



APIDocBooster: An Extract-Then-Abstract Framework Leveraging Large Language Models for Augmenting API Documentation

Journal:	<i>Transactions on Software Engineering and Methodology</i>
Manuscript ID	TOSEM-2023-0432
Manuscript Type:	Paper
Date Submitted by the Author:	21-Nov-2023
Complete List of Authors:	<p>Yang, Chengran ; Singapore Management University, SCIS Liu, Jiakun; Singapore Management University, School of Computing and Information Systems Xu, Bowen; North Carolina State University at Raleigh Treude, Christoph; The University of Melbourne School of Computing and Information Systems, School of Computing and Information Systems Lyu, Yunbo; Singapore Management University, SCIS He, Junda; Singapore Management University, SCIS Li, Ming; Nanjing University Lo, David; Singapore Management University, SCIS</p>
Computing Classification Systems:	<pre>\begin{CCSXML} <ccs2012> <concept> <concept_id>10011007.10011074.10011111.10010913</concept_id> <concept_desc>S</pre>



School of Computing and Information Systems (SCIS)
Singapore Management University (SMU)

Dear Editor,

We are pleased to submit our paper entitled "APIDocBooster: An Extract-Then-Abstract Framework Leveraging Large Language Models for Augmenting API Documentation" co-authored by Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo, for consideration of publication in the journal, ACM Transactions on Software Engineering and Methodology.

Our paper tackles the prevalent issue where existing API documents often fall short in comprehensiveness. Despite some approaches have been proposed to augment API documentation by extracting complementary information from Stack Overflow, these approaches often suffer from inherent readability limitations. On the other hand, through our empirical analysis on leveraging GPT-4 for API documentation augmenting, we have found that while GPT-4 produces coherent and concise summaries, it lacks in informativeness and faithfulness.

To bridge this gap, our work introduces APIDocBooster, an extractive-then-abstractive framework for API document augmentation. This two-stage framework seamlessly fuse the advantages of both extractive (i.e., enabling faithful summaries without length limitation) and abstractive summarization (i.e., producing coherent and concise summaries). APIDocBooster consists of two stages: (1) Context-aware Sentence Section Classification (CSSC) and (2) UPdate SUMmarization (UPSUM). CSSC classifies API-relevant information collected from multiple sources into API documentation sections. UPSUM first generates extractive summaries distinct from the original API documentation and then generates abstractive summaries guided by extractive summaries through in-context learning.

To enable automatic evaluation of APIDocBooster, we construct the first dataset for API documentation augmentation. Our automatic evaluation results reveal that each stage in APIDocBooster outperforms its baselines by a large margin. Our human evaluation also demonstrates the superiority of APIDocBooster over GPT-4.

We look forward to hearing your comments and suggestions on this work.

Yours Sincerely,

Chengran, Jiakun, Bowen, Christoph, Yunbo, Junda, Ming, and David

Singapore Management University

APIDocBooster: An Extract-Then-Abstract Framework Leveraging Large Language Models for Augmenting API Documentation

CHENGRAN YANG and JIAKUN LIU, Singapore Management University, Singapore

BOWEN XU, North Carolina State University, USA

CHRISTOPH TREUDE, University of Melbourne, Australia

YUNBO LYU and JUNDA HE, Singapore Management University, Singapore

MING LI, Nanjing University, China

DAVID LO, Singapore Management University, Singapore

API documentation is often the most trusted resource for programming. Many approaches have been proposed to augment API documentation by summarizing complementary information from external resources such as Stack Overflow. Existing extractive-based summarization approaches excel in producing faithful summaries that accurately represent the source content without input length restrictions. Nevertheless, they suffer from inherent readability limitations. On the other hand, our empirical study on the abstractive-based summarization method, i.e., GPT-4, reveals that GPT-4 can generate coherent and concise summaries but presents limitations in terms of informativeness and faithfulness.

We introduce APIDocBooster, an extract-then-abstract framework that seamlessly fuses the advantages of both extractive (i.e., enabling faithful summaries without length limitation) and abstractive summarization (i.e., producing coherent and concise summaries). APIDocBooster consists of two stages: (1) Context-aware Sentence Section Classification (CSSC) and (2) Update SUMmarization (UPSUM). CSSC classifies API-relevant information collected from multiple sources into API documentation sections. UPSUM first generates extractive summaries distinct from original API documentation and then generates abstractive summaries guided by extractive summaries through in-context learning.

To enable automatic evaluation of APIDocBooster, we construct the first dataset for API documentation augmentation. Our automatic evaluation results reveal that each stage in APIDocBooster outperforms its baselines by a large margin. Our human evaluation also demonstrates the superiority of APIDocBooster over GPT-4 and shows that it improves the informativeness, relevance and faithfulness by 13.89%, 15.15%, and 30.56%, respectively.

CCS Concepts: • **Software and its engineering** → **Documentation**.

Additional Key Words and Phrases: Summarization, Question Retrieval

ACM Reference Format:

Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo. 2023. APIDocBooster: An Extract-Then-Abstract Framework Leveraging Large Language Models for Augmenting

Authors' addresses: Chengran Yang, cryang@smu.edu.sg; Jiakun Liu, jkliu@smu.edu.sg, Singapore Management University, Singapore; Bowen Xu, bxu22@ncsu.edu, North Carolina State University, USA; Christoph Treude, christoph.treude@unimelb.edu.au, University of Melbourne, Australia; Yunbo Lyu, yunbolyu@smu.edu.sg; Junda He, jundahe@smu.edu.sg, Singapore Management University, Singapore; Ming Li, lim@lamda.nju.edu.cn, Nanjing University, China; David Lo, davidlo@smu.edu.sg, Singapore Management University, Singapore.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1049-331X/2023/0-ART0 \$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

0:2 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

API Documentation. *ACM Trans. Softw. Eng. Methodol.* 0, 0, Article 0 (2023), 23 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

The application programming interface (API) is one of the most vital components of modern application development. Software developers typically rely on API reference documentation (API documentation in short) to learn APIs [45, 46, 58]. API documentation is a set of documents indexed by the API name, where each document provides information about a specific API [33]. According to the 2023 Stack Overflow Developer Survey [37], technical documentation is the most trusted resource for programming.

However, creating and maintaining high-quality and readable API documentation still requires significant effort [12]. Existing API documents are often incomplete and not always equally readable [13, 33, 41, 44, 49]. A recent survey of professional developers shows that 60% of the participants have suffered from inadequate API documentation in the last three months [30]. Motivated by this, researchers have proposed many solutions to augment API documentation [11, 45, 47, 50, 56].

The state-of-the-art (SOTA) approaches SISE [45] and DeepTip [50] formulate this task as *update extractive summarization*. Update summarization aims to generate complementary summaries to API documentation, assuming readers are already familiar with the original API documentation [45]. Extractive summarization selects **insight sentences** for the target API from external sources to form extractive summaries. Insight sentences provide API-relevant insights not covered in the API documentation [45]. Nonetheless, extractive summaries tend to show a notable level of redundancy and typically have limited readability [22]. Notably, readability is one of the most important attributes of the desired software documentation [12, 24, 38], highlighting the need to produce concise and coherent summaries [12].

Recently, research in utilizing large language models for abstractive summarization has garnered significant attention [16, 55, 57]. Abstractive summarization generates new sentences, in contrast to extractive summarization which selects sentences from input. The SOTA model GPT-4 demonstrates its potential in generating concise and coherent summaries [16, 55]. However, its effectiveness on API documentation augmentation remains unknown.

In this work, we conduct the first empirical study on GPT-4 for API documentation augmentation and compare it with existing approaches [45, 50] (detailed in Section 2.3). Our empirical study reveals that GPT-4 can generate coherent and concise summaries to augment API documentation. However, GPT-4 gives rise to several drawbacks concerning informativeness and faithfulness. First, the input length limitation of GPT-4 leads to information loss due to the truncation of the prompt. We empirically discover that the token limit of GPT-4 allows for the inclusion of an average of only 6.35 Stack Overflow threads relevant to an API to be embedded into the prompt. Meanwhile, the cost of the GPT-4 API may be a burden for individual developers generating documentation (e.g., each GPT-4 API call costs around \$0.3 if the token number reaches the 8k limit). Furthermore, GPT-4 occasionally generates non-faithful summaries, i.e., summaries that do not accurately reflect the original meaning of API-relevant resources [34]. This phenomenon, known as hallucination [27], has hindered the trustworthy usage of GPT-4 in industries like medicine [31] and law [19]. We also observe that assessing the faithfulness of such hallucinations may require substantial time or, in some cases, might not even be feasible. Since the accuracy of the information in software documentation is considered important for software engineering practitioners [12], the existence of hallucinations also hinders the adoption of GPT-4 in this task. Moreover, our pilot study reveals that the summaries generated by existing extractive summarization approaches [45, 50] are usually more faithful to external resources, which is attributed to direct sentence extraction without modification.

Besides, extractive summaries demonstrate a considerable degree of informativeness and relevance, since there is no limitation on input length.

The success of extractive summaries in terms of informativeness and faithfulness inspires us to use extractive summarization to guide abstractive summarization, addressing the drawbacks of the latter. We leverage an extractive-then-abstractive pipeline including two phases: 1) Extract Phase: extract insight sentences from external resources to form extractive summaries, and 2) Abstract Phase: ask GPT-4 to generate abstractive summaries guided by extractive summaries. The Extract Phase allows input without length limitations, covering more information from external resources and minimizing the costs for API calls. The Abstract Phase ensures that abstractive summaries are aligned with extractive summaries, thereby enhancing faithfulness and facilitating data provenance. We apply GPT-4 as the summarizer for the Abstract Phase.

