

Revisiting, Benchmarking and Exploring API Recommendation: How Far Are We?

Yun Peng[†], Shuqing Li[†], Wenwei Gu[†], Yichen Li[†], Wenxuan Wang[†], Cuiyun Gao^{*‡}, and Michael Lyu[†]

[†]The Chinese University of Hong Kong, Hong Kong, China

[‡]Harbin Institute of Technology, Shenzhen, China

{ypeng, sqli21, wwgu21, ycli21, wxwang, lyu}@cse.cuhk.edu.hk, gaocuiyun@hit.edu.cn

Abstract—Application Programming Interfaces (APIs), which encapsulate the implementation of specific functions as interfaces, greatly improve the efficiency of modern software development. As numbers of APIs spring up nowadays, developers can hardly be familiar with all the APIs and usually need to search for appropriate APIs for usage. So lots of efforts have been devoted to improving the API recommendation task. However, it has been increasingly difficult to gauge the performance of new models due to the lack of a uniform definition of the task and a standardized benchmark. For example, some studies regard the task as a code completion problem, while others recommend relative APIs given natural language queries. To reduce the challenges and better facilitate future research, in this paper, we revisit the API recommendation task and aim at benchmarking the approaches. Specifically, the paper groups the approaches into two categories according to the task definition, i.e., query-based API recommendation and code-based API recommendation. We study 11 recently-proposed approaches along with 4 widely-used IDEs. One benchmark named as APIBench is then built for the two respective categories of approaches. Based on APIBench, we distill some actionable insights and challenges for API recommendation. We also achieve some implications and directions for improving the performance of recommending APIs, including appropriate query reformulation, data source selection, low resource setting, user-defined APIs and query-based API recommendation with usage patterns.

Index Terms—API recommendation, benchmark, empirical study

1 INTRODUCTION

Application Programming Interfaces (APIs) provided by software libraries or frameworks play an important role in modern software development. Almost all programs, even the basic “hello world!” program, include at least one API. However, there are a huge number of APIs from different modules or libraries. For example, Java standard library [70] provides more than 30,000 APIs. It is therefore infeasible for developers to be familiar with all APIs. To address this problem, many approaches are proposed to recommend APIs based on input queries, which describe the programming task in natural language, or surrounding context, i.e., the code already written by developers.

However, a uniform definition of the current API recommendation task is still absent, leading the task hard to be followed by potential researchers. Some studies [11], [44], [61], [82], [84] regard the task as a code completion problem, and recommend any code tokens including APIs. These studies focus on improving the prediction results of all the tokens instead of only APIs. Some studies [31], [37], [49], [76], [78] recommend relative APIs on different levels given natural language queries. Besides, the evaluation results are difficult to be reproduced by future related work. For example, for query-based API recommendation, manual evaluation is generally adopted, so the performances reported by different studies are difficult to be aligned. Comparing with widely-used Integrated Development Environments (IDEs) or search engines is another commonly

adopted yet inconsistent evaluation strategy in previous research. Therefore, to better facilitate future exploration of the API recommendation task, in this paper, we summarize the recent related approaches and build a general benchmark named APIBENCH.

To facilitate the benchmark creation, we group the recent related approaches into two categories according to the task definition: query-based API recommendation and code-based API recommendation:

1) **query-based API recommendation.** Approaches for query-based API recommendation aim at providing related APIs to developers given a query that describes programming requirements in natural language. The approaches can inform developers which API to use for a programming task.

2) **code-based API recommendation.** Approaches for code-based API recommendation aim at predicting the next API given the code surrounding the point of prediction. They can directly improve the efficiency of coding.

Besides the unreproducible evaluation, the two groups of studies face their own challenges. 1) For query-based approaches, high-quality queries play a critical role in accurate recommendation. However, there may exist a knowledge gap between developers and API designers in choosing terms for describing queries or APIs. For example, developers who do not know the term “*heterogeneous list*” in API documents would use other words such as “*list with different types of elements*” in the query. Whether current query reformulation techniques are effective for API recommendation and how effective it is are still remaining

* corresponding author.

unexplored. 2) For code-based approaches, the quality of code before the recommendation point also affects the recommendation performance. Generally, the approaches are evaluated by simulating an actual development, i.e., some parts of a project are removed for imitating a limited context. The APIs to recommend may locate in the front, middle, or back of the code, so exploring the impact of different recommendation points is important for understanding the recommendation capability of existing approaches. Other factors such as whether the APIs are standard or user-defined, lengths of given context, and different domains can also influence the recommendation performance, which have not yet been fully investigated.

To comprehensively understand the above challenges and align the performance of current approaches, we first build a benchmark named APIBENCH. APIBENCH is built on Python and Java, and involves two datasets for evaluation, named as APIBENCH-Q and APIBENCH-C for query-based and code-based approaches, respectively. APIBENCH-Q contains 6,563 Java queries and 4,309 Python queries obtained from Stack Overflow and API tutorial websites. APIBENCH-C contains 1,477 Java projects with 1,229,698 source files and 2,223 Python projects with 414,753 source files obtained from GitHub. Based on APIBENCH, we study the following research questions:

- **RQ1:** How effective are current query-based and code-based API recommendation approaches?
- **RQ2:** What is the impact of query reformulation techniques on the performance of query-based API recommendation?
- **RQ3:** What is the impact of different data sources on the performance of query-based API recommendations?
- **RQ4:** How well do code-based approaches recommend different kinds of APIs?
- **RQ5:** What is the performance of code-based approaches in handling different contexts?
- **RQ6:** How well do code-based approaches perform in cross-domain scenarios?

APIBENCH involves the implementation of the related approaches proposed in the recent five years, specifically including five query-based approaches and five code-based approaches. In RQ1, we compare the performances of the approaches in APIBENCH. To answer RQ2 and RQ3, we apply four popular query reformulation techniques to the queries of APIBENCH-Q and observe the performance of the query-based approaches given reformulated queries. To answer RQ4 to RQ6, we analyze the APIs in APIBENCH-C from different aspects and study the performance of code-based approaches under different experimental settings.

Key Findings. Through the large-scale empirical study, we achieve some findings and summarize the key findings as below.

(1) For query-based API recommendation:

- While current approaches make a good progress on class-level recommendation, recommending the exact API methods is still a challenging task.
- Query reformulation techniques, including query expansion and query modification, are quite effective

in improving the performance of query-based approaches.

- Adding data sources such as Q&A forums and tutorials that are more similar to real-world queries can significantly improve the performance of current approaches.

(2) For code-based API recommendation:

- Recent deep learning models such as Transformers show superior performance on this task. Meanwhile, current IDEs can achieve competitive performance as recent pattern-based and learning-based approaches. They work far more than just recommending APIs based on alphabet orders.
- Current approaches are effective to recommend APIs from standard libraries and popular third-party libraries, but their performance drops a lot when recommending user-defined or project-specific APIs.
- Approaches trained on one single domain face the problem of cross-domain adaptation. Approaches trained on multiple domains achieve satisfying performance when testing on most single domains, and they even outperform those trained on corresponding single domains.

Based on the findings, we conclude some implications and suggestions that would benefit future research. On the one hand, query-based API recommendation approaches should be built along with query reformulation techniques to handle queries with different qualities. We also encourage future work to leverage different data sources and few-shot learning methods to address the low resource challenge in query-based API recommendation. On the other hand, we suggest future code-based API recommendation approaches focus on improving the performance of recommending user-defined APIs as it is currently the major bottleneck.

Contributions. To sum up, our contribution can be concluded as follows.

- To the best of our knowledge, we are the first to systematically study both query-based and code-based API recommendation techniques on two large-scale datasets including Java and Python.
- We build an open-sourced benchmark named APIBENCH to fairly evaluate query-based and code-based approaches.
- We study how different settings can impact the performance of current approaches, including query quality, cross domain adaptation, etc.
- We conclude some findings and implications that would be important for future research in API recommendation.

The rest of this paper is organized as follows. We present the background and regular API recommendation process in Section 2. We describe the details of APIBENCH, current baselines and evaluation metrics in Section 3. Then we introduce the experiment results and potential findings on query-based and code-based API recommendation in Section 4 and Section 5, respectively. Based on the findings, we conclude some implications and future directions in Section 6. Finally, we discuss threats to validity and related work in Section 7

and Section 8, respectively. Section 9 provides our conclusion.

2 BACKGROUND

In this section, we summarize the query-based approaches and code-based approaches, respectively.

2.1 Query-Based API Recommendation Methods

We describe the typical query-based API recommendation process in Figure 1. Given a query “*Calculate int value square root*”, query reformulation techniques first modify the query as “*return int value square root*” or expand it as “*finally calculate int value square root*”. A knowledge base built upon available data sources is also prepared for API candidate selection. Based on the knowledge base, retrieval-based methods or learning-based methods recommend the APIs relevant to the queries.

2.1.1 Query Reformulation Techniques

Input queries can be short in length or vague in semantics. Besides, there may exist a knowledge gap between developers and search engines in query description. For rendering search engines better understand the query semantics, query reformulation is a common pre-processing method. In general, there are two major types of query reformulation approaches:

- 1) query expansion, which adds extra information to the original queries;
- 2) query modification, which modifies, replaces or deletes some words in the original queries.

Query expansion. Query expansion aims at identifying important words that are missing in the input queries. The topic is originally stemmed from the field of natural language processing (NLP). For example, the work [51] utilizes word embeddings to map words in the vector space and finds similar words to enrich the queries. For the API recommendation task, since APIs are encapsulated and organized according to classes and modules, class names and module names are important hints for recommendation. Rahman *et al.* [76], [78] propose to use keyword-API class co-occurrence frequencies and keyword-keyword co-occurrence frequencies to build the relationship between words and API classes, and add the suggested API class for query expansion.

Query modification. Query modification aims at mitigating both the lexical gap and knowledge gap between the user queries and descriptions in knowledge base. The lexical gap, such as mis-spelling, can be easily addressed by spelling correction and synonym search, etc. Recent work focuses on how to mitigate the knowledge gap by replacing inappropriate words in queries. Mohammad *et al.* [4] extract important tokens in code, and Sirres *et al.* [88] leverage discussions and code from Stack Overflow posts to build a knowledge base. Cao *et al.* [12] collect query reformulation history from Stack Overflow and propose a Transformer-based approach to learn how developers change their queries when search engines do not return desired results.

2.1.2 Recommendation with Knowledge Base

Knowledge base. API recommendation approaches generally require a knowledge base that contains all the existing APIs as the search space. There are three primary sources for the knowledge base creation, including: 1) official documentations which contain comprehensive descriptions about the API functionality and structure. 2) Q&A forums, which provide the purposes of APIs and different API usage patterns. Many studies [37], [75] leverage the Q&A pairs from Stack Overflow to select API candidates. 3) Wiki sites, which describe concepts that link different APIs. For example, Liu *et al.* [49] utilizes API concepts from Wikipedia to help build API knowledge graphs.

Retrieval-based methods. Retrieval-based methods retrieve API candidates from the knowledge base and then rank the candidate APIs by calculating the similarities between queries and APIs. For example, Rahman *et al.* [76], [78] utilize the keyword-API occurrence frequencies and API-API occurrence frequencies to find the most relevant APIs. Huang *et al.* [37] first identify the similar posts from Stack Overflow by computing query-documentation similarities and choose the APIs mentioned in posts as candidates. Liu *et al.* [49] build an API knowledge graph to represent relationships between APIs and then calculate the similarities between queries and certain parts of API knowledge graph to rank the APIs.

Learning-based methods. Another type of method is to automatically learn the relationships between queries and APIs based on deep learning techniques. The knowledge base provides query-API pairs as the ground truth. For example, Gu *et al.* [31] formulate the task as a translation problem in which a model is built to translate word sequences into API sequences. They propose an RNN model with an encoder-decoder structure to implement the translation.

2.2 Code-Based API Recommendation Methods

We describe the workflow of code-based API recommendation in Figure 2. Given a target code, context representation is an essential step. Based on the extracted context, pattern-based methods or learning-based methods are adopted by previous studies to recommend the next API.

2.2.1 Context for the Target Code

Most code-based API recommendation methods regard the code before the recommendation point as the context. We name such context as *internal context* since it only considers code in the current source code or current function body. For example, Line 1 ~ 6 of the target code in Figure 2 belongs to internal context. Xie *et al.* [98] find that replacing external APIs in code (such as `Arrays.asList()` in Figure 2) with their implementations can help the identification of common usage patterns. They propose to build a hierarchical context by integrating the implementation out of the current source file. We name the implementation of external APIs as *external context*.

2.2.2 Context Representation

We divide the context representation methods into two types, i.e., pattern-based representation and learning-based representation. **Pattern-based representations** [18], [63],

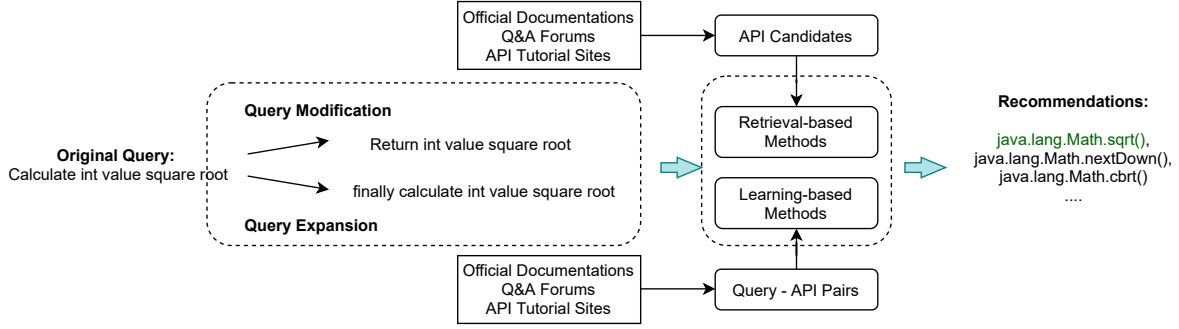


Fig. 1. The typical query-based API recommendation framework.

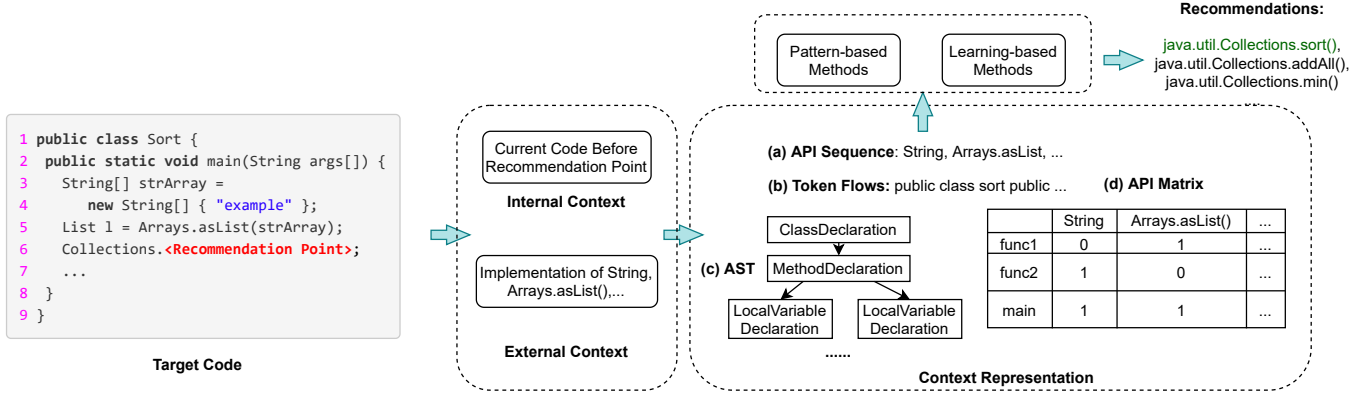


Fig. 2. The typical code-based API recommendation framework.

