

Generate-then-Retrieve: Intent-Aware FAQ Retrieval in Product Search

Zhiyu Chen Jason Choi Besnik Fetahu Oleg Rokhlenko Shervin Malmasi

Amazon.com, Inc. Seattle, WA, USA

{zhiyu, chojson, besnikf, olegro, malmasi}@amazon.com

Abstract

Customers interacting with product search engines are increasingly formulating information-seeking queries. Frequently Asked Question (FAQ) retrieval aims to retrieve common question-answer pairs for a user query with question intent. Integrating FAQ retrieval in product search can not only empower users to make more informed purchase decisions, but also enhance user retention through efficient post-purchase support. Determining when an FAQ entry can satisfy a user’s information need within product search, without disrupting their shopping experience, represents an important challenge. We propose an intent-aware FAQ retrieval system consisting of (1) an intent classifier that predicts when a user’s information need can be answered by an FAQ; (2) a reformulation model that rewrites a query into a natural question. Offline evaluation demonstrates that our approach improves Hit@1 by 13% on retrieving ground-truth FAQs, while reducing latency by 95% compared to baseline systems. These improvements are further validated by real user feedback, where 71% of displayed FAQs on top of product search results received explicit positive user feedback. Overall, our findings show promising directions for integrating FAQ retrieval into product search at scale.

1 Introduction

Product search engines help users find relevant products across large product catalogues and generate sales revenue for e-commerce companies (Grover and Teng, 2001). While such engines are primarily designed to handle keyword searches for products, customer behavior has been changing with an increase in users asking information-seeking service or product related questions (Carmel et al., 2018; Gao et al., 2019). However, most product search engines are not effective at handling non-product search related queries (e.g., “return a package”). Providing correct answers to these Frequently Asked Questions

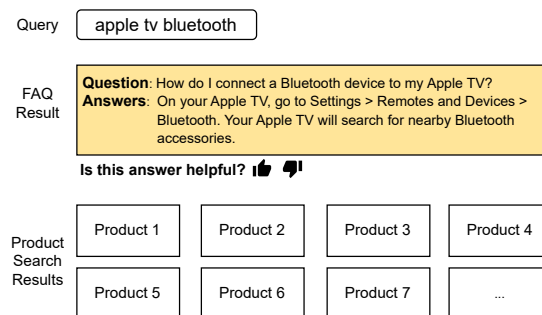


Figure 1: Our proposed aggregated search interface that jointly displays top-1 FAQ and product search results for queries with question intent.

(FAQs) (Gupta and Carvalho, 2019; Mass et al., 2020) is essential to provide a positive pre- and post-purchase experience, which can lead to improved user retention and trust.

Product search and FAQ retrieval are typically powered by independent retrieval systems. This separation is often due to the challenges in combining multiple answering sources (e.g. product details and FAQs) into a holistic retrieval application (Park et al., 2015; Christmann et al., 2022). Furthermore, determining what answering source can satisfy the user’s information need is challenging to perform at scale.

Hence, e-commerce websites tend to isolate FAQ search functionality from product search. For example, Apple offers vertical search where users are required to navigate among different tabs (e.g., product search, support, store location).¹ Such designs require users to navigate multiple links, which can lead to increased user effort and unsatisfactory shopping experiences (Siraj et al., 2020; Nain and Awasthi, 2021; Su et al., 2018). Therefore, we propose to integrate FAQ retrieval into a product search engine, so that users can search products and access FAQs seamlessly from a unified search interface.

¹<https://www.apple.com/us/search/apple-tv-bluetooth>

A potential solution is *aggregated search*, which refers to the task of searching and assembling information from a variety of sources and presenting them in a single unified interface (Murdock and Lalmas, 2008; Wan et al., 2008; Lu et al., 2012). The main challenge here lies in determining **when** and **how** information from multiple *verticals* should be presented effectively and efficiently.

When to Show FAQ Results? Query intent is inherently ambiguous (Krovetz and Croft, 1992; Song et al., 2007; Sanderson, 2008). Figure 1 illustrates an example where the users can use the same query “*apple tv bluetooth*” to retrieve products, or find information about Bluetooth connectivity. In the latter case, the query is intended to express the question, “*How do I connect a Bluetooth device to my Apple TV?*”. We define a query that can be answered by an FAQ entry as having *question intent*. It is important to note that a query with question intent may also have product search intent, as queries can be inherently ambiguous.