We consider SOTA extractive summarization approaches [45, 50] on the API documentation augmentation task as candidates for implementing the Extract Phase. However, we identify four drawbacks of existing approaches, which may hinder their performance in the Extract Phase. Therefore, we build upon existing approaches, extending them to overcome their limitations. Notably, we are particularly concerned about the performance of the Extract Phase as the **extractive summaries directly determine the informativeness of the final API documentation and their relevance to an API**. Specifically, existing approaches **1) Only take input from a single source**. They [45, 50] focus solely on Stack Overflow, neglecting other sources such as tutorial videos and blogs that developers often refer to [37]. **2) Unaware of the API documentation structure**. Existing approaches perform binary classification to identify insight sentences, neglecting the standard structure of API documentation. The structure is important for software documentation [38], which commonly consists of multiple *sections* (e.g., expected behavior, range of values, and cause of exceptions [36]). **3) Neglect contextual dependencies**. Existing approaches analyze each sentence individually and fail to capture semantic dependencies between sentences. Consequently, insight sentences may be considered useless without necessary context, which leads to a 15% output sentence confusion as reported by the existing approach SISE [45]. **4) Suffer from information redundancy**. API documentation augmentation requires generated summaries to be distinct from the original API documentation for readability [45]. However, existing approaches neglect to reduce redundancy between generated summaries and API documentation.

To address the drawbacks of existing approaches and fully leverage the capabilities of GPT-4, we propose APIDocBooster, a novel framework with two stages: **Context-aware Sentence Section Classification (CSSC)** and **UPdate SUMmarization (UPSUM)**. CSSC takes as input documents relevant to a specific API from **multiple sources**. CSSC identifies insight sentences and classifies them into suitable API documentation sections for **structure awareness** while considering the **contextual dependency** of sentences if necessary. UPSUM applies an extract-then-abstract pipeline, which inputs insight sentences concerning a target API and outputs abstractive summaries to augment API documentation. In the Extract phase, we propose an **extractive update summarization algorithm** ExtUP that enables the generated summary to be semantically dissimilar from the API documentation. In the Abstract phase, we leverage GPT-4 to generate summaries guided by the output of ExtUP through in-context learning.

To enable automatic evaluation of this task, we construct the first dataset APISumBench through two-phase labeling. We identify 4,344 API-relevant sentences from multiple sources and linked them to appropriate API documentation sections. Additionally, we produce 48 extractive summaries to augment three sections of API documentation, corresponding to 16 APIs in Java and Python.

To assess APIDocBooster, we conduct automatic evaluation using APISumBench on each stage of APIDocBooster and human evaluation on end-to-end performance. Specifically, our automatic evaluation results reveal that CSSC outperforms the best baseline by 18.18%, 20.31%, and 18.46% in

0:4 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

terms of precision, recall, and F1-scores. Besides, ExtUP outperforms the best baseline by 14.55%, 24.35%, and 16.67% in terms of ROUGE-1, ROUGE-2, and ROUGE-L. Furthermore, our human evaluation results show that APIDocBOOSTER outperforms the GPT-4 baseline by 13.89%, 15.15%, and 30.56% in terms of informativeness (i.e., quantity of insightful information), relevance (i.e., proportion of insightful information in generated summary), and faithfulness (i.e., to what extent the information in summaries aligns with corresponding information in external resources), as well as keeping on par performance on the readability (i.e., fluency and coherence) and the redundancy between the summaries and original API documentation.

The contributions of this paper are the following:

- We conduct the first empirical study on GPT-4 performing abstractive summarization for API documentation augmentation. We identify three main drawbacks of GPT-4 and four drawbacks of existing extractive-based summarization approaches.
- We propose APIDocBOOSTER, a two-stage approach with an extract-then-abstract framework to address these challenges.
- We construct the first dataset APISumBench, which enables automatic evaluation of extractive summarization of API documentation augmentation.
- We evaluate the performance of APIDocBOOSTER via both automatic evaluation and human evaluation. Both evaluation results show that APIDocBOOSTER outperforms the best-performing baseline by a large margin.

2 PRELIMINARY

2.1 Motivation Example

In this section, we show an example of inadequate API documentation, how insightful information from external sources complements the original API documentation, and the necessity for incorporating context information in Table 1.

The first part of Table 1 shows an example of inadequate API documentation. We present the functionality description of an API `TANH`¹ collected from PyTorch [39], a popular python framework for deep learning. The description of the API `TANH` in official API documentation only contains a brief definition of the algorithm used in `TANH` as well as the math formula. Such description discourages developers from better understanding the API due to its incompleteness, e.g., there is a lack of API usage, range of input and output, and information on quality attributes. Those aspects are considered important for API documentation in the literature [33].

The second and third lines of Table 1 show that rich information relevant to the target API from external sources can complement the original API documentation. We take the same API `TANH` as an example. The sentences extracted from the YouTube video caption present the usage scenarios and range of the output for `TANH` API, while the quoted Stack Overflow sentences highlight an essential property of `TANH`. Both sentences can be considered a valuable complement to the API documentation. Furthermore, during our dataset collection process (as detailed in Section 4.1.3), we find that tutorial videos related to the target API often cover detailed operation steps of using the target API. Meanwhile, the Stack Overflow answers to “how-to-do-it” questions shed light on a target API’s purpose and usage scenarios [45]. Note that “how-to-do-it” questions are the most frequent question type on Stack Overflow [20]. These observations inspire us to believe that the information contained in both YouTube videos and Stack Overflow posts has great potential to augment API documentation and complement each other.

Furthermore, we show an example to illustrate the necessity of incorporating contextual information to ensure the comprehensiveness of API insights. Since existing approaches on API documentation

¹<https://pytorch.org/docs/stable/generated/torch.nn.Tanh.html>

Table 1. An Example of Inadequate API Documentation

TANH CLASS torch.nn.Tanh
Functionality Description in Official Documentation : Applies the Hyperbolic Tangent (Tanh) function element-wise. Tanh is defined as: $\text{Tanh}(x) = \frac{\exp(x) - \exp(-x)}{\exp(x) + \exp(-x)}$
Related YouTube Video: https://www.youtube.com/watch?v=lp4skE5U8Cs ... This Tanh is a hyperbolic tangent itself, but in the world of Neural Network, Tanh converts the input value into a non-linear one and keeps it in the range of -1~1. It is used to serve as an activation function. ...
Related Stack Overflow post: https://stackoverflow.com/a/40775144/10143020 ... On the other hand, to overcome the vanishing gradient problem, we need a function whose second derivative can sustain for a long range before going to zero. Tanh is a good function with the above property. ...

augmentation analyze each sentence individually by default, we argue this setting leads to a loss of API-relevant information and hinders readability. We show an example in Table 1. The quoted sentences from Stack Overflow posts are semantically dependent on each other. The first sentence illustrates that a function with a particular property can overcome the vanishing gradient problem. The second sentence explains that TANH has such a property. Presenting either of these two sentences individually proves insufficient in conveying comprehensive insights concerning the target API. Specifically, failing to consider the second sentence as the context of the first one leaves the connection between the first sentence and the target API unclear and obscure. Meanwhile, presenting the second sentence in isolation fails to convey the detailed properties of the target API. This motivates us to incorporate contextual information when performing API documentation augmentation.

2.2 Notions and Task Formulation

This work aims to augment API documentation by generating complementary summaries for each documentation section from multiple online sources. As a initial attempt, we consider two sources, i.e., Stack Overflow and YouTube, as input sources. The rationale is that existing work [43] reveals the potential of combining both platforms for software-related knowledge extraction.

Structuring is one of the most important quality attributes for effective software documentation [38]. API documentation is typically divided into multiple smaller and more manageable sections. We consider three sections (i.e., Functionality, Parameter, and Notes) that appear in the API documentation of popular programming languages (e.g., Python [7] and Java [3]) as well as well-known third-party libraries (e.g., PyTorch [39] and Android [1]). Note that some of the existing approaches [33] categorize the knowledge required for API documentation at a finer granularity. The reason for this coarse categorization is to better fit the structure of real-world API documentation. We leave the fine-grained API documentation augmentation for future work.

0:6 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

We follow the definition of the API documentation section from Java API documentation's writing requirement [36]:

- **Functionality:** Information about the execution, expected behavior, state information and transitions, defined algorithms, and cause of exceptions of the API.
- **Parameters:** Information about the range of valid argument values, return values, and behaviors of passing an invalid argument.
- **Notes:** Information about OS/hardware dependencies, allowed implementation variances, security constraints, and references to any external specifications of the API, also other insightful tips (e.g., API performance, legal concerns) that are helpful for understanding an API [45, 50].

Formally, the input for the API documentation augmentation task includes a target API api , API documentation of the above three sections doc , and a set of sentences extracted from external sources $s = \{s_1, s_2, \dots, s_n\}$. We aim to generate extractive summaries (i.e., an optimal subset $sum = Extract(s|api, doc), sum \subseteq s$), and then abstractive summaries (i.e., $sum' = Abstract(sum|api, doc)$) that satisfy the requirements of readability, relevance of content, and structure of API documentation. These attributes are the most critical to the quality of software documentation [38].

2.3 Empirical Study on GPT-4

We carry out an empirical study to understand the capacity and limitations of using GPT-4 to augment API documentation.

2.3.1 Statistics on External Resources. We begin with exploring how much API-relevant information exists in external sources. Following existing approaches [18, 47], we make an initial attempt on popular Python and Java libraries, i.e., PyTorch and Android. We randomly sample 100 APIs from each library, respectively. As discussed in Section 2.2, we consider Stack Overflow and YouTube as preliminary external sources.

Stack Overflow returns 258.70 posts on average for each API. Corresponding Stack Overflow threads account for 4,597.68 characters on average, which results in an average of approximately 1,150 tokens when tokenized by GPT-4 [5]. The GPT-4 API imposes a maximum context limit of 8,192 tokens, inclusive of both prompts and the subsequent response. We reserve 500 tokens for API responses. The summarization instruction and original API documentation together count 394 tokens on average. Consequently, even if we consider Stack Overflow as the sole external source, only an average of 6.35 threads $((8192 - 500 - 392)/1150)$ can be included within the prompt for standard GPT-4, 27.05 threads for GPT4-32k (i.e., GPT-4 variant with four times the context range and twice the API call cost). Similarly, YouTube returns 80.98 videos on average. The average length of the YouTube video captions is 41,816.49 characters, approximately 10,454 tokens. Even the average number of tokens for a single video exceeds the upper token limit of GPT-4's context.