[64], [97], [98] do not consider all the code tokens. Instead, they only identify APIs to build API usage sequences, as shown in Figure 2 (a), API matrix, as shown in Figure 2 (d), or API dependency graphs to represent the current context. **Learning-based representations** [32], [35], [44], [82], [92] usually represent the context with token flows, as illustrated in Figure 2 (b), or other syntax structures such as Abstract Syntax Trees (ASTs), as illustrated in Figure 2 (c).

2.2.3 Recommendation Based on Context

Pattern-based methods. API recommendation is inherently a recommendation task, so some studies [18], [64] follow the collaborative filtering (*user-item*) methodology of traditional recommendation systems [83]. As shown in Figure 2 (d), they regard the internal context as the *users* and APIs as the *items*. They then calculate the similarities between different *users* to find the most similar API for recommendation. However, the methods do not consider the relationships between APIs. More recent work [97], [98] build API dependency graphs or mines association rules to capture API usage patterns.

Learning-based methods. Hindle *et al.* [35] discover the naturalness of software, rendering it possible to deploy machine learning or deep learning methods on code. Different from pattern-based methods that consider the relationships between API occurrences, learning-based methods regard API as a single code token, and reformulate the code-based API recommendation problem into a *next token prediction* problem. Many statistical language models [62], [80], [82], [92] are proposed to predict the next code token. Besides

using the token sequences, more recent work [32], [44] try to leverage syntax and data flow information for more accurate prediction.

3 METHODOLOGY

In this section, we introduce the scope of the studied APIs, the preparation of benchmark datasets, and implementation details.

3.1 Scope of APIs

To fairly compare the current API recommendation approaches, benchmark datasets should be prepared, during which the scope of studied APIs firstly needs to be defined. In this work, we focus our evaluation on two popular programming languages, i.e., Python and Java.

For facilitating the analysis of the challenges in API recommendation, we divide all APIs into *standard APIs*, *user-defined APIs*, and *popular third-party APIs*. The *standard APIs* refer to the APIs that are clearly defined and built-in in corresponding programming languages while the *user-defined APIs* are defined and used in projects along with *popular third-party APIs*. Following previous work [37], [49], [76], [78] we evaluate query-based API recommendation methods only on the *standard APIs* since currently *standard APIs* have the most comprehensive documentations and extensive discussions to build the knowledge base. We evaluate code-based API recommendation methods on all three kinds of APIs. The details of each kind for different programming languages are depicted below.

TABLE 1
Statistics of Benchmark APIBENCH-C. The data includes both the training set and testing set.

PL	Domain	#Projects	#Files	LOC (per func)	#API (per func)	Total number of APIs (only testset)			LOC Threshold of Short Func	LOC Threshold of Long Func
						Standard	User-defined	Popular		
Python	General	899	230,064	15.24	5.55	1,363,240	1,747,878	54,244	8.875	54.875
	ML	323	46,556	13.89	6.08	629,437	339,821	125,377	12.65	46.05
	Security	126	15,785	18.98	6.72	111,393	64,809	3,613	6	86.5
	Web	568	82,771	14.14	5.05	369,114	241,602	11,832	7.35	51.625
	DL	307	39,577	14.58	6.25	413,295	220,228	76,654	11.675	52.525
Java	General	935	1,056,790	11.16	4.06	5,164,481	3,808,124	36,178	6.26	19.2
	Android	377	87,468	8.24	2.91	517,461	267,141	75,069	7.28	16.8
	ML	52	41,377	12.82	4.77	194,013	136,963	0	7.52	19.74
	Testing	55	23,618	9.93	3.98	105,577	55,241	22	6.44	15.68
	Security	58	20,445	12.35	5.32	125,558	74,471	1,243	6.88	20.78

TABLE 2

Statistics of APIBENCH-Q. Ori. represent the original queries, Exp. represent the expanded queries produced by query expansion techniques, Mod. represent the modified queries produced by query modification techniques.

PL	Stack Overflow			Tutorial Websites		
	Ori.	Exp.	Mod.	Ori.	Exp.	Mod.
Python	1,925	78,157	100,100	2,384	95,360	123,968
Java	1,320	80,343	68,640	5,243	319,783	272,636

(1) **Standard Java APIs.** We choose the version Java 8 for our analysis since it is the most widely-used version in current projects according to the 2020 JVM Ecosystem Report [89]. We collect 34,072 APIs from the Java documentation [70] as *standard APIs*.

(2) **Popular Java third-party APIs.** We choose APIs from the Android library [28] since Android is one of the most popular applications of Java programs. We collect 11,802 APIs from the official documentation of Android in total and regard them as the popular third-party APIs.

(3) **Standard Python APIs.** As Python Software Foundation has stopped the support for Python 2, currently only 6% of developers are still using Python 2, according to the development survey conducted by JetBrains [39]. Considering that APIs of different versions above 3.0 are similar, we choose the newest version 3.9 to ensure the compatibility, and collect 5,241 APIs from Python standard library [72] as the *standard APIs*.

(4) **Popular Python third-party APIs.** Python is well extended by a lot of third-party modules. We choose five widely-used modules with sufficient documentations, including flask [21], django [17], matplotlib [53], pandas [71] and numpy [68]. We collect 215, 700, 4,089, 3,296 and 3,683 APIs from them, respectively.

(5) **User-defined APIs.** For code-based API recommendation, we regard all the functions defined in current projects as *user-defined APIs*. We do not explicitly collect them as a fixed set because they vary across projects. By inspecting the implementations, we can always identify the user-defined APIs.

3.2 Benchmark Datasets

In this section, we describe how we build the benchmark datasets APIBENCH-Q and APIBENCH-C.

3.2.1 Creation of APIBENCH-Q

We build the benchmark dataset APIBENCH-Q by mining Stack Overflow and tutorial websites. Note that we find that currently there is no query-based API recommendation approach specially designed for Python programs, but we still collect the query benchmark for it to facilitate further research investigation.

Mining Stack Overflow. As one of the most popular Q&A forums for developers, Stack overflow contains much discussion about the usage of APIs. Stack Overflow is the primary source for building APIBENCH-Q. We first download all posts from Aug 2008 to Feb 2021 on Stack Overflow (SO) via Stack Exchange Data Dump [20]. Each post is associated with a tag about the related programming language. We filter out the posts not tagged as Java or Python, resulting in 1,756,183 Java posts and 1,661,383 Python posts. Similar to other studies related to Stack Overflow [37], [76], we further filter out the posts based on the following rules:

- To increase the quality of the posts, we remove the posts that do not have endorsed answers.
- We remove the posts that do not contain the HTML tag `<code>`, because we cannot extract any API from them.
- We remove the posts that contain code snippets longer than two lines, since large code snippets in a post usually indicate that the query should be handled by a series of operations, not a single API sequence.
- We check the code in each post and remove the posts that do not contain any APIs involved in this paper, as described in Sec. 3.1.

After the rule-based filtering, we obtain 156,493 Python posts and 148,938 Java posts that contain descriptions about APIs. However, some of the posts are not directly related to API recommendation. For example, some posts only ask about comparing two similar APIs. The unrelated posts are hard to be automatically identified by rules. To ensure the relatedness of the posts in our benchmark dataset, we invite 16 participants with an average of 3-year development

experience in Python or Java for manually checking. For each post, two of the participants are involved to check the following aspects:

- 1) whether the query asks about API recommendation;
- 2) whether the standard APIs recognized by the previous rules are intact, i.e., including the whole class and method names.
- 3) whether the APIs in answers exactly address the query.

If two participants provide the same answers for one post and also one of the above three aspects is not satisfied, we directly remove the post. If the two participants do not reach an agreement, the post will be forwarded to one of the authors to make a final decision.

As the remaining posts are still quite many to be manually checked, we conduct two rounds of annotations. In the first round of annotations, we ask the annotators to label 100 randomly selected posts and conclude the reasons for the cases that they think are not about API recommendation. Then we collect the keywords that frequently appear in these unrelated cases, and remove the posts whose titles contain such keywords. For example, the keyword “error” is mostly used by developers to ask for debugging in Stack Overflow, which is not related to API recommendation. In the second round annotations, we ask annotators to label all the remaining posts. It takes about one month to complete the two-round annotation process. After both rounds of annotations, we manually check 13,775 posts, in which 1,262 posts do not reach an agreement by the annotators and need further check by one of the authors. We use the commonly-used Fleiss Kappa score [22] to measure the agreement degree between the two annotators and the value is 0.77. The result indicates a high agreement between them. Based on the manual check, 3,245 of the 13,775 labeled posts are remained, including 1,925 Python queries and 1,320 Java queries, comprising the first part of our benchmark APIBENCH-Q, as shown in the second column of Table 2.

Mining tutorial websites. API tutorial websites are the second major source of query-API pairs. We choose three popular API tutorial websites GeeksforGeeks [1], Java2s [2] and Kode Java [3] to establish APIBENCH-Q. Different from Stack Overflow that contains discussion on various topics, API tutorial websites focus on providing examples of how to use APIs. Therefore, manually annotating the relatedness of each query to API recommendation is not necessary. We adopt similar rules as mining Stack Overflow to filter out those without code snippets or associated with large code snippets. We finally collect 5,243 Java queries and 2,384 Python queries, which comprise the second part of our benchmark APIBENCH-Q, as shown in the fifth column of Table 2.

3.2.2 Creation of APIBENCH-C

We create the benchmark dataset APIBENCH-C by mining GitHub. GitHub [55] is one of the most popular websites for sharing code and includes large numbers of code repositories on different topics and programming languages.

In order to explore the performance of API recommendation under different domains, we first determine the domains for analysis. According to the JetBrains’ developer survey [39] and topic labels provided by GitHub, we choose




Fig. 3. Distribution of code lines per function at *general* domain (Left: Python, Right: Java).

four popular domains for Python and Java, respectively, as shown in Table 1. For Python, we consider the domains “ML”, “Security”, “Web”, and “DL”; while for Java, we involve domains “Android”, “ML”, “Testing”, and “Security”. For each domain, we collect 500 repositories with the most stars and 500 repositories with the most forks on GitHub¹. Besides the specific domains, we also build a “General” domain which only considers the popularity of repositories. For the “General” domain, we collect 1,000 repositories with the most stars and 1,000 repositories with the most forks on GitHub regardless of the topics.

Not all the collected repositories are applicable for code-based API recommendation. Some popular repositories do not contain enough code, e.g., only including documentations. To remove such repositories, we use *cloc* [6] to scan the code in each repository and filter out the repositories that 1) have fewer than 10 files or 2) have fewer than 1000 lines of code or 3) have code in Python or Java but with the ratio less than 10%. The number of projects, number of files, and average number of code lines for each domain of APIBENCH-C are shown in Table 1.

As most approaches [27], [32], [44], [64], [82] for code-based API recommendation require a training set to learn the API patterns or train the models, we split APIBENCH-C into a training set and a test set with a ratio of 80% and 20%, respectively. Note that we do not split a project both into the training set and test set, but put all the files of the same project into either the training set or test set, because Alon *et al.* [8] and LeClair *et al.* [47] find that code in the same project usually share the same variable names and code patterns, and splitting without considering project can cause data leakage. For the approaches requiring a validation set, we prepare it from the training set.

In order to study the impact of different recommendation points and different lengths of functions on the performance of current approaches, we analyze the average length of functions in each repository. By using Kernel Density Estimation (KDE), we observe that the length distributions of functions in different domains are similar. The distributions of the “General” domain are depicted in Figure 3. From the figure we observe that most Python functions contain 5 ~ 30 lines of code (LOC) and most Java functions contain 5 ~ 20 lines of code. However, there still exist a few extremely short or long functions. These functions are likely

1. The collection was conducted during April 2021.

TABLE 3

The query reformulation techniques and query-based API recommendation approaches involved in the paper. For tools, the years they were last updated are listed. The column name "PL" indicates the applicable programming language.

Approach	Category/ Data Source	Description	PL	Venue	Year
Query Reformulation					
Google Prediction Service	Query expansion, modification	Google Prediction Service [30] is a common query reformulation technique used in Google Search.	Any	-	2021
NLPAUG	Query expansion, modification	NLPAUG [52] is a library that integrates data augmentation techniques in the NLP field. It provides character-level, word-level and sentence-level reformulations to a certain sentence based on DL models or random methods.	Any	-	2021
SEQUER	Query expansion, modification	SEQUER [12] is a Transformer trained for query reformulation based on the query logs from Stack Overflow. We access the tool via the demo link [13] provided by the authors.	Any	ICSE	2021
NLP2API	Query expansion	NLP2API [76] is a query expansion approach that adds predicted API class into the original query. We re-implement it based on the replication package [77] released by the authors.	Java	ICSME	2018
Query-Based API Recommendation					
RACK	Official documentation, Stack Overflow	RACK [78] is a class-level recommendation approach that leverages keyword and API co-occurrence information to make predictions. We re-implement the approach based on the released replication package [79].	Java	ICSE	2016
KG-APISumm	Official documentation, Wikipedia	KG-APISumm [49] is a class-level recommendation approach that builds an API knowledge graph to provide API class recommendations with detailed summaries.	Java	FSE	2019
Naive Baseline	Official documentation	We build an API-level naive baseline that only calculates the similarity score between API descriptions provided by official documentation and queries based on BERTOverflow [46] to recommend API classes. It indicates the basic performance of similarity-based methods.	Any	-	2021
DeepAPI	Official documentation	DeepAPI [31] is an API-level recommendation approach that uses an encoder-decoder model to convert words into APIs. We train this model using the training set released by its authors for 275,000 iterations and evaluate its performance on APIBENCH-Q.	Java	FSE	2016
Lucene	Official documentation	Lucene [23] is a popular java library providing powerful indexing and search features. Given a query, it can be used as a search engine to search related APIs in the official documentation.	Any	-	2021
BIKER	Official documentation, Stack Overflow	BIKER [37] is a method-level recommendation approach that leverages both Stack Overflow posts and API documentations to recommend APIs for a programming task. It fetches the top-k similar Stack Overflow posts to narrow down the scope of API candidates. We re-implement the approach based on the released replication package [38].	Java	ASE	2018

to impact the performance of code-based approaches with extremely long or short contexts. To derive the functions of appropriate lengths for study, we regard the shortest 5% of the functions as extremely short functions, the longest 5% of the functions as extremely long functions, and the middle 90% of the functions as functions with moderate lengths. We then further study the performance of code-based approaches on them in sec. 5.3. The thresholds for distinguishing short and long functions regarding the Line of code are illustrated in the last two columns of Table 1.