Determining when a user query can be answered by either *product* or *FAQ search* is tightly coupled with predicting the user’s information need. Due to query ambiguity, displaying answers from FAQ sources for all searches causes high false positive rates, due to the lack of question intent by the user.

Our analysis (cf. §5.1) from a leading e-commerce site shows that if FAQs would be shown to all queries, 98% of the FAQs would be irrelevant to user’s needs. As users mostly use product search for shopping, injecting FAQ results that are irrelevant or not needed causes significant friction in the user experience. Furthermore, performing FAQ retrieval for every query is inefficient (Tsur et al., 2016) since only a small portion of traffic has question intent (White et al., 2015).² While we cannot disclose the intent distribution of our data for reasons of confidentiality, question intent represent a minor portion of the overall query traffic.

To address the problems above, we train an intent classifier that distinguishes when a query has question intent, and thus, can be answered by an FAQ source. In mixed retrieval scenarios, this allows us to trigger FAQ retrieval and show FAQ results only for queries with question intent, causing less friction for users. In terms of efficiency, running FAQ retrieval only on question intent queries significantly improves latency. Our experiments validate

²In the case of the Bing search engine only 10% of queries were shown to have question intent.

that deploying intent classifier brings a 95% latency reduction compared to baselines without the intent classifier. Lastly, we demonstrate that existing techniques such as upsampling are enough to achieve satisfactory performance in classifying question intent in imbalanced traffic.

How to Show FAQ Results? When a search query has question intent, an interface that jointly displays product search and FAQ retrieval results is required. As in prior work on aggregating web search results (Diaz, 2009), we integrate the top-1 FAQ result alongside product search results, as illustrated in Figure 1, for the following two reasons.

First, since product search is the core functionality of e-commerce search engines, majority of the space is dedicated to the ranked product list. If we consider additional modalities such as mobile search, space constraints are even greater. Displaying more FAQ results comes at the cost of reducing the number of product results, which can lead to reduced revenue (Feng et al., 2007). Second, compared to product search, where users are required to compare multiple options, question intents generally require less exploration since users already have a specific request in their mind.

Given the above reasons, we need to optimize the FAQ retrieval system for **high precision** at the top ranks (i.e., Hit@1). Queries are usually short and consists of several keywords. To achieve a high precision for FAQ retrieval, we propose to rewrite a query with question intent into a more specific natural language question. This rewriting process aims to make the queries semantically and syntactically more similar to the questions found in FAQs than the original queries, inspired by previous studies of query reformulation (Zhao et al., 2011; Zheng et al., 2011; Yu et al., 2020). Our experiments validate that through query reformulation we can achieve significantly higher accuracy in retrieving ground-truth question at first rank (Hit@1) with more than 13% improvement when compared to using original queries for FAQ retrieval.

Contributions We summarize our contributions in this paper as follows:

- To our best knowledge, this is the first work to integrate FAQ retrieval and product search at scale.
- Our proposed intent-aware FAQ retrieval approach is a practical solution that significantly improves performance compared to baseline methods. Our approach achieves a 13% higher preci-

sion in the top-ranked results (Hit@1) and is 95% more efficient than baseline methods.

- To evaluate our design from a user’s perspective, we reviewed feedback from users who interacted with a deployed version of our system. Results showed that 71% of the rendered top-1 FAQ results (at the query level) received explicit positive customer feedback when displayed along with product search results.

2 Related Work

FAQ retrieval. The problem of FAQ retrieval has been extensively studied. Early methods (Whitehead, 1995; Sneyders, 1999) rely on exact keyword matching in FAQs. Karan et al. (2013) propose to derive lexical features such as n-gram overlap and TF-IDF similarity from a query-FAQ pair, and use these features to train a SVM model to classify whether the query is relevant to an FAQ. Karan and Šnajder (2016) combine the scores from BM25, and a classical vector space model to rank the FAQ based on its semantic similarity to the query. More recently, deep learning methods have been applied to FAQ retrieval. Gupta and Carvalho (2019) adapt a sentence matching model (Wang et al., 2017) based on bidirectional LSTM to aggregate question-to-query and question-to-answer similarities. Building on the success of BERT in NLP tasks, Sakata et al. (2019) and Mass et al. (2020) have adopted BERT to rank answers or questions of FAQs using user queries.