2.3.2 How does GPT-4 augment API documentation? We instruct GPT-4 to automatically augment API documentation. We set the parameter *temperature* of GPT-4 as 0 to ensure GPT-4 produces a deterministic response. We consider two settings:

Zero-shot Setting (GPT_{zero}). We ask GPT-4 to generate augmented API documentation without external resources. We craft prompts that include original API documentation, the definition for each section, and an instruction for generation. We follow an existing GPT-based approach on opinion summarization [16] to construct the instruction and prompts.²

Summarization Setting (GPT_{sum}). We ask GPT-4 to summarize insightful information for each section from given external resources. The generated summaries are required to be different from

²We present our prompts in the replication package due to the page limit.

the original API documentation. Specifically, we collect Stack Overflow posts and captions of YouTube videos relevant to each API by calling the Stack Overflow search API and YouTube API (detailed in Section 3.2). The prompt used in this setting also follows existing approaches [16, 55] and consists of original API documentation, the definition for each section, an instruction for summarization, and API-relevant resources.

Table 2. User Study for GPT-4 (5-point Likert scale)

System	Informativeness.	Relevance	Readability	Non-Redundancy	Faithfulness
GPT_{zero}	2.6	2.2	4.3	4.5	-
GPT_{sum}	3.6	3.3	4.4	4.5	3.6
SISE	3.0	3.3	2.3	2.6	5.0

2.3.3 User Study.

User Study Setting. We leverage GPT_{zero} , GPT_{sum} , and SISE (i.e., SOTA extractive-summarization-based approach [45]), to augment API documentation. We randomly sample 10 APIs from PyTorch and Android libraries respectively. As the length of a single video caption exceeds the token limit of GPT-4, we only consider Stack Overflow as external sources for GPT_{zero} and GPT_{sum} . We add each post in the prompt in the same order that they are returned by Stack Overflow, until reaching the context limitation of GPT-4. We ask GPT-4 to generate summaries for all API documentation sections.

Four user study participants, with at least two years of Python or Java programming experience, read all augmented API documentation pages and rate each section on a 5-point Likert scale in terms of **informativeness**, **relevance**, **readability**, **non-redundancy**, and **faithfulness**. Informativeness represents the quantity of insightful information; relevance measures the proportion of insightful information; readability assesses the fluency and coherence of the summaries; non-redundancy evaluates the information redundancy between the summaries and original API documentation; and faithfulness is a measure of the degree to which information in summaries aligns with corresponding information in external resources. Scores 1 to 5 represent a range from extremely low to extremely high, respectively. All of the evaluation metrics are widely used in the software engineering domain [18, 45] and in NLP domain summarization approaches [16, 25]. Participants are also required to highlight sentences that they think should not be in API documentation given the above criteria. Note that, summaries with hallucinated information are given a lower faithfulness score. However, it is impossible to assess the faithfulness of GPT_{zero} as GPT-4 generates the summaries without external resources as reference.

User Study Results. As shown in Table 2, GPT_{sum} performs best in terms of informativeness, readability, and redundancy. SISE performs best in terms of faithfulness and shows on-par relevance compared with GPT_{sum} . The superior performance of GPT_{sum} compared with GPT_{zero} in terms of informativeness ($p < 0.01$) indicates the necessity of leveraging external resources. The lower performance of SISE compared with the GPT-based approach in terms of redundancy and readability ($p < 0.01$) demonstrates the potential of GPT-4 to generate fluent and coherent summaries.

Qualitative Analysis. Two authors of this paper also perform qualitative analysis on the augmented API documentation generated by GPT_{zero} and GPT_{sum} . Specifically, they go through the sentences highlighted by participants of the user study and discuss the rationale behind the unsuitability of these sentences within the API documentation context. Firstly, we observe that GPT_{sum} occasionally produces less informative sentences. For instance, GPT-4 generated the following sentence: “Implementation may vary across different OS/hardware.” for the OS dependency of API LeakyReLU in PyTorch [39]. Upon seeking further explanation from GPT, it clarifies that the implementation differences arise from hardware-specific optimizations of PyTorch on various

0:8 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

platforms. This information applies to most APIs and lacks specificity to the target API, resulting in a sentence with minimal insights to API documentation. We attribute this to the limited availability of high-quality API documentation examples provided to GPT, which hinders GPT's ability to learn from document examples and identify the most insightful information specific to a particular API. However, the input length limitation of GPT-4 makes it challenging to feed more API documentation examples.

We also observe hallucinations, that is information not faithful to given resources, in the generated summaries. Specifically, GPT_{zero} consistently generates incorrect parameter lists, including parameters from APIs with similar names but in different libraries, when original API documentation is not provided. Even when the API documentation is embedded in the prompt, there are instances where the GPT_{zero} still confuses the numerical range of the parameters. We have also noticed a reduction in hallucinations with the use of the GPT_{sum} , although there are occasional occurrences of GPT_{sum} generating information that is not accurate compared with given external resources. On the contrary, SISE does not generate hallucinations since it only extracts sentences from given external resources. Based on the above findings, we argue that providing more external resources can help alleviate hallucinations of GPT-4, enabling GPT-4 to summarize from given information rather than generate information itself.

2.4 Usage Scenario

We propose APIDocBOOSTER as an automated aggregator of API-related information and writer to update API documentation.

More specifically, with the help of APIDocBOOSTER, API documentation's maintenance and update processes become more streamlined and efficient. APIDocBOOSTER simplifies the tedious task of documentation maintenance [12] into an automatic pipeline: 1) automatically collecting API-related information from external sources; 2) automatically generating summaries for each section of the API documentation; and 3) integrating the generated summaries into the original API documentation.

Notably, insights provided by APIDocBOOSTER highlight the characteristics of the target API that developers are concerned about in practice. The focus of the developer community is commonly mirrored in the frequency of corresponding information mentioned in external resources such as Stack Overflow [47]. As the extractive summarization algorithm prioritizes information based on its mention frequency, APIDocBOOSTER enables API-related information with substantial community attention more likely to be involved in the summary. Consequently, APIDocBOOSTER enables API documentation maintainers to focus on updating information that more developers are interested in, which in turn, accommodates the informational needs of a more diverse readership.

It is essential to acknowledge that APIDocBOOSTER's effectiveness relies on the richness of API-related discussions within the development community. APIs with limited discussions might present constraints in APIDocBOOSTER's performance. However, this limitation is deemed acceptable given APIDocBOOSTER's objective: to enhance API documentation that are highly demanded for augmentation. We argue that the existence of extensive community discussions concerning specific APIs implies deficiencies within the original API documentation. The rationale is that technical documentation is the primary way for developers to acquire coding knowledge as indicated by Stack Overflow's survey [37], thereby a sufficiently comprehensive API document would ensure a profound comprehension of API usage for the majority of developers, obviating the necessity for external inquiries. Consequently, an increased frequency of community discourse concerning a specific API indicates a higher probability of inadequacies in its documentation. This, in turn, underscores the imperative to augment the quality of said documentation.




Fig. 1. Overview of APIDocBooster. (1) we train a context identifier and a sentence section classifier to identify the context of each sentence and divide each sentence into either one of the three sections or API-irrelevant, respectively. (2) we generate extractive summaries for each section by using our proposed update summarization algorithm. We then generate abstractive summaries guided by extractive summaries.

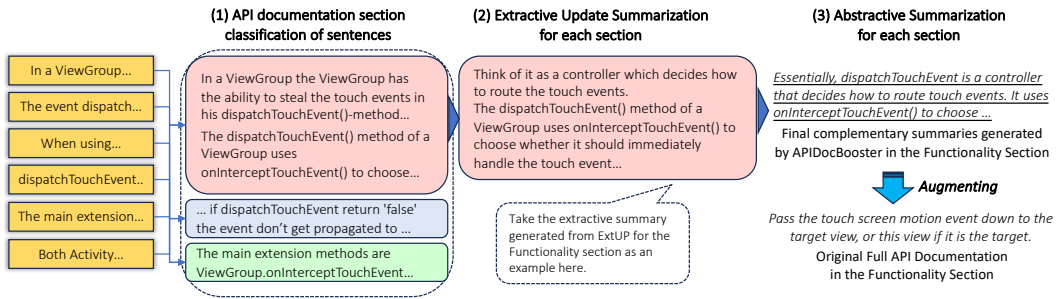


Fig. 2. Example of APIDocBooster's summarization process. (1) Given sentences extracted from external sources, we train a sentence section classifier to classify sentences into either one of three API documentation sections (i.e., **functionality**, **parameter**, **notes**) or API-irrelevant one. Then we leverage our extractive update summarization algorithm to generate extractive summaries for each section. Finally, we ask a large language model to generate abstractive summaries for each section guided by extractive summaries.

3 APPROACH

3.1 Overview

The framework of our approach, namely APIDocBooster, is shown in Figure 1. APIDocBooster takes as input Stack Overflow posts and YouTube videos for a given API. APIDocBooster outputs a set of abstractive summaries, each of them corresponding to a specific API documentation section. APIDocBooster contains two core components: 1) Context-aware Sentence Section Classification (CSSC) and 2) Update SUMmarization (UPSUM). CSSC takes as inputs Stack Overflow posts and YouTube video captions and outputs sets of sentences predicted to be insightful to each section while taking sentence context into account. Given a list of insightful sentences for an API documentation section, UPSUM follows extract-then-abstract pipeline and generates abstractive summaries. In the extract phase, we propose an **Extractive update** summarization algorithm (ExtUP) to generate extractive summaries. Then UPSUM leverage GPT-4 to generate abstractive summaries guided by the output of ExtUP. We also show an example of the processing of data flow in Figure 2.

