In practice, this work has proved effective to substantially reduce the human labeling cost for the QA service in a global commercial search engine, especially for languages with low resources. Our techniques have been deployed in multi-language services.

CCS CONCEPTS

• Information systems → Web search engines; Question answering; • Computing methodologies → Learning from implicit feedback;

KEYWORDS

Question answering; User feedback; Weak supervision

[†]Correspondence author.

[‡]Jian Pei’s research is supported in part by the NSERC Discovery Grant program. All opinions, findings, conclusions and recommendations in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

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](https://permissions.acm.org).
KDD ’20, August 23–27, 2020, Virtual Event, CA, USA

© 2020 Association for Computing Machinery.
ACM ISBN 978-1-4503-7998-4/20/08...\$15.00
<https://doi.org/10.1145/3394486.3403343>

ACM Reference Format:

Linjun Shou, Shining Bo, Feixiang Cheng, Ming Gong, Jian Pei, and Daxin Jiang. 2020. Mining Implicit Relevance Feedback from User Behavior for Web Question Answering. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD ’20)*, August 23–27, 2020, Virtual Event, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3394486.3403343>

1 INTRODUCTION

Question answering (QA) has become a de facto feature in *search result pages* (SERP for short) in most commercial search engines. For a query bearing some question intent, such as a noun phrase like symptoms of coronavirus, a search engine can extract the most relevant passage from web documents and put it in an individual block at the top of a SERP. Figure 1 shows the screenshot of the QA feature of a commercial search engine, where the query is “normal temperature for children in Celsius”.

Typically, a QA block is composed of a question, the passage to answer the question, the URL of the source web document from which the passage is extracted, and the links to collect user feedback (e.g., “Was the QA block helpful?”). Clearly, a well-designed QA block can deliver informative answers to search engine users in an intuitive and straightforward manner, save user time, and improve user experience. QA blocks have become even more popular on mobile devices as voice search is adopted by more and more users.

No doubt the magic behind a QA block is empowered by various machine learning algorithms, including the latest deep neural networks [13, 15, 20, 24, 28, 40, 42]. While machine learning algorithms have attracted extensive attention, people often overlook a critical challenge in making QA blocks industry scale commercial products – the need of huge amounts of training data. In practice, a commercial search engine receives extremely diverse questions in open domain at web scale. To handle such a complex and huge question space, the QA models for search engines often have to involve tens of millions of parameters, which cause the models easily overfitting to small training data. Consequently, we usually have to use millions of training examples to train a model in order to overcome overfitting and biases.

It is well recognized that obtaining large amounts of high quality training data is a bottleneck for commercial search engines. Using human judges to label training data is very expensive in financial



Figure 1: Example QA features in web search engines.

cost and time. To make the challenge even tougher, a commercial search engine often provides services in global markets with various languages. It is unrealistic to manually label millions of training samples for each language.

A practical approach for search engines to collect massive training data for search tasks is to exploit implicit relevance feedback from user behavior mined from search logs. There exists a rich body of literature [1, 4, 14, 16, 19, 37]. Can we simply extend the existing best practice to collect user implicit relevance feedback data to train QA models? Unfortunately, all existing works target at the relevance of web *documents*, rather than *passages*. Collecting and understanding implicit relevance feedback for QA blocks is much more sophisticated and demands dramatic innovation beyond any existing approaches.

While we will develop our novel approach and present our best engineering practice later in the paper, let us illustrate some challenges in collecting and understanding implicit relevance feedback for QA block using a real example.

Figure 2 shows two QA examples. In the first case, “*What’s the normal body temperature for child?*”, there is no click in the QA block. In the second case, “*What’s the normal body temperature for adult?*”, user clicks on the URL in the QA block are observed. We further examine these two cases, and find that the passage in the first case perfectly answers the question. Therefore, a user can obtain satisfactory information by simply going through the content of the passage. No follow-up action is needed.

For the second case, the information in the passage (about child body temperature) does not accurately match the user intent (for adult body temperature). A user may have to explore more information in the source page from which the passage is extracted. In this case, the title “human body temperature” of the source page may trigger a user’s interest to click on the URL and read more in that page. This example illustrates a unique characteristic of QA. As the content of the passage is already presented to users in a QA block, users may not need to click on the URL to get the answer. Consequently, the correlation between user clicks and passage relevance may be much weaker than the correlation between user clicks and page relevance in web search results. We will provide more insights in Section 3.3.2.

Another major difference between QA block and web document results is the number of results presented in SERP. Given a user question, a search engine usually returns a list of web documents, but only a single QA block. Most previous click models leverage the relative rank order of documents to gain more reliable implicit

Question (a):	<i>What’s the normal body temperature for child?</i>
Passage (a):	<i>The average normal body temperature for children is about 37 degree. A child’s temperature usually averages from around 36.3 degree in the morning to 37.6 degree in the afternoon.</i>
URL (a):	<i>Human Body Temperature: Fever - Normal - Low https://www.disabled-world.com/calculators-charts/degrees.php</i>
Label:	<i>Relevant</i>
User behavior:	<i>No Click</i>
Question (b):	<i>What’s the normal body temperature for adult?</i>
Passage (b):	<i>The average normal body temperature for children is about 37 degree. A child’s temperature usually averages from around 36.3 degree in the morning to 37.6 degree in the afternoon.</i>
URL (b):	<i>Human Body Temperature: Fever - Normal - Low https://www.disabled-world.com/calculators-charts/degrees.php</i>
Label:	<i>Irrelevant</i>
User behavior:	<i>Click</i>

Figure 2: Examples of user behavior for web QA.

feedback. However, this idea cannot be used to QA blocks, since SERP contains only one QA block per question.

In this paper, we investigate user behavior in QA blocks and propose a novel approach to mine implicit relevance feedback from noisy user behavior data for passage relevance. To the best of our knowledge, this is the first systematic study submitted for publication to address the data collection challenges for QA blocks. We make the following contributions.

First, we capture three types of user behavior when users interact with QA blocks, namely *click behavior*, *re-query behavior*, and *browsing behavior*. By analyzing the aggregated sequences of user actions in the context of complete search sessions, we obtain interesting insights about the correlation between user behavior and passage relevance.

Second, we examine several possible methods that automatically extract user feedback signals from user behavior data. With a small amount of human labeled data as ground truth, we reveal strong correlation between extracted feedback signals and passage relevance, and further assess the feasibility of learning implicit feedback with reasonable accuracy.

Third, we incorporate implicit feedback mined from user behavior data into a weakly-supervised approach for QA model training, and carry out extensive experiments on several QA datasets in English. The experimental results clearly show our approach greatly improves the QA performance on all datasets, especially under the low-resource conditions.

Last, we deploy our approach in a commercial search engine in two non-English markets. We find users speaking different languages uniformly follow similar behavior patterns when they interact with QA blocks. Consequently, the implicit relevance feedback model trained in en-US (english) market can be successfully transferred to foreign markets without any tuning. In de-DE (German)

and fr-FR (French) markets, our approach significantly improves the QA service by around 3.0% in the AUC metric. Moreover, this approach can automatically refresh the QA model by continuously collecting relevance feedback from users, which further saves the labeling cost. We expect our approach can save millions of dollars of labeling cost when scaling out to more markets.

