

# Doc2Token: Bridging Vocabulary Gap by Predicting Missing Tokens for E-commerce Search

Kaihao Li<sup>1</sup>, Juexin Lin<sup>1</sup> and Tony Lee<sup>1</sup>

<sup>1</sup>Walmart Global Technology, USA

## Abstract

Addressing the “vocabulary mismatch” issue in information retrieval is a central challenge for e-commerce search engines, because product pages often miss important keywords that customers search for. Doc2Query [1] is a popular document-expansion technique that predicts search queries for a document and includes the predicted queries with the document for retrieval. However, this approach can be inefficient for e-commerce search, because the predicted query tokens are often already present in the document. In this paper, we propose Doc2Token, a technique that predicts relevant tokens (instead of queries) that are missing from the document and includes these tokens in the document for retrieval. For the task of predicting missing tokens, we introduce a new metric, “novel ROUGE score”. Doc2Token is demonstrated to be superior to Doc2Query in terms of novel ROUGE score and diversity of predictions. Doc2Token also exhibits efficiency gains by reducing both training and inference times. We deployed the feature to production and observed significant revenue gain in an online A/B test, and launched the feature to full traffic on Walmart.com.

## Keywords

Document Expansion, Information Retrieval, E-commerce Search

## 1. Introduction

The vocabulary gap problem in e-commerce search is a central challenge, as it arises from discrepancies between the vocabulary used by customers and sellers when describing products. Customer queries are often short and ambiguous, while product descriptions tend to be more detailed and explicit. For instance, a customer might search for “small building set”, intending to find a set that offers simpler building experiences for young children. However, in the product catalog, those products are often characterized by piece count and target age group, which do not align directly with this search query.

Different approaches have been proposed to address the vocabulary mismatch issue. In the context of lexical retrieval, query expansion [2, 3, 4, 5, 6] and document expansion [1, 7, 8] are two effective techniques. Query expansion enriches user queries with additional terms or synonyms to better capture the user’s intent, while document expansion enriches product information with additional keywords or phrases. Doc2Query [7] is a document expansion technique that predicts and indexes search queries for documents. Recently, embedding-based dense retrieval models [9] have demonstrated their ability to align queries and documents by projecting them into a representation space to learn their semantic similarity. Although dense retrieval with approximate nearest neighbor search has shown impressive results, lexical


---

*eCom’24: ACM SIGIR Workshop on eCommerce, July 18, 2024, Washington, D.C., USA*

✉ kaihao.li@walmart.com (K. Li); juexin.lin@walmart.com (J. Lin); tony.lee@walmart.com (T. Lee)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



**Figure 1:** Doc2Token Overview.

retrieval remains an important component of e-commerce search due to its desirable properties, such as interpretability, scalability, as well as handling of rare words and numerical tokens.

In this paper, we propose a new document expansion technique, called Doc2Token, which generalizes Doc2Query for application in e-commerce search, as depicted in Figure 1. Our task is to, given a product, generate relevant keywords that are absent from the product’s indexed metadata to ensure that the product is retrieved when customers search using these keywords. We call these “novel tokens”. We observed that Doc2Query’s predicted queries often contain tokens already in the product metadata instead of novel tokens, which makes it inefficient for our task. In contrast, Doc2Token is designed to predict novel tokens (instead of queries). The approach is to prepare a dataset of pairs of product and novel token, then train a seq2seq model on that dataset. By design, Doc2Token efficiently generates tokens with a high probability of being novel, rather than producing long and redundant sequences like in Doc2Query. Using the product shown in Figure 1 as an example, Doc2Query predicts “6 year old boy toy”, “3 in 1 creator”, “sea animal toy”, and “building toy for boy”. On the other hand, Doc2Token provides a more diversified set of tokens, including “small”, “kit”, and “tank”. To incorporate the Doc2Token predictions into the search system, we added them to the product’s metadata and indexed for retrieval matching and ranking.