In order to study the performance of current approaches on different kinds of APIs, we convert source files of each repository into ASTs and extract all the function calls in them. We label a function call as a *standard* API or *popular* third-party API if it matches one of the APIs collected in

Sec. 3.1. We label a function call as a *user-defined* API if its implementation can be found in the current repository via import analysis. The average number of APIs per function, number of *standard* APIs and number of *user-defined* APIs are shown in columns 6 ~ 8 of Table 1, respectively.

3.3 Implementation Details

In this section, we describe the details of each approach involved in the benchmark and the metrics for evaluation.

Query reformulation techniques. We choose four popular query reformulation techniques, including Google Prediction Service [30], NLPAUG [52], SEQUER [12], and NLP2API [76]. The detailed description of each technique is illustrated in Table 3. Google prediction service is included

TABLE 4

The code based API recommendation baselines included in this empirical study. For tools we list the year of its most recent update time. The column name "PL" indicates the applicable programming language.

Approach	Representation	Description	PL	Venue	Year
Practical IDE					
PyCharm	Code tokens	PyCharm [42] is a widely-used Python IDE released by JetBrains. We use its basic code completion feature [40] to recommend APIs given an incomplete piece of code.	Python	-	2021
Visual Studio Code	Code tokens	Visual Studio Code [56] is a widely-used code editor released by Microsoft. We use its popular <i>Python</i> extension (ID: ms-python.python) that includes IntelliSense (Pylance) to recommend APIs given an incomplete piece of code.	Python	-	2021
Eclipse	Code tokens	Eclipse [24] is a widely-used Java IDE released by Eclipse Foundation. We use its default code completion feature to recommend APIs given an incomplete piece of code.	Java	-	2021
IntelliJ IDEA	Code tokens	IntelliJ IDEA [41] is a widely-used Java IDE released by JetBrains. We use the default code completion feature to recommend APIs given an incomplete piece of Java code.	Java	-	2021
Approach in Academia					
TravTrans	AST	TravTrans [44] encodes AST node sequences generated by pre-order traversal into Transformer and predicts the masked AST nodes that correspond to some certain tokens in the source code. We re-implement it based on the replication package [45] provided by the authors.	Python	ICSE	2021
PyART	Token flow Data flow	PyART [32] leverages a predictive model trained on data flow information, token similarity and token co-occurrence based on the Random Forest algorithm to recommend APIs. We re-implement it based on the replication package [33] provided by the authors.	Python	ICSE	2021
Deep3	AST, DSL	Deep3 [80] is a general approach for building probabilistic models of code based on learning decision trees. The decision trees are represented as program in a domain specific language (DSL) called TGEN which is agnostic to programming languages. We re-implement it based on the replication package [81] provided by the authors.	Python	ICML	2016
FOCUS	API Matrix	FOCUS [64] leverages the idea of traditional recommendation systems and uses a 3D matrix to represent the relationship between different projects, functions and APIs. It recommends APIs and usage patterns from the most similar projects. We re-implement it based on the released replication package [65].	Java	ICSE	2019
PAM	API sequence	PAM [25] builds a probabilistic model to mine the most "interesting" API patterns in certain projects. We re-implement it based on the released replication package [26].	Java	FSE	2016
PAM-MAX	API sequence	As PAM is originally a context-insensitive API recommendation approach, we build PAM-MAX by selecting the most correct API recommendation among all API patterns mined in the training set. It indicates the theoretical maximum performance a context-insensitive approach can reach when it can always find the most related projects to mine API patterns.	Java	FSE	2016

as one of the most effective approaches in practice, while SEQUER [13] is the state-of-the-art approach. NLPAUG [52] is considered since it is widely used for query reformulation in many NLP studies [43], [60], [74], [101]. We also include NLP2API [77] since it differs from major reformulation methods by first predicting the API class related to the query and then adding the predicted API class into the query.

Query-based API recommendation approaches. We choose five query-based API recommendation approaches published by recent top conferences, including KG-APISumm [49], BIKER [37], RACK [78], and DeepAPI [31], along with a popular search library Lucene [23]. The de-

tailed description of each baseline is shown in Table 3. We reproduce the five approaches based on the replication packages released by the authors. Besides, we build a naive baseline that recommends APIs by computing the similarities between queries and API descriptions based on BERTOverflow [46]. The native baseline serves as an indicator of the basic performance of similarity-based models. We also notice that different sources are adopted by the approaches for creating the knowledge base. For example, the naive baseline and DeepAPI only consider official documentation, while BIKER and RACK also involve the Q&A forum – Stack Overflow. We list the knowledge source of

each approach in Table 3. During implementation, we do not align the sources of the approaches, since the sources are claimed as contributions in the original papers. Instead, we design a separate RQ to study the impact of knowledge sources on the performance of API recommendations.

During studying the impact of query reformulation on the recommendation performance, we implement all the four query reformulation techniques for each of the six API recommendation baselines because all the baselines do not integrate query reformulation techniques in the original papers.

Code-based API recommendation approaches. We choose four IDEs and five approaches published on recent top conferences as our code-based API recommendation baselines. A detailed description of each baseline is shown in Table 4. For the IDEs and some of the approaches such as TravTrans [44] and Deep3 [80], they can predict any code tokens besides API tokens. In this paper, we focus on evaluating their performance in recommending APIs. Following prior research [27], [32], [44], [64], [82], we use the training set of APIBENCH-C to train each of the approaches in academia for a fair comparison.

PAM [27] is the only context-intensive approach, primarily designed for intra-project API pattern mining. In this paper, we also extend the approach to cross-project recommendation by selecting the best API from projects in the training set for each test case. The extended version of PAM is named as PAM-MAX, which indicates the theoretical maximum performance the context-insensitive approach can achieve.

Evaluation metrics. Since both query-based and code-based API recommendation baselines output a ranked list of candidate APIs, we adopt the commonly-used metrics in recommendation tasks for evaluation. Table 5 shows the details of each metric.

The Mean Reciprocal Rank (MRR), Mean Average Precision (MAP), and Normalized Discounted Cumulative Gain (NDCG) metrics are widely adopted by previous API recommendation studies [37]. In this study, we also involve a new metric Success Rate. The Success Rate@k is defined to evaluate the ability of an approach in recommending correct APIs based on the top-k returned results regardless of the orders. To determine the relevance score in NDCG calculation, we use a relevance score of 1 if an approach hits the correct API class, and a relevance score of 2 if the correct API method is hit. Therefore, we can align the performance of class-level and method-level approaches.

4 EMPIRICAL RESULTS OF QUERY REFORMULATION AND QUERY BASED API RECOMMENDATION

In this section, we study the RQ 1-3 discussed in Sec. 1 and provide the potential findings concluded from the empirical experiments. Since currently no query-based API recommendation approach is specially designed for Python APIs, we focus on studying query-based API recommendation approaches for Java.

TABLE 5
The evaluation metrics used in this empirical study.

Metric	Calculation	Description
Success Rate@k	$\frac{\sum_{i=1}^N \text{ContainCorrect}_k(q_i)}{N}$	ContainCorrect _k (q) returns 1 if the top k results returned for query q contain the correct API, otherwise it returns 0.
MAP@k	$\text{MAP@k} = \frac{\sum_{i=1}^N \text{AveP}_k(q_i)}{N}$ $\text{AveP@k} = \frac{\sum_{i=1}^N P_k(i) \times \text{rel}(i)}{m}$	rel(i) returns 1 if the i _{th} result is the correct API, otherwise it returns 0.
MRR	$\frac{\sum_{i=1}^N 1/\text{firstpos}(q_i)}{N}$	firstpos(q) returns the position of the first correct API in the results, if it cannot find a correct API in results, it returns +∞.
NDCG@k	$\text{NDCG@k} = \frac{\sum_{i=1}^N \frac{\text{DCG@k}(q_i)}{\text{IDCG@k}(q_i)}}{N}$ $\text{DCG@k} = \sum_{t=1}^k \frac{\text{rel}_t(q_i)}{\log_2(t+1)}$	rel _t (q) returns 2 if t _{th} result exactly matches one correct API, and it returns 1 if t _{th} result matches the API class but fails to match the API, otherwise it returns 0. IDCG@k is the best DCG@k by re-arranging the order of current results.

4.1 Effectiveness of Query-Based API Recommendation Approaches (RQ1-1)

To answer RQ1, we evaluate the six query-based API recommendation baselines listed in Table 3 by using the original queries in our benchmark APIBENCH-Q. The evaluation results are illustrated in Table 6.

Class-level v.s. Method-level. Regarding the **class-level** recommendation, as shown in Table 6, we can find that BIKER achieves the highest Success Rate, e.g., 0.67 for Success Rate@10, indicating that BIKER is more effective in finding the correct API class in the top-10 returned results for 60%~70% of cases. Unsurprisingly, the naive baseline shows the worst performance for all the metrics. Even so, the naive baseline can successfully predict the correct API class for around 20% of cases. However, with respect to the **method-level** recommendation, all the approaches show obvious declines. For example, the Success Rate@10 of BIKER is only 0.37, decreasing by 44.8% compared to the class-level recommendation. The Success Rates@10 of DeepAPI and Lucene are only around 0.10, which is far from the requirement of practical development. On average, the approaches fail to give the exact methods for 57.8% APIs that they give the correct classes in top-10 returned recommendations. Thus, recommending method-level APIs still remains a great challenge.

TABLE 6

The basic performance of query-based API recommendation baselines without applying any query reformulation techniques at different metrics (Top-1,3,5,10). Note that we define NDCG as a uniform metric to evaluate class level and method level together, so the NDCG scores listed in two levels have the same values. The red numbers indicate the best performance achieved in top-10 results.

Baseline	Level	Success Rate@k				MAP@k				MRR	NDCG@k			
		Top-1	Top-3	Top-5	Top-10	Top-1	Top-3	Top-5	Top-10		Top-1	Top-3	Top-5	Top-10
RACK	Class	0.17	0.30	0.35	0.41	0.17	0.23	0.24	0.24	0.25	0.17	0.24	0.26	0.28
KG-APISumm	Class	0.19	0.33	0.40	0.50	0.19	0.25	0.26	0.27	0.28	0.19	0.24	0.27	0.31
Naive Baseline	Class	0.07	0.13	0.16	0.21	0.07	0.10	0.10	0.10	0.11	0.07	0.09	0.10	0.13
	Method	0.02	0.03	0.04	0.05	0.01	0.02	0.03	0.03	0.03	0.07	0.09	0.10	0.13
DeepAPI	Class	0.19	0.27	0.29	0.30	0.19	0.22	0.23	0.23	0.23	0.17	0.22	0.23	0.24
	Method	0.05	0.09	0.10	0.11	0.05	0.07	0.07	0.07	0.07	0.17	0.22	0.23	0.24
Lucene	Class	0.15	0.21	0.24	0.29	0.15	0.17	0.18	0.17	0.19	0.12	0.15	0.16	0.20
	Method	0.04	0.08	0.10	0.14	0.04	0.06	0.06	0.06	0.07	0.12	0.15	0.16	0.20
BIKER	Class	0.33	0.51	0.59	0.67	0.33	0.41	0.41	0.39	0.44	0.27	0.32	0.35	0.42
	Method	0.12	0.23	0.29	0.37	0.12	0.16	0.18	0.18	0.19	0.27	0.32	0.35	0.42

Finding 1: Existing approaches fail to predict 57.8% method-level APIs that could be successfully predicted at the class level. The performance achieved by the approaches is far from the requirement of practical usage. Accurately recommending the method-level APIs still remains a great challenge.

Retrieval-based methods v.s. Learning-based methods.

By comparing learning-based methods, such as DeepAPI and naive baseline, with the other retrieval-based methods, we can observe that learning-based methods achieve relatively lower performance regarding the Success Rate@10 metric. For example, on average, retrieval-based methods can accurately predict 46.8% class-level and 25.5% method-level APIs among all the cases in the top-10 returned results, respectively, while learning-based methods can only successfully recommend 25.5% class-level and 8% method-level APIs. A possible reason may be the insufficient training data for the learning-based methods in this task domain. Since there are more than 30,000 APIs from the official documentation, learning-based methods require a large number of query-API pairs for training. However, even the largest Q&A forum, Stack Overflow, contains only about 150,000 posts after our pre-processing, which is not enough for model training.

Finding 2: Learning-based methods do not necessarily outperform retrieval-based methods in recommending more correct APIs. The insufficient query-API pairs for training limit the performance of learning-based methods.

Performance in API ranking. From Table 6, we find that there exist obvious gaps between the scores of Success Rate@k and the metrics for evaluating API ranking, such as MAP@k and NDCG@k. For example, RACK achieves Success Rate@10 score at 0.41, but its MAP@10 score is only 0.24. This indicates that although the approaches are able to find the correct APIs, they cannot well rank them ahead in the returned results. The low MRR scores, e.g., 0.11 ~ 0.44 for class-level API recommendation and 0.03 ~ 0.19 for method-level API recommendation, and NDCG scores also show the poor ranking performance of the approaches.

The results manifest that API ranking is still challenging for current approaches.

Finding 3: Current approaches cannot well rank the correct APIs, considering the huge gap between the scores of Success Rate and the other ranking metrics.

To sum up, accurately recommending method-level APIs and ranking candidate APIs still remain great challenges. Besides, the insufficient data for training hinder the performance of current learning-based approaches.

4.2 Effectiveness of Query Reformulation Techniques (RQ2)

Original queries can be short in length or contain vague terms. Query reformulation aims at changing original queries for facilitating downstream tasks. In this RQ, we explore the impact of query reformulation on the performance of query-based API recommendation.

We implement the four query reformulation techniques, as listed in Table 3, for the original queries. We name the queries reformulated by query expansion techniques and query modification techniques as expanded queries and modified queries, respectively. For each original query, we conduct the reformulation 10 times, producing 10 expanded or modified queries, with the statistics shown in Table 2. Note that NLPAUG [52] is a comprehensive data augmentation library for general NLP tasks. We choose the popular word-level insertion and substitution methods designed for manipulating single sentences based on five models, including BERTOverflow [46], Google News Word2vec [29], Stack Overflow Word2vec [93], WordNet [57], and Random model, in the library to generate expanded and modified queries.

The queries output by the query reformulation techniques are not ranked in order, and may impact the downstream API recommendation performance variously. To explore the maximum potential effect brought by query reformulation techniques, we evaluate the API recommendation approaches on each reformulated query and choose the best result for analysis. For example, SEQUER [12] generates




Fig. 4. The maximum improvement of Success Rate@10 by all query reformulation techniques on **class-level** query-based API recommendation baselines. We do not evaluate the performance of RACK and KG-APISumm in NLP2API reformulated queries as they are only class-level recommendation approaches while NLP2API directly give the predicted API classes. Note that we include Google Prediction Service and SEQUER as expansion techniques here because they expand the queries in most cases.