The rest of the paper is organized as follows. We first review the related work in Section 2. We then present our approach in Section 3. We report the extensive experimental results in Section 4, and conclude the paper in Section 5.

2 RELATED WORK

Our study is mainly related to the previous work on QA and learning from user feedback. We provide a brief review on those topics here.

2.1 Question Answering (QA)

The purpose of web QA is to offer users an efficient information access mechanism by directly presenting an answer passage to the web search engine users [2, 6, 8]. There are various methods for web QA in literature. For example, Moldovan *et al.* [26] proposed a window-based word scoring technique to rank potential answer pieces for web QA. Cui *et al.* [11] learned transformations of dependency path from questions to answers to improve passage ranking. Yao *et al.* [41] tried to fulfill the matching using minimal edit sequences between dependency parse trees. AskMSR [5], a search-engine based QA system, used Bayesian Neural Network relying on data redundancy to find short answers.

In recent years, deep neural networks have achieved excellent performance in QA [8, 34]. Convolutional Neural Networks (CNN) and Recurrent Neural Networks (RNN), such as the Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU), were applied to learn the representations of questions and answers [31, 32]. Attention mechanisms were employed to model the interaction between questions and answers [35, 38], which led to better performance than simply modeling query and answer separately. Most recently, deep pre-trained models, such as BERT [36], XLNet [39], have become the new state-of-the-art approaches of QA models.

To tackle web-scale open-domain question answering, the statistical machine learning models require large amounts of training data. In this paper, we do not target at developing another QA model. Instead, we aim to find a model-agnostic approach for training data collection.

2.2 Learning from User Feedback

User feedback has been intensively explored in web page ranking to improve search quality [17, 18]. There are two types of user feedback. *Explicit (or shallow) feedback* is that a user takes an extra effort to proactively express her satisfaction with the search results, e.g., through a simple up-voting or down-voting button. *Implicit feedback* is the inference of user satisfaction from a user’s search and/or browse sessions, without extra efforts from the users.

Rocchio [29] is a pioneer to leverage relevance feedback for information retrieval by explicitly gathering feedback through a button for up-voting or down-voting. Another means of collecting explicit feedback was through side-by-side comparisons [3, 33]. In practice, the chances of receiving explicit feedback from users

are very low, since explicit feedback disturbs users in their normal interaction with search engines.

Compared with explicit feedback, implicit feedback can be collected at much lower cost and in much larger quantity, without putting any burden on users of search systems [19]. Various features have been extracted from user behavior data, such as click-through information, average dwell time, and number of page visits in post-search browsing sessions [1, 4]. For example, Joachims *et al.* [19] derived relative preferences from click-through information. Agichtein *et al.* [1] explored page clicks and page visits as features to improve ordering of top results in web search. Gao *et al.* [?] and Huang *et al.* [16] used click-through data for deep semantic model training to learn semantic matching between queries and documents in page ranking. A major challenge in exploiting implicit feedback is that it is inherently noisy or even biased [18]. To address the challenge, various methods have been proposed. For example, Craswell *et al.* [10] proposed four simple hypotheses about how position bias may arise. Dupret and Piwowarski [12] proposed a set of assumptions on user browsing behavior in order to estimate the probability a document was to be viewed. Chapelle and Zhang [7] proposed a dynamic Bayesian Network model to indicate whether a user was satisfied with a clicked document and then left the page.

Although user feedback for web page ranking has been well studied, there is little work on user feedback for web QA. The closest work to our study is by Kratzwald and Feuerriegel [22], who designed feedback buttons to explicitly ask users to assess the overall quality of the QA result. Different from them [22], our work mainly focuses on the mining of implicit relevance feedback for web QA. To the best of our knowledge, it is the first study in this frontier.

3 OUR APPROACH

Our goal is to derive training data from online user behavior to train QA models. To achieve this goal, our basic idea is to learn an implicit relevance feedback model. Given a question q and the passage p served by the search engine, the feedback model extracts features from user behavior and predicts the relevance r between q and p . The predicted results (q, p, r) are then used as training data to train QA models.

Based on this idea, we first conduct a comprehensive analysis on user behavior in QA sessions in Section 3.1 and propose a systematic categorization to cover all types of user behavior. We then design a rich set of user behavior features in Section 3.2 to make sure we do not miss any useful implicit feedback signals. In Section 3.3, we carefully compare various algorithms that learn implicit feedback models, and apply the best model to a huge volume of user behavior data to derive a large amount of training data. Section 3.4 elaborates how we leverage the derived training data in a weakly-supervised approach for the training of QA model.

3.1 Taxonomy of User Behavior

We propose a taxonomy of user behavior summarized in Table 1. At the higher level, we distinguish two types of user behavior, which correspond to explicit and implicit feedback to web QA systems. In the following, we first show an empirical study of explicit feedback in a commercial search engine, and explain why it is not efficient

Table 1: Taxonomy of user behavior in web QA system

Type	Type	behavior
Explicit	Click	Up-vote/Down-vote
	Re-query	Reformulation
Implicit	Click	Answer Click
		Answer Expansion Click
		Outside Answer Click
		Related Click
	Browsing	Browse

or effective to collect training data through explicit feedback. We then describe user implicit feedback in detail.

To collect explicit feedback, commercial search engines, such as Google and Bing, provide links at the bottom of a QA block, as illustrated in Figure 1. However, only a very small fraction of users send their explicit feedback. In a real commercial web QA system, the coverage of explicit feedback, i.e., clicking on the feedback links, is only less than 0.001% of the total QA impressions. Moreover, we find that users strongly tend to send negative feedback – the positive over negative ratio is about 1:17. To form a balanced training dataset, we have to sample almost equal amounts of positive and negative examples from the skewed label distribution. This further reduces the size of valid training data that can be derived from explicit feedback. Consequently, the explicit feedback may not be a good source to collect training data for web QA model.

We then cast our attention to implicit feedback. Basically, all actions related to QA blocks recorded in search logs are either queries and clicks. Therefore, the first two categories in our taxonomy correspond to these two types of actions. We further refine the categories of actions related to QA blocks into sub-groups. For example, we distinguish the types of clicks based on the components that are clicked on. Finally, we also model the information about general SERP browsing that may be useful to QA blocks, which is the last category in our taxonomy. The details of the taxonomy are introduced in the following.

Re-query Behavior: we consider the sequence of user queries in a session and particularly note whether a user issues a new query by modifying the previous one in the session. Interchangeably we also call this behavior *reformulation*.

Click Behavior: we distinguish four types of clicks, depending on the components being clicked on (see Figure 3 for illustration).

- **Answer Click:** a user clicks on the source page URL of the answer passage (indicated by ① in Figure 3).
- **Answer Expansion Click:** a user clicks on a special button (indicated by ② in Figure 3) to expand the folded QA answer due to the maximum length limit for display.
- **Outside Answer Click:** a user clicks on the links to the web documents in the SERP (indicated by ③ in Figure 3) other than the source page URL for the web QA passage.
- **Related Click:** a user clicks on the related queries (indicated by ④ in Figure 3) to explore more information.