3.2 Pre-processing

In the pre-processing step, we Stack Overflow posts and YouTube video captions relevant to the given API.

Searching for API-relevant Resources. To extract Stack Overflow posts relevant to specific API, we call the Stack Exchange API [8]. We search with the query *library_name + API_name* and

0:10 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

consider returned answer posts with at least one up vote as posts relevant to the target API. We also filter out Stack Overflow answer without textual content. To obtain the YouTube videos relevant to the target API, we use the official YouTube API [9] to search for videos with the same query as for Stack Overflow Exchange API search.

Sentence Extraction. We apply NLTK’s sentence tokenizer [?] to split the textual content from Stack Overflow posts and captions of YouTube videos into sentences.³ Furthermore, since video caption has much more API-irrelevant content compared to Stack Overflow answers, we denoise the caption data. As described in Section 2.3.1, YouTube video captions have 9.10 times more characters than Stack Overflow posts on average. However, the majority of the sentences are extraneous to the specific APIs. In our manually labeled dataset (detailed in Section 4.1), the average number of insight sentences for the target API, among 553 sentences per video, is 11.9. To reduce noise in input, inspired by existing multi-doc summarization for question and answering [52], we leverage the BM25 [40] algorithm to coarsely filter out video captions sentences that are irrelevant to API. BM25 is a ranking function that is used to estimate the relevance between a given query and a document. Formally, given a query Q , the relevance score is given by:

$$Score(Q, d) = \sum_i^n IDF_{q_i} \frac{h(q_i, d) * (k_1 + 1)}{h(q_i, d) + k_1 * (1 - b + b * \frac{|D|}{avgdl})} \quad (1)$$

where q_i is the i^{th} word in the query Q , $h(q_i, d)$ indicates the frequency of q_i occurring in the document d , $avgdl$ is the **average document length** in the text collection, k_1 and b are adjustable parameters (by default $k_1 = 1.2$ and $b = 0.75$, following Robertson et al. [40]), and $|D|$ is the length of the document in terms of words. The $IDF(q_i)$ is the IDF value of q_i . We observe that sentences, along with their context, which contain name entities within API documentation (e.g., API name and API parameters), tend to be insightful to target API. Thus, we set each caption sentence as d and manually construct the query for each API as “*API name, hyper-parameter names, input and output*”. Following existing work [52], we consider sentences with positive BM25 scores as well as their context (i.e., preceding and following sentence) as input for CSSC. We consider preceding and following sentences as context since adjacent sentences are more likely to have semantic dependencies.

3.3 Context-aware Section Classification

CSSC trains a BERT-based classifier to classify sentences from relevant Stack Overflow answers and YouTube video captions as either one of the three section types (i.e., function, parameters, notes) or API-irrelevant. In addition, we implement a context awareness mechanism to improve the classification performance of context-dependent sentences.

3.3.1 Sentence Section Classifier. We fine-tune BERT to create a four-class classifier to divide each input sentence into either one of the three API documentation sections or irrelevant. We leverage BERT, a pre-trained language model that achieves promising performance on relevant software engineering tasks [18, 53], to fine-tune a sentence section classifier. Previous approaches [45, 50] treat the task as a binary classification that identifies insight sentences to the target API. Unlike previous approaches, we aim to further classify insight sentences into a finer granularity, i.e., sections of the API document. Thus we perform multi-class classification. By doing so, each sentence is mapped to a certain document section, enabling augmented API documentation at a more fine-grained level.

³YouTube caption files contain the text of what is said in the video. It also contains timestamps for when each line of text should be displayed.

Algorithm 1: Context-Awareness Mechanism**Input:** api : target API; S : List of the API-relevant sentences**Output:** S_label : section label of each sentence

```

1 Initialization:  $context\_flag = []$ ;  $l^{update} = ""$ 
2 for  $s_{i-1}, s_i, s_{i+1}$  in  $S$  do
3   if  $context\_identifier(s_{i-1}, s_i)$  then
4      $context\_flag\_preceding[i] = True$ 
5      $context\_score\_preceding[i] = context\_identifier(s_{i-1}, s_i)$ 
6   else
7      $context\_flag[i] = False$ 
8   end
9   if  $context\_identifier(s_i, s_{i+1})$  then
10     $context\_flag\_following[i] = True$ 
11     $context\_score\_following[i] = context\_identifier(s_i, s_{i+1})$ 
12  else
13     $context\_flag[i] = False$ 
14  end
15 end
16 for  $s_i$  in  $S$  do
17    $s_{context} = s[i]$ 
18   if  $context\_flag\_preceding[i] == True$  then
19      $s_{context} = s_{context} + s_{i-1}$ 
20   end
21   if  $context\_flag\_following[i] == True$  then
22      $s_{context} = s_{context} + s_{i+1}$ 
23   end
24   if  $context\_flag\_preceding[i] == True$  or  $context\_flag\_following[i] == True$  then
25      $l^{origin} = section\_classify(s_i, api)$ 
26      $l^{context} = section\_classify([s_i, s_{context}], api)$ 
27      $l^{update} = update\_section\_prediction(l^{origin}, l^{context})$ 
28   end
29 end
30  $sent\_labels = argmax(l^{update})$ 

```

In our task, an input instance is in the form of a pair of API and corresponding sentences, i.e., $\langle API, S \rangle$; thus, we feed the pairs of sentences and APIs to the BERT model for fine-tuning. Specifically, for each pair $\langle S, API \rangle$, we concatenate the sentence S and the target API API into a sequence with padding if necessary, i.e., $\langle [CLS], S, [SEP], API, [SEP] \rangle$, to serve as input to the classifier. Both $[CLS]$ and $[SEP]$ tokens are used to form the BERT input [26]. The $[CLS]$ vector extracted from the classifier decoder is fed into the final classification layer to produce the prediction. Following standard practice, we use cross entropy as the loss function.

3.3.2 Context-Awareness Mechanism. As we discussed in Section 2.1, we argue that the necessary context of a given sentence plays a key role in the understanding of this sentence and delivering complete information. Sentences requiring contextual information for the comprehension of API insights are denoted as context-dependent sentences. However, not all sentences need context to be understandable, where unnecessary context is considered as noise. Thus, we propose a context-awareness mechanism to 1) determine context-dependent sentences, 2) concat context-dependent

0:12 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

sentences with necessary context, and 3) modify the section classification results for context-dependent sentences based on their context. We detail the algorithm of the context-awareness mechanism shown in Algorithm 1 as follows.

Training Context Identifier. Firstly, we train a context identifier to identify context-dependent sentences. To train such a context identifier, we fine-tune a binary BERT-based classifier to model the context dependence for a pair of consecutive sentences. The classifier inputs a pair of consecutive sentences and outputs whether they are contextually dependent on each other. To train such a model, we reuse the training data of our dataset APISumBench (described in Section 4.1.2). In APISumBench, annotators label consecutive sentences with context-dependence as the same semantic unit. We reuse the context information inferred by the semantic unit label to construct positive examples. Precisely, we extract pairs of consecutive sentences in the same semantic unit as positive samples (i.e., label the pair with 1). To construct the negative pairs (i.e., label the pair with 0), we sample pairs of sentences randomly selected from different Stack Overflow answers or YouTube videos. We fine-tune the BERT-based classifier with the above dataset. We apply the same fine-tuning strategy as our sentence section classifier. Following standard practice, we use cross entropy as the loss function.

Context Dependency Prediction (Line 2-15). After we obtain our context identifier model, we predict if each sentence input in CSSC is context-dependent. For a given sentence, we only consider its preceding and following sentence as the candidates of context-dependent. The rationale is that semantic dependencies are more probable between adjacent sentences. Next, a sentence is predicted to be context-dependent if at least one of its preceding and following sentences has contextual dependence with this sentence. Besides, we pick the positive class' prediction probability \mathbf{c} as each pair's context-dependent score, indicating the extent to which the pair of sentences are semantically dependent on each other.

Context-aware Prediction Adjustment (Line 16-29). Once a sentence s_i is predicted as context-dependent, we update its prediction logit $l = [l_1, l_2, l_3, l_4]$, i.e., raw prediction score for API documentation section classification, by considering the prediction that the input sentence is concatenating with context. The value l_i is a raw prediction score (i.e., the score before normalization) of each class. Firstly, we obtain the prediction logits l^{origin} of s_i by feeding s_i and corresponding the API api to the section classifier:

$$l^{origin} = section_classify(s_i, api) \quad (2)$$

where $section_classify$ denotes the section classifier. Then, assuming s_{i-1} and s_{i+1} are contextual sentences of the current sentence s_i , we obtain the context-aware prediction logits of s_i by concatenating s_{i-1} , s_i , and s_{i+1} and feeding them to the section classifier:

$$l^{context} = section_classify([s_{i-1}, s_i, s_{i+1}], api) \quad (3)$$

Next, we update the prediction logits l^{update} . Specifically, in each dimension of l^{update} , the value l_j^{update} is obtained by weighting both context-aware prediction logits $l_j^{context}$ and original prediction logits l_j^{origin} :

$$l_j^{update} = \frac{\sum((1-c) \times l_j^{origin} + c \times l_j^{context})}{num(context)} \quad (4)$$

where $num(context)$ denotes the number of contextual sentences of the current sentence. \mathbf{c} represents the extent of semantic dependency between the current sentence and its contextual sentences. Finally, given update prediction logits $l^{update} = [l_1^{update}, l_2^{update}, l_3^{update}, l_4^{update}]$, we obtain the update predicted section y' by applying the $argmax$ function:

$$y' = argmax(l^{update}) \quad (5)$$

3.4 Update Summarization

Given the output of CSSC (i.e., a set of insightful sentences for each API documentation section) and a budget of summarization length (i.e., number of sentences) as input, UPSUM outputs a set of summaries for augmenting API documentation sections.

