

# When Language Model Meets Private Library

Daoguang Zan<sup>1,2\*</sup>, Bei Chen<sup>3</sup>, Zeqi Lin<sup>3</sup>, Bei Guan<sup>2,4</sup>, Yongji Wang<sup>2,4,5</sup>, Jian-Guang Lou<sup>3</sup>

<sup>1</sup>Cooperative Innovation Center, Institute of Software, Chinese Academy of Sciences

<sup>2</sup>University of Chinese Academy of Sciences; <sup>3</sup>Microsoft Research Asia

<sup>4</sup>Integrative Innovation Center, Institute of Software, Chinese Academy of Sciences

<sup>5</sup>State Key Laboratory of Computer Science, Institute of Software, Chinese Academy of Sciences

{daoguang@, guanbei@, ywang@itechs.}iscas.ac.cn

{beichen, zeqi.lin, jlou}@microsoft.com

## Abstract

With the rapid development of pre-training techniques, a number of language models have been pre-trained on large-scale code corpora and perform well in code generation. In this paper, we investigate how to equip pre-trained language models with the ability of code generation for private libraries. In practice, it is common for programmers to write code using private libraries. However, this is a challenge for language models since they have never seen private APIs during training. Motivated by the fact that private libraries usually come with elaborate API documentation, we propose a novel framework with two modules: the APIRetriever finds useful APIs, and then the APICoder generates code using these APIs. For APIRetriever, we present a dense retrieval system and also design a friendly interaction to involve users. For APICoder, we can directly use off-the-shelf language models, or continually pre-train the base model on a code corpus containing API information. Both modules are trained with data from public libraries and can be generalized to private ones. Furthermore, we craft three benchmarks for private libraries, named TorchDataEval, MonkeyEval, and BeatNumEval. Experimental results demonstrate the impressive performance of our framework.<sup>1</sup>

## 1 Introduction

Code generation, automatically generating code snippets based on user descriptions, is one of the long-standing challenges in the software engineering and artificial intelligence communities. With the rapid development of pre-training techniques, a number of language models are pre-trained on large-scale code corpora and able to generate decent code snippets, for example, Codex (Chen et al., 2021), AlphaCode (Li et al., 2022), CODEGEN (Nijkamp et al., 2022), and InCoder (Fried et al., 2022).

\*Work done as an intern at Microsoft Research Asia.

<sup>1</sup>Our work is available at <https://github.com/microsoft/PyCodeGPT/tree/main/apicoder>.

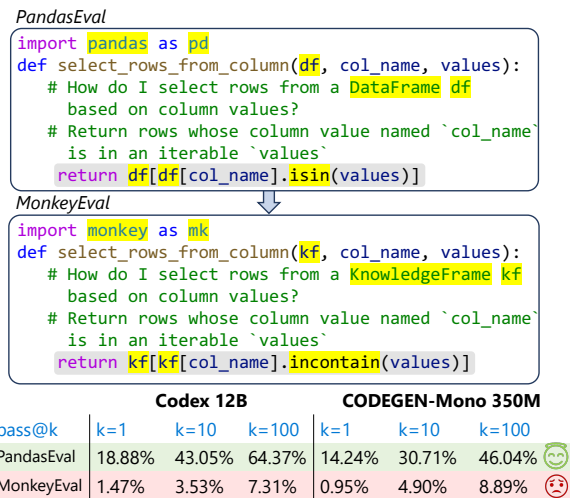


Figure 1: A practical example of converting PandasEval (public) to MonkeyEval (private). The changed parts are highlighted in yellow. The performance of Codex 12B and CODEGEN-MONO 350M is shown at the bottom.

They bring fresh energy to code generation and improve coding efficiency (Vaithilingam et al., 2022). Although making remarkable progress, these models may be biased towards generating code that is similar to the training distribution (Chen et al., 2021). What if one wants to generate code beyond the training distribution? A real-world scenario for programmers is to write code using a private library, which is very common in practice. For example, for security and functionality reasons, companies often build private libraries for internal use only. Private libraries provide a number of private APIs that have not been seen by the language models and are also not publicly available on any code hosting platform like GitHub. Therefore, it is worth exploring whether and how pre-trained language models can generate code using private libraries.