Browsing Behavior: a user reads the content of the QA passage or other components in the SERP, and does not give any input to the search engine.

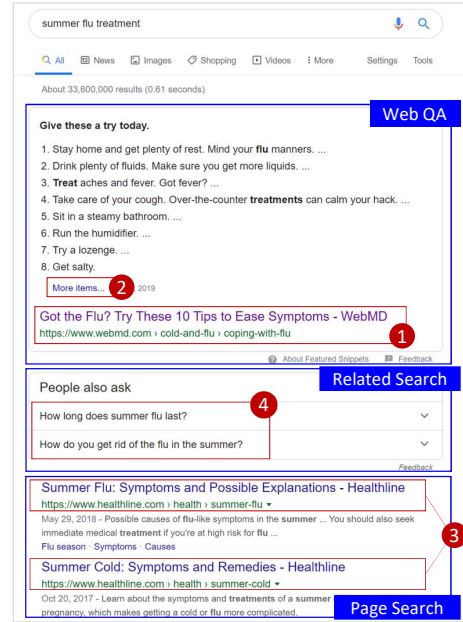


Figure 3: Illustration of user click behavior, including Answer Click, Answer Expansion Click, Outside Answer Click, and Related Click.

Table 2: User behavior features. Here “Answer” means the source URL for the passage in web QA.

Name	Type	Description
$RFRate$	Re-query	rate of re-query
$AnswerCTR$	Click	CTR of answer
$AnswerOnlyCTR$	Click	CTR with only click on answer
$AnswerSatCTR$	Click	satisfied CTR of answer
$AnswerExpRate$	Click	CTR of answer expansion
$OTAnswerCTR$	Click	CTR outside of answer
$OTAnswerOnlyCTR$	Click	CTR with only click outside of answer
$OTAnswerSatCTR$	Click	satisfied CTR outside of answer
$BothClickCTR$	Click	CTR of both click on/outside of answer
$RelatedClickRate$	Click	CTR of related queries
$NoClickRate$	Browsing	no click rate
$AbandonRate$	Browsing	abandonment rate
$AvgSourcePageDwellTime$	Browsing	average source page dwell time
$AvgSERPDwellTime$	Browsing	average SERP dwell time

3.2 Feature Extraction from User Behavior

To learn implicit feedback models, we need to design user behavior features, which should be sensitive and robust in capturing relevance feedback signals from users. To meet this goal, we follow two principles in our feature design. First, our feature set exhaustively covers all types of user behavior discussed in Section 3.1, and makes sure we do not miss any useful relevance signals. Second, we design aggregated features that summarize the behavior of multiple users in multiple sessions. Aggregated features can effectively reduce noise and biases in individual users and individual actions. The features are listed in Table 2. Most features are straightforward.

Please refer to the ‘‘Description’’ column for the meaning of the features. We only select several features to explain as follows.

- *AvgSourcePageDwellTime*: the average time from the user clicking into the source page of the web QA answer to the user leaving the source page.
- *AvgSERPDwellTime*: the average time from SERP loaded successfully to the completion of the search session.
- *AbandonRate*: the percentage of sessions with no click on the SERP. In those sessions, a user just browses SERP and leaves the search sessions.

The click-through rate (*CTR*) for a component, which can be either a QA passage, an answer expansion, a related search, or a web document outside of QA block, is defined as

$$CTR = \frac{N_{click}}{N_{impression}} \quad (1)$$

where $N_{impression}$ denotes the total number of impressions of the component and N_{click} the number of clicks on the component.

A satisfied click (*SatClick*) on a component is a click on the component followed by the corresponding Dwell Time greater than or equal to a pre-defined threshold. The satisfied click-through rate (*SatCTR*) is then defined as

$$SatCTR = \frac{N_{SatClick}}{N_{impression}} \quad (2)$$

where $N_{SatClick}$ denotes the number of *SatClicks* on the component.

3.3 Implicit Feedback Modeling

In this section, we target at building an effective implicit feedback model to predict the relevance between a question and a passage based on the features designed in Section 3.2. We first prepare a dataset with 18k QA pairs, where each QA pair is augmented with the set of user behavior features as well as a human judged label (Section 3.3.1). Intuitively, a single feature, such as *AnswerCTR* or *AnswerSatCTR*, may be too weak a signal to infer user satisfaction. We verify this assumption as a baseline in Section 3.3.2. We then consider various machine learning models, including Logistic Regression (LR) [25], Decision Tree (DT) [30], Random Forest (RT) [23], and Gradient Boost Decision Tree (GBDT) [21], and conduct an empirical study to choose the best model (Section 3.3.3). We also analyze the feature importance and derive rules from the learned models, which help us gain insights into users’ decision process when they interact with a QA block.

3.3.1 Dataset. We create a dataset which consists of 18k QA pairs, where each QA pair is augmented with the set of user behavior features as well as a human judged label. To be more specific, the dataset is a table, where each row is in the form of $\langle Question, Passage, User\ behavior\ features, Label \rangle$. The QA pairs are sampled from a half-year log from January to June 2019) of one commercial search engine. We sample the QA pairs whose number of impressions in the log is around 50. This number is set due to two considerations. First, we want to aggregate multiple users’ behavior to reduce the noise from individual users. Second, we want to avoid too popular queries, which tend to be short and too easy to answer. Each QA pair is sent to three crowd sourcing judges and the final label is derived




Figure 4: Precision-recall curves of *AnswerCTR*(a) and *AnswerSatCTR*(b).

based on a majority voting. This 18k dataset is further randomly split into 14k/2k/2k as training, dev and test sets, respectively. We plan to make this data set public if this paper is accepted.

3.3.2 Baseline. Click-through rate (*CTR*) and satisfied click-through rate (*SatCTR*) have been widely adopted in existing works as indicators for the relevance of a web page with respect to a given user query. Analogously in our study of passage relevance, we start with users’ clicks on the source URL of the answer passage. We first investigate the feature of *AnswerCTR* in Table 2 by plotting a precision-recall curve in Figure 4(a).

- For QA pairs whose $CTR > 0$, the recall is less than 0.33. In other words, when the passage is relevant to the question, there are more than two thirds of cases where users do not make any single click on the source URL. Please note that the number of clicks is counted throughout all the impressions for that question-passage pair. This is a very different observation compared to the cases for page ranking. However, when we consider the feature of question answering, this result is not surprising since users may simply browse the content of the passage and get the information. No further click into the source URL is needed at all.
- The highest precision is less than 0.77, across the full range of recall values. This suggests that clicking into the source URL does not necessarily suggest a relevant passage. We find in most clicked cases, the passage is partially relevant to the question. Therefore, users may want to click into the source URL to explore more information.

We further investigate the correlation between *SatCTR* and passage relevance. Similarly, we plot the precision-recall curves in Figure 4(b). CTR_t means we consider the clicks followed by dwell time on the source page for longer than t seconds as a satisfied click. We experiment with dwell time thresholds 5s, 15s, and 25s, and observe similar trend as in Figure 4(a).