Specifically, UPSUM employs an extract-then-abstract summarization pipeline, in which the abstractive summarization is guided by the output of extractive summarization. Extractive summarization involves the direct selection of sentences from the source text, ensuring consistency of information and establishing a clear lineage between the source and the generated summaries [22]. Conversely, abstractive summarization entails generating summary content from the model, resulting in a fluent summary but potentially containing words not present in the source and opinions contradicting the source (i.e., hallucination) [22]. The purpose of the extract-then-abstract pipeline is threefold: 1) to link the generated content in abstractive summaries with corresponding external sources, 2) to take advantage of both extractive (to enhance faithfulness and mitigate hallucination) and abstractive summarization (to enhance readability), and 3) to alleviate the length constraints inherent in abstractive summarization models.

Precisely, in the extractive summarization phase, UPSUM performs extractive update summarization. Update summarization is a double-objective summarization task [45], which aims to generate a summary by minimizing redundancy 1) across generated content and 2) between generated summaries and original API documentation. We achieve this goal by utilizing our proposed extractive update summarization algorithm ExtUP for update summarization. In the abstractive summarization phase, LLM performs in-context learning and generates abstractive summaries guided by the extractive summaries.

3.4.1 Extractive Update Summarization. We propose an extractive update summarization approach, ExtUP. Specifically, given a target API, ExtUP inputs a set of sentences that are insightful to the specific API documentation section. ExtUP outputs an optimal subset of sentences as the summary that represents the main ideas of the input while minimizing the information redundancy between the summary and original API documentation. We modify a simple yet effective summarization approach, TextRank [35], for update summarization.

TextRank is an unsupervised extractive summarization approach based on the PageRank algorithm. PageRank is interpreted as the stationary distribution of a random walk on a graph that restarts from a uniform location in the node sets at each time. The restart probability is $1 - \phi$ while the probability of moving to a neighbor node is ϕ . In TextRank, all input sentences are constructed as a sentence graph with undirected edges. Each node represents a sentence, and the initial weight of an edge is based on the token similarity between two nodes. TextRank recursively computes the similarity of each sentence to all candidates as follows:

$$C(S_i) = (1 - \phi) + \phi \times \sum_{V_j \in In(V_i)} \frac{1}{Out(S_j)} C(S_j) \quad (6)$$

To incorporate the update mechanism in TextRank, we add a penalty factor p_{doc} to deprioritize sentences that are semantically similar to original API documentation when performing summarization. This penalty factor is inspired by Biased-TextRank [28]. Biased-TextRank is designed for query-focused summarization, which is biased toward query-oriented sentences. Specifically, our penalty factor changes the way random restart probabilities ($1 - \phi$) are assigned in the sentence graph. In TextRank, all nodes have an equal restart probability of $1 - \phi$. However, in our scenario, sentence nodes that are semantically similar to the API documentation are expected to have a lower restart probability. Nodes with lower random restart likelihood are less likely to be reached when randomly jumping, and thus achieve lower Summarization scores. The penalty factor is defined as

0:14 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

follows:

$$p_{doc} = 1 - \text{cosine_similarity}(\text{node}, \text{API}doc) \quad (7)$$

where node and $\text{API}doc$ denote the sentence representation in node graph and original API documentation, respectively. We apply a software-engineering domain text representation model [18] to obtain their representations. The formulation of calculating sentences' ranking is defined below, which is followed by a normalization:

$$C(S_i) = p_{doc} \times (1 - \phi) + \phi \times \sum_{V_j \in \text{In}(V_i)} \frac{1}{\text{Out}(S_j)} C(S_j) \quad (8)$$

where $\text{In}(S_i)$ represents all nodes that point to S_i and $\text{Out}(S_i)$ represents all nodes that S_i points to. $C(S_i)$ denotes the TextRank summarization score of a node sentence. ϕ is the damping factor, defaulting as 0.85 [35].

We also replace the way of assigning the edge weight in the sentence graph from calculating token overlaps to calculating the cosine similarity between two sentences' representations. We use the same text representation model [18] as calculating the penalty factor p_{doc} . Finally, given the budget of summary length, i.e., k sentences, we select the top- k sentences with the highest scores as the extractive summary.

3.4.2 Abstractive Summarization. We leverage GPT-4 to generate abstractive summaries guided by the extractive summaries. Specifically, we ask GPT to select insightful information from the extractive summaries (i.e., S_e) concerning the specification of API sections, paraphrase and synthesize the insightful information, and form abstractive summaries (i.e., S_a). We follow existing GPT-based approaches for opinion and news summarization to construct the prompts [16, 55], including extractive summaries S_e , the summarization instruction i , original API documentation d , and context pairs $\{(S_{a1}, S_{e1}) \dots (S_{an}, S_{en})\}$.⁴ We link the sentences in extractive summaries and abstractive summaries by greedily calculating the cosine similarity between sentence representations [54]. Formally, the summarization process is represented as:

$$\hat{S}_a = \underset{S_a}{\text{argmax}} P_{GPT}(S_a | S_e, d, i, \{(S_{a1}, S_{e1}) \dots (S_{an}, S_{en})\}) \quad (9)$$

As we observe that GPT-4 tends to generate less informative content in Section 2.3, we perform in-context learning to enhance the informativeness of generated summaries. In-context learning can make LLM rapidly adapt to or recognize the desired task by providing example input-output pairs $\{(S_{a1}, S_{e1}) \dots (S_{an}, S_{en})\}$ in the prompt for few-shot learning [17]. We provide pairs of ground-truth extractive summaries and informative abstractive summaries, aiming to make GPT-4 more adept at producing informative summaries. Specifically, we sample human-crafted extractive summaries from our dataset APISumBench (detailed in Section 4.1) and ask GPT-4 to generate augmented API documentation following the same prompt. We then manually revise the augmented API documentation to remove less informative content.

4 EXPERIMENT SETTING

4.1 APISumBench

We craft APISumBench, the first dataset which enables automatic evaluation of CSSC and ExtUP. In APISumBench, we perform two-stage labeling and acquire two phases of data. In the first phase, we invite annotators to label sentences from Stack Overflow posts and YouTube videos relevant to the target API with API documentation sections. The first phase consists of 4,344 sentences, which scale is on par with previous API review summarization work [47]. The second phase comprises

⁴Detailed prompts are provided in our repository due to the page limit.

48 extractive summaries, generated from the data gathered in the first phase, and corresponds to three API documentation sections from 16 different APIs. The first phase data enables automatic evaluation on CSSC, i.e., sentence classification for API documentation section. The second phase data enables automatic evaluation on ExtUP, i.e., extractive summarization for API documentation augmentation.

4.1.1 Data Collection. We collect candidate APIs as well as relevant Stack Overflow posts and YouTube videos. Following our empirical study in Section 2.3, we randomly sample 10 PyTorch APIs and 10 Android APIs as candidates. To reduce the labor cost, following existing approaches [18], We collect the top 10 Stack Overflow posts returned by Stack Overflow search engine for each API. We notice that numerous Stack Overflow posts, which are identified as relevant to the target API by the Stack Overflow search engine, solely contain API in the code snippets. Those posts lack insight sentences towards the target API. Thus, we filter out posts containing target API only in code snippets. For YouTube video captions, we observe that automatically generated captions have no punctuation, thus causing those captions to be unable to be naturally split into sentences. To improve the data quality in APISumBench, we only automatically extract videos with manually crafted captions. We apply the youtube-transcript-api tool [10] to automatically extract videos with manual captions. In this step, 189 Stack Overflow posts and 84 YouTube videos are collected, i.e., 11.72 Stack Overflow posts and 4.2 YouTube videos for each API on average. We then use NLTK [4] to break the collected Stack Overflow posts and YouTube videos into sentences. For YouTube video captions, we follow our pre-processing step in Section 3.2 to reduce the noise. Finally, we obtain 5,414 sentences related to 20 APIs.

4.1.2 Labeling Process. Following prior works [15, 18], we perform two-phase labeling. Four participants are involved in the labeling process. None of them participated in the user study detailed in Section 2.3. All annotators have at least two years of programming experience in Java and Python.

Phase 1: API Documentation Section Classification. We ask annotators to label both API documentation sections and the context dependence of each sentence [45]. Specifically, annotators cluster consecutive sentences that are context-dependent as a cluster. Annotators consider a cluster as a single sentence in the following labeling. Then they label each sentence with API documentation sections. Sentences in the same cluster are labeled with the same section label. The whole labeling process consists of 3 steps:

1) Iterative Guideline Refinement. We iteratively refine the guideline for annotators before full labeling. Specifically, we ask annotators to label the data for one API following the guideline. We then calculate their labeling agreement in terms of Kappa value⁵. Next, we refine the guideline based on the annotators' feedback. We repeat the process until the Kappa value reaches a substantial level of agreement (i.e., 0.62), which agreement surpasses common summarization datasets in software engineering domain [18] and NLP domain [14]. Finally, the data for four APIs used in the refinement step are dropped due to the low agreement.

2) Full Labeling. Annotators then follow the refined guideline to label the remaining 16 APIs in the full labeling process. Each annotator independently labels sentences for 4 APIs. In total, we label 4,344 sentences with API documentation sections in the first labeling stage.

3) Validation Labeling. We perform validation labeling to assess the quality of the data. We randomly sample 654 sentences from the full labeling data (i.e., 15%) and the first author labels the sampled sentences independently. The Cohen's Kappa score between the annotators and the first

⁵We interpret kappa values 0 as poor agreement, 0.01–0.20 as slight agreement, 0.21–0.40 as fair, 0.41–0.60 as moderate, 0.61–0.80 as substantial, and 0.81–1.00 as almost perfect agreement [21].

0:16 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

author achieves a substantial level of agreement (i.e., 0.61), which is higher than that of existing summarization datasets in the NLP domain [14] and Software Engineering domain [18].