To assess performance on the novel-token-prediction task, we introduce a new evaluation metric called “novel ROUGE score”, denoted by “nROUGE”, to measure the ROUGE score [10] specifically for novel tokens. Our results indicate that Doc2Token surpasses Doc2Query in terms of nROUGE score. Regarding efficiency, Doc2Token is capable of predicting more diverse results while significantly reducing training and inference times. The effectiveness of this approach is further demonstrated through online relevance evaluation and A/B testing, confirming that the novel tokens generated are not only novel but also relevant to products, ultimately driving customer engagement.

Our contributions are summarized as follows:

- We propose a novel technique, Doc2Token, for document expansion in e-commerce search, encompassing both the training setup and the modification of the loss function.
- We introduce a new metric, “novel ROUGE score”, to evaluate the performance of predicting novel tokens.
- We demonstrate that Doc2Token achieves improvements in both effectiveness and efficiency compared to Doc2Query.

## 2. Methodology

We define the task of novel token prediction in this section. For a product  $i$ , let  $p_i$  and  $T_i$  represent the product text and the set of tokens extracted from  $p_i$ , respectively. The goal is to predict novel tokens, i.e., tokens that are absent from  $T_i$ . To achieve this, we first collect a list of relevant queries (section 3.1) for the product based on historical search logs. Next, for each product, we assemble a set of unique tokens, disregarding their sequence, through the subsequent process. We concatenate all queries, divide the concatenated sequence into individual tokens, count their frequencies, and exclude tokens already present in product  $i$ . As a result, for product  $i$ , we have a target token set  $\bigcup_{t_{ik} \notin T_i} (t_{ik}, f_{ik})$ , where  $t_{ik}$  represents the  $k$ th unique token from target queries for product  $i$ , and  $f_{ik}$  denotes the frequency of token  $t_{ik}$ . Instead of using all tokens as one training target, we divide them into  $k$  training instances. We train a seq2seq generative language model, T5 [11], with an encoder-decoder structure. It takes the product text  $p_i$  as input and outputs a sequence of tokens  $\hat{t}_i$  in an autoregressive manner. More formally,

$$\hat{t}_i = \text{Decoder}(\text{Encoder}(p_i)). \quad (1)$$

To account for the token frequency, we modify the T5 loss as follows. The loss for a product-token pair is:

$$\text{loss}_{\text{weighted}}(t_{ik}, \hat{t}_i) = (f_{ik})^\alpha * \text{loss}_{T5}(t_{ik}, \hat{t}_i), \quad (2)$$

where  $\alpha$  is the smoothing factor, set to 0.5 in our implementation. This value is chosen to balance the contribution of token frequency to the overall loss. In practice, we use the cross-entropy loss for  $\text{loss}_{T5}$ .

For model inference, we employ beam search [12] to generate the top N predictions. Since T5 tokenizes words into subtokens [13], the target token  $t_{ik}$  and predicted token  $\hat{t}_i$  may consist of multiple subtokens, although they are always single words. We utilize the beam score [12] to determine the confidence of the prediction, and we only retain predictions with scores greater than a predetermined cutoff value (more details are discussed in Section 3.4).

## 3. Experiments

### 3.1. Datasets

We sampled the product-query pairs from user engagement data on Walmart.com with at least a certain number of add-to-cart (ATCs) over a two-year period. Then we did the following preprocessing steps. The product information used throughout the experiment includes product title, product type, brand, color, gender and description.

Preprocessing	product-query pairs	product-token pairs	products
None	81.8M		8.7M
RF	66.4M		8.0M
RF + FMF + PTF (Doc2Query)	14.9M		3.1M
RF + FMF + PTF + tokenization + OTF (Doc2Token)		10.3M	3.1M

**Table 1**

Summary statistics for each data preprocessing step.

