

Singapore Management University

## Institutional Knowledge at Singapore Management University

---

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

---

5-2023

### Stop words for processing software engineering documents: Do they matter

Yaohou FAN

Chetan ARORA

Christoph TREUDE

Singapore Management University, [ctreude@smu.edu.sg](mailto:ctreude@smu.edu.sg)

Follow this and additional works at: [https://ink.library.smu.edu.sg/sis\\_research](https://ink.library.smu.edu.sg/sis_research)



Part of the [Software Engineering Commons](#)

---

#### Citation

FAN, Yaohou; ARORA, Chetan; and TREUDE, Christoph. Stop words for processing software engineering documents: Do they matter. (2023). *Proceedings of the 2nd Workshop on Natural Language-based Software Engineering, 2023 May 20*. 40-47.

Available at: [https://ink.library.smu.edu.sg/sis\\_research/8912](https://ink.library.smu.edu.sg/sis_research/8912)

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email [cherylds@smu.edu.sg](mailto:cherylds@smu.edu.sg).

# Stop Words for Processing Software Engineering Documents: Do they Matter?

Yaohou Fan  
The University of Melbourne  
Australia  
yaohouf@student.unimelb.edu.au

Chetan Arora  
Monash University  
Australia  
chetan.arora@monash.edu

Christoph Treude  
The University of Melbourne  
Australia  
christoph.treude@unimelb.edu.au

**Abstract**—Stop words, which are considered non-predictive, are often eliminated in natural language processing tasks. However, the definition of uninformative vocabulary is vague, so most algorithms use general knowledge-based stop lists to remove stop words. There is an ongoing debate among academics about the usefulness of stop word elimination, especially in domain-specific settings. In this work, we investigate the usefulness of stop word removal in a software engineering context. To do this, we replicate and experiment with three software engineering research tools from related work. Additionally, we construct a corpus of software engineering domain-related text from 10,000 Stack Overflow questions and identify 200 domain-specific stop words using traditional information-theoretic methods. Our results show that the use of domain-specific stop words significantly improved the performance of research tools compared to the use of a general stop list and that 17 out of 19 evaluation measures showed better performance.

Online appendix: <https://zenodo.org/record/7865748>

**Index Terms**—Software Engineering Documents, Natural Language Processing (NLP), Stop Words.

## I. INTRODUCTION

With the development of natural language processing technology, the amount of text data has exploded in various fields. This phenomenon is also occurring in the field of software engineering, and text analysis using natural language processing (NLP) techniques has become increasingly prevalent in the field of software engineering [1], [2]. However, without effective pre-processing methods, raw text corpora often pose computational and analytical barriers to NLP tasks such as indexing, text classification and information retrieval [3]. Because of the large number of uninformative words present in the raw text, these words are defined as ‘stop words’. They appear frequently in text mining tasks and carry little information, simply connecting words in a sentence, such as the words ‘the’, ‘and’, ‘to’, and so on. Theoretically, the removal of stop words can improve the statistical significance of important terms [4], [5]. Therefore, stop word removal is an important step in text pre-processing.

Using a standard list of stop words to remove these high-frequency and uninformative words has become the standard in research and industry. Researchers often use off-the-shelf generic stop word lists, such as those provided by the Natural Language Processing Toolkit (NLTK), to remove noise [6]. The use of stop word lists is controversial because it is impossible to have a uniform list of stop words and the

semantic importance of each word depends on the task and usage. Blind use of stop word lists may result in the loss of important information, which may adversely affect the accuracy of text mining algorithms [7], [8]. However, stop words are difficult to define rigorously, so there are few papers that systematically investigate the impact of stop word removal on algorithm performance. This has led to the current disagreement on whether stop words should be retained or removed.

More reliable lists of stop words might be generated from a specific domain, as common words are distinct in different domains. The technical language used in software engineering texts is different from others; therefore, there might be specific stop words in the texts. Using existing stop lists may result in many useless terms remaining in the data. In fact, experts often maintain an ad hoc stop list [9], [10] manually. In contrast, few studies have captured the distribution of stop words in a specific domain, i.e., limiting the corpus to a certain domain to extract stop words.

Inspired by this, we address two research questions (RQs):

- RQ1 Does removing stop words from general stop lists improve the performance of algorithms that have been previously published?
- RQ2 Can stop lists generated from the software engineering (SE) domain positively affect downstream tasks in their domain?

To answer RQ1, this study replicated the tools proposed in three previous papers in the SE domain. After this, stop lists of different lengths were used to pre-process the raw text data and reproduce all tools. The analysis of the impact on the output of the tools was completed by looking at the proportion of words removed in the text data and measures of the performance of the tools used in the original papers.