10 queries for the original query “Get min value between two double type values.” We choose the one achieving the best performance for analysis, i.e., using the reformulated query “Get min value between two double type values in java” for the recommendation. We study the impact of query reformulation on API recommendation from the following two aspects:

- 1) whether query reformulation techniques can help predict more correct APIs;
- 2) whether query reformulation can improve the API ranking performance.

4.2.1 Influence on predicting more correct APIs

With query reformulation v.s. Without query reformulation. The Success Rate metric reflects the proportion of the APIs an approach can correctly predict. The results of implementing the reformulation techniques on API recommendation approaches are illustrated in Figure 4 (class-level) and Figure 5 (method-level). From the figures, we observe that query reformulation can increase the performance of API recommendation in most cases. Only for a few cases, the performance drops, which can be attributed to the inefficiency of some query reformulation techniques. For example, NLPAUG (WordNet) and NLPAUG (Random)

tend to poorly modify the original queries for recommendation, as shown in Figure 4 (b) and Figure 5 (b). Overall, on average the process improves the class-level and method-level recommendation by 0.11 and 0.08, which is a corresponding boost of 27.7% and 49.2% compared with the basic performance on original queries.

Finding 4: Query reformulation techniques, including query expansion and query modification, are quite effective in helping query-based API recommendation approaches give the correct API by adding an average boost of 27.7% and 49.2% on class-level and method-level recommendations, respectively.

Query expansion v.s. Query modification. By comparing the class-level and method-level recommendation results of query expansion and query modification in Figure 4 and Figure 5, respectively, we observe that all the query expansion techniques improve the API recommendation performance, but not all the query modification techniques benefit the recommendation. For example, NLPAUG (WordNet) and NLPAUG (Random) generally decrease the performance of current approaches both in class-level and method-level recommendations. This indicates that query expansion techniques bring more stable improvement than query modification techniques. Furthermore, on average,

query expansion techniques improve the performance by 0.13 and 0.10 on class-level and method-level recommendation, which is much higher than the improvement of 0.09 and 0.06 achieved by query modification techniques. This also suggests that query expansion techniques are more effective than query modification techniques.

Finding 5: Query expansion is more stable and effective to help current query-based API recommendation approaches give correct APIs than query modification.

Comparing different query expansion techniques. As shown in Figure 4 (a) and Figure 5 (b), NLP2API and NLPAUG (BERT) present the largest improvement on the performance of query-based API approaches at both class level and method level. For analyzing the improvement, we use two examples to illustrate the query expansion results of NLP2API and NLPAUG (BERT), respectively. In both examples, the most effective approach BIKER fails to predict the API based on the original queries but succeeds given the reformulated queries.

Example 1: Query Expansion

TECHNIQUE	NLP2API
ORIGINAL QUERY	Returns a new Document instance
PROCESSED QUERY	DocumentBuilderFactory Returns a new Document instance

In the first example, NLP2API expands the query by adding a predicted API class *DocumentBuilderFactory* that is related to the original query. With such an explicit hint, the recommendation approach can narrow down the search scope and pinpoint the requested API method.

Example 2: Query Expansion

TECHNIQUE	NLPAUG (BERT)
ORIGINAL QUERY	Java reverse string
PROCESSED QUERY	java reverse character string

In the second example, the query is looking for the API *java.lang.StringBuilder.reverse()*, whose description in official documentation is “Causes this character sequence to be replaced by the reverse of the sequence”. NLPAUG (BERT) adds a relevant word *character* to enrich the semantics of the original query.

Example 3: Query Expansion

TECHNIQUE	NLPAUG (W2V-SO)
ORIGINAL QUERY	Convert from Radians to Degrees in Java
PROCESSED QUERY	Convert from AV Radians to Degrees in Long Java

Comparing NLPAUG (W2V) with NLPAUG (BERT) and NLP2API in Figure 4 and Figure 5, we find that the NLPAUG (W2V) is much less effective. To obtain a possible reason for such a difference, we give the third example below. As shown in this example, we find that NLPAUG (W2V) adds two irrelevant words into the original query, which negatively impacts the prediction results of BIKER. This also indicates that contextual embeddings such as BERT are more effective than traditional word embeddings.

Finding 6: Add predicted API class names or relevant words into original queries are effective query expansion methods for improving the performance of API recommendation.

Comparing different query modification techniques.

Among all query modification techniques, NLPAUG (BERT) presents the biggest improvement on all the baselines at both class level and method level. Example 4 illustrates how NLPAUG (BERT) modifies words in the original query. In the example, the original query asks about ways to calculate the time difference between two dates and the correct API is *java.time.Period.between()*. The description of the API in its official documentation is “obtains a period consisting of the number of years, months, and days between two dates”. However, the word “difference” used in the original query does not clearly describe the functional request. NLPAUG (BERT) modifies the word into “months” which exactly appears in the official description. Based on the modifications, the correct API is recommended.

Example 4: Query Modification

TECHNIQUE	NLPAUG (BERT)
ORIGINAL QUERY	How do I calculate difference between two dates
PROCESSED QUERY	how do they I calculate months difference between two dates

From the second and fourth examples above, we find that BERT-based models show great performance on both query expansion and query modification to help improve the performance of current query-based API recommendation approaches. This indicates that even though the current data source limits the performance of these models to directly predict the correct APIs, they can be used to improve the query quality as query reformulation techniques.

Finding 7: BERT-based data augmentation shows superior performance in query modification compared with other query modification techniques.

4.2.2 Influence on the performance of API ranking

In this section, we analyze the impact of query reformulation techniques on the performance of API ranking. Since the ideal case is that the correct APIs rank first in the returned results, we use the metric NDCG@1 which considers both class-level and method-level recommendation performance. We compute the changes of NDCG@1 scores for the query-based API recommendation approaches before and after query reformulation. Besides, to focus our analysis on the performance of API ranking instead of the overall recommendation accuracy, the computation is performed only on the cases that are correctly predicted with and without query reformulation.

The results are illustrated in Figure 6. As can be seen in Figure 6 (a), most query expansion techniques also improve the ranking results of the query-based recommendation approaches. Among all the query expansion techniques, SEQUER, NLPAUG (BERT), RACK and NLP2API can improve the ordering performance relatively better than the others. The biggest improvement of 0.14 (32% boost) is achieved by NLP2API on the Lucene approach. We also find that

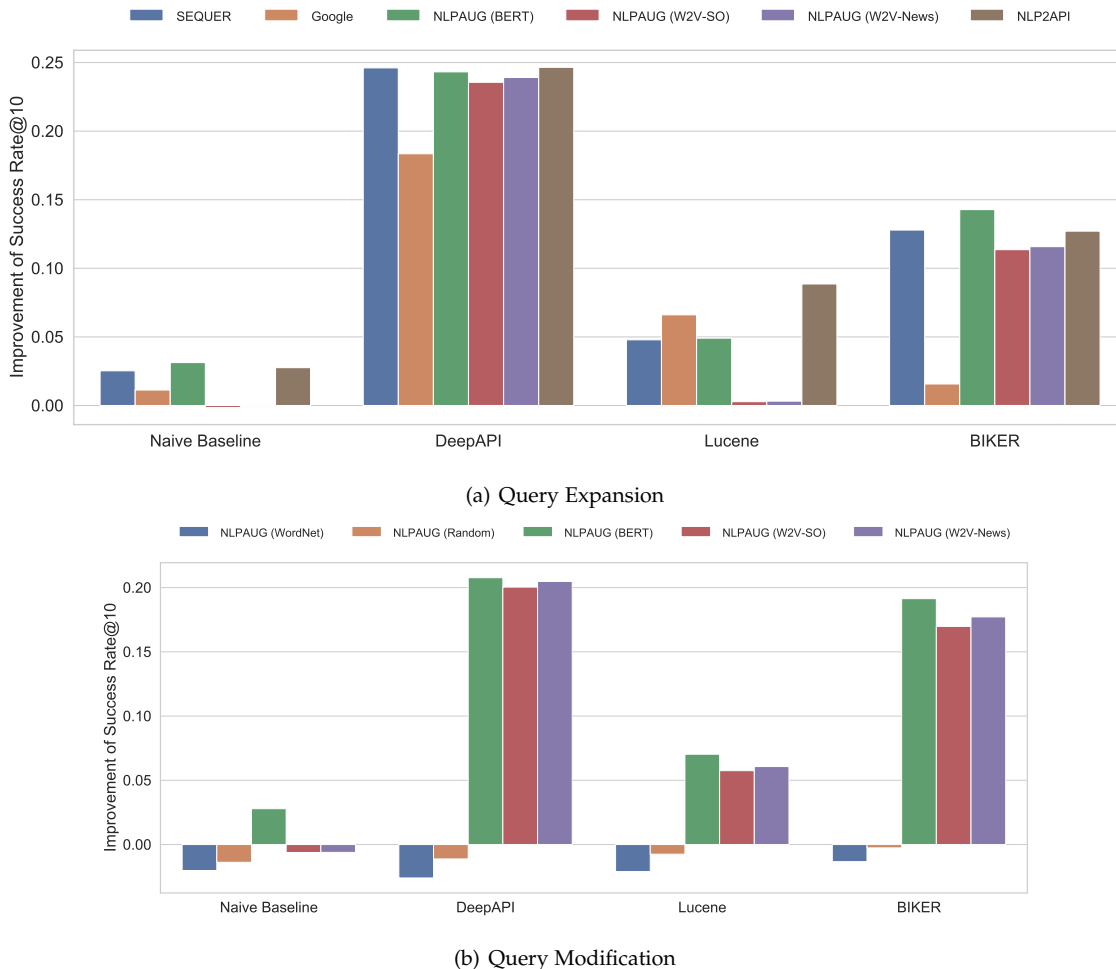


Fig. 5. The maximum improvement of Success Rate@10 by all query reformulation techniques on **method-level** query-based API recommendation baselines.

on average query expansion also improves MRR by 0.09 (36% boost) and 0.08 (89% boost) on class-level and method-level recommendation, respectively, which indicates that the correct APIs are ranked much higher based on reformulated queries.

According to Figure 6 (b), compared with query expansion techniques, query modification techniques are much less effective in improving the API ranking performance. For example, the average improvement of NDCG@1 brought by query modification is 0.01 (4% boost), which is 0.06 (14% boost) for query expansion techniques. Comparing different data augmentation methods, we also find that WordNet and random methods tend to negatively impact the ranking results, leading to 24% and 14% drop in terms of NDCG@1, respectively. The results indicate that inappropriate query modification will reduce the ranking performance of the query-based recommendation approaches.

Finding 8: Expanding queries or modifying queries with appropriate data augmentation methods can improve the ranking performance of the query-based API recommendation techniques.

To sum up, query reformulation, especially query expansion, can not only help current approaches recommend more correct APIs, but also improve the ranking performance. However, the reformulation step is generally ignored by

current studies. Future work is suggested to involve such a step for more accurate API recommendation.

4.2.3 A special Query Modification Method: Word Deletion

In previous subsections, we compare and evaluate different query expansion and modification techniques. They aim at enriching the original queries by adding, replacing or modifying some words without deleting words. In this section, we focus on studying the impact of word deletion, a special query modification method, on the performance of query-based API recommendation approaches. Different from the previous query reformulation techniques which rely on external data sources, the word deletion method we studied does not leverage any extra knowledge. Our goal is to explore whether original queries contain meaningless or noisy words. Specifically, we randomly delete some words from the original query every time and produce ten different modified queries for one original query.

The maximum and average Success Rate@10 scores based on the modified queries are illustrated in Figure 7. As can be seen, the average performance of the query-based API recommendation approaches, denoted as the orange bar, decreases by 0.05 (13% drop) at class level and 0.03 (18% drop) at method level. The results are not surprising, and indicate that most words in the original

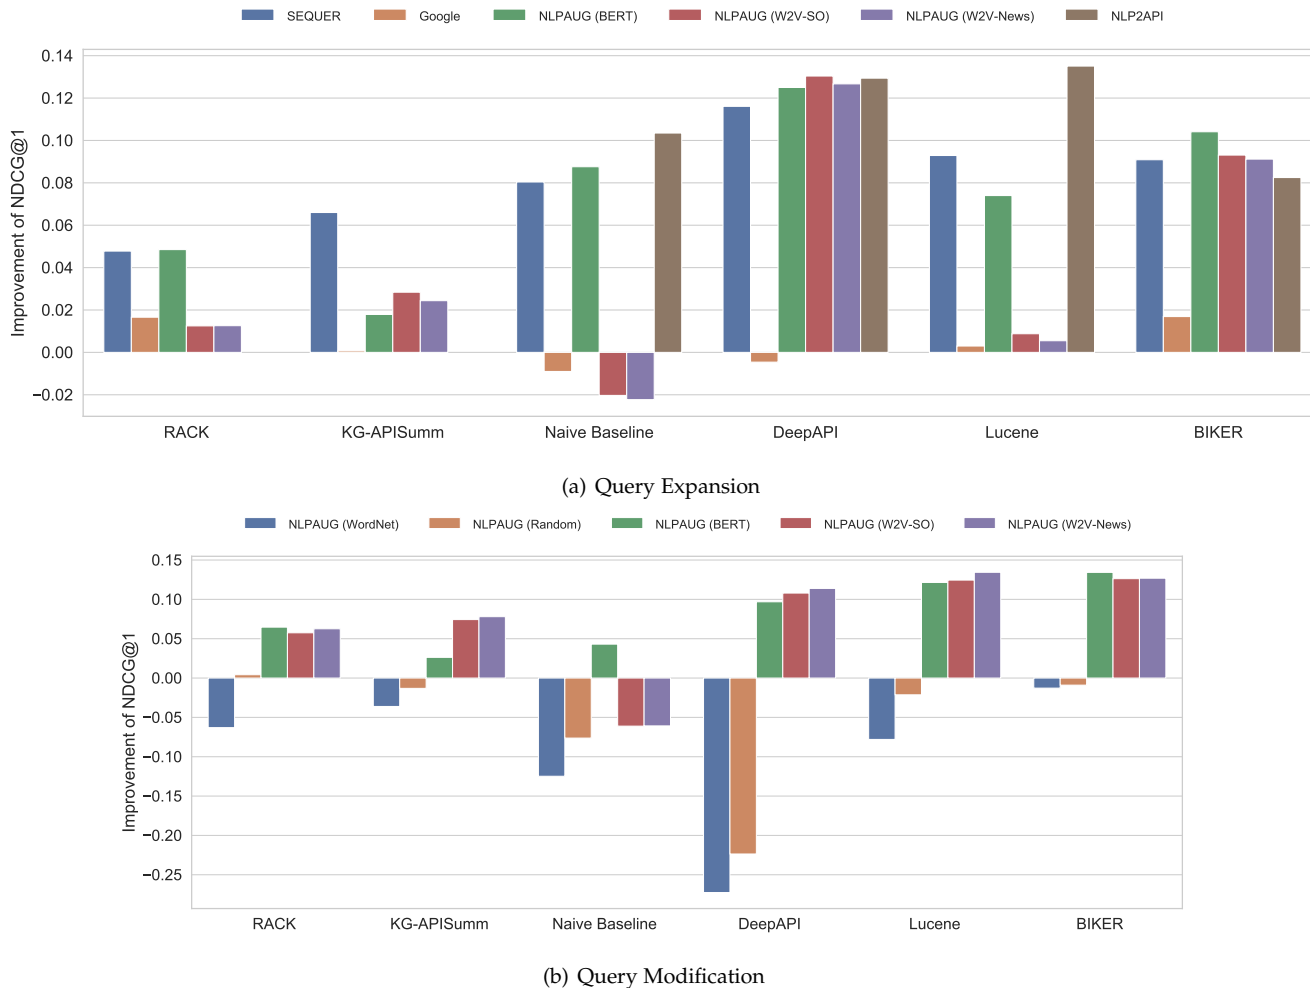


Fig. 6. The maximum improvement of NDCG@1 by all query reformulation techniques on query-based API recommendation baselines under original successful cases.

queries are helpful for the recommendation. However, the maximum scores, denoted as the green bar, all show that word deletion improves the recommendation performance with an average boost of 38% and 64% for class level and method level, respectively. The improvement demonstrates that the original queries contain noisy words that can bias the recommendation results, although most of the words are useful for recommendation.