Phase 2: Update Summarization. Following existing work [18], in the second phase, given a set of sentences associated with each API documentation section, annotators select at most five sentences to form a summary for each section, considering the insights not covered in the target API's API documentation [45] and the redundancy among all sentences. Phase 2 employs the same guideline refinement step and full labeling step as phase 1. Finally, 48 (16APIs \times 3sections) extractive summaries are produced.

4.1.3 Dataset Statistics. In APISumBench, the first phase data consists of 4,344 labeled sentences labeled with API documentation sections or irrelevant to the API. The average number of sentences input into the second stage for each section (i.e., function, parameter, and notes) is 18.76, 4.33, and 11.08, respectively. The average number of sentences in the generated summaries of each section is 4.69, 2.77, and 4.31. On average, 20.51% of sentences in the generated summaries are from YouTube video captions, while 79.49% are from Stack Overflow answers. We empirically observe that sentences from YouTube videos often discuss the operational steps involved in utilizing the target API, while those operation details are not commonly included in Stack Overflow answers. This indicates that the information in YouTube videos can effectively supplement Stack Overflow in the API documentation augmentation task.

4.2 Baselines

4.2.1 Baselines For CSSC.

Baseline #1: SISE is the state-of-the-art insight sentence extractor [45] that identifies insight Stack Overflow sentences for a given API to augment API documentation. We modify the classifier in SISE from binary to multi-class.

Baseline #2: DeepTip is the state-of-the-art approach to identify tip sentences to answer API-related questions from Stack Overflow answers based on textCNN model [50]. We reuse the architecture of DeepTip with one hot encoding but modify the classifier from a binary classifier to a multi-class classifier.

4.2.2 Baselines For ExtUP.

Baseline #1: TextRank is a widely used graph-based unsupervised summarization approach [35]. Since our UPSUM is built on top of TextRank, we consider TextRank as our baseline.

Baseline #2: LexRank [23] is an unsupervised graph-based text summarization approach that is similar to TextRank. Differently, LexRank leverages inverse document frequency-based cosine similarity to measure the similarity between two sentence nodes.

Baseline #2: TechSumBot is the SOTA approach for question-answering summarization [18] in software engineering. It consists of three sub-modules: 1) the usefulness ranking module to filter out useless sentences to the query, 2) the centrality estimation module to extract summative sentences, and 3) the redundancy removal module to remove redundant sentences. We remove the usefulness ranking module of TechSumBot because all the input sentences in this stage are considered useful for API understanding.

Baseline #3: TechSumBot++ Since the original TechSumBot only considers redundancies among input sentences, we propose a new variant of TechSumBot by integrating redundancy information between selected sentences and API documentation. To achieve this goal, we modify TechSumBot's redundancy-reducing component to enable update summarization. In the redundancy reduction component, TechSumBot iteratively calculates the similarity of each candidate sentence to the one selected in the summary and removes redundant sentences. Our modification further calculates the similarity between the candidate sentences and the original API document. We then remove

1 APIDocBooster: An Extract-Then-Abstract Framework Leveraging Large Language Models for Augmenting API
2 Documentation 0:17

3
4 redundant candidate sentences from the API documentation documentation in each iteration. We
5 keep the remaining settings of TechSumBot the same.

6
7 **4.2.3 End-To-End Baseline.** We set GPT_{sum} (detailed in Section 2.3.2) as the end-to-end baseline
8 for APIDocBOOSTER.

9 **4.3 Implementation Details**

10
11 **4.3.1 CSSC.** We implement a BERT-based API section classifier and context identifier using
12 *Hugging Face* [2], a popular deep-learning library for Transformers. We use the ‘bert-base-uncased’
13 version of the model. Following the fine-tuning strategy in the BERT paper [26], we tune the same
14 key hyper-parameters, i.e., learning rate, batch size, and epoch numbers.

15
16 **4.3.2 UPSUM.** We obtain the sentence vector representation using the in-domain sentence repre-
17 sentation model proposed by TechSumBot. Since there is no clear rule about how many sentences
18 should be considered in the summary of each API section [45], we set the abstractive summary
19 length, representing the output sentence length of APIDocBOOSTER, to be equivalent to the length
20 of the ground-truth extractive summaries created in APISumBench, i.e., they have the same number
21 of sentences. Moreover, we set the length of the extractive summary (i.e., the output of ExtUP) to
22 be twice that of the abstractive summary, thereby providing ExtUP a certain margin for error. The
23 information in extractive summaries which is irrelevant to the target API, can subsequently be
24 dropped in the abstractive summary by leveraging GPT’s comprehensive understanding of the API
25 documentation learned from in-context learning. We use the GPT-4 checkpoint (i.e., ‘gpt-4-0613’).
26 For in-context learning, we set sample pairs N as 3 since existing summarization approaches work
27 best on this setting [55]. No summaries sampled for in-context learning exists in the test set.

28 **4.4 Automatic Evaluation Metrics**

29 Following previous work on identifying API-relevant sentences [45, 50], we use weighted precision,
30 weighted recall, and weighted F1 for the sentence section classification stage.

31 Following existing summarization approaches [18, 28], we employ ROUGE as our main automatic
32 evaluation metrics for ExtUP [32]. We report the F1 score for ROUGE-1, ROUGE-2, and ROUGE-L.
33 ROUGE-1 and ROUGE-2 evaluate the unigram and bigram overlap between the generated and
34 ground-truth summaries. ROUGE-L evaluates the overlap of the longest common sub-sequence
35 between the generated and ground-truth summaries. We use the ROUGE calculation code in the
36 pyrouge library [6].

37 **5 EVALUATION**

38
39 This section presents our experiment results as well as the corresponding analysis to answer the
40 following three research questions:

41 **RQ1:** How effective is CSSC for sentence section classification?

42 **RQ2:** How effective is ExtUP for update extractive summarization?

43 **RQ3:** How effective is APIDocBOOSTER in augmenting API documentation?

44 **5.1 RQ1: Effectiveness in the Sentence Section Classification Stage**

45
46 **Experiment Setting.** To answer this research question, we compare the performance of CSSC and
47 the multi-class classification baselines described in Section 4.2.1 for sentence section classification
48 task. We use the first-stage result of APISumBench, i.e., a set of sentences labeled with API sections,
49 to evaluate the performance of all approaches. Moreover, we consider the following two evaluation
50 settings:
51

0:18 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

- **Cross-API Setting:** We randomly select 75% of the APIs in APISumBench, in which the corresponding sentences are used to fine-tune the models. The remaining 25% of the APIs are used for evaluation.
- **Within-API Setting:** We randomly select 75% of the sentences in APISumBench for fine-tuning models. The remaining 25% of the sentences are used for evaluation.

Table 3. Effectiveness in Section Classification Stage

Setting	System	Precision	Recall	F1
Cross-API Setting	SISE	0.44	0.41	0.43
	DeepTip	0.52	0.51	0.51
	CSSC w/o context	0.58	0.57	0.57
	CSSC	0.63	0.62	0.62
Within-API Setting	SISE	0.62	0.61	0.61
	DeepTip	0.66	0.64	0.65
	CSSC w/o context	0.74	0.72	0.73
	CSSC	0.78	0.77	0.77

Result & Analysis. We report the performance of CSSC and selected baselines in Table 3. CSSC *w/o context* refers to our sentence section classifier without the context-awareness mechanism. **Overall, our CSSC outperforms all baselines in both cross-API and within-API settings.** Specifically, in the cross-API setting, our CSSC outperforms the best-performing baseline (i.e., DeepTip) by 21.15%, 21.57%, and 21.57% in terms of precision, recall, and F1 score, respectively. In the within-API setting, our CSSC outperforms the best-performing baseline (i.e., DeepTip) by 18.18%, 20.31%, and 18.46% in terms of precision, recall, and F1 score, respectively.

By comparing CSSC and CSSC *w/o context*, we observe that the context-awareness mechanism improves the performance of CSSC in terms of all evaluation metrics. Specifically, the context-awareness mechanism improves the performance of CSSC by 8.62%, 8.77%, and 7.02% in terms of precision, recall, and F1 in the cross-API setting; 5.40%, 6.94%, and 5.48% in the within-API setting.

5.2 RQ2: Effectiveness in the Update Extractive Summarization Stage

Experiment Setting. To answer this research question, we compare the performance of our extractive update summarization algorithm ExtUP with considered extractive summarization baselines described in Section 4.2.2. We report ROUGE scores to assess the token overlapping between the generated extractive summaries and the ground-truth summaries in APISumBench. We consider the following two evaluation settings:

- **Summarization from ground-truth API-relevant sentences:** To assess the theoretical summarization ability of each summarization approach, we consider the output of the first stage in APISumBench, i.e., ground-truly insightful sentences associated with specific API documentation sections, as the input of all summarization approaches. Each approach outputs summaries corresponding to API sections. In this setting, there is no training data since all approaches are unsupervised.
- **Summarization from CSSC output:** To simulate real-world application scenarios, we evaluate the performance of ExtUP by summarizing the output of CSSC. Since there is no existing baseline, we combine the best performer in CSSC stage and the best baseline of ExtUP as the baseline. We follow the same cross-API setting of RQ1 to split the dataset, train, and test the models.

Table 4. Effectiveness of Summarizing from Ground-Truth API-Relevant Sentences

System	ROUGE-1	ROUGE-2	ROUGE-L
LexRank	0.606	0.468	0.580
TextRank	0.640	0.528	0.620
TechSumBot	0.672	0.549	0.653
TechSumBot++	0.694	0.579	0.674
ExtUP	0.795	0.720	0.783

Result & Analysis. In Table 4, we report the result for ExtUP summarizing from ground-truthly API-relevant sentences. **The result shows that ExtUP consistently outperforms all the baselines** in terms of ROUGE-1, ROUGE-2, and ROUGE-L. In particular, compared with the best-performing baseline, i.e., TechSumBot++, our ExtUP achieves better performance by 14.55%, 24.35%, and 16.67% in terms of ROUGE-1, ROUGE-2, and ROUGE-L. Note that the performance of TechSumBot++ is higher than TechSumBot for all the evaluation metrics, which justifies our adjustment of TechSumBot.