We explored the SE-related corpus and used two automated methods to statistically count stop words to answer RQ2. These domain-specific stop words are examined in the same way in the replicated tools. We found that using stop lists generated from the public domain had a positive effect on the algorithms’ output in some evaluation metrics, but a negative effect in others. SE domain-specific stop word lists performed better than general lists, but it is important to take context into account before removing stop words as blind use can have a negative impact. We conclude the paper with a

systematic comparison and analysis of all stop lists. All stop word lists and scripts are available in our online appendix: <https://zenodo.org/record/7865748>

## II. RELATED WORK

Research on stop words dates back to 1957, when H.P. Luhn first introduced the concept of stop words [11]. Subsequent studies have systematically studied stop words and summarised seven properties [12]. Stop words are generally a single set of words, and most stop lists contain definite articles and coordinating conjunctions. Researchers have worked to extract stop words from general knowledge bases. The Brown Corpus is the most widely used corpus of general language, which contains about one million words. Many stop lists for general language have been generated from this corpus [13], [14]. In addition, Montemurro and Zanette [15] learned linguistic structures in a large corpus of written languages and exported a stop list.

Static stop lists drawn from a public knowledge base are seen by many practitioners as the standard and the default choice. However, Lo et al. [16] pointed out that ‘standard’ stop lists ignore domain knowledge properties. These ‘standard’ stop lists have become outdated over time, with the first English stop lists published in the 1970s [17]. The usage of some popular expressions has changed, due to social factors such as technological and cultural shifts. It is vital to update and revise the stop lists for current use. Other studies have similarly identified the drawbacks of ‘standard’ stop words [18], [19]. Experts have proposed many heuristics to overcome these limitations. The methods commonly used in the industry can be divided into the two categories discussed in the following.

### A. Methods based on Zipf’s Law

The study of word frequency in texts first appeared in the last century when George K. Zipf customised the rule of thumb to rank lists of words using the frequency of terms (tf) as a measure [20]. A study borrowed this concept by statistically calculating the relative importance of individual words and phrases [11], further developing the hypothesis that the most discriminatory words should appear in the middle of the frequency rank. Based on these frequency descriptions, Zipf’s law approach can be summarised as:

- Term-frequency High (TF-high): remove the most frequently occurring words.
- Term-frequency 1 (TF1): remove words that occur once.
- Inverse document frequency (IDF): remove words that occur frequently in multiple documents.

Myerson [21] combined TF-high and statistical methods to obtain 500 stop words from a Chinese corpus. The first step of their method was to filter the list of candidate stop words using the TF-high method. After this, the statistical correlation (Chi-squared statistic) between the candidate words and the text categories was calculated, and words with low correlation were identified as stop words. Their results showed that 500 stop words effectively removed 43% of the words from the corpus and improved the accuracy of the classification task

by 7%. Salton and Buckley [22] combined the TF-high and IDF methods to propose a method called TF-IDF, which is effective in representing the weight of terms in an arbitrary setting, where the weight is also the usefulness.

### B. Machine learning-based methods

The main objective of machine learning methods is to reduce the dimensionality of the document vector representation in the feature space [23]. They use a method of ranking terms to eliminate words that do not contain information. All term ranking methods are based on three components: (a) calculating importance values for each word based on statistical indicators; (b) ranking features in decreasing order of importance; and (c) establishing a threshold to filter the ranked list. Commonly used statistical indicators are:

- Information gain (IG) [24]: IG is the difference between the information entropy of a feature (word) before and after it appears in the text, and it considers the information load of the feature in the text. The smaller the load, the less important the feature becomes.
- Odds ratio [25]: It is a statistical measure that quantifies the strength of the association between two events (word and text). When the presence or absence of a word does not change the semantics of the text, they are independent of each other, so the word can be classified as a stop word.
- Expected cross entropy [26]: Cross entropy measures the information about the variability between two probability distributions, and in the calculation of stop lists, it indicates how difficult a word is for text classification. Therefore, words that do not contribute to text classification are evidently stop words.

Asubiaro [27] proposed an entropy-based method to determine the probability that Yoruba words are stop words. Their experiment calculated entropy for all words with non-noun lexical forms, and those with values greater than 0.6 were identified as stop words. The results show that a stop list of length 256 removes 66% of the words in the corpus, and they indicated that after translating the list into English, only 69.1% of the stop words appeared in the common English stop list.

## III. REPLICATED TOOLS AND THEIR EVALUATION METRICS

This paper replicates three tools implemented using NLP techniques relevant to the field of software engineering: an automated app review classifier [28], a natural language query recommendation tool [29], and a requirements change impact analysis tool [30], with the aim of answering RQ1 and RQ2 using different stop lists in these three downstream tasks. We selected these tools based on the availability and operability of their replication packages.

### A. Tool 1 - App Reviews Classification

Numerous studies have investigated the classification of texts related to software engineering, such as Application Programming Interface (API) documents [31] and README files [32]. Maalej et al. used machine learning algorithms to

build classifiers of user reviews, thus helping developers filter and process useful information from reviews [28].

- Data: (1) Their experimental data consists of 4400 reviews from the Apple App Store and Google Play Store. These reviews were divided into four categories, 737 User Experience (UE), 378 Bug Report (BP), 299 Feature Request (FR), and 2721 Rating (RT). (2) List of stop words used to purify text: NLTK stop words (179 words).
- Model: A Naïve Bayes binary classifier was constructed for each of the four categories, and the results were then aggregated to obtain the final classification output.
- Evaluation Metrics:
  - Precision: refers to the proportion of samples classified as positive that are actually positive.
  - Recall: is the fraction of the true positive sample that was retrieved.
  - F1 measure: Harmonised mean of precision and recall.