To understand what kinds of words are noisy for the accurate recommendation, we manually check 545 out of 6,563 queries for which the recommendation approaches perform better after word deletion. We summarize the cases as below:

1) 349 (64%) queries contain unnecessary or meaningless words.

Example 5: Word Deletion

TECHNIQUE	Random Deletion
ORIGINAL QUERY	Standard way to iterate over a StringBuilder in java
PROCESSED QUERY	Standard way to iterate over a StringBuilder in java

In Example 5, the phrases “Standard way to” and “in java”

are not beneficial for pinpointing the correct API. Stop word removal also has a limited effect on eliminating these words.

2) 156 (29%) queries contain too detailed words for explanation.

Example 6: Word Deletion

TECHNIQUE	Random Deletion
ORIGINAL QUERY	converts a color into a string like 255,0,0
PROCESSED QUERY	converts a color into a string like 255,0,0

In Example 6, the phrase “like 255,0,0” is used to explain the “string”. However, such phrases never appear in the official documentation and the specific number adversely impacts the recommendation results.

3) 34 (6%) queries contain extreme long descriptions.

Example 7: Word Deletion

TECHNIQUE	Random Deletion
ORIGINAL QUERY	how to add progress bar to zip utility while zipping or extracting in java
PROCESSED QUERY	how to add progress bar to zip utility while zipping or extracting in java

In Example 7, the words after *while* actually describe nothing about the task. The long descriptions can decrease

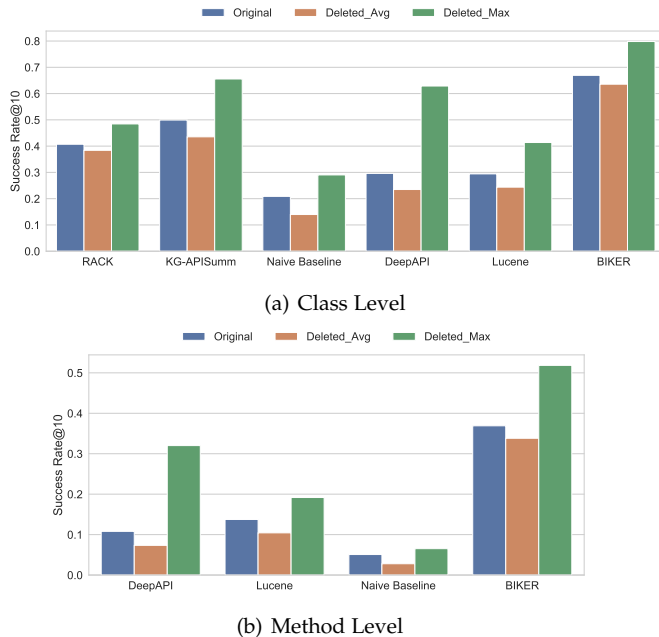


Fig. 7. The maximum and average Success Rate@10 on all baselines when randomly deleting some words in original queries.

the weight of useful words in the queries thus confusing API recommendation approaches.

Finding 9: *Original queries raised by users usually contain noisy words which can bias the recommendation results, and query reformulation techniques should consider involving noisy-word deletion for a more accurate recommendation.*

4.3 Data Sources (RQ3)

In RQ1-1, we highlight that insufficient data greatly limits the performance of current learning-based methods. In this section, we conduct a deep analysis on the influence of different data sources on the recommendation results. From Table 3, we can observe that current approaches generally leverage three different data sources: official documentation, Q&A forums, and tutorial websites. For analysis, we choose two methods, Lucene and naive baseline, which are flexible to incorporate different data sources. Specifically, we evaluate the methods on the part of queries from the tutorial websites collected in APIBENCH-Q, and the method training is conducted based on the following knowledge base:

- 1) only official documentation,
- 2) only Stack Overflow posts, and
- 3) both official documentation and Stack Overflow posts.

The experiment results are shown in Figure 8. As can be seen, training on Stack Overflow posts achieves much better performance than on official documentation at both class and method levels. For example, Lucene achieves a 29% boost in class-level and an 169% boost in method-level recommendation when searching based on Stack Overflow than on official documentation; and the naive baseline even achieves a 71% boost in class-level and a 602% boost in method-level recommendation. The advantage of leveraging Stack Overflow posts may be attributed that the discussion on Stack Overflow is more natural and similar to




Fig. 8. The Success Rate@10 of Lucene and Naive Baseline under three data source settings.

user queries, compared with the descriptions in the official documentation. Besides, the extended usage of some APIs is rarely mentioned in official documentation but is widely discussed in Stack Overflow. An example is used to illustrate the influence of different data sources.

Example 8: Data Source

BASELINE	Lucene
ORIGINAL QUERY	Compute the md5 hash of a File
CORRECT API	<code>java.security.MessageDigest.digest()</code> , <code>java.security.MessageDigest.getInstance()</code>
API DESCRIPTION	Completes the hash computation by performing final operations such as padding
SIMILAR SO POST	How can I generate an MD5 hash in Java?

In Example 8, the query asks about the API for generating an MD5 hash of a file. However, there is no standard API specially designed to generate the MD5 hash, so Lucene focuses on two words “hash” and “file” for recommendation. But the official description of ground truth API `java.security.MessageDigest.digest()` does not contain the word “file” since it is a general API that not merely handles files. Under this circumstance, Lucene recommends a more relevant but wrong API `java.nio.file.attribute.FileTime.hashCode()`. When involving Stack overflow posts, as there already exists discussion on how to generate the MD5 hash, Lucene can easily pinpoint and recommend the correct API in the posts.

The advantage of leveraging Stack Overflow for recommendation is also demonstrated by the BIKER approach [37], which is the most effective approach in Section 4.1. Our finding is consistent with the claim in the work [37] that Stack Overflow posts can mitigate the semantics gap between user queries and official descriptions.

Finding 10: *Apart from official documentation, using other data sources such as Stack Overflow can significantly improve the performance of query-based API recommendation approaches.*

5 EMPIRICAL RESULTS OF CODE-BASED API RECOMMENDATION

In this section, we study the RQ1 and RQ 4 ~ 6 discussed in Sec 1. To study RQ1, RQ4 and RQ5, we evaluate the performance of all the code-based API recommendation approaches on the “General” domain of our benchmark APIBENCH-C, as shown in Table 1, since the “General” domain includes code with different topics and can reflect

TABLE 7

The performance of code-based API recommendation baselines at different metrics (Top-1,3,5,10). All baselines are trained and tested on the full dataset from “General” domain of APIBENCH-C except for PyART. Since PyART takes months to train and test on our full dataset, we randomly sampled 20% of original training and testing testset to evaluate it. The “PL” column indicates the programming language the baselines target. The **red** number indicates the best performance.

PL	Baseline	Success Rate@k				MAP@k				MRR	NDCG@k			
		Top-1	Top-3	Top-5	Top-10	Top-1	Top-3	Top-5	Top-10		Top-1	Top-3	Top-5	Top-10
Python	TravTrans	0.45	0.57	0.59	0.62	0.45	0.50	0.51	0.51	0.51	0.45	0.52	0.53	0.54
	Deep3	0.21	0.34	0.37	0.43	0.20	0.27	0.28	0.28	0.28	0.21	0.29	0.30	0.32
	PyART	0.29	0.38	0.46	0.60	0.29	0.33	0.35	0.37	0.37	0.29	0.34	0.37	0.41
Java	FOCUS	0.01	0.03	0.04	0.06	0.01	0.02	0.02	0.03	0.03	0.01	0.02	0.03	0.04
	PAM	0.01	0.02	0.03	0.05	0.01	0.02	0.02	0.02	0.02	0.01	0.02	0.02	0.03
	PAM-MAX	0.22	0.32	0.36	0.45	0.22	0.26	0.27	0.28	0.28	0.22	0.27	0.29	0.32

TABLE 8

The performance of code-based API recommendation baselines along with 4 widely used IDEs tested on 500 cases sampled from the testset of all domains in APIBENCH-C. The “PL” column indicates the programming language the baselines target. The **red** number indicates the best performance. The rows with **gray** background indicates the performance of IDEs.

PL	Baseline	Success Rate@k				MAP@k				MRR	NDCG@k			
		Top-1	Top-3	Top-5	Top-10	Top-1	Top-3	Top-5	Top-10		Top-1	Top-3	Top-5	Top-10
Python	TravTrans	0.38	0.46	0.48	0.50	0.38	0.42	0.42	0.43	0.43	0.38	0.43	0.44	0.44
	Deep3	0.19	0.26	0.31	0.38	0.19	0.22	0.23	0.24	0.24	0.19	0.23	0.25	0.28
	PyCharm	0.31	0.42	0.47	0.49	0.31	0.36	0.37	0.37	0.37	0.31	0.38	0.40	0.40
	VSCode	0.05	0.15	0.21	0.35	0.05	0.09	0.11	0.13	0.13	0.05	0.11	0.14	0.18
Java	FOCUS	0.02	0.04	0.05	0.07	0.02	0.03	0.03	0.04	0.04	0.02	0.03	0.04	0.04
	PAM	0.01	0.02	0.05	0.07	0.01	0.02	0.02	0.03	0.03	0.01	0.02	0.03	0.04
	PAM-MAX	0.27	0.38	0.43	0.56	0.27	0.31	0.33	0.34	0.34	0.27	0.33	0.35	0.39
	Eclipse	0.28	0.42	0.49	0.60	0.28	0.34	0.35	0.37	0.37	0.28	0.36	0.39	0.42
	IntelliJ IDEA	0.42	0.58	0.65	0.67	0.42	0.49	0.51	0.51	0.51	0.42	0.51	0.54	0.55

the overall performance of baselines. For studying the ability of cross-domain adaptation in RQ6, we evaluate the performance of the approaches on all the five domains of our APIBENCH-C.

5.1 Effectiveness of Existing Approaches (RQ1-2)

According to Table 4, three approaches for Python and three approaches for Java are evaluated on the “General” domain of APIBENCH-C. The results are depicted in Table 7. We can observe that the learning-based method TravTrans obtains the best performance on the Python dataset, achieving 0.62 and 0.54 for Success Rate@10 and NDCG@10, respectively. The results mean that TravTrans can successfully recommend 62% of APIs in our benchmark and well predict the API rankings. However, the traditional statistical method Deep3 only achieves 0.43 and 0.32 for Success Rate@10 and NDCG@10, respectively, while the pattern-based method FOCUS and PAM achieve less than 0.10 for both Success Rate@10 and NDCG@10. This suggests that learning-based methods obtain superior performance in code-based API recommendation, which is quite different from query-based API recommendation. The possible reason is that lots of well-organized public code repositories provide sufficient data for training code-based API recommendation models.

We also find that FOCUS and PAM show low recommendation accuracy, with all the metric values lower than 0.1. The low performance is attributed to the context representation of the approaches. PAM is a context-insensitive approach, which only mines the top-N APIs that are most likely to be used in the training set and directly recommends them for each file in the test set; while FOCUS takes one step further by extracting the APIs in the test set and building a matrix to match the APIs in the training set. Such coarse-grained context representation or context-insensitive representation do not well capture the relations between APIs. PAM-MAX shows the theoretical best performance context-insensitive methods can achieve. However, the performance of PAM-MAX is still lower than TravTrans and PyART which consider fine-grained code features such as code tokens and data flows. The results indicate the effectiveness of fine-grained approaches for code-based API recommendation.

Besides the recent code-based API recommendation approaches, we also compare the widely-used IDEs. Since it is hard to automatically evaluate IDEs’ recommendation performance, we sampled 500 APIs from the original large test set of APIBENCH-C based on the distribution shown in Table 1. We then conduct a manual evaluation by imitating




Fig. 9. The Success Rate@10 of baselines on three categories of APIs at the “General” domain of APIBENCH-C.

the behaviors of developers on the 500 sampled APIs. We show the results on the sampled test set in Table 7. As can be seen, for Python, Pycharm achieves the Success rate@10 at 0.49 and NDCG@10 at 0.40, which is truly competitive to the performance of TravTrans, with Success Rate@10 and NDCG@10 at 0.50 and 0.44, respectively. For Java, IDEs also show competitive performance compared with the baseline approaches. The results demonstrate that the widely-used IDEs are generally effective in API recommendation and far from relying on alphabet orders for recommendation.

Finding 11: *DL models such as TravTrans show superior performance on code-based API recommendation by achieving a Success Rate@10 of 0.62, while widely-used IDEs also obtain satisfying performance by achieving a Success Rate@10 of 0.5 ~ 0.6.*

5.2 Capability to Recommend Different Kinds of APIs (RQ4)

Exploring which kinds of APIs tend to be wrongly predicted is essential for understanding the bottleneck of current approaches and for providing clues for further improvement. In Section 3, we have classified all APIs into standard APIs, popular third-party APIs and user-defined APIs. In this RQ, we study the performance of current baselines for different kinds of APIs. Specifically, we evaluate TravTrans, Deep3, FOCUS, PAM and PAM-MAX on the full test set of the “General” domain, with results shown in Figure 9.

As can be seen in Figure 9, all the approaches achieve a very high Success Rate@10 on standard APIs. For example, TravTrans even successfully recommends more than 90% of standard APIs in the test set. The approaches also present relatively good performance for the popular third-party libraries, e.g., TravTrans achieves a Success Rate@10 of more than 0.8. As standard APIs and popular APIs from third-party libraries are widely used in real-world projects, data-driven methods can achieve superior performance. However, the approaches are hard to correctly recommend the user-defined APIs and fail to predict 35.3% ~ 91.3% more of user-defined APIs comparing to the prediction of standard APIs.

Finding 12: *Although current approaches achieve good performance on recommending standard and popular third-party libraries, they face the challenges of correctly predicting the user-defined APIs.*




Fig. 10. The Success Rate@10 of baselines on extremely short, normal and extremely long contexts at the “General” domain of APIBENCH-C.