However, all of the experiments from early work assume that the FAQ retrieval system is deployed independently and all input queries have a question intent. In our work, we evaluate our proposed search interface in a more realistic setting by simulating search traffic queries with both product search and questions intents (with significant imbalance). We argue that this setting is more suitable for studying the benefits of aggregated search interfaces (Murdock and Lalmas, 2008) in large-scale e-commerce businesses.

Keyword-to-Question Generation. Generating questions from search queries is an example of query rewriting that is first used in community-based question answering websites (e.g., Yahoo! Answers, and Quora) to retrieve related questions. Zhao et al. (2011) propose a template-based method to rewrite keyword-based queries into questions by first extracting a set of query-question

pairs from search engine logs. Then for each input query, the most relevant templates are retrieved to generate questions. A similar method is proposed by Zheng et al. (2011), but the difference is allowing users to refine the generated question with generated refinement keywords. Ding and Balog (2018) use a statistical model to synthesize keyword-question pairs which are then used to train a neural model (Gu et al., 2016). However, the synthesized queries are noisy and additional filtering mechanism has to be used to improve performance. Recently, Iovine et al. (2022) proposed a bidirectional keyword-question rewriting model that leverages non-parallel data through cycle training. Their experiments showed that sequence-to-sequence text generation models can perform the keyword-to-question task with high accuracy, and improve retrieval results in various scenarios.

Inspired by earlier keyword-to-question generation approaches, we utilize the state-of-the-art generation models to reformulate keyword queries into questions for FAQ retrieval. Experiments show that our reformulation model trained from human annotated query-question pairs significantly improves Hit@1 by 13% compared to using original query.

3 Method

Our intent-aware FAQ retrieval approach, shown in Figure 2, consists of two main components: (1) an intent classifier that takes a user query as input and determines whether it has question intent and can be answered by an FAQ; (2) a query reformulation model which reformulates queries with question intent into a natural language question that is used for FAQ retrieval. Regardless of query intent, the product search is always performed.³

3.1 Question Intent Classification

Unlike prior work on FAQ retrieval, we do not assume that all input queries have question intent. Instead, we train a binary intent classifier that takes an input query and predict its intent into: (1) non-question intent; (2) question intent. The intent classifier corresponds to a fine-tuned RoBERTa model (Liu et al., 2019) trained for the binary classification task. To handle class imbalance, we over-sample the minority class (question intent) to approximate a balanced class distribution.

³Product search is beyond the scope of our work.

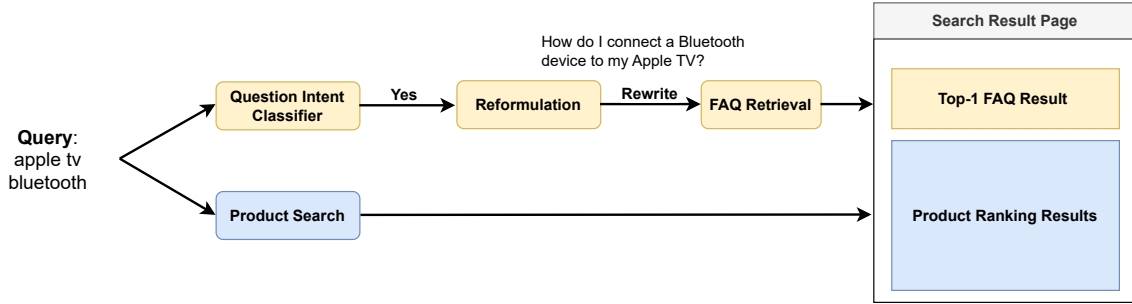


Figure 2: An overview of our proposed intent-aware FAQ retrieval approach. While product search is performed by default, FAQ retrieval is triggered only for queries with question intent.

3.2 Query Reformulation

Once a query is classified as having question intent, the query is reformulated into a natural language question. We train a sequence-to-sequence Transformer model (Vaswani et al., 2017) that reformulates the query into a question. The natural language question is used for FAQ retrieval, which are discussed in details in §4.1. The assumption behind our method is that generated questions are syntactically closer to FAQ questions than the original keyword queries, which can bring additional improvements in FAQ retrieval.