Note that these reproduced models did not perform exactly the same as reported in previous work [28], since we were unable to access their source code. We refer readers to Maalej et al. [28] for the comparison. Only the replicated model from our own implementation was used for subsequent analysis and evaluation.

### B. Tool 2 - Query Recommendation Tool (RACK)

API retrieval is another application of NLP techniques in the field of software engineering. API retrieval tools accept natural language queries from developers as input and then return a set of related API classes or methods automatically. A query recommendation tool, called RACK, proposed by Rahman and Lo [29], first captures the intention of queries from code annotations in the IDE and then converts these queries into relevant API classes.

- Model: Such a conversion is achieved in two steps with the Keyword-API association in Stack Overflow and the Query Reformulation:
  - Keyword-API Mapping Database: 344K pairs of questions and answers from the Stack Overflow website were analysed to generate intrinsic associations between the keywords in question and the API classes, and the associations were used to build the Keyword-API Mapping Database.
  - Query Reformulation: In this step, the raw unstructured natural language query input is pre-processed (i.e., tokenization, stop word removal, stemming) using NLP techniques, and a keyword vector is obtained; RACK accesses the database to collect candidate API classes and sorts them heuristically to output a sorted list of relevant API classes.
- Data: The authors provide 175 natural language queries with ground truth. This study used these queries, then evaluated the accuracy of the recommendations by varying the stop list used in the pre-processing phase. In

addition, a stop list of length 797 is used to reformulate the queries.

- Evaluation Metric:
  - Top-10: Refers to the percentage of search queries in which at least one API class was correctly recommended by the recommendation technology in the Top-10 results.
  - Mean Reciprocal rank@10 (MRR@10): Reciprocal Rank@10 is the multiplicative reciprocal of the first relevant API class ranking in the Top-10 results returned by the tool. Mean Reciprocal Ranking@10 (MRR@10) is the average of such metrics across all search queries in the dataset.
  - Mean Average Precision@10 (MAP@10): Precision@10 is the precision of all relevant APIs that appear in the Top-10 recommendation list. Average Precision@10 (AP@10) is the average Precision@10 for relevant APIs recommended for a natural language (NL) query. Mean Average Precision@10 is the average of the AP@10 for all NL queries.
  - Mean Recall@10 (MR@10): Recall@10 represents the percentage of Top-10 results for which this tool correctly recommends all relevant APIs for an NL query. Mean Recall@10 (MR@10) is the average of these measures for all queries.

### C. Tool 3 - Requirements Change Impact Analysis

During the development of software, requirements are frequently changed to ensure that the final product meets expectations. Changing one requirement in a requirements specification requires the entire specification to be modified to keep it correct and consistent. Manually analysing the impact of a change to one requirement on other requirements is time consuming. A natural language processing approach has been proposed to deal with the impact of requirement changes [30], which takes into account the phrase structure of requirement statements and captures the change propagation conditions to achieve a more accurate analysis of the impact of changes. The proposed approach is given by Arora et al. [30], [33].

- Data: The original paper tested the proposed approach using 14 scenarios from the requirements specifications of two industrial cases (case A and case B). For this study, we only tested three scenarios from case B due to resource limitations: query 2, query 4, and query 5.
- Stop list: Generated from Brown Corpus
- Evaluation metrics: Area under the curve (AUC): Commonly used in statistics and machine learning to assess the performance of a binary classification model by measuring the ranking performance of the model. Ranking performance refers to the probability of ranking a positive sample ahead of a negative sample in classification. In this experiment, it assesses the quality of the ranking of the impacted requirements.

#### IV. STOP WORD LIST CONSTRUCTION

To answer RQ2, we used a corpus from Stack Overflow, a programming question and answer website. Today, the field of software engineering is made up of a wide variety of technologies, languages, and platforms, and developers are required to have a broad skill set. Even experienced senior software engineers find it difficult to keep up with the pace of evolution of technology. Thus, Stack Overflow leverages the expertise of its users to provide answers to technical concerns, and over time, it has become a repository of software engineering knowledge. Meanwhile, the question-and-answer text on the site is generated by engineers, and much of the text consists of technical terms. It is therefore one of the largest corpora in software engineering. Some studies have mined the textual information contained therein, e.g., [34].

We crawled the Stack Overflow question API with four popular tags ('javascript', 'css', 'python2' and 'deep learning'), focusing on questions that have been answered. We sort these questions by the number of times they have been viewed, taking the top 2500 from each category, for a total of 10,000 questions crawled as a corpus for the software engineering domain.

##### A. Data processing

So far, the natural language text in the corpus is unstructured and cleaning the original text using regular expressions is necessary to improve the accuracy of stop word identification. The most common pre-processing steps are applied: (1) removing mentions, (2) removing URLs, (3) removing hashtags, (4) removing numbers, (5) removing text in brackets, (6) removing non-English words, (7) removing extraneous space.

After this, the size of the corpus is reduced by lowercasing and lemmatizing all words. For grammatical reasons, the original corpus contains different forms of a word, such as see, saw, and seen, but they have the same semantic meaning. Lemmatization analyses the morphology of the word, thus removing the flexion endings to return the basic form of the word.