5.3 Capability to Handle Different Contexts (RQ5)

As context representation is an important part of the current code-based API recommendation shown in Figure 2, it is worthwhile to study the impact of different contexts on the performance of current approaches. In this RQ, we explore the impact of the following two different types of context.

- lengths of functions, which evaluates the capability of current approaches to handle different lengths of contexts;
- different recommendation points, since different recommendation points affect how much context an approach can be aware of before recommendation.

Capability to handle different lengths of functions. In Section 3 and Table 1 we classify all functions of APIBENCH-C into extremely short functions, functions of moderate lengths, or extremely long functions by sampling the first 5%, middle 90% and last 5% according to the distribution of function lengths. As code-based API recommendation is often based on the context in a function, the length of function can represent the length of context that an approach needs to handle. We study the performance of current baselines on functions of different lengths and show the results of TravTrans, Deep3, FOCUS, PAM, and PAM-MAX in Figure 10.

From Figure 10, we find that most baselines share similar performance distributions on functions of different lengths. They present the best performance on functions with moderate lengths and suffer from performance drops on extremely long or short functions. To be more specific, the performance drops by 7.1% for extremely long functions and 10.6% for extremely short functions on average. The results indicate that context length can affect the performance of current approaches. Besides, the approaches are more difficult to recommend correct APIs for the functions of extremely short lengths than those of extremely long lengths.

Finding 13: *Context length can impact the performance of current approaches in API recommendation. The approaches perform poorly for the functions with extremely short or long lengths, and accurate recommendation for the extremely short functions is more challenging.*

Capability to handle different recommendation points. Similar to the previous work [64], we first define three locations of recommendation points. Suppose that the LOC of a function is n and the total number of APIs used in the function is m . we define a recommendation point that is on the a_{th} line of the function and is the b_{th} API in the function locates on




Fig. 11. The Success Rate@10 of baselines on three categories of recommendation points at the *general* domain of APIBENCH-C.

- 1) the front of function if $a/n < 1/4$ and $b/m < 1/4$, or
- 2) the middle of function if $1/4 < a/n < 3/4$ and $1/4 < b/m < 3/4$, or
- 3) the back of function if $a/n > 3/4$ and $b/m > 3/4$.

For all the APIs in the test set of the “General” domain, we replace them with placeholders of the above three types of recommendation points for evaluation. We also remove APIs in extremely long or short functions (according to the thresholds shown in Table 1) to alleviate the influence of function lengths. We show the results of TravTrans, Deep3, FOCUS, PAM and PAM-MAX in Figure 11.

From Figure 11, we observe that most approaches obtain the best performance on the back recommendation points and acquire the worst performance on the front recommendation points. To be more specific, on average five approaches achieve 0.316, 0.347 and 0.348 for Success Rate@10 at the front, middle and back recommendation points, respectively. The results demonstrate that the current approaches tend to be less effective on the front recommendation points, which is reasonable since the context information is much limited before the front recommendation points.

Finding 14: *The location of recommendation points can affect the performance of current approaches. Current approaches are more effective in handling back recommendation points than front recommendation points.*

To sum up, different contexts can affect the performance of current code-based API recommendation approaches. Among them, the extremely short contexts and front recommendation points bring the most challenges for the accurate recommendation.

5.4 Adaptation to Cross-Domain Projects (RQ6)

We have divided APIBENCH-C into five different domains in Section 3. In this section, we aim at studying the adaption capability of current approaches for cross-domain projects. We train the approaches in one domain and evaluate them in other different domains. We choose the approaches TravTrans, Deep3, and PyART, which are all designed for Python, for analysis. We do not involve the approaches FOCUS, PAM, or PAM-MAX, since they use coarse-grained context representations or context-insensitive feature, and are difficult to incorporate project-specific information. The first four rows of Table 9 list the cross-domain Success Rate@10 of TravTrans, Deep3 and PyART, respectively.

According to Table 9, the approaches trained on one domain generally perform best on the test set of the same

domain. For example, when trained on data from the “Security” domain, TravTrans, Deep3 and PyART obtain the best scores at 0.54, 0.51 and 0.48 on the test set of the same domain, respectively, in terms of Success Rate@10. However, their performance drops by 2.1% ~ 43.1% when recommending APIs from different domains.

Finding 15: *Current approaches using fine-grained context representation are sensitive to the domain of the training data and suffer from performance drop when recommending cross-domain APIs.*

We also analyze the cross-domain performance of the approaches when training on multiple domains instead of on one single domain. Such analysis is worthwhile to explore whether different domains can complement each other. Then we train the approaches on the projects from the “General” domain of APIBENCH-C and evaluate them on the other four different domains. We show the results in the last row of table 9.

From the table, we can see that the approaches trained on the “General” domain generally show the best performance when evaluating on different domains. For example, TravTrans trained on the “General” domain achieves the Success Rate@10 of 0.72, 0.76, 0.78 and 0.74 on ML, Security, Web and DL domains, respectively, which is significantly higher than the corresponding best scores obtained by TravTrans trained on single one domain. We observe an average boost of 14% for the performance of the approaches when trained on multiple domains than on a single domain. The results indicate that training approaches on multiple domains greatly improve the recommendation performance.

Finding 16: *Training on multiple domains helps the current approaches to recommend APIs in different single domains, and the performance is generally better than only training on a single domain.*

6 DISCUSSION AND FUTURE WORK

6.1 Query Reformulation for Query-based API Recommendation

In Section 4.2, we find that query reformulation techniques can not only help current query-based API recommendation approaches find more correct APIs but also improve the ranking performance. Based on query reformulation, BIKER can even achieve a Success Rate@10 of 0.80 in class-level and 0.51 in method-level API recommendation. The results demonstrate that query quality has a great impact on the recommendation results and suggest that query reformulation should become a common pre-processing technique used before query-based API recommendation. We also discover that some query reformulation techniques, such as adding predicted API class names or relevant words, can improve the performance of query-based API recommendation approaches. However, to the best of our knowledge, few studies have considered integrating these techniques, which could be one major reason that current approaches achieve limited performance.

By implementing a random deletion strategy, in Section 4.2.3 we find that user queries usually contain noisy words, which can bias the recommendation results. We

TABLE 9

The cross-domain Success Rate@10 of Python code-based API recommendation baselines. The rows list the domains where three baselines are trained and the columns list the domains where three baselines are evaluated. The red number indicates the best performance an approach achieves when trained on one domain (The largest number in each row). The numbers with gray background indicates the best performance achieved on a specific testing domain (The largest number in each column).

Training Domain	TravTrans				Deep3				PyART			
	ML	Security	Web	DL	ML	Security	Web	DL	ML	Security	Web	DL
ML	0.64	0.58	0.53	0.71	0.42	0.41	0.36	0.48	0.39	0.35	0.40	0.40
Security	0.40	0.54	0.54	0.39	0.31	0.51	0.42	0.29	0.36	0.48	0.47	0.36
Web	0.54	0.63	0.64	0.51	0.33	0.42	0.46	0.31	0.42	0.47	0.50	0.40
DL	0.66	0.58	0.50	0.68	0.44	0.39	0.33	0.44	0.43	0.36	0.38	0.45
General	0.72	0.76	0.78	0.74	0.55	0.65	0.62	0.57	0.44	0.44	0.46	0.46

summarize three kinds of cases in which a query contains noisy words. However, there exists very little work that aims to detect and eliminate the irrelevant words for recommendation systems, which poses a great challenge for current approaches to be robust when handling various user queries. Although a random deletion strategy reduces the overall performance on average, the positive improvement of deletion on some specific words indicates the potential benefits of noisy word deletion.

Implication 1: *Current query-based API recommendation approaches should be integrated with query reformulation techniques to be more effective.*

6.2 Data Sources for Query-based API Recommendation

In Section 4.1, we point out that current query-based API recommendation approaches face the problem of building a comprehensive knowledge base due to the lack of enough data such as query-API pairs. In Section 4.3, we further discover that there is a semantic gap between user queries and descriptions from the official documentation. Both the lack of enough data for knowledge base creation and the semantic gap increase the difficulty of accurate API recommendation based on only official documentation. Such challenges can not be easily solved by improving learning-based models or pattern-based models. One effective way to mitigate the difficulty is to involve Stack Overflow posts, as analyzed in Section 4.3. While Stack Overflow is only one type of data source, our analysis demonstrates that adding appropriate data sources can improve the performance of query-based API recommendation approaches.

Implication 2: *Apart from query reformulation, adding appropriate data sources provides another solution to bridge the gap between queries and APIs.*

6.3 Low Resource Setting in Query-based API Recommendation

In Section 4.1, we find that current learning-based methods do not necessarily outperform traditional retrieval-based

methods. We attribute the results to the limited data such as query-API pairs in the query-based API recommendation task, which is a low-resource scenario [16], [34]. We also discover that pre-trained models such as BERT show superior performance in query reformulation in Section 4.2. This indicates that current pre-trained models can implicitly mitigate the semantic gap between user queries and official descriptions of APIs. Future work is suggested to explore how to make the best use of pre-trained models for query-based API recommendation based on limited available data.

Implication 3: *Few-shot learning with powerful pre-trained models can be a solution to further improve the performance of query-based API recommendation.*

6.4 User-defined APIs

In Section 5.2, we find that current code-based API recommendation approaches, no matter pattern-based or learning-based models, all face the challenge of recommending user-defined APIs. User-defined APIs have become the major bottleneck to further improve the performance of current code-based API recommendation approaches. However, as user-defined APIs usually do not appear in the training set, they can hardly be learned by machine learning methods or be mined by pattern-based methods. A possible solution used by current approaches [37], [45] is to regard the API as a code token and predict the token based on previous contexts. However, this solution also fails if the API token never appears in previous context. Thus, accurately predicting user-defined APIs should be one major direction of code-based API recommendation in future work.

Implication 4: *User-defined API recommendation is one major bottleneck for improving the performance of current code-based API recommendation approaches and remains unsolved.*

6.5 Query-based API Recommendation with Usage Patterns

In this paper, we only focus on testing whether an approach can recommend the correct APIs, but we believe developers

can always benefit more from detailed information about how to use the recommended APIs. A common method is to provide summaries such as the signature and constraints extracted from official documentation along with the recommended APIs. For example, KG-APISumm proposed by Liu *et al.* [49] provides a detailed summary of the recommended API class. However, official documentation sometimes cannot provide enough usage information about an API, which may cause API misuse. For instance, a fresh developer may search “*how to read a file*” in Python and the recommended API should be `fileObject.read()`, but without sufficient experience to use file operations, the developers may forget to close the file after reading it.

A possible solution to complement official documentation and avoid possible API misuse is to provide usage patterns from other developers. In the above example, a common usage pattern `open()`, `fileObject.read()`, `fileObject.close()` can prevent dangerous file operations. As there exist some pattern mining approaches on code, we can combine query-based API recommendation with code-based pattern mining methods for better providing the usage pattern.

Implication 5: *Code-based API recommendation approaches can provide usage patterns to enrich the results returned by query-based API recommendation approaches.*

7 THREAT TO VALIDITY

In this section, we describe the possible threats we may face in this study and discuss how we mitigate them.

7.1 Internal Validity

Our research may face the following internal threats:

Baseline Re-implementation. In this paper, we re-implemented several baselines according to the code or replication packages released by their authors. However, as some baselines are not primarily designed for API recommendation, we slightly modified their code and adapted them into our task and our benchmark. For example, we limit the prediction scope of code completion baselines to only API tokens. Such adaptations may cause the performance of baselines to be slightly different from those in the original papers. To mitigate this threat and validate the correctness of our re-implementation, we refer to some related work that cites these baselines and confirm our experiment results with them.

Data Format. In this paper, we align the performance of all code-based API recommendation baselines with a uniform evaluation process. However, different approaches may require different input formats and different preprocessing methods. For example, TravTrans requires the inputs of ASTs so we need to transform the code in APIBENCH into ASTs. To mitigate the possible impacts brought by such transformations, we make sure that except for the original source code we do not involve extra information and knowledge in this modification process.

Data Quality. We build APIBENCH-Q by manually selecting and labeling API-related queries from Stack Overflow and some tutorial websites. This process involved

some human checks so that some subjective factors may influence the quality of our datasets. To mitigate this threat, we involve at least two persons to label one case and let one of our authors further check if the previous two persons give different opinions to the case. We also implement some rules to automatically filter out the cases that are explicitly unrelated to API recommendation.

7.2 External Validity

Our research may face the following external threats:

Data Selection. To the best of our knowledge, APIBENCH is the largest benchmark in API recommendation task. We try to make it more representative by selecting real-world code repositories from the most popular domains at GitHub and real-world queries from the largest Q&A forum StackOverflow according to several developer surveys [39], [89]. All findings in this empirical study are based on this dataset. However, there may still be slight differences when adapting our findings into other domains and datasets that we do not discuss in this paper.

Programming Language. Our study focuses on the API recommendation on Python and Java, and the findings included in this study may not be generalized to other programming languages that have different API call patterns from Python and Java. However, we believe the impacts of programming languages should not be significant as Python and Java are the most representative dynamically typed and statically typed languages, respectively, while most programming languages can be classified into these two categories.

8 RELATED WORK

In this section, we list the work related to query reformulation, API recommendation and API-related empirical study, respectively.

8.1 Query Reformulation

The effectiveness of query-based approaches highly depends on the input natural language queries, which motivates many previous query reformulation works. Query expansion and query modification are two major reformulation methods, one of which adds several missing words and phrases which are relevant to the queries, the other one replaces or modifies words and phrases. Existing works show that even some minor textual changes influence the final results a lot. To better evaluate the API recommendation methods, we utilize multiple existing methods to reformulate the original queries.

Rahman *et al.* propose ACER [4]. It takes the initial query as the input and identifies appropriate search terms from source code by using a novel term weight named CodeRank, and then it suggests the best reformulation of the original query by using document structures, query quality analysis and machine learning techniques. Rahman *et al.* propose NLP2API [76], which automatically identifies the relevant API classes for a given query and then uses these API classes to reformulate the original query. Lu *et al.* propose to expand the queries by using synonyms. They extract the natural phrases from code identifiers, match the expanded

queries with the identifiers, and then sort the methods in the code base by the matched identifier [51]. Sirres *et al.* present COCABU [88] to resolve the vocabulary mismatch problem. Their approach uses common developer questions and expert answers to augment the user queries and improves the relevance of returned code examples. Nie *et al.* propose QECK [66] to solve the term mismatch problem. QECK retrieves relevant question-answer pairs collected from Stack Overflow as the Pseudo Relevance Feedback (PRF) documents and identifies the specific words from the documents to expand the original query.

There are also some query reformulation techniques in industry or open-source communities. For example, the Google search engine utilizes its own prediction service [30] to reformulate the natural language queries provided by users. The open-source library NLPAUG [52] provides NLP utilities to conduct data augmentation.