3.3 Proposed Intent-Aware FAQ Retrieval

As illustrated in Figure 2, once a query is identified with question intent and is reformulated into a natural language question, the FAQ retrieval component takes it as input and returns the top-1 FAQ result. If a query does not contain a question intent, we do not initiate the FAQ retrieval process and only return product search results.

Our retrieval component ranks FAQ results solely based on questions, without utilizing the associated answers. The rationale behind this is that a well-reformulated question should closely match the ground-truth question of an FAQ entry, allowing a simple ranking component to accurately position the correct FAQ at the top position using only the question as input. We claim that our intent-aware FAQ approach can satisfy users’ information needs whether they are looking for FAQs, qualified products, or both, and therefore provides a more convenient pre-purchase and post-purchase experience.

4 Experimental Setup

In this section, we first discuss datasets used, followed by implementation details of retrieval baselines, experimental settings and evaluation metrics.

4.1 Datasets

Intent Classification Dataset As the majority of search queries issued to e-commerce websites are focused on product search, annotations on random samples of queries yield only a tiny fraction of queries with question intent. This is not suitable for our needs, as we require a training set that includes a balanced distribution of intents. This challenge is addressed by applying several cycles of (1) training an intent classifier; (2) generating predictions on a set of unseen queries; (3) select question-intent queries with high probability (>0.9) for the next annotation; (4) correct model predictions through human annotations.

For the first cycle, when human annotations were not yet available, we utilized an existing keyword extraction algorithm (RAKE) (Rose et al., 2010) on a publicly available product question corpus (Rozen et al., 2021) to generate pairs of (question, query with question intent). For example, given a product question "How do I connect a Bluetooth device to my Apple TV", we can extract keywords "connect Bluetooth device Apple TV" as the corresponding query with question intent. For queries with shopping intent, we randomly sampled queries from our traffic and applied a simple filter (e.g., removing queries that start with a question word). Based on these initial training samples, we prototyped our first model and improved it through subsequent human annotations. More details on intent classifier training are summarized in §4.2.

Overall, our annotations resulted in 18,972 queries including 5,562 question-intent queries and 13,410 non-question queries. We allocate 50% of the data as the training set, 25% as the validation set, and the remaining 25% as the testing set. Note that since there are multiple cycles of training and annotation, the distribution of intent in this set does not reflect the distribution of the real traffic.

Query Reformulation Dataset From the intent classification dataset, we randomly sample 1,500 question intent queries from the training set and 500 question intent queries from the testing set. For each query, we ask annotators to reformulate it into a natural language question. The resulting 1,500 query-question pairs form the training set is used for training reformulation models, while the remaining 500 pairs are divided evenly into validation and testing sets.

FAQ Corpus For FAQ retrieval, we annotate each query in the query reformulation test set with a ground-truth FAQ, ensuring full coverage for the test set in our experiments.⁴

4.2 Implementation Details for Our Method

Intent Classifier We fine-tuned RoBERTa large model (Liu et al., 2019) on the intent classification dataset (cf. §4.1). Due to significant class imbalance, we upsampled (with repetition) the question-intent queries until balanced distribution of intents are satisfied (we cannot disclose the exact sampling ratio). Standard cross entropy (CE) loss is adopted. This model is trained for 8 epochs with $2e-6$ learning rate and 256 batch size, which is distributed evenly on 8 NVIDIA K80 GPUs. We use Adam as our optimizer and early stopping (*patience* = 10) is used to prevent overfitting.

Reformulation Model We train two reformulation models using BART-base (Lewis et al., 2019) and T5-base (Raffel et al., 2020), respectively. Both models take a user query with question intent as input, and output its reformulated question. The CE loss is adopted to train the models by maximizing the likelihood of generating the human’s reformulation. This model is trained for 10 epochs with $1e-5$ learning rate and a single Tesla A100 GPU. We use Adam as our optimizer, and batch size of 16. The training is halted using early stopping (*patience* = 3). All intent classifier and reformulation models are implemented using HuggingFace.⁵

4.3 Implementation Details for FAQ Retrieval

We evaluate our approach for FAQ retrieval on the following ranking models of different complexity:

BM25 (Robertson et al., 2009) We use BM25 as an unsupervised FAQ retrieval model. All the

FAQ questions are indexed using Lucene.⁶ For a user query, we retrieve top-50 documents based on BM25 scores regardless of the query intent. It is possible that some queries with non-question intent may not return any results.