Besides, ExtUP outperforms TextRank by 24.22%, 36.36%, and 26.29% in terms of ROUGE-1, ROUGE-2, and ROUGE-L, respectively. ExtUP is a variant of the TextRank algorithm. Considering that ExtUP deprioritizes sentences that are semantically similar to the original API documentation, the results demonstrate that ExtUP can effectively reduce the information redundancy between the generated summaries and API documentation.

Table 5. Effectiveness of Summarizing from CSSC Output

System	ROUGE-1	ROUGE-2	ROUGE-L
CSSC w/o context & TechSumBot++	0.345	0.272	0.319
CSSC & ExtUP	0.417	0.361	0.398

We also report the performance of ExtUP summarizing from CSSC output in Table 5. Specifically, **CSSC & ExtUP outperforms our construct baseline** by 20.87%, 32.72%, and 24.76% in terms of ROUGE-1, ROUGE-2, and ROUGE-L, respectively. The superior performance of ExtUP over the baseline demonstrates the effectiveness of our ExtUP in generating extractive update summaries.

5.3 RQ3: End-to-End Effectiveness

Experiment Setting. To answer this research question, we compare the abstractive summarization performance of APIDocBooster with GPT_{sum} detailed in Section 2.3. We conduct a user study to examine the quality of generated summaries. We follow the user study setting and evaluation metrics used in our empirical study on GPT-4 (detailed in Section 2.3). Specifically, we still ask participants to read all augmented API documentation and rate each of them on a 5-point Likert scale in terms of **informativeness**, **relevance**, **readability**, **redundancy**, and **faithfulness**. The user study involves four participants, all are PhD students with at least two years of programming experience in Java or Python. None of them participated in the benchmark construction process. We augment APIs in the test set of APISumBench following the cross-API setting of RQ1.

Result & Analysis. We show the user study results in Table 6. We observe **APIDocBooster significantly outperforms GPT_{sum} in terms of informativeness, relevance, and faithfulness** ($p < 0.01$) by 13.89%, 15.15%, and 30.56%, respectively. Specifically, APIDocBooster is 'very

0:20 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo

Table 6. End-to-end Effectiveness (5-point Likert scale)

System	Informativeness.	Relevance	Readability	Non-Redundancy	Faithfulness
<i>GPT_{sum}</i>	3.6	3.3	4.4	4.5	3.6
APIDocBOOSTER	4.1	3.8	4.5	4.5	4.7

informative in terms of informativeness; between ‘relevant’ to ‘very relevant’ in terms of relevance; and ‘very faithful’ in terms of faithfulness.

We also observe **APIDocBOOSTER and *GPT_{sum}* achieves on-par performance in terms of redundancy and readability.** The negligible difference in terms of redundancy suggests that our extractive update summarization algorithm, ExtUP, performs on par with GPT-4 itself in identifying and eliminating redundant information. Note that, we employ update summarization in the Extract Phase rather than the Abstract Phase, which could reduce the risk that GPT-4 potentially alters the semantics of generated summaries and produces hallucinations during update summarization. Meanwhile, it is reasonable that APIDocBOOSTER performs comparably to *GPT_{sum}* in terms of readability. This can be attributed to the fact that both approaches leverage the abstractive summarization capabilities of GPT-4.

6 THREATS TO VALIDITY

Threats to internal validity are related to APIDocBOOSTER and the implementation errors of the baselines. We mitigate this threat by double-checking our code and reusing the replication packages of the baselines. Meanwhile, we acknowledge that the current version of our dataset may still suffer from the subjective bias of the annotators. We plan to expand APISumBench to mitigate this threat in the future. We also acknowledge that APIDocBOOSTER still has threats to presenting factually inaccurate information, which stem from errors in the external resources. Nevertheless, the implementation of our data provenance mechanism allows effective verification of factualness by tracing back to the original input.

Threats to external validity are related to the generalizability of our benchmark and experimental results. Specifically, the APISumBench dataset is built on Python and Java. They are two of the most popular programming languages. In addition, the design of APIDocBOOSTER is programming-language and API agnostic. APIDocBOOSTER can easily be extended to other programming languages and API types after training with the appropriate data.

Threats to construct validity are related to evaluation metrics. We use ROUGE, a widely used automatic evaluation metric for text summarization approaches [18, 28]. Furthermore, we use informativeness, relevance, readability, diversity, and faithfulness in our human evaluation, which are widely used to evaluate SE domain and NLP domain summarization tasks [16, 18, 25, 51].

7 RELATED WORK

API Documentation Issues. Researchers often conduct surveys about software document issues with software developers. Uddin et al. [49] summarized and explored the frequency and severity of ten commonly encountered document issues. They highlighted that the most pressing problem is document content issues. Aghajani et al. [13] built a taxonomy of software document issues, including 162 issue types. In their following work [12], they investigated 1) the types of document considered useful for each specific development task and 2) the types of documentation issues considered more relevant to developing practitioners. Furthermore, researchers also employ mining-based strategies to capture software document issues that developers [42] or end users [29] frequently discuss. Our work is inspired by previous empirical studies on API documentation issues and aims to augment inadequate API documentation.

1 APIDocBooster: An Extract-Then-Abstract Framework Leveraging Large Language Models for Augmenting API
2 Documentation 0:21

3
4 **API Document Augmentation.** Some approaches enhance the API documentation by integrating
5 usage code examples and corresponding descriptions [48, 56]. Uddin et al. [48] treat pairs of API
6 usage examples and corresponding comments extracted from Stack Overflow as a new form of API
7 documentation. Zhang et al. [56] proposed ADECK to extract code examples from Stack Overflow
8 and map them to a specific API usage scenario. Unlike prior works that focus on providing code
9 examples, our work targets augmenting the natural language part of API documentation, which is
10 considered a complementary approach to the above approaches.

11 Many works aim to enrich the textual part of API documentation. Treude et al. [45] try to identify
12 insightful sentences to specific API from Stack Overflow to augment Java API documentation. Then
13 DeepTip [50] extract Stack Overflow sentences as API tips by training a CNN-based classifier. All
14 previous works position the task as a binary classification to classify insightful sentences solely
15 from Stack Overflow. Differently, our approach positions the task as a multi-class classification to
16 map each sentence to an API document section, followed by an updating summarization task to
17 improve the readability of API documentation. In addition, we consider the context of the sentence
18 as well as the structure of API documentation. Finally, we are the first to leverage the abstraction
19 summarization approaches and GPT in this task.

20 21 8 CONCLUSION AND FUTURE WORK

22 To augment API documentation, this work proposes a two-stage framework APIDocBooster to au-
23 tomatically generate complementary summaries to API documentation based on multiple resources
24 at the section level. APIDocBooster leverages a BERT-based classifier named CSSC to classify
25 sentences from multiple sources into the corresponding API documentation sections. Then CSSC
26 implements a context-awareness mechanism to adjust the prediction results of context-dependent
27 sentences considering their context. Next, UPSUM proposes an extractive update summarization
28 algorithm to generate extractive summaries semantically different from the original API docu-
29 mentation, which is followed by GPT generating abstractive summaries guided by the extractive
30 summaries. Comprehensive evaluation results on our proposed dataset APISumBench and human
31 evaluation show that 1) both CSSC and UPSUM outperform their baselines by a large margin, and
32 2) APIDocBooster performs better than the strong baseline GPT.

33 In the future, we plan to develop APIDocBooster in the form of a browser plug-in to assist
34 API users for better understanding the target API. In addition, we plan to explore the potential of
35 other online sources, e.g., API tutorials and technical blogs, to further improve the performance of
36 APIDocBooster.

37 38 9 DATA AVAILABILITY

39 The replication package of APIDocBooster and APISumBench is at [https://github.com/autumn-](https://github.com/autumn-city/APIDocBooster)
40 [city/APIDocBooster](https://github.com/autumn-city/APIDocBooster).

41 42 REFERENCES

- 43 [1] [n. d.]. *Android Documentation*. <https://developer.android.com/reference>
- 44 [2] [n. d.]. *HuggingFace*. <https://huggingface.com/>
- 45 [3] [n. d.]. *Java Documentation*. <https://docs.oracle.com/en/java/>
- 46 [4] [n. d.]. *NLTK*. <https://www.nltk.org/>
- 47 [5] [n. d.]. *OpenAI Tokenizer Calculator*. <https://platform.openai.com/tokenizer>
- 48 [6] [n. d.]. *Pyrouge*. <https://github.com/bheinzerling/pyrouge>
- 49 [7] [n. d.]. *Python Documentation*. <https://docs.python.org/>
- 50 [8] [n. d.]. *Stack Exchange API*. <https://api.stackexchange.com/>
- 51 [9] [n. d.]. *YouTube Data API*. <https://developers.google.com/youtube/v3>
- 52 [10] [n. d.]. *YouTube Transcript API*. <https://developers.google.com/youtube/v3>