The experiments with *CTR* and *SatCTR* verify that the single feature of user clicks into the source URL is not a good indicator for passage relevance. Clicks do not necessarily indicate relevant passages, and vice versa. Thus, we have to consider more complex models to combine the sequences of user actions in search sessions.

3.3.3 Machine Learning Models. We apply machine learning models to combine the various types of user behavior features. The training target is to fit the human judged labels. We evaluate the

Table 3: Comparison of feedback modeling methods. The best results are highlighted in bold, and the secondary best results are marked with underline.

Method	AUC	ACC	F1
<i>AnswerCTR</i>	58.28	53.87	18.11
<i>AnswerSatCTR</i> _{5s}	57.73 (-0.55)	53.50 (-0.37)	16.52 (-1.59)
<i>AnswerSatCTR</i> _{15s}	57.75 (-0.53)	52.73 (-1.14)	13.43 (-4.68)
<i>AnswerSatCTR</i> _{25s}	57.79 (-0.49)	52.63 (-1.24)	12.23 (-5.88)
LR	<u>71.41 (+13.13)</u>	<u>66.00 (+12.13)</u>	66.78 (+48.67)
DT	61.75 (+3.37)	63.43 (+9.56)	60.92 (+42.81)
RF	71.14 (+12.86)	<u>67.47 (+13.60)</u>	65.66 (+47.55)
GBDT	73.69 (+15.41)	68.00 (+14.13)	66.08 (+47.97)




Figure 5: Relative weights of top 8 features in the GBDT model.

model performance by common binary classification metrics, including area under the curve (AUC), accuracy (ACC), and F1 score. **Machine Learning Models Considered.** We apply various models and evaluate the results, including Logistic Regression (LR), Decision Tree (DT), Random Forest (RT), and Gradient Boost Decision Tree (GBDT).

Results. As summarized in Table 3, the machine learning approach significantly outperforms baseline methods (i.e. *AnswerCTR* and *AnswerSatCTR*) on all metrics. In terms of AUC and ACC, the GBDT model achieves the best performance. In terms of F1 score, the performance of GBDT model (66.08) is very close to the best result (66.78). Overall, we consider GBDT as the best model.

Model Interpretation. To get more insights about user behavior on QA, we first investigate the impact of individual feature based on the best model GBDT. The top 8 features of the model are shown in Figure 5.

The results indicate that *AvgSERPDwellTime* and *OTAnswerOnlyCTR* have the highest feature importance, followed by *AnswerOnlyCTR*, *AnswerSatCTR*, *AnswerExpRate* and *AbandonRate*. Reformulation related features like *RFRate* as well as *RelatedClickRate* have relatively low importance. We can also obtain the following insights about user behavior on web QA.

- Click features related to web QA answer itself, such as *AnswerOnlyCTR* and *AnswerSatCTR*, are not the most important features. This aligns well with our previous observations.
- SERP Dwell Time suggests the time period during which a user lingers on the search result page. Since the content of passage is presented in the QA block (in the SERP) as the answer to the user’s question, the length of SERP Dwell time may be a good indicator of the relevance of the passage.

To further reveal users’ decision process in a search session, we examine the decision tree model DT in Table 3. Some interesting insights gained from the paths on the tree are listed in the following.

- If “*AvgSERPDwellTime* is long \wedge *NoClickRate* is large”, i.e., the SERP is abandoned, the passage usually has good relevance. Users may just browse the QA passage for a while and then get the information needed.
- “*AnswerCTR* is small \wedge *OTAnswerCTR* is large” is often a strong signal that the passage has poor relevance. In such cases, users may not be satisfied with the passage answer, and then click more on other documents.
- “*AnswerOnlyCTR* is large \wedge *AvgSERPDwellTime* is long” is also a positive signal for passage relevance. The passage is relevant to the question, but due to the space limit of the QA block, the displayed content cannot fully answer user’s question. Therefore, the user clicks on the source URL.
- “*NoClickRate* is large \wedge *RelatedClickRate* is large” suggests that the passage is not relevant to the user question. The user revises the query to express her information need.

Summary: Unlike the cases of page relevance, user clicks (including satisfied clicks) are not a good indicator for passage relevance in web QA. However, using machine learning models to combine various user behavior, it is still feasible to extract relevance feedback from the search sessions.

3.4 Pre-training QA models

Through an implicit feedback model, we can derive a large amount of training data. However, user behavior data may be very noisy. Although we develop aggregated features to reduce biases in individual sessions, the prediction accuracy of the best model is only 68% (see Table 3). How to leverage such noisy training data becomes an interesting problem.

Inspired by the pre-training idea in learning deep neural networks, we apply large-scale automatically derived implicit feedback data to pre-train QA models, such as BiLSTM and BERT, as weak supervision. Intuitively, although the derived training data is noisy, it still contains valuable relevance signals and can roughly guide the parameters of QA models to a region close to the optimal solutions. In the second stage, we apply the human labeled data to fine-tune the parameters to get the final model. As verified in our experiments (Section 4), the strategy of pre-training plus fine-tuning is remarkably better than training models with only human labeled data without pre-training.

Technically, let $Q = \{w_1, w_2, w_3, \dots, w_{|Q|}\}$ be a question with $|Q|$ words (or word pieces), $P = \{w_1, w_2, w_3, \dots, w_{|P|}\}$ be a passage with $|P|$ words (or word pieces). We use cross-entropy (CE) as our loss function for both pre-training and fine-tuning, defined by

$$y = F_{QAModel}(\langle Q, P \rangle) \quad (3)$$

$$L_{CE} = -\frac{1}{n} \sum_i [\hat{y}_i \log(y_i) + (1 - \hat{y}_i) \log(1 - y_i)] \quad (4)$$

where $F_{QAModel}(\cdot)$ are the QA models, which will be described in Section 4.2, y_i represents our QA relevance model output, \hat{y}_i represents the true label, and n is the number of training samples.

Table 4: Statistics of experiment datasets.

Dataset	Train	Dev	Test	Labels
FeedbackQA _{log}	22M	-	-	-
FeedbackQA _{ctr, gbd}	4M	10k	10k	50%+/50%-
DeepQA	30k	2k	2k	57.6%/42.4%-
MS Marco	10k	1k	1k	50%+/50%-
WikiPassageQA	10k	1k	1k	50%+/50%-
FrenchGermanQA	50k	2k	2k	50%+/50%-

4 EXPERIMENTS

In this section, we report extensive experiments to verify our proposed approach using real data in a commercial search engine.

4.1 Datasets and Metrics

We conduct experiments on several datasets as follows, with their statistics shown in Table 4. More detailed description of the following datasets as well as one extra dataset is presented in Appendix A. AUC and ACC are used as the evaluation metrics for the QA relevance models.

FeedbackQA_{log}: An English QA dataset collected from the latest half year’s web QA system log in a commercial search engine. Each item of the log consists of a tuple \langle query, passage, user behavior \rangle .

FeedbackQA_{ctr, gbd}: For each QA pair in FeedbackQA_{log}, the feedback models in Table 3 are employed to predict a relevance label. The subscript indicates the model employed.

DeepQA: An English QA dataset where each case consists of three parts, i.e. question, passage, and a binary label (i.e. 0 or 1) by crowd sourcing human judges. The queries and passages are sampled from a commercial search engine.