##### B. TF-IDF based method for identifying stop words

There is a long history of using the TF-IDF method and its variants to separate stop words and content words from the corpus. Extensive experience has shown it to be an effective method [13], [35], [36]. TF-IDF is the product of two quantities TF and IDF. Based on information theory, TF estimates the probability of occurrence of a term when normalised by the total frequency in the document or set of documents. The calculation of term frequency of the  $i$ -th term in the  $j$ -th document is given by:

$$tf_{ij} = \log(1 + freq(t, d)) \quad (1)$$

The inverse document frequency is calculated by giving the ratio of the total number of documents  $|D|$  to the amount of documents that contain the term  $t$  in logarithmic form, and

hence the importance  $t$  in the document is denoted by the following expression:

$$idf = \log\left(\frac{|D|}{count(d \in D : t \in d)}\right) \quad (2)$$

The importance of a word is expressed as the product of the probability of its occurrence and the amount of information it represents. This paper identifies words that occur in many documents (high  $df$ ) and words that occur less frequently (low  $tf$ ) as stop words and selects the top 200 words for the stop list based on TF-IDF scores.

##### C. Stop words recognition based on Poisson distribution

On the other hand, Church and Gale [37] have studied the statistical distribution of words and have proposed the hypothesis that stop words follow a Poisson distribution. We refer to this statistical method of computing stop word lists as a Poisson stop list. There have been studies that validate this unsupervised method by creating a stop list in a corpus of Polish [38]. Under their assumptions, the document frequency ( $df$ ) of words can be approximated from the term frequency ( $tf$ ) and the number of documents ( $N$ ), by using probability theory:

$$\frac{\text{Estimated } df}{N} = 1 - P(\text{term does not occur}) \quad (3)$$

Assuming a Poisson distribution for stop words, with  $\mu$  as the average number of occurrences of a word in each document, the probability of the word occurring  $k$  times is:

$$\mu = \frac{tf}{N} \quad (4)$$

$$P(k, \mu) = \frac{e^{-\mu} * \mu^k}{k!} \quad (5)$$

Randomly distributed words (stop words) should yield an estimated  $df/df$  close to 1, while highly cluttered words (key words) show an increasing trend in value. In this paper, the words in the corpus with an estimated  $df/df$  of 1 were calculated as candidate stop words, and after this a stop list of 200 length was randomly selected in order to keep the lengths of the stop lists consistent.

Figure 1 illustrates the process of generating a stop list from an SE domain-specific corpus for this paper. First, the raw data is pre-processed and then the lists are ranked using statistical metrics (TF-IDF and Poisson distribution). Following this, the first author leveraged his domain expertise to manually eliminate unstructured terms, as the regular expressions did not fully clean the text. Finally, a final list related to the software engineering domain is completed.

TABLE I  
OTHER FREELY AVAILABLE STOP WORD LISTS USED IN THIS EXPERIMENT

ID	Stop word list	Length		Source URL
1	Ranks NL (Google)	32	(very small)	http://www.ranks.nl/stopwords
2	MongoDB	174	(small)	https://github.com/mongodb/mongo/blob/master/src/mongo/db/fts/stop_words_english.txt
3	Azure Gallery	311	(medium)	https://gallery.azure.ai/Experiment/How-to-modify-default-stopword-list-1
4	Alir3z4	1298	(large)	https://github.com/Alir3z4/stopwords/blob/master/english.txt
5	Technology Domain	87		https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0254937




Fig. 1. The flowchart for generating a domain-specific stop list

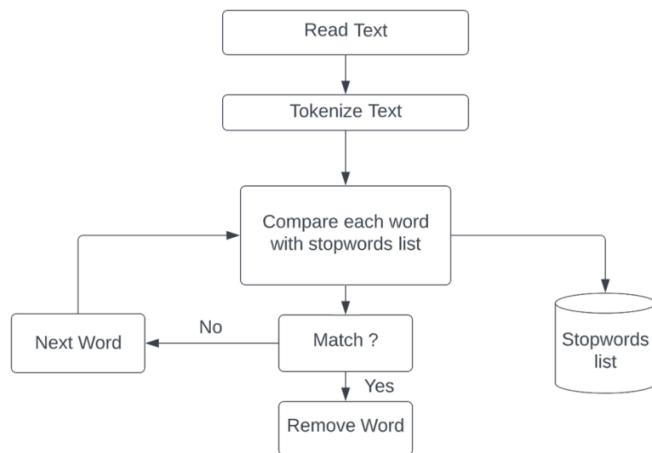


Fig. 2. Classic workflow of stop word removal

#### D. Other stop word lists used in the experiment

To answer RQ1, we collected five existing stop lists that are publicly available, as shown in Table I, four of which are generated from public knowledge bases and widely used in NLP pre-processing tasks. The last stop list is a stop list in a technical domain [39].

Stop word lists 1-4 are not related to a specific domain; they are generated from language in common use and range in length from 32 to 1298. We sort the lengths of stop word lists from small to large and define them as ‘very small’, ‘small’, ‘medium’, and ‘large’. We note that there is no existing research to divide the lengths of stop word lists. The reason for doing this is to facilitate subsequent comparisons.