It is challenging for existing language models to generate code that uses a private library directly. A practical evidence is shown in Figure 1. We built a pseudo private library named Monkey based on a

public one named Pandas. PandasEval (Zan et al., 2022) is a benchmark consisting of 101 Pandas programming problems. We convert all Pandas-related keywords in PandasEval into the new version and construct MonkeyEval (details in Section 4). As seen in Figure 1, Codex 12B and CODEGEN-MONO 350M show a significant drop in performance on the private MonkeyEval compared to their performance on the public PandasEval. For example, Codex 12B drops from 18.88% to 1.47% on pass@1, showing the inadequacy of the language models in code generation for private libraries.

To meet the challenge, we propose a framework to equip pre-trained language models with the ability to generate code that uses private libraries. As is known, private libraries usually come with elaborate API documentation, which motivates our main idea to mimic the process of a programmer learning to write code using private libraries. This process is also known as API practices in the software engineering field (Snodgrass and Winnie, 2019): first learning the private API documentation and then invoking the APIs to implement the needed functionalities. Analogically, there are two modules in our framework: an APIRetriever first retrieves the useful APIs based on the programming problem and the API documentation, and then an APICoder uses these APIs to generate code. For APIRetriever, we train a dense retriever and also design a friendly interaction to involve users in the loop optionally. For APICoder, we can directly use existing language models of code generation, such as CODEGEN, to invoke the private APIs; furthermore, to better teach a language model how to invoke APIs, we also continually pre-train the base model on a code corpus containing API information from public libraries, and obtain our reinforced model called CODEGENAPI. Since we only have access to the data of public libraries during training, we expect that APIRetriever and APICoder can be generalized to private libraries via learning.

To evaluate the code generation for private libraries, we craft three benchmarks, named TorchDataEval, MonkeyEval, and BeatNumEval. TorchDataEval includes 50 programming problems using the TorchData library. The last two are adapted from PandasEval and NumpyEval (Zan et al., 2022), respectively, each consisting of 101 programming problems. Extensive experiments on the three benchmarks have revealed that our framework effectively improves the performance of pre-trained

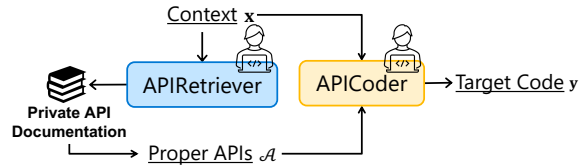


Figure 2: The overview of our proposed framework.

language models on code generation for private libraries. We also provide a thorough analysis to facilitate progress in this direction.

## 2 Framework

First, we would like to define the task of code generation formally. Given *context*, the task aims to generate *target code*. In Figure 1, context and target code are shown in white and grey backgrounds, respectively. Context consists of a comment, which is a natural language description of the programming problem, and a code snippet including import statements, function header, etc. Target code solves the programming problem in context. We denote the context by  $\mathbf{x}$ . Code generation model  $\mathcal{M}$  outputs target code  $\mathbf{y}$  based on  $\mathbf{x}$ . For the task of code generation for private library, the context  $\mathbf{x}$  contains the instruction for using a private library, such as an import statement. The target code  $\mathbf{y}$  contains the calls of the corresponding private library APIs.

As mentioned in Section 1, private libraries are usually equipped with elaborate API documentation. As a technical reference manual outlining how to use the library, API documentation typically includes a quick start guide, tutorials, and an instruction for each API (e.g., API name, signature, description, parameters, and examples). To take advantage of the API documentation, we propose to mimic the generic process of a programmer coding with private APIs, and design a framework to generate code that can invoke private APIs. The framework consists of APIRetriever and APICoder with the overview shown in Figure 2. Given the context, APIRetriever  $\mathcal{M}_R$  aims to retrieve possible used APIs from the API documentation; and APICoder  $\mathcal{M}_C$  is dedicated to generating code using the retrieved APIs. The process can be formalized as  $\mathcal{A} = \mathcal{M}_R(\mathbf{x})$  and  $\mathbf{y} = \mathcal{M}_C(\mathcal{A}; \mathbf{x})$ , where  $\mathcal{A}$  represents the set of information of all proper APIs, and each  $\mathbf{a} \in \mathcal{A}$  is the information of an API. In our implementation, we design the API information to include the API name, signature and description. Note that we only use the first sentence of the API