SentTrans (Reimers and Gurevych, 2019) We adopt a sentence encoder model⁷ that is trained on Google’s Natural Question corpus (Kwiatkowski et al., 2019) to compute the similarity between queries reformulations and FAQ questions. For retrieval, we rank FAQ questions based on their cosine similarity against either the query or its rewritten question.

BERT (Devlin et al., 2019) We fine-tune a point-wise ranker using BERT on our query reformulation dataset to rank a query against all FAQs. For a query, we treat the ground-truth reformulation as the positive sample and randomly sample 100 reformulations of other queries as negative samples. Finally, the FAQ questions are ranked based on the classification score w.r.t the query. Hinge loss function is applied to train the model for 10 epochs with a batch size of 25.

BERT-Rerank (Dai and Callan, 2019) Directly ranking all FAQs with BERT is computationally expensive. A more efficient approach is to rerank the top- k results of BM25 using BERT. In our experiments, we test k with 10 and 50.

4.4 Simulating Intent Classification with Retrieval Baselines

Although the FAQ retrieval approaches discussed in §4.3 were not originally designed for intent classification, we include them as baselines to approximate false matches on real traffic. Specifically, we evaluate the probability of FAQ results shown to users who do not have the intent of seeking FAQs (reflected by precision), as well as the likelihood of FAQ results not appearing for users who are specifically searching for FAQs (reflected by recall), in the presence and absence of our intent classifier. For BM25, we consider a query to have question intent if it meets both requirements:

1. The number of returned FAQ results is more than a threshold x ;
2. BM25 score of the top-1 result is larger than a threshold y .

⁴We cannot disclose the performance on the full internal FAQ database.

⁵<https://huggingface.co/>

⁶<https://lucene.apache.org/>

⁷<https://www.sbert.net/docs/pretrained-models/nq-v1.html>

By default, we set $x = 1$ and $y = 0$, which means at least one FAQ is returned for a question intent query. To obtain optimal BM25 results, we use the validation set for fine-tuning the thresholds ($x = 40$, $y = 5$). Similarly for SentTrans, we find the optimal cosine similarity threshold (0.6) based on the validation set, and classify queries with question intent if an FAQ is retrieved above the threshold.

4.5 Evaluation Metrics

To measure the intent classification performance, we report results on precision, recall and F1. We report mean reciprocal rank (MRR) and Hit@1 to evaluate the FAQ retrieval performance. Hit@1 is the most critical metric since only the top-1 retrieved FAQ result will be displayed to users.

5 Experiments and Results

We study the following research questions:

RQ1: How much does using the question intent classifier as a filtering step benefit the integrated FAQ retrieval pipeline?

RQ2: How effective is query-to-question reformulation for FAQ retrieval?

RQ3: How efficient is our intent-aware FAQ retrieval approach when integrated with a product search engine?

5.1 Question Intent Classification

The intent classifier determines when to trigger FAQ retrieval and display the FAQ results alongside the product search results. Table 1 shows the evaluation results on intent classification. For reasons of confidentiality, we report results as relative differences with the baseline method (BM25). The details on approximating precision and recall for baselines are summarized in §4.4. Since the evaluation set is highly imbalanced, BM25 has the lowest F1 score with the majority of queries being classified as question intent queries, leading to extremely high recall and low precision. Even with optimal thresholds, BM25 obtains only an F1 score that is 48% lower than our method. Although SentTrans outperforms BM25, it still falls significantly behind our method. To answer **RQ1**, intent classifier improves precision by 26% and recall by 51% compared to the strongest SentTrans baseline.

These results indicate that question intent classifier is required because other baselines cannot effectively distinguish queries with question intent.

Methods	Precision	Recall	F1
BM25	0.00	0.00	0.00
BM25 (Optimal)	+0.41	-0.60	+0.36
SentTrans (Optimal)	+0.54	-0.54	+0.46
Our Method	+0.80	-0.03	+0.84

Table 1: Results of question intent classification. Scores are relative to BM25.