Sarica and Luo [39] collected natural language texts that describe technologies at all levels to identify stop words in technical language texts and constructed a database. They combined algorithms such as TF, TF-IDF, and information entropy to automatically identify candidate stop words. Human experts then evaluated the importance of these candidates in filtering out the real stop words.

#### E. Stop word Removal

We used the classic stop word removal method to sanitise the raw data. The classic algorithm accepts an English text as input and outputs the text after all stop words have been removed. The process is illustrated in Figure 2.

TABLE II

NUMBER OF STOP WORDS REMOVED FOR APP REVIEW CLASSIFICATION TOOL

Stop list	Length	words remaining	% removed
Very small	32	7849	18.70%
Small	174	7734	19.89%
Medium	311	7656	20.70%
Large	1298	7257	24.84%
Technology Domain	87	7823	18.97%
SE Domain (TF-IDF)	200	7617	21.10%
SE Domain (Poisson)	200	7612	21.16%

## V. EXPERIMENT ANALYSIS

### A. Tool 1 - App Reviews Classification

1) *Stop list Evaluation*: Eliminating stop words usually results in a significant reduction in the number of words in the corpus, allowing for a more accurate approach to text mining. This paper summarises how many stop words were removed from the raw text data using different stop lists for this replicated tool. Note that after lowercasing and lemmatizing of the raw text data, there are 9655 unique words remaining.

Table II shows that the number of words retained after pre-processing is negatively correlated with the length of the stop list, and a large stop list significantly removes more stop words than a small stop list. Additionally, the stop lists generated in the SE domain resulted in fewer words remaining compared

to the longer stop lists (medium), which implicitly indicates that they contain more accurate domain-related stop words.

2) *Model Performance Evaluation:* Tables III and IV represent the performance of the models obtained after pre-processing with different stop lists, respectively. The small and large stop lists are the best general lists for Tool 1, as their models perform better overall, outperforming the replicated model on 11 of 12 measures. However, it is interesting to note that as the length of the general stop list grows, the precision metric in the RT classification tends to decrease, when all other metrics are getting better. The opposite trend is observed in the Recall indicator, where recall values increase with the length of the stop list in almost all four categories.

It is also clear that stop words extracted from the SE domain perform better than those from the technical domain, which is to be expected since there is SE-related text in user application reviews; for example, some SE domain lexicons appear in feature request reviews. Furthermore, domain-specific stop lists performed no better than generic ones of the same volume in this downstream task. This is because most of the vocabulary in the user review text data is derived from the knowledge base of the end user, meaning that the body of the text is still composed of general language.

#### B. Tool 2 - Query Recommendation Tool (RACK)

1) *Stop list Evaluation:* Table V shows the status of the lexicon in the corpus after the removal of stop words using different stop lists. For the API recommendation tool, stop lists identified from the general knowledge base perform poorly, and the longest stop list removes only 15.76% of the words. In contrast, the SE domain-specific stop list with only 200 words removed about 20% of the words in the original text. The stop list generated from the technology domain does not contain relevant stop words, as it removed almost no words from the text.

2) *Model Performance Evaluation:* From the experimental results (Table VI and Table VII), the size of the stop word list has very little effect on the performance of the model, with all general stop lists giving similar performance. The model even performs stable with no stop word removal. However, other stop lists, such as the large stop list, have a negative impact on the model behaviour. We speculate that the large stop list removed useful information from the data, resulting in a weaker performance of the model. In addition, the smaller stop lists (very small and small) gave the same results as those without the stop list. This suggests that, in the context of this task, these lists did not contain the appropriate stop words and did not reformulate the query well. This validates our previous hypothesis that the definitions of stop words are different between domains and that stop words generated from generic knowledge bases would not maximize the improvement of the algorithm in domain-specific tasks.

The stop list from the SE domain was superior to any other list, and the TF-IDF method outperformed the original performance on all four metrics. Furthermore, the Poisson stop list exceeded the original performance only for the Top-10 and

MR@K metrics. On the contrary, the model obtained using the stop list of the technology domain performed the same as if the stop words were not removed, indicating that this stop list does not contain the appropriate stop words of the SE domain.

#### C. Tool 3 - Requirements Change Impact Analysis

Table VIII shows a very clear trend in the experimental outcomes, with all of the common stop word lists performing very similarly, except for the medium stop list and the large one. The medium stop list appeared to be weak in Query 5, only obtaining an AUC of 0.447 down from the original performance of 0.561. However, it performed the best of all stop lists in Query 2, achieving an AUC close to 0.6. The large stop list performed poorly in all three scenarios, which is consistent with the experimental finding for Tool 2, where some larger stop lists removed some valuable words. Furthermore, we found that the model performed surprisingly better in Query 5 without stop word removal.

Domain-specific stop lists improved the performance of the model (Table IX), and all lists outperformed the original model in two of the three cases. SE domain-specific stop words had the most significant effect on the model performance, producing an AUC value of 0.602 in Query 5 compared to the original performance of 0.561.