*Relevance filter (RF).* The engaged products are not always relevant to the search query, because users’ decisions are influenced by factors other than relevance, such as price, visual appeal, ranking, etc. Additionally, the minimal match criteria [14, 15] for our lexical retrieval is not always 100%. As a result, customers may be shown products that don’t fully match their search terms. For instance, a customer may search for “vanilla ice cream” but end up buying chocolate ice cream. To mitigate such noise, we removed product-query pairs predicted to be irrelevant by applying a relevance model. (The relevance model is a BERT [16] cross-encoder model with a classifier head that takes the query and product information as input and outputs a relevance score. It was trained on manually annotated relevance data.)

*Full match filter (FMF).* For the purpose of predicting novel tokens, we focused on product-query pairs with a vocabulary mismatch. This implies that at least one query token was not found in the product information. Thus, we removed pairs where all query tokens were in the product information.

*Price token filter (PTF).* Customer queries sometimes include price and deal intent (e.g., “under \$500” or “on sale”). Such phrases are not very useful for our task, since price and deals can fluctuate rapidly, so it does not make sense to include them as training labels. We utilized regular expressions to identify and eliminate these phrases from the query.

*Overlapping token filter (OTF).* This step excludes all query tokens that are present in the product information. (This is a stronger extension of the full match filter.) After this step, only novel tokens remain.

Detailed data statistics can be found in Table 1. In later sections, we show results for both Doc2Query and Doc2Token. For the Doc2Query dataset, we applied the first three filters, resulting in 14.9M product-query pairs. For the Doc2Token dataset, we built upon the Doc2Query dataset by further dividing the queries into tokens and applying the overlapping token filter, resulting in 10.3M product-token pairs. For each dataset, we partitioned it by product into training, validation, and test sets in an 8:1:1 ratio to ensure no product overlap between the sets.

### 3.2. Metrics

To evaluate model performance, we utilize the standard ROUGE score, a widely-used evaluation metric for summarization tasks [17]. The ROUGE score assesses the quality of generated text by comparing it to a reference text, considered as ground truth, using n-gram overlaps. Unlike the summarization task, which may require complex evaluations based on n-grams, lexical retrieval

primarily relies on unigrams. Therefore, we measure the ROUGE scores in terms of unigrams. In our context, we formulate the ROUGE score as follows.

$$ROUGE_{Precision} = \frac{1}{N} \sum_i \frac{\sum_{token \in y_i} count_{match}(token, \hat{y}_i)}{count(\hat{y}_i)}, \quad (3)$$

$$ROUGE_{Recall} = \frac{1}{N} \sum_i \frac{\sum_{token \in y_i} count_{match}(token, \hat{y}_i)}{count(y_i)}, \quad (4)$$

where  $y_i$ ,  $\hat{y}_i$  represent the reference text and predicted text for product  $i$ , respectively.  $count$  denotes the number of tokens, and  $count_{match}$  denotes the number of co-occurrences of a reference token in the predicted text.  $N$  is the number of products.

However, a higher ROUGE score does not necessarily indicate better performance at predicting novel tokens. We observed that text predictions with high ROUGE scores often exhibit substantial token overlap with the product information. To assess the performance for novel tokens, we introduce a new metric, ‘‘novel ROUGE score’’ (nROUGE), where the reference text consists solely of novel tokens. Formally, the nROUGE score is defined as follows.

$$nROUGE_{Precision} = \frac{1}{N} \sum_i \frac{\sum_{token \in y_i^*} count_{match}(token, \hat{y}_i)}{count(\hat{y}_i)}, \quad (5)$$

$$nROUGE_{Recall} = \frac{1}{N} \sum_i \frac{\sum_{token \in y_i^*} count_{match}(token, \hat{y}_i)}{count(y_i^*)}, \quad (6)$$

where  $y_i^*$  represents the novel reference text for product  $i$ , i.e., the tokens in the reference text but not in the product information.

### 3.3. Experiment Setup

For the model training, we fine-tuned the public T5-base model using eight A-100 Nvidia GPUs, learning rate of 1e-4, batch size of 64, maximum input sequence of 256, and maximum output sequence of 32. For model inference, we employed a top-k beam search strategy with a beam size of 10. The input of the model is a text of product information consisting of the product title and some attributes, such as brand, color, etc.