Retrieval Model	Input Query Type	MRR	Hit@1
BM25	Original	0.00	0.00
	Reformulation (BART)	+0.07	+0.10
	Reformulation (T5)	+0.11	+0.17
SentTrans	Original	+0.01	+0.01
	Reformulation (BART)	+0.08	+0.12
	Reformulation (T5)	+0.13	+0.19
BERT (pointwise)	Original	+0.02	+0.02
	Reformulation (BART)	+0.04	+0.06
	Reformulation (T5)	+0.07	+0.09
BERT-Rerank (top-10)	Original	+0.05	+0.07
	Reformulation (BART)	+0.10	+0.15
	Reformulation (T5)	+0.14	+0.20
BERT-Rerank (top-50)	Original	+0.08	+0.08
	Reformulation (BART)	+0.10	+0.13
	Reformulation (T5)	+0.14	+0.19

Table 2: Results of FAQ retrieval by different retrieval models and query types. Scores are relative to BM25 using the original queries.

5.2 FAQ Retrieval

For each retrieval system being compared, there are three input types: (1) original query, (2) BART reformulated query, and (3) T5 reformulated query. Table 2 summarizes the FAQ retrieval results. Similar to Section 5.1, we only report numbers relative to the BM25 baseline where original queries are used (first row in Table 2).

First, we observe that using the reformulated query improves the performance of all FAQ retrieval systems in terms of MRR and Hit@1, which answers **RQ2**. T5 query reformulations consistently show better results than BART query reformulations. Even for the most competitive BERT-Rerank baselines, T5 can further improve Hit@1 by more than 10%. The results indicate that our reformulation method has the advantage of improving the precision at top ranks. Although our reformulation models are not designed to generate exact questions from FAQs,⁸ they generate questions that are sufficiently similar to the FAQs in our corpus. As a result, they significantly improve the retrieval performance.

⁸This is a challenging annotation task because workers must match queries against our entire FAQ corpus.

Second, Table 2 shows that reformulations allow BM25 to achieve comparable results with strong BERT-Rerank baselines. Notably, BM25 with T5’s reformulated queries outperforms two BERT-Rerank baselines using the original queries. Yet, when using T5’s reformulations, BERT-Rerank (top-10) achieves the best MRR and Hit@1 among all methods, outperforming the original query by 13% in Hit@1, and 9% in MRR. Considering the model’s complexity, the combination of BM25 and the reformulation method is a promising FAQ solution in industry settings.

5.3 Efficiency Comparison

An important consideration in industry settings is the efficiency of the proposed solution. We assess the impact of our proposed solution in terms of computational cost. As BERT Rerank (top-10) with T5 reformulation yields the best results, we compare its inference time with and without an intent classifier. We measure the speed on a single *p3.8xlarge* instance,⁹ with a batch size of 16.

Table 3 shows the inference time¹⁰ normalized by the baseline that does not use an intent classifier. For **RQ3**, our proposed intent-aware FAQ retrieval system that reranks top-10 results using BERT and T5 reformulations can save on average 95% of inference time, which is a significant benefit for real-world applications. We also observed that despite adding an extra layer of inference from reformulation model, the increase in latency is negligible since majority of queries are already filtered out and only a small fraction of queries are reformulated.

Pipeline	Query	Inference Time
BERT-Rerank (top 10)	Original	1.0
	Original	0.0404 (↓95.96%)
Ours	Reformulation (BART)	0.0446 (↓95.54%)
	Reformulation (T5)	0.0461 (↓95.39%)

Table 3: Inference speed comparison on full query traffic (including question intent and non-question intent queries) with and without intent classifier. The inference time is normalized based on the time taken by the method that does not utilize an intent classifier.

6 Online Deployment

Our intent-aware FAQ retrieval solution is deployed and integrated into the product search interface of

⁹<https://aws.amazon.com/ec2/instance-types/>

¹⁰The raw inference time cannot be disclosed.

a leading global e-commerce website. The live version uses a much larger FAQ corpus than the ones used here. When FAQ results are displayed, we collect optional explicit user feedback on whether the answer is helpful or not, as demonstrated in Figure 1. Over one month of traffic was collected from the US marketplace, showing that 71% of the rendered FAQ results received explicit positive customer feedback. The feedback results were aggregated at the query level, allowing each query to receive multiple positive feedback responses. These findings demonstrate that our practical solution is not only effective in offline evaluation, but also helpful to real users.