#### D. Overall Stop list Performance Comparison

The overall performance of the stop lists generated based on a total of 19 metrics (i.e., all evaluation metrics used by the three replicated tools) used in the experiment is given in Table X. The fact that 12 of the metrics obtained worse results without the stop list is an indication that the stop list can have a positive effect on the model output. We see that among the non-SE domain-specific stop lists, the small stop list with 174 stop words gives the best performance, improving the performance of the tools on 13 metrics. The other stop lists have less of a positive effect on the performance of the tools, as some indicators show worse performance. The stop lists generated from the corpus of other domains (technology domain) performed negatively. Finally, the SE domain-related stop lists gave the best performance, especially the one generated using TF-IDF, which had a positive impact on the model for 17 of the 19 metrics.

## VI. CONCLUSION

This paper evaluates the performance of existing tools by removing stop words in three downstream tasks related to the SE domain using various stop lists. Combining the experimental results and analysis, we can provide answers to our two Research Questions in this section.

The experimental results in Table X answer RQ1, whether general stop lists have a positive effect on the performance of algorithms' evaluation measures and tasks. While using stop lists generated from the public domain had a positive effect on the algorithms' output in some of the evaluation metrics, other metrics showed that the model gave worse results. For example, a larger stop list increased the F1 measure value for

TABLE III

PERFORMANCE OF GENERAL STOP LIST FOR APP REVIEW CLASSIFICATION TOOL. GREEN INDICATES BETTER PERFORMANCE THAN THE ORIGINAL, ORANGE WORSE PERFORMANCE THAN THE ORIGINAL.

	PD (bug report)			RT (rating)			FR (feature request)			UE (user experience)		
	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1
Original (NLTK stop list)	9.9%	29.5%	14.8%	73.8%	72.9%	73.3%	5.7%	17.9%	8.7%	11.6%	27.4%	16.3%
No stop list	10.8%	28.6%	15.7%	75.6%	71.0%	73.3%	5.3%	15.5%	7.9%	8.7%	18.8%	11.9%
Very small	12.2%	35.7%	18.2%	74.9%	74.8%	74.9%	5.7%	19.0%	8.8%	10.3%	23.9%	14.4%
Small	12.7%	50.0%	20.3%	72.1%	78.2%	75.0%	6.8%	28.6%	10.9%	12.2%	34.0%	18.0%
Medium	12.7%	52.7%	20.5%	71.0%	81.0%	75.6%	7.8%	34.5%	12.7%	11.2%	32.0%	16.6%
Large	10.4%	56.2%	17.5%	68.6%	83.1%	75.1%	7.5%	48.8%	13.0%	12.5%	45.2%	19.6%

TABLE IV

PERFORMANCE OF THE DOMAIN-SPECIFIC STOP LIST FOR THE APP REVIEW CLASSIFICATION TOOL

	PD (bug report)			RT (rating)			FR (feature request)			UE (user experience)		
	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1	Pre	Rec	F1
Original	9.9%	29.5%	14.8%	73.8%	72.9%	73.3%	5.7%	17.9%	8.7%	11.6%	27.4%	16.3%
Technology domain	10.5%	30.4%	15.6%	75.4%	71.2%	73.2%	4.4%	14.3%	6.7%	12.2%	34.0%	18.0%
SE domain (Poisson)	10.0%	37.5%	15.8%	72.1%	78.0%	74.9%	7.1%	29.8%	11.5%	11.6%	32.0%	17.0%
SE domain (TF-IDF)	10.7%	40.2%	16.9%	72.2%	78.2%	75.1%	7.9%	33.3%	12.8%	11.7%	32.5%	17.2%

TABLE V

NUMBER OF STOP WORDS REMOVED FOR QUERY RECOMMENDATION TOOL

Stop list	Length	words remaining	% removed
Very small	32	389	4.18%
Small	174	380	6.40%
Medium	311	369	9.85%
Large	1298	342	15.76%
Technology Domain	87	400	1.48%
SE Domain (TF-IDF)	200	328	19.21%
SE Domain (Poisson)	200	326	20.44%

TABLE VI

PERFORMANCE OF THE GENERAL STOP LIST FOR QUERY RECOMMENDATION TOOL

	Top-10	MRR@10	MAP@10	MR@K
Original	83.43%	52.29%	45.73%	54.07%
No stop words	81.14%	50.72%	44.94%	51.10%
Very Small	81.14%	50.72%	44.94%	51.10%
Small	81.14%	50.72%	44.94%	51.10%
Medium	81.14%	50.97%	44.83%	51.91%
Large	80.57%	49.71%	44.07%	51.03%

TABLE VII

PERFORMANCE OF DOMAIN-SPECIFIC STOP LIST FOR QUERY RECOMMENDATION TOOL

	Top-10	MRR@10	MAP@10	MR@K
Original	83.43%	52.29%	45.73%	54.07%
Technology domain	81.14%	50.72%	44.94%	51.10%
SE domain (Poisson)	83.85%	52.29%	43.27%	54.47%
SE domain (TF-IDF)	84.17%	53.20%	45.82%	56.18%

TABLE VIII

PERFORMANCE (NORMALIZED AUC) OF THE GENERAL STOP LIST FOR TOOL 3

	Original	None	V. Small	Small	Medium	Large
Query 2	0.585	0.585	0.588	0.588	0.595	0.538
Query 4	0.981	0.981	0.981	0.981	0.981	0.961
Query 5	0.561	0.579	0.579	0.577	0.447	0.462