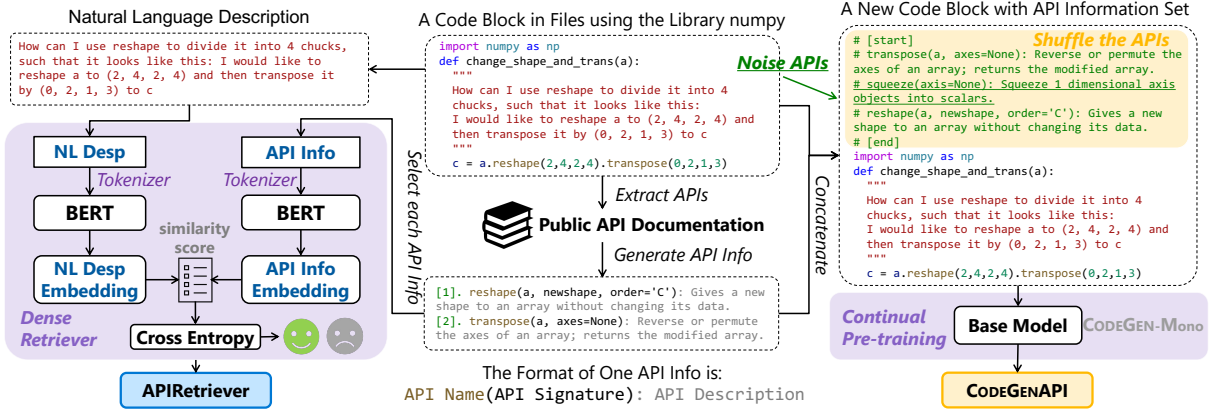


Figure 3: The training process of APIRetriever and CODEGENAPI.

description since it is sufficient to summarize.

### 3 Methodology

We have introduced our framework that provides pre-trained models a fantastic way to deal with private libraries. In this section, we present the data collection, followed by the detailed design of our APIRetriever and APICoder.

#### 3.1 Data Collection

We collect API information and code files of public libraries due to the fact that we can only access data from public libraries. Then we train the models based on the public data with the expectation that the model can be generalized to private libraries. For API information, we consider the 31 most popular public libraries in Python (e.g., Pandas, NumPy, and scikit-learn) according to the popularity ranking on StackOverFlow<sup>2</sup>. For each of the libraries, we crawled its API documentation and extracted detailed information about each API, including the API name, signature, description, parameters, usage examples, and so on. Please refer to Appendix A for the details of the 31 public libraries. For code files, we first collected a 330GB corpus from GitHub containing 60.6M python files and then extracted those files that involved one or more API calls from the 31 public libraries. After a bunch of pre-processing strategies, for example, de-duplicating, cleaning, and formatting, we finally obtained 4.54M python files, denoted by  $\mathcal{D}$ .

#### 3.2 APIRetriever

APIRetriever aims to find the proper APIs based on the description of a programming problem. We

regard it as a dense retrieval task (Qu et al., 2020; Xiong et al., 2020; Santhanam et al., 2021; Formal et al., 2022) and design a simple dual-encoder model (Karpukhin et al., 2020) to retrieve the possible used APIs for each programming problem. To further boost the retrieval performance, a friendly interaction approach is designed to involve users.

**Training.** To train APIRetriever, we need a large amount of pairwise data, natural language description and API information. We first segment each python file  $d \in \mathcal{D}$  into  $K$  code blocks  $(d_1, d_2, \dots, d_K)$  using the pip-tools, i.e., redbaron, autopep8, and docformatter, where each code block is a relatively well-rounded code fragment, such as a function or a class. For each code block  $d_i$ , we extract all API names and obtain the corresponding API signatures and descriptions by searching our collected 31 API documentations<sup>3</sup>. The information of an API includes its name, signature and description, denoted by  $\mathbf{a} \in \mathcal{A}$ . Each  $\mathbf{a}$  and the natural language description  $\mathbf{p}$  extracted from the same code block  $d_i$  are regarded as a positive training sample. For the negative training sample, we randomly sample an API  $\hat{\mathbf{a}}$  that is unrelated to  $d_i$  from the same library. In total, we obtained 40.3M  $(\mathbf{p}, \mathbf{a}, \hat{\mathbf{a}}_1, \hat{\mathbf{a}}_2, \dots)$  sets as training samples. As in Figure 3, the left part shows the training process of APIRetriever. Our APIRetriever is a dual-encoder model. The two dense encoder,  $E_p(\cdot)$  and  $E_a(\cdot)$ , map  $\mathbf{p}$  and  $\mathbf{a}$  to  $z$ -dimensional vectors, respectively. Then, we use the dot product of their vectors to calculate the similarity score formalized as  $E_p(\mathbf{p})^\top E_a(\mathbf{a})$ , where  $E_p(\cdot)$  and  $E_a(\cdot)$  are implemented by two independent BERT (Devlin et al.,

<sup>2</sup><https://stackoverflow.com/tags?tab=popular>

<sup>3</sup>If an API name matches more than one candidate API, we randomly pick one.