7 Conclusion

We proposed to integrate FAQ retrieval with product search to address challenges in aggregated search from an e-commerce perspective. Our approach first classifies queries with question intent, which then reformulates into natural language questions that are used to retrieve FAQs. Offline experimental results show that on the best-performed FAQ retrieval system (i.e., BERT-Rerank (top 10)), the proposed intent classifier saves a substantial amount of inference costs (95%) and also improved retrieval performance through query reformulation by 13% on Hit@1. These improvements are also reflected in online evaluation: over one month of user feedback demonstrated that about 71% of the rendered FAQ results were considered to be helpful. Overall, the findings in this work suggest promising directions for e-commerce platforms to support FAQ retrieval at scale without disrupting customer’s shopping experience.

Limitations and Future Work

One limitation of our approach is that we do not display or rank multiple reformulations. It is possible that a query can be reformulated into multiple possible questions. For example, the query “*apple tv bluetooth*” can be reformulated into “*How do I connect a Bluetooth device to my Apple TV*” or “*Does apple TV support Bluetooth*”. In our future work, we aim to explore the integration of multiple reformulations into the FAQ retrieval process to further enhance the overall user experience. Another limitation is that we do not train an end-to-end FAQ retrieval model. In the future, we plan to train the FAQ retrieval model using the reformulations so that the original query can directly be used.

References

- David Carmel, Liane Lewin-Eytan, and Yoelle Maarek. 2018. Product question answering using customer generated content-research challenges. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 1349–1350.
- Philipp Christmann, Rishiraj Saha Roy, and Gerhard Weikum. 2022. Conversational question answering on heterogeneous sources. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 144–154.
- Zhuyun Dai and Jamie Callan. 2019. Deeper text understanding for ir with contextual neural language modeling. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*, pages 985–988.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Fernando Diaz. 2009. Integration of news content into web results. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, pages 182–191.
- Heng Ding and Krisztian Balog. 2018. [Generating Synthetic Data for Neural Keyword-to-Question Models](#). In *Proceedings of the 2018 ACM SIGIR International Conference on Theory of Information Retrieval*, pages 51–58, Tianjin China. ACM.
- Juan Feng, Hemant K Bhargava, and David M Pennock. 2007. Implementing sponsored search in web search engines: Computational evaluation of alternative mechanisms. *INFORMS Journal on Computing*, 19(1):137–148.
- Shen Gao, Zhaochun Ren, Yihong Zhao, Dongyan Zhao, Dawei Yin, and Rui Yan. 2019. Product-aware answer generation in e-commerce question-answering. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 429–437.
- Varun Grover and James TC Teng. 2001. E-commerce and the information market. *Communications of the ACM*, 44(4):79–86.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Sparsh Gupta and Vitor R Carvalho. 2019. Faq retrieval using attentive matching. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 929–932.
- Andrea Iovine, Anjie Fang, Besnik Fetahu, Jie Zhao, Oleg Rokhlenko, and Shervin Malmasi. 2022. [CycleKQR: Unsupervised bidirectional keyword-question rewriting](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11875–11886, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Mladen Karan and Jan Šnajder. 2016. Faqir—a frequently asked questions retrieval test collection. In *Text, Speech, and Dialogue: 19th International Conference, TSD 2016, Brno, Czech Republic, September 12-16, 2016, Proceedings 19*, pages 74–81. Springer.
- Mladen Karan, Lovro Žmak, and Jan Šnajder. 2013. Frequently asked questions retrieval for croatian based on semantic textual similarity. In *Proceedings of the 4th Biennial International Workshop on Balto-Slavic Natural Language Processing*, pages 24–33.
- Robert Krovetz and W Bruce Croft. 1992. Lexical ambiguity and information retrieval. *ACM Transactions on Information Systems (TOIS)*, 10(2):115–141.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. 2019. Natural questions: a benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. [BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). *CoRR*, abs/1910.13461.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Wei Lu, Qi Wang, and Birger Larsen. 2012. Simulating aggregated interfaces. In *Workshop on aggregated search*, pages 24–28.
- Yosi Mass, Boaz Carmeli, Haggai Roitman, and David Konopnicki. 2020. Unsupervised faq retrieval with question generation and bert. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 807–812.
- Vanessa Murdock and Mounia Lalmas. 2008. Workshop on aggregated search. In *ACM SIGIR Forum*, volume 42, pages 80–83. ACM New York, NY, USA.