MS Marco: An open source QA dataset [27], which contains questions generated from real anonymized Bing user queries. To evaluate the effectiveness of our approach under low-resource setting (i.e., only a small amount of labeled data are available), the dataset is sub-sampled to form a positive/negative balanced set with 10k/1k/1k as training, dev and testing sets, respectively.

WikiPassageQA: A Wikipedia based open source QA dataset [9], targeting non-factoid passage retrieval tasks. To evaluate our approach in low resource setting, the dataset is sub-sampled to form a positive/negative balanced set with 10k/1k/1k as training, dev and testing sets, respectively.

FrenchGermanQA: This dataset is collected in a similar process as DeepQA. The main difference is that this dataset targets at French and German languages.

4.2 Baselines and Models

To evaluate the effectiveness of our approach to mine implicit feedback, we set the following two baselines.

- **Original**: Only the human labeled data is used to train the QA model.
- **FBQA_{ctr}**: The FeedbackQA_{ctr} data is used for pre-training the QA model at the first stage. At the second stage, the QA model is further fine-tuned using the human labeled data.

In our approach, the best performing feedback model GBDT in Table 3 is used to auto-label large scale pre-training data (i.e. FeedbackQA_{gbd}) for pre-training the QA model. At the second

stage, the QA model is further fine-tuned using the human labeled data. This approach is referred to as **FBQA_{FA}** in the later experiment results.

We build the QA relevance models based on two popular deep neural networks, BiLSTM and BERT_{base}^{*}. The detailed description of these two models as well as the experimental setting are presented in Appendix B.

4.3 Results and Discussions

4.3.1 Overall Comparison Results. Table 5 shows the experimental results across all settings. We observe the following.

- Compared with the two baselines Original and FBQA_{ctr}, our implicit feedback approach FBQA_{FA} achieves significant improvements over different pre-training data size $\{0.5m, 1m, 4m\}$ and different QA fine-tuning data size $\{5k, 10k, 20k, 30k\}$. When the size of feedback pre-training data reaches 4m, our model gets the best results on the experiment set: for BiLSTM, there is about 5 AUC points increase on average; for BERT, 6 AUC points increase on average.
- Especially for low resource settings such as 5k and 10k QA fine-tuning data, our approach shows excellent results to save the labeling cost. Take BERT setting as an example. When the size of pre-training data equals to 4m and fine-tuning data equals to 5k, our model can get 76.53 AUC metric, which is even higher than the Original result on 30k fine-tuning data. In other words, with only 1/6 of the human labeled data, our model can still outperform the model trained on the full dataset. This experiment verifies the effectiveness of our approach to save large labeling cost.
- When we increase the size of implicit feedback pre-training data from 0.5m to 1m and 4m, our model is able to get consistent gains in all experimental settings. While for FBQA_{ctr}, the gains are not consistent: increasing pre-training data size does not necessarily improve the metrics, which aligns with our findings in Section 3.3.
- Our approach shows substantial gains over the baselines with both the BiLSTM and BERT models, which verifies the model agnostic characteristic of our approach. It is expected that BERT based QA models outperform BiLSTM based models since BERT benefits from large scale unsupervised pre-training as well as a large size of model parameters. It is interesting to find that even on top of the powerful deep pre-trained model such as BERT, further significant gains can be obtained. This demonstrates the huge potential of the inexpensive, abundant implicit feedback derived from large-scale user behavior data as complementary data sources to the expensive human labeled data of a relatively small size.

4.3.2 Effect of Pre-training Data Size. To further analyze the effect of user implicit feedback on improving web QA, we explore the model performance with respect to the size of the feedback data employed in the pre-training stage. The experiments are conducted on the DeepQA dataset using BERT_{base} models. Pre-training data size is set to $\{0, 1, 2, 3, 4, 5, 6\}$ millions. The results are shown in

^{*}Our goal is to verify the effectiveness of approach, so we do not use BERT_{large}, which is time and resource consuming.

Table 5: Performance comparison between our methods and baselines on the DeepQA dataset. All ACC and AUC metrics in the table are in percentage, where the sign % are omitted.

Model	Method	Pre-training Data Size	Performance on Different Fine-tuning Data Size (AUC/ACC)			
			5k	10k	20k	30k
BiLSTM	Original	-	60.45/58.21	61.30/59.92	61.55/61.99	62.40/61.74
		0.5m	59.90/57.60 (-0.55/-0.61)	61.25/58.25 (-0.05/-1.67)	61.40/60.50 (-0.15/-1.49)	60.65/59.29 (-1.75/-2.45)
	FBQA _{ctr}	1.0m	60.25/58.45 (-0.20/-0.24)	61.35/58.12 (+0.05/-1.80)	62.65/57.43 (+1.10/-4.56)	61.35/60.69 (-1.05/-1.05)
		4.0m	60.50/56.99 (+0.05/-1.22)	59.75/58.39 (-1.55/-1.53)	60.90/59.15 (-0.65/-2.84)	62.25/61.73 (-0.15/-0.01)
	FBQA _{FA}	0.5m	61.95/59.66 (+1.50/+1.45)	62.50/60.96 (+1.20/+1.04)	62.85/62.74 (+1.30/+0.75)	64.23/62.50 (+1.83/+0.76)
		1.0m	62.80/60.44 (+2.35/+2.23)	63.20/61.20 (+1.90/+1.28)	63.45/63.00 (+1.90/+1.01)	65.57/63.05 (+3.17/+1.31)
		4.0m	64.13/62.15 (+3.68/+3.94)	65.45/63.33 (+4.15/+3.41)	65.46/64.17 (+3.91/+2.18)	67.35/64.35 (+4.95/+2.61)
BERT	Original	-	69.31/64.86	71.81/67.76	72.47/67.07	75.28/68.26
		0.5m	67.35/62.76 (-1.96/-2.10)	72.96/66.66 (+1.15/-1.10)	75.11/68.26 (+2.64/+1.19)	77.76/71.07 (+2.48/+2.81)
	FBQA _{ctr}	1.0m	72.33/67.06 (+3.02/+2.20)	73.76/67.36 (+1.95/-0.40)	76.16/69.16 (+3.69/+2.09)	77.42/68.26 (+2.14/+0.00)
		4.0m	72.19/65.66 (+2.88/+2.90)	73.92/67.96 (+2.11/+0.20)	76.81/67.96 (+4.34/+0.89)	77.94/69.36 (+2.66/+1.10)
	FBQA _{FA}	0.5m	72.26/65.27 (+2.95/+0.41)	76.03/68.87 (+4.22/+1.11)	77.79/69.47 (+5.32/+2.40)	77.92/69.47 (+2.34/+1.21)
		1.0m	73.53/66.37 (+4.22/+1.51)	76.29/68.97 (+4.48/+1.15)	78.63/68.77 (+6.16/+1.70)	79.82/70.17 (+4.54/+1.91)
		4.0m	76.53/68.57 (+7.22/+3.71)	78.17/68.57 (+6.36/+0.81)	79.79/71.17 (+7.32/+4.10)	81.03/71.57 (+5.78/+3.31)




Figure 6: Performance on dataset DeepQA with different FeedbackQA pre-training data size.