### 3.4. Results

In this section, we compare our proposed Doc2Token model to the baseline Doc2Query model from both effectiveness and efficiency perspectives. To assess effectiveness, we measure performance based on the ROUGE and nROUGE scores. In terms of efficiency, we report the resources used for model training and inference. We evaluated on four models, including both Doc2Query and Doc2Token models with and without full-match filter in data preprocessing step.

Model	Cutoff	ROUGE			nROUGE			# of predicted tokens		
		Precision	Recall	F1	Precision	Recall	F1	Total	Novel	%
Doc2Query w/o FMF	0.00	0.505	0.704	0.533	0.409	0.606	0.427	28.9	9.0	31
	0.34	0.535	0.680	0.545	0.434	0.574	0.438	25.5	8.3	33
Doc2Token w/o (FMF + OTF)	0.00	0.448	0.696	0.596	0.346	0.636	0.393	10.0	10.0	100
	0.15	0.534	0.664	0.548	0.453	0.593	0.464	8.2	8.2	100
	0.18	0.564	0.640	0.556	0.479	0.560	0.469	7.5	7.5	100
Doc2Query w/ FMF	0.00	0.418	0.797	0.496	0.338	0.755	0.402	31.8	6.2	19
	0.51	0.558	0.673	0.554	0.497	0.622	0.481	15.2	3.2	21
Doc2Token w/ (FMF + OTF)	0.00	0.214	0.364	0.240	0.205	0.767	0.286	10.0	9.8	98
	0.29	0.517	0.281	0.329	0.507	0.623	0.496	3.2	3.2	100
	0.33	0.553	0.261	0.321	0.554	0.583	<b>0.500</b>	2.6	2.6	100

**Table 2**

Results to compare Doc2Query models and Doc2Token models.

### 3.4.1. Offline evaluation results

In Table 2, we present the results for each model based on the top 10 predictions with various beam score cutoffs. The predictions were chosen if their beam scores exceeded the respective cutoff value. These cutoff values were tuned, for each model, to achieve the optimal nROUGE F1 score. Additionally, to assess models’ efficiency in predicting novel tokens, we concatenated the predictions and calculated the total number of predicted tokens and the number of predicted novel tokens. For the Doc2Query models, we present the result without any beam score cutoff, as well as the result that achieves the highest nROUGE F1 score. For the Doc2Token models, we reported three results: one without any cutoff, one with the optimal nROUGE F1 score, and a result generating a similar number of novel tokens as the optimal Doc2Query model.

To evaluate the effectiveness of the full-match filter in data preprocessing, we trained the models without incorporating that step. We observed no apparent impact on the ROUGE F1 score. However, there was a substantial improvement in the nROUGE F1 score for both Doc2Query (from 0.438 to 0.481) and Doc2Token (from 0.469 to 0.500) at the optimal cutoff values. This is expected, as with the full-match filter, our training data primarily relies on labels containing novel tokens.

In our comparison between the Doc2Query and Doc2Token models, we observed that the Doc2Query models tend to achieve higher ROUGE F1 scores than the Doc2Token models. However, Doc2Token excels in achieving superior nROUGE F1 scores. Comparing the models with a similar number of predicted novel tokens, for example, the Doc2Query model with a cutoff of 0.51 and Doc2Token model with a cutoff of 0.29, the Doc2Token model outperforms the Doc2Query model in both nROUGE precision and nROUGE recall, yielding a higher nROUGE F1 score. With optimal cutoffs, the Doc2Token model shows the superior performance compared to the Doc2Query model by achieving a 3.95% higher nROUGE F1 score from 0.496 to 0.500. This improvement is statistically significant, with a 95% confidence interval for the Doc2Token F1 score of (0.498, 0.501) obtained through bootstrap resampling. Moreover, the Doc2Token model is more efficient in generating novel tokens, achieving nearly 100% of predicted tokens being novel. In contrast, Doc2Query produces only 20% novel tokens, indicating a higher degree

Model	Preprocessing	Train Time	Inference Time
Doc2Query	w/o FMF	1002	136
Doc2Token	w/o (FMF + OTF)	1416	76
Doc2Query	all	166	141
Doc2Token	all	<b>108</b>	<b>76</b>

**Table 3**

Results to compare model train and inference time. Train time and inference time are measured at minutes per one epoch and minutes per 100k products, respectively.

of redundancy. This is expected, as the Doc2Token model is designed to predict more diverse novel tokens.