- A Nain and A Awasthi. 2021. E-tourism: A study of tourist satisfaction. *International Journal of Trend in Scientific Research and Development*, 5(6).
- Seonyeong Park, Soonchoul Kwon, Byungsoo Kim, Sangdo Han, Hyosup Shim, and Gary Geunbae Lee. 2015. [Question answering system using multiple information source and open type answer merge](#). In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 111–115, Denver, Colorado. Association for Computational Linguistics.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 3(4):333–389.
- Stuart Rose, Dave Engel, Nick Cramer, and Wendy Cowley. 2010. Automatic keyword extraction from individual documents. *Text mining: applications and theory*, pages 1–20.
- Ohad Rozen, David Carmel, Avihai Mejer, Vitaly Mirkis, and Yftah Ziser. 2021. Answering product-questions by utilizing questions from other contextually similar products. *arXiv preprint arXiv:2105.08956*.
- Wataru Sakata, Tomohide Shibata, Ribeka Tanaka, and Sadao Kurohashi. 2019. [FAQ retrieval using query-question similarity and bert-based query-answer relevance](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1113–1116.
- Mark Sanderson. 2008. Ambiguous queries: test collections need more sense. In *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 499–506.
- Ahsan Siraj, J Guo, MW Kamran, Q Li, and Y Zhu. 2020. Characteristics for e-satisfaction in e-retailing-evidence from chinese e-commerce. *International Journal of Innovation, Creativity and Change*, 12(5):375–376.
- Eriks Sneiders. 1999. Automated faq answering: Continued experience with shallow language understanding. In *Question Answering Systems. Papers from the 1999 AAAI Fall Symposium*, pages 97–107.
- Ruihua Song, Zhenxiao Luo, Ji-Rong Wen, Yong Yu, and Hsiao-Wuen Hon. 2007. Identifying ambiguous queries in web search. In *Proceedings of the 16th international conference on World Wide Web*, pages 1169–1170.
- Ning Su, Jiyin He, Yiqun Liu, Min Zhang, and Shaoping Ma. 2018. User intent, behaviour, and perceived satisfaction in product search. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 547–555.
- Gilad Tsur, Yuval Pinter, Idan Szpektor, and David Carmel. 2016. Identifying web queries with question intent. In *Proceedings of the 25th International Conference on World Wide Web*, pages 783–793.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Stephen Wan, Cecile Paris, and Alexander Krumpholz. 2008. From aggravated to aggregated search: Improving utility through coherent organisations of an answer space. In *Proceedings of the SIGIR 2008 Workshop on Aggregated Search*, volume 3. Citeseer.
- 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-17*, pages 4144–4150.
- Ryen W White, Matthew Richardson, and Wen-tau Yih. 2015. Questions vs. queries in informational search tasks. In *Proceedings of the 24th international conference on World Wide Web*, pages 135–136.
- Steven D Whitehead. 1995. Auto-faq: An experiment in cyberspace leveraging. *Computer Networks and ISDN Systems*, 28(1-2):137–146.
- Shi Yu, Jiahua Liu, Jingqin Yang, Chenyan Xiong, Paul Bennett, Jianfeng Gao, and Zhiyuan Liu. 2020. Few-shot generative conversational query rewriting. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, pages 1933–1936.
- Shiqi Zhao, Haifeng Wang, Chao Li, Ting Liu, and Yi Guan. 2011. Automatically generating questions from queries for community-based question answering. In *Proceedings of 5th international joint conference on natural language processing*, pages 929–937.
- Zhicheng Zheng, Xiance Si, Edward Y. Chang, and Xiaoyan Zhu. 2011. [K2Q: generating natural language questions from keywords with user refinements](#). In *Fifth International Joint Conference on Natural Language Processing, IJCNLP 2011, Chiang Mai, Thailand, November 8-13, 2011*, pages 947–955. The Association for Computer Linguistics.