TABLE IX

PERFORMANCE (NORMALIZED AUC) OF THE DOMAIN-SPECIFIC STOP LIST FOR TOOL 3

	Original	Technology domain	SE domain (TF-IDF)	SE domain (Poisson)
Query 2	0.585	0.588	0.588	0.588
Query 4	0.981	0.981	0.981	0.981
Query 5	0.561	0.579	0.602	0.602

TABLE X  
OVERALL PERFORMANCE

Stop word list	better	worse	same
No stop words	4	12	3
Very Small	10	7	2
Small	13	5	1
Medium	11	7	1
Large	11	8	0
Technology domain	9	9	1
SE domain (Poisson)	12	5	2
SE domain (TF-IDF)	17	1	1



Tool 1 but led to a decrease in recall. On the other hand, experiments conducted in Tool 2 and Tool 3 showed that all ‘standard’ stop lists had little effect on the algorithm and the model performed similarly to that without the removal of stop words. For RQ2, the SE domain-related stop list showed a positive effect for all three tools and led to the best overall performance. Stop lists generated using TF-IDF outperformed the Poisson method, with 17 of 19 indicators performing better than the original. Although the Poisson method performs similarly to the small stop list, it is superior to other general lists. Therefore, we argue that SE domain-specific stop word lists performed better than general lists.

An overall conclusion can be drawn that it is not possible to have a standard list of stop words. The semantics of each word are different in each domain. The context of the task needs to be taken into account before removing stop words, and the use of domain-specific stop lists can have a positive effect. However, blind use of a stop list may have a negative impact on the results of the algorithm.

## VII. LIMITATIONS AND FUTURE IMPROVEMENT

Future studies can contribute in two directions, domain-specific corpora and algorithms for identifying stop words. The text in the domain-specific corpus used for this paper comes from Stack Overflow. However, such a corpus is biased because there is a wide variety of programming languages used in SE. Therefore, a corpus with more data and broader coverage specific to the SE domain needs to be constructed to support more accurate stop word generation.

On the other hand, we did not use more advanced algorithms to generate stop words to exclude variants, and only the most commonly used methods were employed. These traditional methods are flawed because they are based on some naive assumptions that do not hold in all cases. The current experiment has demonstrated the positive effect of domain-specific stop words in SE using a baseline approach. Future work could investigate stop words in their domain further using more advanced algorithms as well as the impact of different stop word lists on different software engineering tasks.

## VIII. DATA AVAILABILITY

All stop word lists and scripts are available in our online appendix: <https://zenodo.org/record/7865748>