Figure 6. By increasing the size of implicit feedback data in pre-training from 0 to 4 millions, the model performance improves accordingly. However, when the data size reaches a certain scale, e.g. 4 millions in our experiments, the AUC metric on test set slowly flattens out. This suggests that the noise in the implicit feedback data may limit the growth of the improvement.

4.3.3 Results on the MS Marco and WikiPassageQA datasets. We further apply our pre-trained model (trained on 4m FeedbackQA_{gdbt}) from the implicit relevance feedback data to two open benchmark QA datasets, MS Marco and WikiPassageQA. The results are reported in Table 6. We find that the queries in the MS Marco dataset are simpler than those in the DeepQA dataset, consequently, with 10k human labeled data and the BERT model, the AUC can reach as high as 94.01%. Our approach also shows improvement on this dataset, although the gain is not as large as that on the DeepQA dataset. For the BiLSTM model, the gain is around 2 points on average in terms of AUC over 2k, 5k, and 10k human labeled training data in the fine-tuning stage. For the BERT model, since the baseline is already very strong, the improvement is less than 1 point.

Table 6: Comparison Results on MS Marco and WikiPassageQA datasets (all AUC metrics are percentage numbers with % omitted).

Model	Method	MS Marco			WikiPassageQA		
		2k	5k	10k	2k	5k	10k
BiLSTM	Original	64.70	64.25	65.61	55.39	58.37	60.72
	FBQA _{ctr}	62.52	65.50	66.46	53.80	58.95	61.16
	FBQA _{FA}	65.65	66.24	68.66	57.23	60.38	62.30
BERT	Original	87.93	93.02	94.01	78.03	83.70	84.70
	FBQA _{ctr}	88.70	93.42	94.47	78.04	81.18	85.66
	FBQA _{FA}	88.75	94.02	94.81	80.10	84.14	86.16

On the WikiPassageQA dataset, our approach also shows gains, with around 1.3 points AUC gains for BERT and around 1.8 points AUC gains for BiLSTM on average over 2k, 5k and 10k labeled fine-tuning data.

4.4 Applications to non-English QA

We further apply our approach to several non-English markets in a commercial search engine. We find user behavior in different countries is very consistent. Consequently, the implicit relevance feedback model trained in en-US market can be successfully transferred to foreign markets without any tuning. The results are shown in Table 7.

In the de-DE (German) and the fr-FR (French) markets, our approach significantly improves the QA service in AUC metric, saving a huge amount of human labeling cost. Take fr-FR as example, our approach shows around 3.2 consistent AUC gains across all training data sizes. Meanwhile, the FBQA_{FA} model is able to get 76.43 AUC with only 5k training data while the Original model needs 30k training data to get similar results.

5 CONCLUSION AND FUTURE WORK

This paper proposes a novel framework of mining implicit relevance feedback from user behavior. The implicit feedback models are further applied to generate weak supervised data to train QA models.

Table 7: Results on French and German QnA.

Model	AUC of fr-FR & de-DE			
	5k	10k	30k	50k
Original	73.05/71.46	73.99/73.15	76.23/75.84	76.82/77.11
FBQA _{FA}	76.43/76.64	77.26/76.22	79.28/78.83	80.31/79.76

Our extensive experiments demonstrate the effectiveness of this approach in improving the performance of QA models and thus reducing the human labeling cost.

Mining implicit feedback from user behavior data for web QA task is an interesting area to explore. In this study, we mainly focus on users' search behavior. As future work, we may combine users' search behavior with browse behavior. Moreover, we may also conduct deeper analysis on the question types and compare the effectiveness of implicit feedback on different types of queries. Understanding when to trigger QA block from user feedback is another interesting problem. Finally, the application of our approach to more languages is also in our future plan.