- 0:22 Chengran Yang, Jiakun Liu, Bowen Xu, Christoph Treude, Yunbo Lyu, Junda He, Ming Li, and David Lo
- [11] Emad Aghajani, Gabriele Bavota, Mario Linares-Vásquez, and Michele Lanza. 2019. Automated documentation of android apps. *IEEE Transactions on Software Engineering* 47, 1 (2019), 204–220.
 - [12] Emad Aghajani, Csaba Nagy, Mario Linares-Vásquez, Laura Moreno, Gabriele Bavota, Michele Lanza, and David C Shepherd. 2020. Software documentation: the practitioners’ perspective. In *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*. IEEE, 590–601.
 - [13] Emad Aghajani, Csaba Nagy, Olga Lucero Vega-Márquez, Mario Linares-Vásquez, Laura Moreno, Gabriele Bavota, and Michele Lanza. 2019. Software documentation issues unveiled. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*. IEEE, 1199–1210.
 - [14] Stefanos Angelidis, Reinald Kim Amplayo, Yoshihiko Suhara, Xiaolan Wang, and Mirella Lapata. 2021. Extractive opinion summarization in quantized transformer spaces. *Transactions of the Association for Computational Linguistics* 9 (2021), 277–293.
 - [15] Stefanos Angelidis and Maria Lapata. 2018. Summarizing Opinions: Aspect Extraction Meets Sentiment Prediction and They Are Both Weakly Supervised. In *2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 3675–3686.
 - [16] Adithya Bhaskar, Alex Fabbri, and Greg Durrett. 2023. Prompted Opinion Summarization with GPT-3.5. In *Findings of the Association for Computational Linguistics: ACL 2023*. 9282–9300.
 - [17] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
 - [18] Yang Chengran, Bowen Xu, Ferdian Thung, Yucen Shi, Ting Zhang, Zhou Yang, Xin Zhou, Jieke Shi, Junda He, DongGyun Han, et al. 2022. Answer Summarization for Technical Queries: Benchmark and New Approach. In *2022 37nd IEEE/ACM International Conference on Automated Software Engineering (ASE)*.
 - [19] Amy B Cyphert. 2021. A human being wrote this law review article: GPT-3 and the practice of law. *UC Davis L. Rev.* 55 (2021), 401.
 - [20] Lucas BL De Souza, Eduardo C Campos, and Marcelo de A Maia. 2014. Ranking crowd knowledge to assist software development. In *Proceedings of the 22nd International Conference on Program Comprehension*. 72–82.
 - [21] DK Donker, A Hasman, and HP Van Geijn. 1993. Interpretation of low kappa values. *International journal of bio-medical computing* 33, 1 (1993), 55–64.
 - [22] Wafaa S El-Kassas, Cherif R Salama, Ahmed A Rafea, and Hoda K Mohamed. 2021. Automatic text summarization: A comprehensive survey. *Expert systems with applications* 165 (2021), 113679.
 - [23] Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research* 22 (2004), 457–479.
 - [24] Golará Garousi, Vahid Garousi-Yusifoglu, Guenther Ruhe, Junji Zhi, Mahmoud Moussavi, and Brian Smith. 2015. Usage and usefulness of technical software documentation: An industrial case study. *Information and software technology* 57 (2015), 664–682.
 - [25] Dandan Huang, Leyang Cui, Sen Yang, Guangsheng Bao, Kun Wang, Jun Xie, and Yue Zhang. 2020. What Have We Achieved on Text Summarization?. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 446–469.
 - [26] Devlin Jacob, Chang Ming-Wei, Lee Kenton, and Toutanova Kristina. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT*. 4171–4186.
 - [27] Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *Comput. Surveys* 55, 12 (2023), 1–38.
 - [28] Ashkan Kazemi, Verónica Pérez-Rosas, and Rada Mihalcea. 2020. Biased TextRank: Unsupervised Graph-Based Content Extraction. In *Proceedings of the 28th International Conference on Computational Linguistics*. 1642–1652.
 - [29] Hammad Khalid, Emad Shihab, Meiyappan Nagappan, and Ahmed E Hassan. 2014. What do mobile app users complain about? *IEEE software* 32, 3 (2014), 70–77.
 - [30] Junaed Younus Khan, Md Tawkat Islam Khondaker, Gias Uddin, and Anindya Iqbal. 2021. Automatic detection of five api documentation smells: Practitioners’ perspectives. In *2021 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE, 318–329.
 - [31] Peter Lee, Sebastien Bubeck, and Joseph Petro. 2023. Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine. *New England Journal of Medicine* 388, 13 (2023), 1233–1239.
 - [32] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*. 74–81.
 - [33] Walid Maalej and Martin P Robillard. 2013. Patterns of knowledge in API reference documentation. *IEEE Transactions on Software Engineering* 39, 9 (2013), 1264–1282.
 - [34] Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan McDonald. 2020. On Faithfulness and Factuality in Abstractive Summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. 1906–1919.

1
2 APIDocBooster: An Extract-Then-Abstract Framework Leveraging Large Language Models for Augmenting API
3 Documentation 0:23

- 4 [35] Rada Mihalcea and Paul Tarau. 2004. Texttrank: Bringing order into text. In *Proceedings of the 2004 conference on*
5 *empirical methods in natural language processing*. 404–411.
- 6 [36] Oracle. [n. d.]. *API Specification for JAVA documentation*. [https://www.oracle.com/java/technologies/javase/api-](https://www.oracle.com/java/technologies/javase/api-specifications.html#class)
7 [specifications.html#class](https://www.oracle.com/java/technologies/javase/api-specifications.html#class)
- 8 [37] Stack Overflow. [n. d.]. *Stack Overflow Developer Survey 2023*. [https://survey.stackoverflow.co/2023/#learning-to-code-](https://survey.stackoverflow.co/2023/#learning-to-code-learn-code-learn)
9 [learn-code-learn](https://survey.stackoverflow.co/2023/#learning-to-code-learn-code-learn)
- 10 [38] Reinhold Plösch, Andreas Dautovic, and Matthias Saft. 2014. The value of software documentation quality. In *2014*
11 *14th International Conference on Quality Software*. IEEE, 333–342.
- 12 [39] PyTorch. [n. d.]. *PyTorch Documentation*. <https://pytorch.org/docs/stable/>
- 13 [40] Stephen Robertson, Hugo Zaragoza, et al. 2009. The probabilistic relevance framework: BM25 and beyond. *Foundations*
14 *and Trends® in Information Retrieval* 3, 4 (2009), 333–389.
- 15 [41] Martin P Robillard and Robert DeLine. 2011. A field study of API learning obstacles. *Empirical Software Engineering* 16,
16 6 (2011), 703–732.
- 17 [42] Christoffer Rosen and Emad Shihab. 2016. What are mobile developers asking about? a large scale study using stack
18 overflow. *Empirical Software Engineering* 21, 3 (2016), 1192–1223.
- 19 [43] Agus Sulistya, Gede Artha Azriadi Prana, Abhishek Sharma, David Lo, and Christoph Treude. 2020. Sieve: Helping
20 developers sift wheat from chaff via cross-platform analysis. *Empirical Software Engineering* 25 (2020), 996–1030.
- 21 [44] Christoph Treude, Justin Middleton, and Thushari Atapattu. 2020. Beyond accuracy: assessing software documentation
22 quality. In *Proceedings of the 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on*
23 *the Foundations of Software Engineering*. 1509–1512.
- 24 [45] Christoph Treude and Martin P Robillard. 2016. Augmenting api documentation with insights from stack overflow. In
25 *2016 IEEE/ACM 38th International Conference on Software Engineering (ICSE)*. IEEE, 392–403.
- 26 [46] Gias Uddin, Olga Baysal, Latifa Guerrouj, and Foutse Khomh. 2019. Understanding how and why developers seek and
27 analyze API-related opinions. *IEEE Transactions on Software Engineering* 47, 4 (2019), 694–735.
- 28 [47] Gias Uddin and Foutse Khomh. 2017. Automatic summarization of API reviews. In *2017 32nd IEEE/ACM International*
29 *Conference on Automated Software Engineering (ASE)*. IEEE, 159–170.
- 30 [48] Gias Uddin, Foutse Khomh, and Chanchal K Roy. 2021. Automatic api usage scenario documentation from technical
31 q&a sites. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 30, 3 (2021), 1–45.
- 32 [49] Gias Uddin and Martin P Robillard. 2015. How API documentation fails. *Ieee software* 32, 4 (2015), 68–75.
- 33 [50] Shaohua Wang, NhatHai Phan, Yan Wang, and Yong Zhao. 2019. Extracting API tips from developer question and
34 answer websites. In *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*. IEEE, 321–332.
- 35 [51] Bowen Xu, Zhenchang Xing, Xin Xia, and David Lo. 2017. AnswerBot: Automated generation of answer summary to
36 developers’ technical questions. In *ASE 2017*. IEEE, 706–716.
- 37 [52] Yumo Xu and Mirella Lapata. 2020. Coarse-to-fine query focused multi-document summarization. In *Proceedings of the*
38 *2020 Conference on empirical methods in natural language processing (EMNLP)*. 3632–3645.
- 39 [53] Chengran Yang, Bowen Xu, Junaed Younus Khan, Gias Uddin, Donggyun Han, Zhou Yang, and David Lo. 2022.
40 Aspect-Based API Review Classification: How Far Can Pre-Trained Transformer Model Go?. In *2022 IEEE International*
41 *Conference on Software Analysis, Evolution and Reengineering (SANER)*. IEEE Computer Society.
- 42 [54] Chengran Yang, Bowen Xu, Ferdian Thung, Yucen Shi, Ting Zhang, Zhou Yang, Xin Zhou, Jieke Shi, Junda He,
43 Donggyun Han, et al. 2022. Answer Summarization for Technical Queries: Benchmark and New Approach. In
44 *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*. 1–13.
- 45 [55] Haopeng Zhang, Xiao Liu, and Jiawei Zhang. 2023. Extractive summarization via chatgpt for faithful summary
46 generation. *arXiv preprint arXiv:2304.04193* (2023).
- 47 [56] Jingxuan Zhang, He Jiang, Zhilei Ren, Tao Zhang, and Zhiqiu Huang. 2019. Enriching API documentation with code
48 samples and usage scenarios from crowd knowledge. *IEEE Transactions on Software Engineering* 47, 6 (2019), 1299–1314.
- 49 [57] Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B Hashimoto. 2023.
50 Benchmarking large language models for news summarization. *arXiv preprint arXiv:2301.13848* (2023).
- 51 [58] Hao Zhong and Hong Mei. 2017. An empirical study on API usages. *IEEE Transactions on Software Engineering* 45, 4
52 (2017), 319–334.