## REFERENCES

- [1] J. Murphy, K. Fu, K. Otto, M. Yang, D. Jensen, and K. Wood, “Function based design-by-analogy: a functional vector approach to analogical search,” *Journal of Mechanical Design*, vol. 136, no. 10, 2014.
- [2] M. W. Glier, D. A. McAdams, and J. S. Linsey, “Exploring automated text classification to improve keyword corpus search results for bio-inspired design,” *Journal of Mechanical Design*, vol. 136, no. 11, 2014.
- [3] J. Kaur and P. K. Buttar, “A systematic review on stopword removal algorithms,” *International Journal on Future Revolution in Computer Science & Communication Engineering*, vol. 4, no. 4, pp. 207–210, 2018.
- [4] G. Manco, E. Masciari, M. Ruffolo, and A. Tagarelli, “Towards an adaptive mail classifier,” in *Proc. of Italian Association for Artificial Intelligence Workshop*, 2002.
- [5] M. Choy, “Effective listings of function stop words for twitter,” *arXiv preprint arXiv:1205.6396*, 2012.
- [6] S. Bird, E. Klein, and E. Loper, *Natural language processing with Python: analyzing text with the natural language toolkit*. ” O’Reilly Media, Inc.”, 2009.
- [7] M. P. Sinka and D. W. Corne, “Towards modernised and web-specific stoplists for web document analysis,” in *Proceedings IEEE/WIC Int’l. Conf. on Web Intelligence*. IEEE, 2003, pp. 396–402.
- [8] M. Makrehchi and M. S. Kamel, “Automatic extraction of domain-specific stopwords from labeled documents,” in *European Conference on Information Retrieval*. Springer, 2008, pp. 222–233.
- [9] S. Sarica, J. Luo, and K. L. Wood, “Technet: Technology semantic network based on patent data,” *Expert Systems with Applications*, vol. 142, p. 112995, 2020.
- [10] K. Seki and J. Mostafa, “An application of text categorization methods to gene ontology annotation,” in *Proc. 28th annual international ACM SIGIR conference on research and development in information retrieval*, 2005, pp. 138–145.
- [11] H. P. Luhn, “A statistical approach to mechanized encoding and searching of literary information,” *IBM Journal of research and development*, vol. 1, no. 4, pp. 309–317, 1957.
- [12] J. Kaur and P. K. Buttar, “Stopwords removal and its algorithms based on different methods,” *International Journal of Advanced Research in Computer Science*, vol. 9, no. 5, pp. 81–88, 2018.
- [13] C. Fox, “A stop list for general text,” in *Acm sigir forum*, vol. 24, no. 1-2. ACM New York, NY, USA, 1989, pp. 19–21.
- [14] G. V. Maverick, “Computational analysis of present-day american english,” 1969.
- [15] M. A. Montemurro and D. H. Zanette, “Towards the quantification of the semantic information encoded in written language,” *Advances in Complex Systems*, vol. 13, no. 02, pp. 135–153, 2010.
- [16] R. T.-W. Lo, B. He, and I. Ounis, “Automatically building a stopword list for an information retrieval system,” in *Journal on Digital Information Management: Special Issue on the 5th Dutch-Belgian Information Retrieval Workshop*, vol. 5, 2005, pp. 17–24.
- [17] C. Van Rijsbergen, “Information retrieval: theory and practice,” in *Proc. Joint IBM/University of Newcastle upon Tyne Seminar on Data Base Systems*, vol. 79, 1979.
- [18] S. Chakrabarti, B. Dom, R. Agrawal, and P. Raghavan, “Using taxonomy, discriminants, and signatures for navigating in text databases,” in *VLDB*, vol. 97, 1997, pp. 446–455.
- [19] C. Silva and B. Ribeiro, “The importance of stop word removal on recall values in text categorization,” in *Proc. International Joint Conference on Neural Networks, 2003.*, vol. 3. IEEE, 2003, pp. 1661–1666.
- [20] G. Zipf, “Selective studies and the principle of relative frequency in language,” *Human Behavior and the Principle of Least-Effort*, 1949.
- [21] R. B. Myerson *et al.*, “Fundamentals of social choice theory,” *Quarterly Journal of Political Science*, vol. 8, no. 3, pp. 305–337, 2013.
- [22] G. Salton and C. Buckley, “Term-weighting approaches in automatic text retrieval,” *Information processing & management*, vol. 24, no. 5, pp. 513–523, 1988.
- [23] A. Aizawa, “An information-theoretic perspective of tf-idf measures,” *Information Processing & Management*, vol. 39, no. 1, pp. 45–65, 2003.
- [24] Y. Yang and X. Liu, “A re-examination of text categorization methods,” in *Proc. 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, 1999, pp. 42–49.
- [25] D. Mladenić, “Feature subset selection in text-learning,” in *European conference on machine learning*. Springer, 1998, pp. 95–100.
- [26] D. Koller and M. Sahami, “Hierarchically classifying documents using very few words,” Stanford InfoLab, Tech. Rep., 1997.
- [27] T. V. Asubiaro, “Entropy-based generic stopwords list for yoruba texts,” *International Journal of Computer and Information Technology*, vol. 2, no. 5, 2013.
- [28] W. Maalej, Z. Kurtanović, H. Nabil, and C. Stanik, “On the automatic classification of app reviews,” *Requirements Engineering*, vol. 21, no. 3, pp. 311–331, 2016.
- [29] M. M. Rahman, C. K. Roy, and D. Lo, “Rack: Code search in the ide using crowdsourced knowledge,” in *2017 IEEE/ACM 39th Int’l. Conf. on Software Engineering Companion*. IEEE, 2017, pp. 51–54.
- [30] C. Arora, M. Sabetzadeh, A. Goknil, L. C. Briand, and F. Zimmer, “Change impact analysis for natural language requirements: An nlp approach,” in *2015 IEEE 23rd International Requirements Engineering Conference*. IEEE, 2015, pp. 6–15.
- [31] G. Petrosyan, M. P. Robillard, and R. De Mori, “Discovering information explaining api types using text classification,” in *2015 IEEE/ACM 37th*

- IEEE Int'l. Conf. on Software Engineering*, vol. 1. IEEE, 2015, pp. 869–879.
- [32] G. A. A. Prana, C. Treude, F. Thung, T. Atapattu, and D. Lo, “Categorizing the content of github readme files,” *Empirical Software Engineering*, vol. 24, no. 3, pp. 1296–1327, 2019.
- [33] C. Arora, M. Sabetzadeh, A. Goknil, L. C. Briand, and F. Zimmer, “Narcia: an automated tool for change impact analysis in natural language requirements,” in *Proc. 2015 10th Joint Meeting on Foundations of Software Engineering*, 2015, pp. 962–965.
- [34] C. Treude and M. P. Robillard, “Augmenting api documentation with insights from stack overflow,” in *2016 IEEE/ACM 38th Int'l. Conf. on Software Engineering*. IEEE, 2016, pp. 392–403.
- [35] A. Aizawa, “The feature quantity: an information theoretic perspective of tfidf-like measures,” in *Proc. 23rd annual international ACM SIGIR conference on Research and development in information retrieval*, 2000, pp. 104–111.
- [36] M. Konchady, *Text mining application programming*. Charles River Media, Inc., 2006.
- [37] K. W. Church and W. A. Gale, “Poisson mixtures,” *Natural Language Engineering*, vol. 1, no. 2, pp. 163–190, 1995.
- [38] M. Jungiewicz and M. Łopuszyński, “Unsupervised keyword extraction from polish legal texts,” in *Int'l. Conf. on natural language processing*. Springer, 2014, pp. 65–70.
- [39] S. Sarica and J. Luo, “Stopwords in technical language processing,” *Plos one*, vol. 16, no. 8, p. e0254937, 2021.