REFERENCES

- [1] Eugene Agichtein, Eric Brill, and Susan T. Dumais. 2018. Improving Web Search Ranking by Incorporating User Behavior Information. *SIGIR Forum* 52, 2 (2018), 11–18.
- [2] David Ahn, Valentin Jijkoun, Gilad Mishne, Karin Müller, Maarten de Rijke, Stefan Schlobach, M Voorhees, and L Buckland. 2004. Using Wikipedia at the TREC QA Track. In *TREC*. Citeseer.
- [3] Kamal Ali and Chi Chao Chang. 2006. On the relationship between click rate and relevance for search engines. *WIT Transactions on Information and Communication Technologies* 37 (2006).
- [4] Mikhail Bilenko and Ryan W. White. 2008. Mining the Search Trails of Surfing Crowds: Identifying Relevant Websites from User Activity. In *WWW*. ACM, New York, NY, USA, 51–60.
- [5] Eric Brill, Susan T. Dumais, and Michele Banko. 2002. An Analysis of the AskMSR Question-Answering System. In *EMNLP*.
- [6] Davide Buscaldi and Paolo Rosso. 2006. Mining knowledge from wikipedia for the question answering task. In *Proceedings of the International Conference on Language Resources and Evaluation*. 727–730.
- [7] Olivier Chapelle and Ya Zhang. 2009. A dynamic bayesian network click model for web search ranking. In *WWW*. 1–10.
- [8] Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading wikipedia to answer open-domain questions. *arXiv preprint arXiv:1704.0051* (2017).
- [9] Daniel Cohen, Liu Yang, and W. Bruce Croft. 2018. WikiPassageQA: A Benchmark Collection for Research on Non-factoid Answer Passage Retrieval. In *SIGIR*. 1165–1168.
- [10] Nick Craswell, Onno Zoeter, Michael J. Taylor, and Bill Ramsey. 2008. An experimental comparison of click position-bias models. In *WSDM*. 87–94.
- [11] Hang Cui, Renxu Sun, Keya Li, Min-Yen Kan, and Tat-Seng Chua. 2005. Question answering passage retrieval using dependency relations. In *SIGIR*. 400–407.
- [12] Georges Dupret and Benjamin Piwowarski. 2008. A user browsing model to predict search engine click data from past observations. In *SIGIR*. 331–338.
- [13] Abdessamad Echihabi, Ulf Hermjakob, Eduard Hovy, Daniel Marcu, Eric Melz, and Deepak Ravichandran. 2008. How to Select an Answer String? In *Advances in open domain question answering*. Springer, 383–406.
- [14] Jianfeng Gao, Kristina Toutanova, and Wen tau Yih. 2011. Clickthrough-based latent semantic models for web search. In *SIGIR '11*.
- [15] Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. 2019. Unicoder: A Universal Language Encoder by Pre-training with Multiple Cross-lingual Tasks. In *EMNLP/IJCNLP*.
- [16] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry P. Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *CIKM*. 2333–2338.
- [17] Thorsten Joachims. 2002. Optimizing search engines using clickthrough data. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, July 23–26, 2002, Edmonton, Alberta, Canada. 133–142.
- [18] Thorsten Joachims, Laura A. Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately interpreting clickthrough data as implicit feedback. In *SIGIR 2005: Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, Salvador, Brazil, August 15–19, 2005. 154–161.
- [19] Thorsten Joachims, Laura A. Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2017. Accurately Interpreting Clickthrough Data as Implicit Feedback. *SIGIR Forum* 51, 1 (2017), 4–11. <https://doi.org/10.1145/3130332.3130334>
- [20] Michael Kaisser and Tilman Becker. 2004. Question Answering by Searching Large Corpora With Linguistic Methods. In *TREC*.
- [21] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. In *Advances in Neural Information Processing Systems*. 3146–3154.
- [22] Bernhard Kratzwald and Stefan Feuerriegel. 2019. Learning from On-Line User Feedback in Neural Question Answering on the Web. In *The World Wide Web Conference*. ACM, 906–916.
- [23] Andy Liaw, Matthew Wiener, et al. 2002. Classification and regression by randomForest. *R news* 2, 3 (2002), 18–22.
- [24] Yosi Mass, Haggai Roitman, Shai Erera, Or Rivlin, Bar Weiner, and David Konopnicki. 2019. A Study of BERT for Non-Factoid Question-Answering under Passage Length Constraints. *CoRR abs/1908.06780* (2019). arXiv:1908.06780
- [25] Scott Menard. 2002. *Applied logistic regression analysis*. Vol. 106. Sage.
- [26] Dan I. Moldovan, Sanda M. Harabagiu, Marius Pasca, Rada Mihalcea, Roxana Girju, Richard Goodrum, and Vasile Rus. 2000. The Structure and Performance of an Open-Domain Question Answering System. In *ACL*.
- [27] Tri Nguyen, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng. 2016. MS MARCO: A Human Generated Machine Reading Comprehension Dataset. In *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016)*, Barcelona, Spain, December 9, 2016.
- [28] Dragomir Radev, Weiguo Fan, Hong Qi, Harris Wu, and Amardeep Grewal. 2002. Probabilistic question answering on the web. In *WWW*. ACM, 408–419.
- [29] Joseph Rocchio. 1971. Relevance feedback in information retrieval. *The Smart retrieval system-experiments in automatic document processing* (1971), 313–323.
- [30] S Rasoul Safavian and David Landgrebe. 1991. A survey of decision tree classifier methodology. *IEEE transactions on systems, man, and cybernetics* 21, 3 (1991), 660–674.
- [31] Ming Tan, Cicero Nogueira dos Santos, Bing Xiang, and Bowen Zhou. 2016. Improved Representation Learning for Question Answer Matching. In *ACL*.
- [32] Ming Tan, Bing Xiang, and Bowen Zhou. 2015. LSTM-based Deep Learning Models for non-factoid answer selection. *CoRR abs/1511.04108* (2015). arXiv:1511.04108
- [33] Paul Thomas and David Hawking. 2006. Evaluation by comparing result sets in context. In *Proceedings of the 15th ACM international conference on Information and knowledge management*. ACM, 94–101.
- [34] Shuohang Wang, Mo Yu, Xiaoxiao Guo, Zhiguo Wang, Tim Klinger, Wei Zhang, Shiyu Chang, Gerry Tesaro, Bowen Zhou, and Jing Jiang. 2018. R³: Reinforced Ranker-Reader for Open-Domain Question Answering. In *AAAI, New Orleans, Louisiana, USA, February 2-7, 2018*. 5981–5988.
- [35] Zhiguo Wang, Wael Hamza, and Radu Florian. 2017. Bilateral Multi-Perspective Matching for Natural Language Sentences. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*. 4144–4150.
- [36] Zhiguo Wang, Patrick Ng, Xiaofei Ma, Ramesh Nallapati, and Bing Xiang. 2019. Multi-passage BERT: A Globally Normalized BERT Model for Open-domain Question Answering. *CoRR abs/1908.08167* (2019). arXiv:1908.08167
- [37] Ryan W. White, Mikhail Bilenko, and Silviu Cucerzan. 2007. Studying the Use of Popular Destinations to Enhance Web Search Interaction. In *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '07)*. ACM, New York, NY, USA, 159–166.
- [38] Liu Yang, Qingyao Ai, Jiafeng Guo, and W. Bruce Croft. 2016. aNMM: Ranking Short Answer Texts with Attention-Based Neural Matching Model. In *CIKM, Indianapolis, IN, USA, October 24-28, 2016*. 287–296.
- [39] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. XLNet: Generalized Autoregressive Pretraining for Language Understanding. *CoRR abs/1906.08237* (2019). arXiv:1906.08237
- [40] Ze Yang, Linjun Shou, Ming Gong, Wutao Lin, and Daxin Jiang. 2020. Model Compression with Two-stage Multi-teacher Knowledge Distillation for Web Question Answering System. *Proceedings of the 13th International Conference on Web Search and Data Mining* (2020).
- [41] Xuchen Yao, Benjamin Van Durme, Chris Callison-Burch, and Peter Clark. 2013. Answer Extraction as Sequence Tagging with Tree Edit Distance. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 9-14, 2013, Westin Peachtree Plaza Hotel, Atlanta, Georgia, USA*. 858–867.
- [42] Fei Yuan, Linjun Shou, Xuanyu Bai, Ming Gong, Yaobo Liang, Nan Duan, Yan Fu, and Daxin Jiang. 2020. Enhancing Answer Boundary Detection for Multilingual Machine Reading Comprehension. *ArXiv abs/2004.14069* (2020).

A MORE DESCRIPTION OF DATA SETS

In the following, we present more details about the datasets in Section 4.1. We add one extra dataset called DeepQA_{factoid}. For clarity, the original data DeepQA is renamed as DeepQA_{general}. The statistics are summarized in Table 8.

- **DeepQA_{general}**: An English QA dataset where each case consists of three parts, i.e. question, passage, and a binary label (i.e. 0 or 1) by crowd sourcing human judges. The data collection process is as follows. First, for each question Q , all the passages P extracted from the top 10 relevant web documents returned by the search engine are collected to form a candidate set of $\langle Q, P \rangle$ pairs. Next, each $\langle Q, P \rangle$ pair is sent to three crowd sourcing judges, and a label (1 for relevance and 0 otherwise) is derived based on a majority voting.
- **DeepQA_{factoid}**: This dataset is collected in a similar process as DeepQA_{general}. The main difference is that the queries in DeepQA_{factoid} are mainly factoid queries, i.e., queries asking about “what”, “who”, “where”, and “when”.
- **MS Marco**: An open source QA dataset [27], which contains questions generated from real anonymized Bing user queries. Each question is associated with multiple passages extracted from the Bing web search results. Well trained judges read the question and its related passages, if there is an answer present, the supporting passages are annotated as relevant, while others are labeled as irrelevant. To evaluate the effectiveness of our approach under low-resource setting (i.e., only a small amount of labeled data are available), the dataset is further sub-sampled to form a positive/negative balanced set with 10k/1k/1k as training, dev and testing sets, respectively.
- **WikiPassageQA**: An open source QA dataset [9] based on Wikipedia, targeting non-factoid passage retrieval tasks. It contains thousands of questions with annotated answers. Each question is associated with multiple passages in the documents. In order to obtain a dataset of low resource setting, the dataset is further sub-sampled to form a positive/negative balanced set with 10k/1k/1k as training, dev and testing sets, respectively.