### 3.4.2. Model efficiency

Table 3 presents the results for training and inference times. The training time is primarily affected by the size of the training data. Without the full-match filter, splitting queries into tokens explodes the data size, resulting in a longer training time for Doc2Token compared to Doc2Query. However, with the full-match filter, the situation changes: the Doc2Token strategy significantly reduces the dataset size, leading to shorter training times than Doc2Query. For inference time, we sampled 100,000 products from the test dataset and conducted model inference on the top 10 results with a batch size of 16 using a single K80 GPU machine. The inference time for Doc2Token is faster than that for Doc2Query, as the output of Doc2Token is generally shorter. The results are in agreement with the efficiency discussions from Table 2.

### 3.5. Examples

Product Input	title: Toddler Floaties, Swim Vest for Boys and Girls Age 2-7 Years Old, 20-50 Pounds Children Water Wings Arm Floaties in Puddle/Sea/Pool/Beach (Dinosaur) brand: Dark Lightning color: Blue gender: Unisex
Doc2Query	“swimming vest for <b>kid</b> ”, “toddler boy swim vest”, “swim vest for <b>kid</b> ”, “boy floaty”, “ <b>kid</b> floaty”
Doc2Token	“ <b>float</b> ”, “ <b>kid</b> ”, “floaty”, “ <b>floater</b> ”, “ <b>salvavida</b> ”, “ <b>swimmy</b> ”, “ <b>baby</b> ”, “ <b>floatation</b> ”, “children”, “ <b>life</b> ”
Product Input	title: Hanno Muller-Brachmann - North German Poets - Classical - CD brand: Artists color: white
Doc2Query	“classical cds”, “germany cd”, “ <b>country</b> music cd”, “ <b>north</b> germany cd”, “ <b>west</b> germany cd”
Doc2Token	“ <b>music</b> ”, “ <b>country</b> ”, “ <b>5</b> ”, “ <b>b</b> ”, “classical”, “ <b>christmas</b> ”, “ <b>soundtrack</b> ”

**Table 4**

Two product examples for Doc2Query and Doc2Token: A good example & A case for improvement

Table 4 showcases two example products along with its corresponding Doc2Query and Doc2Token predictions. The Doc2Query model produces the top 5 queries, while the Doc2Token model generates the top 10 tokens. The novel tokens, after the process of stemming [14, 15], are bold. In general, the Doc2Query model produces queries containing tokens that are already



present in the product information, while the Doc2Token model does not. In contrast, all tokens produced by the Doc2Token model are relevant and absent from the product. In the first positive example, the Doc2Token model is capable of predicting a Spanish word “salvavida” (“lifeguard” in English), indicating its ability to handle Spanish queries. Queries in Spanish are commonly observed in US e-commerce search. The second example, a product from northern Germany, illustrates some bad predictions from the models. The predicted tokens “country”, “west”, “christmas” are irrelevant. This is mainly due to a lack of media-related data in our training set. Replacing T5 with a more knowledgeable LLM could potentially address this issue.

## 4. Implementation and online tests

We implemented the Doc2Token model in production because of its superior performance compared to the Doc2Query model as shown in Section 3.4. We ran the model inference on all products in our catalog using cost-effective K80 GPUs and a batch size of 16. We predicted the top 10 tokens and retained the predictions with scores above 0.33. The inference process is conducted offline on a daily basis. For online usage, the Doc2Token predictions serve as an additional text matching field in Solr [14], which is an enterprise search platform built on Apache Lucene [15] utilized for the search retrieval at Walmart.com.