8.2 API Recommendation

8.2.1 Query Based API Recommendation

Multiple existing works [15], [31], [37], [48], [49], [54], [58], [73], [78], [90], [91], [99], [100], [102], [103], [104], [109], [109] explore the possibility to provide developers with concrete API recommendation, using the natural language queries as input. Most of these works utilize open-source code bases, and some also use the knowledge in crowd-sourcing forums and wiki websites for augmentation.

Retrieval-based methods. Portfolio [54], proposed by McMillan *et al.*, recommends relevant APIs by utilizing several NLP techniques and indexing approaches with spreading activation network (SAN) algorithms as well as PageRank [10]. Zhang *et al.* supplement the call graph with control flow analysis and design Flow-Augmented Call Graph (FACG) to utilize it for API recommendation [103]. Chan *et al.* model API invocations as API graphs and design subgraph search algorithm to recommend APIs [15]. Rahman *et al.* collect the crowdsourced knowledge on Stack Overflow to extract keyword-API correlations and find several relevant API classes based on them [78]. Huang *et al.* propose BIKER [37] to bridge the lexical gap and knowledge gap that previous approaches faced during API recommendation. BIKER obtains API candidates from Stack Overflow, and uses the similarity between queries and documentations as well as Stack Overflow posts to recommend API methods. Liu *et al.* propose KG-APISumm [49], which is the first knowledge graph designed for API recommendation. KG-APISumm sorts the APIs through similarity calculation between queries and relevant parts of the constructed knowledge graph to recommend API classes.

Other than coding problems encountered by developers, there is another source of natural language functionality descriptions for APIs: feature requests from product managers or users. Thung *et al.* propose a method to recommend APIs based on the feature requests by learning from other modifications of the projects [91]. Xu *et al.* propose MULAPI [100], which takes feature locations, project repositories and API libraries into consideration when recommending APIs.

Learning-based methods. DeepAPI [31], proposed by Gu *et al.*, is the first approach that combines deep learning

with API recommendation. It reformulates the API recommendation task as a query-API translation problem and uses an RNN Encoder-Decoder model to recommend API sequences. Xiong *et al.* propose to use representation learning to recommend web-based smart service [99]. Ling *et al.* propose GeAPI [48] based on graph embedding to provide more semantic information about and between APIs. GeAPI utilizes projects' source code to automatically construct API graphs and leverages graph embedding techniques for API representation. Given a query, it searches relevant subgraphs on the original graph and recommends them to developers. Zhou *et al.* propose BRAID [109] and utilize approaches such as active learning as well as learning-to-rank based on the feedback of users to further improve the performance.

8.2.2 Code Based API Recommendation

Pattern-based methods. Zhong *et al.* propose MAPO [108] to mine API usage patterns and then recommends the relevant usage patterns to developers. Schäfer *et al.* propose Pythia [87] to utilize static pointer analysis and usage-based property inference to recommend APIs for JavaScript. Wang *et al.* propose UP-Miner [96] and use source code to extract succinct usage patterns to recommend APIs. Nguyen *et al.* propose APIREC [61], which uses fine-grained code changes and the corresponding changing contexts to recommend APIs. D'Souza *et al.* propose PyReco [18]. It first extracts API usages from open-source projects and uses such information to rank the API recommendation results by utilizing nearest neighbor classifier techniques. Fowkes *et al.* propose PAM [25] to tackle the problem that the recommended API lists are large and hard to understand. PAM mines API usage patterns through an almost parameter-free probabilistic algorithm and uses them to recommend APIs. Niu *et al.* propose another API usage pattern mining approach, which segments the data using the co-existence relationship of object usages to mine API usage patterns [67]. Liu *et al.* propose RecRank [50] to improve the top-1 accuracy based on API usage paths. Nguyen *et al.* propose FOCUS [64], which mines open-source repositories and analyzes API usages in similar projects to recommend APIs and API usage patterns based on context-aware collaborative-filtering techniques. Wen *et al.* propose FearS [97], which mines open-source repositories and extracts API sequences that are implemented together in the same tasks frequently to recommend APIs.

Learning-based methods. Hindle *et al.* adopt the n-gram model, a widely-used statistical language model, on the code of software [35], and develop a code suggestion tool based on the n-gram model. Tu *et al.* propose to enhance the n-gram model by adding a cache component [92]. Raychev *et al.* propose to extract API sequences from open-source projects and index them into statistical language models to recommend APIs [82]. Nguyen *et al.* propose a graph-based language model GraLan [62]. Based on GraLan, they design an AST-based language model named ASTLan to recommend APIs. Raychev *et al.* propose a probabilistic model with decision trees named TGEN [80] to predict code tokens. Several recent works try to utilize syntax and data flow information for more accurate recommendation, besides focusing on token sequences [32], [44]. He *et al.* pro-

pose PyART [32], which utilizes a predictive model along with data-flow, token similarity and token co-occurrence to recommend APIs. Kim *et al.* leverage Transformer-based techniques to learn the syntactic information from source code [44].

8.3 Empirical Study on APIs

There are several empirical studies focusing on different aspects of APIs. Zhong *et al.* conduct an empirical study on how to use diversified kinds of APIs [105]. Hora *et al.* study API evolution together with its influences on large software ecosystem [36]. Zhong *et al.* study API parameter rules [106]. Monperrus *et al.* conduct an empirical study on rules in API documentation [59]. Ajam *et al.* study API topic issues on Stack Overflow [7], and Linares Vásquez *et al.* focus on how API changes affect activities on Stack Overflow [95]. Moreover, other work with respect to cross-language API mapping relations [107], API migration [9], API usability [19], API deprecation [85], API learning obstacles [86], usage of specific types of APIs [5], [69], the correlation between APIs and software quality [94], as well as order of API usage patterns [14], are also explored. Among these studies, our paper is the first to systematically evaluate both the query-based and code-based API recommendation techniques and provides the first comprehensive benchmark in API recommendation.

9 CONCLUSION

In this paper, we present an empirical study on the API recommendation task. We classify current work into query-based and code-based API recommendation, and build a benchmark named APIBENCH to align the performance of different recommendation approaches. We conclude some findings based on the empirical results of current approaches.

For query-based API recommendation approaches, we find that 1) recommending method-level APIs is still challenging; 2) query reformulation techniques have great potential to improve the quality of user queries thus they can help current approaches better recommend APIs. What's more, user queries also contain some meaningless and verbose words and even a simple word deletion method can improve the performance; 3) approaches built upon different data sources have quite different performances. Q&A forums such as Stack Overflow can greatly help mitigate the gap between user queries and API descriptions.

For code-based API recommendation, we emphasize the superior performance of current deep learning models such as Transformer. However, they still face the challenge of recommending user-defined APIs. We also find different contexts, such as different location of recommendation points and context length, can impact the performance of current approaches. Besides, current approaches suffer from recommending cross-domain APIs.

Based on the findings, we summarize some future directions on improving the performance of API recommendation. For query-based approaches, we encourage researchers to integrate query reformulation techniques with query-based API recommendation approaches to obtain better performance, but how to choose the best query reformulation

strategy still remains as future work. We also believe some few-shot learning methods and different data sources can bridge the gap between user queries and knowledge base under low resource scenarios. For code-based approaches, we recommend future work to focus on improving the performance of user-defined API recommendation and train the approach on multiple domains instead of a single domain.

We released our benchmark APIBENCH and all experimental results at Github². We hope this empirical study can remove some barriers and motivate future research on API recommendation.

². The address of benchmark is at <https://github.com/JohnnyPeng18/APIBench>

REFERENCES

- [1] Geeks4geeks website. <https://www.geeksforgeeks.org/>, 2021.
- [2] Java2s website. <http://www.java2s.com/>, 2021.
- [3] Kode java website. <https://kodejava.org/>, 2021.
- [4] Mohammad Masudur Rahman 0001 and Chanchal K. Roy. Improved query reformulation for concept location using coderank and document structures. *PeerJ PrePrints*, 5, 2017.
- [5] Rabe Abdalkareem, Olivier Nourry, Sultan Wehaibi, Suhaib Mujahid, and Emad Shihab. Why do developers use trivial packages? an empirical case study on npm. In Eric Bodden, Wilhelm Schäfer, Arie van Deursen, and Andrea Zisman, editors, *Proceedings of the 2017 11th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2017, Paderborn, Germany, September 4-8, 2017*, pages 385–395. ACM, 2017.
- [6] AlDanial. Count lines of code. <https://github.com/AlDanial/cloc>, 2021.
- [7] George Ajam, Carlos Rodríguez, and Boualem Benatallah. API topics issues in stack overflow q&as posts: An empirical study. In *XLVI Latin American Computing Conference, CLEI 2020, Loja, Ecuador, October 19-23, 2020*, pages 147–155. IEEE, 2020.
- [8] Uri Alon, Shaked Brody, Omer Levy, and Eran Yahav. code2seq: Generating sequences from structured representations of code, 2019.
- [9] Hussein Alrubaye. How does API migration impact software quality and comprehension? an empirical study. *CoRR*, abs/1907.07724, 2019.
- [10] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. *Comput. Networks*, 30(1-7):107–117, 1998.
- [11] Marcel Bruch, Martin Monperrus, and Mira Mezini. Learning from examples to improve code completion systems. In *Proceedings of the 7th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on the foundations of software engineering*, pages 213–222, 2009.
- [12] Kaibo Cao, Chunyang Chen, Sebastian Baltes, Christoph Treude, and Xiang Chen. Automated query reformulation for efficient search based on query logs from stack overflow. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pages 1273–1285. IEEE, 2021.
- [13] Kaibo Cao, Chunyang Chen, Sebastian Baltes, Christoph Treude, and Xiang Chen. The demo link for sequer. <https://sequer-tpz.novfjxa-uc.a.run.app/?query=>, 2021.
- [14] Ervina Çergani, Sebastian Proksch, Sarah Nadi, and Mira Mezini. Investigating order information in api-usage patterns: A benchmark and empirical study. In Leszek A. Maciaszek and Marten van Sinderen, editors, *Proceedings of the 13th International Conference on Software Technologies, ICSoft 2018, Porto, Portugal, July 26-28, 2018*, pages 91–102. SciTePress, 2018.
- [15] Wing-Kwan Chan, Hong Cheng, and David Lo. Searching connected api subgraph via text phrases. In *Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering*, pages 1–11, 2012.
- [16] Mona Diab. Data paucity and low resource scenarios: Challenges and opportunities. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, page 3612, New York, NY, USA, 2020. Association for Computing Machinery.
- [17] Django. Django api reference. <https://docs.djangoproject.com/en/3.2/ref/>, 2021.
- [18] Andrea Renika D'Souza, Di Yang, and Cristina V Lopes. Collective intelligence for smarter api recommendations in python. In *2016 IEEE 16th International Working Conference on Source Code Analysis and Manipulation (SCAM)*, pages 51–60. IEEE, 2016.
- [19] Stefan Endrikat, Stefan Hanenberg, Romain Robbes, and Andreas Steffik. How do API documentation and static typing affect API usability? In Pankaj Jalote, Lionel C. Briand, and André van der Hoek, editors, *36th International Conference on Software Engineering, ICSE '14, Hyderabad, India - May 31 - June 07, 2014*, pages 632–642. ACM, 2014.
- [20] Stack Exchange. Stack exchange data dump. <https://archive.org/details/stackexchange>, 2021.
- [21] Flask. Flask api reference. <https://flask.palletsprojects.com/en/2.0.x/api/>, 2021.
- [22] J. Fleiss. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76:378–382, 1971.
- [23] Apache Software Foundation. Lucene. <https://lucene.apache.org/>, 2021.
- [24] Eclipse Foundation. The eclipse ide. <https://www.eclipse.org/eclipseide/>, 2021.
- [25] Jaroslav Fowkes and Charles Sutton. Parameter-free probabilistic api mining across github. In *Proceedings of the 2016 24th ACM SIGSOFT international symposium on foundations of software engineering*, pages 254–265, 2016.
- [26] Jaroslav Fowkes and Charles Sutton. The replication package of pam. <https://github.com/mast-group/api-mining>, 2021.
- [27] Jaroslav M. Fowkes and Charles Sutton. Parameter-free probabilistic API mining across github. In Thomas Zimmermann, Jane Cleland-Huang, and Zhendong Su, editors, *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2016, Seattle, WA, USA, November 13-18, 2016*, pages 254–265. ACM, 2016.
- [28] Google. Android api reference. <https://developer.android.com/reference>, 2021.
- [29] Google. The google news word2vec model. <https://code.google.com/archive/p/word2vec/>, 2021.
- [30] Google. The interface of google prediction service. <http://suggestions.google.com/complete/search?>, 2021.
- [31] Xiaodong Gu, Hongyu Zhang, Dongmei Zhang, and Sunghun Kim. Deep api learning. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 631–642, 2016.
- [32] Xincheng He, Lei Xu, Xiangyu Zhang, Rui Hao, Yang Feng, and Baowen Xu. Pyart: Python api recommendation in real-time. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pages 1634–1645. IEEE, 2021.
- [33] Xincheng He, Lei Xu, Xiangyu Zhang, Rui Hao, Yang Feng, and Baowen Xu. The replication package of pyart. <https://github.com/PYART0/PyART>, 2021.
- [34] Michael A. Hedderich, Lukas Lange, Heike Adel, Jannik Strötgen, and Dietrich Klakow. A survey on recent approaches for natural language processing in low-resource scenarios, 2021.
- [35] Abram Hindle, Earl T. Barr, Zhendong Su, Mark Gabel, and Premkumar T. Devanbu. On the naturalness of software. In Martin Glinz, Gail C. Murphy, and Mauro Pezzè, editors, *34th International Conference on Software Engineering, ICSE 2012, June 2-9, 2012, Zurich, Switzerland*, pages 837–847. IEEE Computer Society, 2012.
- [36] André C. Hora, Romain Robbes, Nicolas Anquetil, Anne Etien, Stéphane Ducasse, and Marco Tulio Valente. How do developers react to API evolution? the pharo ecosystem case. In Rainer Koschke, Jens Krinke, and Martin P. Robillard, editors, *2015 IEEE International Conference on Software Maintenance and Evolution, ICSME 2015, Bremen, Germany, September 29 - October 1, 2015*, pages 251–260. IEEE Computer Society, 2015.
- [37] Qiao Huang, Xin Xia, Zhenchang Xing, David Lo, and Xinyu Wang. Api method recommendation without worrying about the task-api knowledge gap. In *2018 33rd IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 293–304. IEEE, 2018.
- [38] Qiao Huang, Xin Xia, Zhenchang Xing, David Lo, and Xinyu Wang. The replication package of biker. <https://github.com/tkdsheep/BIKER-ASE2018>, 2021.
- [39] JetBrains. Python developer survey conducted by jetbrains and python software foundation. <https://www.jetbrains.com/lp/python-developers-survey-2020/>, 2020.
- [40] JetBrains. The basic code completion feature of pycharm ide. https://www.jetbrains.com/help/pycharm/auto-completing-code.html#basic_completion, 2021.
- [41] JetBrains. The intellij idea ide. <https://www.jetbrains.com/idea/>, 2021.
- [42] JetBrains. The pycharm ide. <https://www.jetbrains.com/pycharm/>, 2021.
- [43] Elise Jing, Kristiana Schneck, Dennis Egan, and Scott A. Waterman. Identifying introductions in podcast episodes from automatically generated transcripts, 2021.
- [44] Seohyun Kim, Jinman Zhao, Yuchi Tian, and Satish Chandra. Code prediction by feeding trees to transformers. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pages 150–162. IEEE, 2021.
- [45] Seohyun Kim, Jinman Zhao, Yuchi Tian, and Satish Chandra. The replication package of travtrans. <https://github.com/facebookresearch/code-prediction-transformer>, 2021.
- [46] lanwuwei. A pre-trained bert on stackoverflow corpus. <https://github.com/lanwuwei/BERTOverflow>, 2021.