Table 8: Statistics of experiment datasets.

Dataset	Train	Dev	Test	Labels
DeepQA _{factoid}	30k	2k	2k	55.7%/44.3%-
DeepQA _{general}	30k	2k	2k	57.6%/42.4%-
MS Marco	10k	1k	1k	50%/50%-
WikiPassageQA	10k	1k	1k	50%/50%-

B MORE DETAILS ON EXPERIMENTAL SETTING

We build the QA relevance models based on the following two popular deep neural networks.

- **BiLSTM**. It consists of three parts. The first is an embedding layer, which maps each token to a vector with a fixed dimension. The second is multi-layer Bidirectional LSTM, which

encodes both questions and passages based on the token embeddings, i.e. $H^q = BiLSTM_1(Q)$ and $H^p = BiLSTM_2(P)$, where H^q and H^p are the representations for a question and a passage respectively. The parameters of $BiLSTM_1$ and $BiLSTM_2$ are not shared. Following the BiLSTM layer is a prediction layer, which includes a combination layer to concatenate H^q and H^p , and then a fully connection layer to predict the relevance of the passage and the question.

- **BERT_{base}**: It contains 12 bidirectional transformer encoders. We concatenate the question text and the passage text together as a single input of the BERT encoder. We then feed the final hidden state of the first token ($[CLS]$ token embedding) from the input into a two-layer feed-forward neural network. The final output is the relevance score between the input question and the passage. In all cases, the hidden size is set to 768. The number of self-attention heads is set to 12, and the feed-forward filter size is set to 3072.

For the experiments with BiLSTM, we use two BiLSTM layers with the hidden size of BiLSTM cell set as 128. We set the maximum length of questions and maximum length of passages to be 30 and 200, respectively, for data sets DeepQA, MS Marco, WikiPassageQA and FrenchGermanQA. We set the batch size to 256 and the dimension of word embedding to 300. During pre-training, the learning rate is set to $\{1e^{-5}, 3e^{-5}, 5e^{-5}\}$, the dropout rate to $\{0.1, 0.3, 0.5\}$ and the max epoch to 50. We choose the best model for fine-tuning based on the evaluation metric on the dev set. During fine-tuning, the learning rate is set to $\{1e^{-5}, 3e^{-5}, 5e^{-5}\}$, the dropout rate to 0.5 and the max epoch to 50. The best model is selected based on the evaluation metric on the dev set as well.

For the experiments with BERT_{base}, we use the huggingface version pre-trained BERT_{base} model[†]. We set the maximum sequence length to 200 for datasets DeepQA, MS Marco, WikiPassageQA and FrenchGermanQA. We set the batch size to 128, the number of gradient accumulation steps to 2 and the learning rate warmup ratio to 0.1. During pre-training, the learning rate is set to $\{1e^{-5}, 3e^{-5}, 5e^{-5}\}$ and the max epoch to 3. We choose the best model for fine-tuning based on the evaluation metric on the dev set. During fine-tuning, the learning rate is set to $\{1e^{-5}, 3e^{-5}, 5e^{-5}\}$, and the max epoch to 3. The best model is selected based on the evaluation metric on the dev set as well.

C MORE EXPERIMENTAL RESULTS AND DISCUSSIONS

We conduct further experiments on DeepQA_{factoid} dataset. The results are shown in Table 9. Similar to DeepQA_{general} dataset, our implicit feedback approach FBQA_{FA} achieves significant improvements over different pre-training data size $\{0.5m, 1m, 4m\}$ and different QA fine-tuning data size $\{5k, 10k, 20k, 30k\}$. For comparison, click based FBQA_{ctr} model is not able to get consistent gains compared with Original baseline, which proves that click does not necessarily suggest relevance in terms of QA task.

The results further demonstrate the effectiveness of our proposed implicit feedback QA approach by leveraging large-scale user behavior data as complementary data sources to save the human labeling cost.

[†]<https://github.com/huggingface/pytorch-transformers>

Table 9: Results on DeepQA_{factoid}. All ACC and AUC metrics in the table are in percentage , where the sign % are omitted.

Model	Method	Pre-training Data Size	Performance on Different Fine-tuning Data Size (AUC/ACC)			
			5k	10k	20k	30k
BiLSTM	Original	-	64.02/60.80	64.73/59.00	65.18/60.80	64.03/58.95
		0.5m	63.62/60.35 (-0.40/-0.45)	65.51/59.00 (+0.78/+0.00)	64.56/60.40 (-1.24/-0.40)	61.79/58.05 (-2.24/-0.90)
	FBQA _{ctr}	1.0m	63.67/60.30 (-0.35/-0.50)	64.86/58.65 (+0.15/-0.35)	65.53/60.40 (+0.35/-0.40)	61.35/57.40 (-2.68/-1.55)
		4.0m	62.83/60.20 (-1.19/-0.60)	65.09/59.05 (+0.36/+0.05)	64.67/59.70 (-0.51/-1.10)	65.24/60.95 (+1.21/+2.00)
	FBQA _{FA}	0.5m	65.66/62.05 (+1.64/+1.25)	67.87/63.00 (+3.14/+4.00)	68.84/63.80 (+3.66/+3.00)	69.80/64.75 (+5.77/+5.80)
		1.0m	66.36/64.00 (+2.34/+3.2)	67.81/63.45 (+3.08/+4.45)	69.23/64.70 (+4.05/+3.90)	72.16/66.40 (+8.13/+7.45)
		4.0m	67.22/65.25 (+3.20/+4.45)	70.26/65.10 (+5.53/+6.10)	70.79/64.95 (+5.61/+4.15)	73.31/67.15 (+9.28/+8.20)
BERT	Original	-	68.88/65.46	71.43/66.77	73.87/67.06	78.37/69.56
		0.5m	71.44/66.27 (+2.56/+0.81)	75.17/67.06 (+3.74/+0.29)	77.78/70.78 (+3.91/+3.72)	80.00/72.87 (+1.63/+3.31)
	FBQA _{ctr}	1.0m	71.73/66.17 (+2.85/+0.71)	74.19/67.26 (+2.76/+0.49)	77.01/69.46 (+3.14/+2.40)	78.30/72.07 (-0.07/+2.51)
		4.0m	70.55/66.06 (+1.67/+0.60)	75.51/68.26 (+4.08/+1.49)	78.10/69.77 (+4.23/+2.71)	79.54/70.97 (+1.17/+1.41)
	FBQA _{FA}	0.5m	74.00/67.77 (+5.12/+2.31)	76.57/68.87 (+5.14/+2.10)	78.78/70.67 (+4.91/+3.61)	80.91/72.37 (+2.54/+2.81)
		1.0m	74.33/68.67 (+5.45/+3.21)	75.77/68.77 (+4.34/+2.00)	78.40/71.17 (+4.53/+4.11)	79.31/71.77 (+0.94/+2.21)
		4.0m	76.13/69.77 (+7.25/+4.31)	77.93/71.27 (+6.50/+4.50)	79.94/72.77 (+6.07/+5.71)	81.21/73.27 (+2.84/+3.71)