We evaluated the performance of the Doc2Token feature from both relevance and engagement perspectives. For the relevance evaluation, we enlisted the human annotators to assess the top 10 ranked products from impacted queries. These assessments were based on a three-point scale (exact match, substitute, irrelevant), considering factors such as product title, image, and product page at Walmart.com. We then computed the NDCG@10 based on this 3-point scale, showing a 0.49% lift (p-value=0.066). For engagement assessment, we conducted a two-week A/B test for the feature on live traffic. The test revealed a statistically significant 0.28% lift in revenue (p-value = 0.013). While the NDCG@10 improvement is statistically marginal, the statistically significant revenue increase demonstrates the effectiveness of the Doc2Token feature. By introducing relevant products in the retrieval process, the Doc2Token feature is able to enhance the end-to-end search results, assisting customers in finding what they are searching for.

## 5. Conclusions

In this study, we present Doc2Token, a novel document expansion technique for e-commerce search engines. We introduce the novel ROUGE score, a new metric crafted to evaluate the efficacy of document expansion endeavors. Our analysis has demonstrated that Doc2Token surpasses Doc2Query in terms of efficiency and effectiveness in addressing the vocabulary mismatch challenge. The Doc2Token feature has been deployed and evaluated online, resulting in a significant improvement in both relevance and revenue.



## References

- [1] R. Nogueira, W. Yang, J. Lin, K. Cho, Document expansion by query prediction, arXiv preprint arXiv:1904.08375 (2019).
- [2] J. J. Rocchio, Relevance feedback in information retrieval, in: G. Salton (Ed.), *The Smart retrieval system - experiments in automatic document processing*, Englewood Cliffs, NJ: Prentice-Hall, 1971, pp. 313–323.
- [3] G. A. Miller, Wordnet: a lexical database for english, *Communications of the ACM* 38 (1995) 39–41.
- [4] Y. Lv, C. Zhai, A comparative study of methods for estimating query language models with pseudo feedback, in: *Proceedings of the 18th ACM conference on Information and knowledge management*, 2009, pp. 1895–1898.
- [5] L. Liu, M. Li, J. Lin, S. Riedel, P. Stenetorp, Query expansion using contextual clue sampling with language models, arXiv preprint arXiv:2210.07093 (2022).
- [6] L. Wang, N. Yang, F. Wei, Query2doc: Query expansion with large language models, arXiv preprint arXiv:2303.07678 (2023).
- [7] R. Nogueira, J. Lin, A. Epistemic, From doc2query to docttttquery, *Online preprint* 6 (2019) 2.
- [8] T. Formal, B. Piwowarski, S. Clinchant, Splade: Sparse lexical and expansion model for first stage ranking, in: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021, pp. 2288–2292.
- [9] L. Xiong, C. Xiong, Y. Li, K.-F. Tang, J. Liu, P. Bennett, J. Ahmed, A. Overwijk, Approximate nearest neighbor negative contrastive learning for dense text retrieval, arXiv preprint arXiv:2007.00808 (2020).
- [10] C.-Y. Lin, Rouge: A package for automatic evaluation of summaries, in: *Text summarization branches out*, 2004, pp. 74–81.
- [11] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *The Journal of Machine Learning Research* 21 (2020) 5485–5551.
- [12] Y. Yang, L. Huang, M. Ma, Breaking the beam search curse: A study of (re-) scoring methods and stopping criteria for neural machine translation, arXiv preprint arXiv:1808.09582 (2018).
- [13] T. Kudo, J. Richardson, Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing, arXiv preprint arXiv:1808.06226 (2018).
- [14] D. Shahi, *Apache solr*, Springer, 2016.
- [15] Apache lucene, <http://lucene.apache.org>, 2019.
- [16] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, *CoRR abs/1810.04805* (2018).
- [17] A. See, P. J. Liu, C. D. Manning, Get to the point: Summarization with pointer-generator networks, arXiv preprint arXiv:1704.04368 (2017).