- [47] Alexander LeClair, Siyuan Jiang, and Collin McMillan. A neural model for generating natural language summaries of program subroutines. *ICSE '19*, page 795–806. IEEE Press, 2019.
- [48] Chun-Yang Ling, Yan-Zhen Zou, Ze-Qi Lin, and Bing Xie. Graph embedding based api graph search and recommendation. *Journal of Computer Science and Technology*, 34(5):993–1006, 2019.
- [49] Mingwei Liu, Xin Peng, Andrian Marcus, Zhenchang Xing, Wenkai Xie, Shuangshuang Xing, and Yang Liu. Generating query-specific class api summaries. In *Proceedings of the 2019 27th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 120–130, 2019.
- [50] Xiaoyu Liu, LiGuo Huang, and Vincent Ng. Effective api recommendation without historical software repositories. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, pages 282–292, 2018.
- [51] Meili Lu, Xiaobing Sun, Shaowei Wang, David Lo, and Yucong Duan. Query expansion via wordnet for effective code search. In Yann-Gaël Guéhéneuc, Bram Adams, and Alexander Serebrenik, editors, *22nd IEEE International Conference on Software Analysis, Evolution, and Reengineering, SANER 2015, Montreal, QC, Canada, March 2-6, 2015*, pages 545–549. IEEE Computer Society, 2015.
- [52] Edward Ma. Nlp augmentation. <https://github.com/makcedward/nlpaug>, 2019.
- [53] Matplotlib. Matplotlib api reference. <https://matplotlib.org/stable/api/index.html>, 2021.
- [54] Collin McMillan, Mark Grechanik, Denys Poshyvanyk, Qing Xie, and Chen Fu. Portfolio: finding relevant functions and their usage. In *Proceedings of the 33rd International Conference on Software Engineering*, pages 111–120, 2011.
- [55] Microsoft. Github website. <https://github.com/>, 2021.
- [56] Microsoft. The visual studio code editor. <https://code.visualstudio.com/>, 2021.
- [57] George A. Miller. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, November 1995.
- [58] Martin Monperrus, Michael Eichberg, Elif Tekes, and Mira Mezini. What should developers be aware of? an empirical study on the directives of api documentation. *Empirical Software Engineering*, 17(6):703–737, 2012.
- [59] Martin Monperrus, Michael Eichberg, Elif Tekes, and Mira Mezini. What should developers be aware of? an empirical study on the directives of API documentation. *Empir. Softw. Eng.*, 17(6):703–737, 2012.
- [60] Benjamin Newman, Prafulla Kumar Choubey, and Nazneen Rajani. P-adapters: Robustly extracting factual information from language models with diverse prompts, 2021.
- [61] Anh Tuan Nguyen, Michael Hilton, Mihai Codoban, Hoan Anh Nguyen, Lily Mast, Eli Rademacher, Tien N Nguyen, and Danny Dig. Api code recommendation using statistical learning from fine-grained changes. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 511–522, 2016.
- [62] Anh Tuan Nguyen and Tien N Nguyen. Graph-based statistical language model for code. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, volume 1, pages 858–868. IEEE, 2015.
- [63] Anh Tuan Nguyen, Tung Thanh Nguyen, Hoan Anh Nguyen, Ahmed Tamrawi, Hung Viet Nguyen, Jafar Al-Kofahi, and Tien N Nguyen. Graph-based pattern-oriented, context-sensitive source code completion. In *2012 34th International Conference on Software Engineering (ICSE)*, pages 69–79. IEEE, 2012.
- [64] Phuong T Nguyen, Juri Di Rocco, Davide Di Ruscio, Lina Ochoa, Thomas Degueule, and Massimiliano Di Penta. Focus: A recommender system for mining api function calls and usage patterns. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, pages 1050–1060. IEEE, 2019.
- [65] Phuong T Nguyen, Juri Di Rocco, Davide Di Ruscio, Lina Ochoa, Thomas Degueule, and Massimiliano Di Penta. The replication package of focus. <https://github.com/crossminer/FOCUS>, 2021.
- [66] Liming Nie, He Jiang, Zhilei Ren, Zeyi Sun, and Xiaochen Li. Query expansion based on crowd knowledge for code search. *IEEE Trans. Serv. Comput.*, 9(5):771–783, 2016.
- [67] Haoran Niu, Iman Keivanloo, and Ying Zou. Api usage pattern recommendation for software development. *Journal of Systems and Software*, 129:127–139, 2017.
- [68] Numpy. Numpy api reference. <https://numpy.org/doc/stable/reference/>, 2021.
- [69] Semih Okur and Danny Dig. How do developers use parallel libraries? In Will Tracz, Martin P. Robillard, and Tevfik Bultan, editors, *20th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE-20), SIGSOFT/FSE'12, Cary, NC, USA - November 11 - 16, 2012*, page 54. ACM, 2012.
- [70] Oracle. Java se 8 api documentation. www.oracle.com/technetwork/java/javase/documentation/jdk8-doc-downloads-2133158.html, 2021.
- [71] Pandas. Pandas api reference. <https://pandas.pydata.org/docs/reference/index.html>, 2021.
- [72] Python. Python standard library. <https://docs.python.org/3/library/>, 2021.
- [73] Lianyong Qi, Qiang He, Feifei Chen, Wanchun Dou, Shaohua Wan, Xuyun Zhang, and Xiaolong Xu. Finding all you need: web apis recommendation in web of things through keywords search. *IEEE Transactions on Computational Social Systems*, 6(5):1063–1072, 2019.
- [74] Maithra Raghu and Eric Schmidt. A survey of deep learning for scientific discovery, 2020.
- [75] Mohammad Masudur Rahman and Chanchal Roy. Effective reformulation of query for code search using crowdsourced knowledge and extra-large data analytics. 06 2018.
- [76] Mohammad Masudur Rahman and Chanchal Roy. Nlp2api: Query reformulation for code search using crowdsourced knowledge and extra-large data analytics. In *2018 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, pages 714–714, 2018.
- [77] Mohammad Masudur Rahman and Chanchal Roy. The replication package for nlp2api. <https://github.com/masud-technop e/NLP2API-Replication-Package>, 2021.
- [78] Mohammad Masudur Rahman, Chanchal K Roy, and David Lo. Rack: Automatic api recommendation using crowdsourced knowledge. In *2016 IEEE 23rd International Conference on Software Analysis, Evolution, and Reengineering (SANER)*, volume 1, pages 349–359. IEEE, 2016.
- [79] Mohammad Masudur Rahman, Chanchal K Roy, and David Lo. The replication package of rack. <https://github.com/masud-technop e/RACK-Replication-Package>, 2021.
- [80] Veselin Raychev, Pavol Bielik, and Martin Vechev. Probabilistic model for code with decision trees. In *Proceedings of the 2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA 2016*, page 731–747, New York, NY, USA, 2016. Association for Computing Machinery.
- [81] Veselin Raychev, Pavol Bielik, and Martin Vechev. The replication package of deep3. <https://github.com/eth-sri/ModelsPHOG>, 2021.
- [82] Veselin Raychev, Martin Vechev, and Eran Yahav. Code completion with statistical language models. In *Proceedings of the 35th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 419–428, 2014.
- [83] Francesco Ricci, Lior Rokach, and Bracha Shapira. *Introduction to Recommender Systems Handbook*, pages 1–35. Springer US, Boston, MA, 2011.
- [84] Romain Robbes and Michele Lanza. Improving code completion with program history. *Automated Software Engineering*, 17(2):181–212, 2010.
- [85] Romain Robbes, Mircea Lungu, and David Röthlisberger. How do developers react to API deprecation?: the case of a smalltalk ecosystem. In Will Tracz, Martin P. Robillard, and Tevfik Bultan, editors, *20th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE-20), SIGSOFT/FSE'12, Cary, NC, USA - November 11 - 16, 2012*, page 56. ACM, 2012.
- [86] Martin P. Robillard and Robert DeLine. A field study of API learning obstacles. *Empir. Softw. Eng.*, 16(6):703–732, 2011.
- [87] Max Schäfer, Manu Sridharan, Julian Dolby, and Frank Tip. Effective smart completion for javascript. *Technical Report RC25359*, 2013.
- [88] Raphael Sirres, Tegawendé F. Bissyandé, Dongsun Kim, David Lo, Jacques Klein, Kisub Kim, and Yves Le Traon. Augmenting and structuring user queries to support efficient free-form code search. In *Proceedings of the 40th International Conference on Software Engineering, ICSE '18*, page 945, New York, NY, USA, 2018. Association for Computing Machinery.

- [89] snyk. Jvm ecosystem report 2020. <https://snyk.io/blog/developers-dont-want-to-leave-java-8-as-64-bit-hold-firm-on-their-preferred-release/>, 2021.
- [90] Xiaobing Sun, Congying Xu, Bin Li, Yucong Duan, and Xintong Lu. Enabling feature location for api method recommendation and usage location. *IEEE Access*, 7:49872–49881, 2019.
- [91] Ferdian Thung, Shaowei Wang, David Lo, and Julia L. Lawall. Automatic recommendation of API methods from feature requests. In Ewen Denney, Tevfik Bultan, and Andreas Zeller, editors, *2013 28th IEEE/ACM International Conference on Automated Software Engineering, ASE 2013, Silicon Valley, CA, USA, November 11–15, 2013*, pages 290–300. IEEE, 2013.
- [92] Zhaopeng Tu, Zhendong Su, and Premkumar Devanbu. On the localness of software. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 269–280, 2014.
- [93] Efstathiou Vasiliki, Chatzilenas Christos, and Spinellis Diomidis. Word Embeddings for the Software Engineering Domain, March 2018.
- [94] Mario Linares Vásquez, Gabriele Bavota, Carlos Bernal-Cárdenas, Massimiliano Di Penta, Rocco Oliveto, and Denys Poshyvanyk. API change and fault proneness: a threat to the success of android apps. In Bertrand Meyer, Luciano Baresi, and Mira Mezini, editors, *Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering, ESEC/FSE’13, Saint Petersburg, Russian Federation, August 18–26, 2013*, pages 477–487. ACM, 2013.
- [95] Mario Linares Vásquez, Gabriele Bavota, Massimiliano Di Penta, Rocco Oliveto, and Denys Poshyvanyk. How do API changes trigger stack overflow discussions? a study on the android SDK. In Chanchal K. Roy, Andrew Begel, and Leon Moonen, editors, *22nd International Conference on Program Comprehension, ICPC 2014, Hyderabad, India, June 2–3, 2014*, pages 83–94. ACM, 2014.
- [96] Jue Wang, Yingnong Dang, Hongyu Zhang, Kai Chen, Tao Xie, and Dongmei Zhang. Mining succinct and high-coverage api usage patterns from source code. In *2013 10th Working Conference on Mining Software Repositories (MSR)*, pages 319–328. IEEE, 2013.
- [97] Fengcai Wen, Emad Aghajani, Csaba Nagy, Michele Lanza, and Gabriele Bavota. Siri, write the next method. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pages 138–149. IEEE, 2021.
- [98] Rensong Xie, Xianglong Kong, Lulu Wang, Ying Zhou, and Bixin Li. Hirec: Api recommendation using hierarchical context. In *2019 IEEE 30th International Symposium on Software Reliability Engineering (ISSRE)*, pages 369–379. IEEE, 2019.
- [99] Wei Xiong, Zhihui Lu, Bing Li, Bo Hang, and Zhao Wu. Automating smart recommendation from natural language api descriptions via representation learning. *Future Generation Computer Systems*, 87:382–391, 2018.
- [100] Congying Xu, Xiaobing Sun, Bin Li, Xintong Lu, and Hongjing Guo. MULAPI: improving API method recommendation with API usage location. *J. Syst. Softw.*, 142:195–205, 2018.
- [101] Le Xue, Mingfei Gao, Zeyuan Chen, Caiming Xiong, and Ran Xu. Robustness evaluation of transformer-based form field extractors via form attacks, 2021.
- [102] Weizhao Yuan, Hoang H Nguyen, Lingxiao Jiang, Yuting Chen, Jianjun Zhao, and Haibo Yu. Api recommendation for event-driven android application development. *Information and Software Technology*, 107:30–47, 2019.
- [103] Qirun Zhang, Wujie Zheng, and Michael R. Lyu. Flow-augmented call graph: A new foundation for taming API complexity. In Dimitra Giannakopoulou and Fernando Orejas, editors, *Fundamental Approaches to Software Engineering - 14th International Conference, FASE 2011, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2011, Saarbrücken, Germany, March 26–April 3, 2011. Proceedings*, volume 6603 of *Lecture Notes in Computer Science*, pages 386–400. Springer, 2011.
- [104] Wujie Zheng, Qirun Zhang, and Michael R. Lyu. Cross-library API recommendation using web search engines. In Tibor Gyimóthy and Andreas Zeller, editors, *SIGSOFT/FSE’11 19th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE-19) and ESEC’11: 13th European Software Engineering Conference (ESEC-13)*, Szeged, Hungary, September 5–9, 2011, pages 480–483. ACM, 2011.
- [105] Hao Zhong and Hong Mei. An empirical study on API usages. *IEEE Trans. Software Eng.*, 45(4):319–334, 2019.
- [106] Hao Zhong, Na Meng, Zexuan Li, and Li Jia. An empirical study on API parameter rules. In Gregg Rothermel and Doo-Hwan Bae, editors, *ICSE ’20: 42nd International Conference on Software Engineering, Seoul, South Korea, 27 June - 19 July, 2020*, pages 899–911. ACM, 2020.
- [107] Hao Zhong, Suresh Thummalapenta, and Tao Xie. Exposing behavioral differences in cross-language API mapping relations. In Vittorio Cortellessa and Dániel Varró, editors, *Fundamental Approaches to Software Engineering - 16th International Conference, FASE 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16–24, 2013. Proceedings*, volume 7793 of *Lecture Notes in Computer Science*, pages 130–145. Springer, 2013.
- [108] Hao Zhong, Tao Xie, Lu Zhang, Jian Pei, and Hong Mei. Mapo: Mining and recommending api usage patterns. In *European Conference on Object-Oriented Programming*, pages 318–343. Springer, 2009.
- [109] Yu Zhou, Xinying Yang, Taolue Chen, Zhiqiu Huang, Xiaoxing Ma, and Harald C. Gall. Boosting API recommendation with implicit feedback. *CoRR*, abs/2002.01264, 2020.