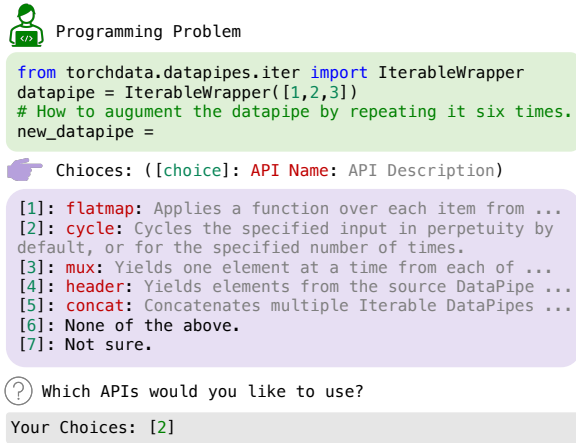


Figure 4: Friendly interaction interface for users.

2019) with base-uncased version. We use BERT instead of CodeBERT (Feng et al., 2020) as most tokens in  $\mathbf{p}$  and  $\mathbf{a}$  are natural language rather than programming language.

**Inference.** After the training phase, we can use APIRetriever to retrieve private APIs for each programming problem description. In detail, we apply  $E_a$  to all the APIs and index them by FAISS (Johnson et al., 2019) offline. Given a new programming problem description  $\mathbf{p}$  at run-time, we only need to produce its embedding  $v_p = E_p(\mathbf{p})$  and recall the top- $k$  APIs with the embeddings closest to  $v_p$ .

**Human Interaction with APIRetriever.** In order to further increase the accuracy of API retrieval, we provide a friendly interaction interface to allow humans in the loop with APIRetriever, as shown in Figure 4. In the interaction interface, we give the programming problem and the top-5 APIs retrieved by APIRetriever, and let users choose one or more APIs that may be used in the target code. Note that we only provide API names and descriptions to users, as we find in our empirical experiments that providing API signatures has a negative effect on making the correct choice.

### 3.3 APICoder

APIRetriever finds useful APIs for a programming problem, and then APICoder aims to generate code that solves the problem with these APIs. We make use of the most straightforward way for APICoder: prompting API information set  $\mathcal{A}$  in front of the context  $\mathbf{x}$ . Formally, the APICoder can be written as  $\mathbf{y} = \mathcal{M}_C(\text{Concat}(\mathcal{A}, \mathbf{x}))$ , where  $\text{Concat}(\mathcal{A}, \mathbf{x})$  means to concatenate the API information set and the context. Examples can be found

in Figure 3. Each API information is in the form of “name(signature):description”. This is to mimic programmers learning the APIs properly before writing code using them.

Technically speaking, the off-the-shelf code generation models, such as CodeT5, CodeGPT, CodeClippy, CodeParrot, CODEGEN, and Codex, can be applied directly to land APICoder. Although these base models can achieve gains in correctly invoking APIs, they have not learned how to use them as an explicit training task. To better use the APIs, we devised a fantastic idea of continually pre-training the base models using code files with API information inserted.

In practice, we use CODEGEN-MONO 350M (Nijkamp et al., 2022) as our base model, based on which we continually pre-train and obtain our reinforced model called CODEGENAPI. CODEGEN is a GPT-based model skilled at generating code. We choose it because it is by far the most popular and publicly available model. As for the training corpus, we use the collected python files  $\mathcal{D}$  mentioned in Section 3.1. Firstly, as done for APIRetriever, each file  $\mathbf{d} \in \mathcal{D}$  is split into  $K$  code blocks  $(d_1, d_2, \dots, d_K)$ . For each code block  $d_i$ , we obtain the set of API information  $\mathcal{A}_i$ . Then, the  $K$  code blocks and sets of API information are cross-merged to output a new file  $\hat{\mathbf{d}} = (\mathcal{A}_1, d_1, \mathcal{A}_2, d_2, \dots, \mathcal{A}_K, d_K)$ . This mimics the process of API information as a prompt for each block. Then, we continually pre-train the base model on the new code files, teaching the model to write code based on the prompted APIs. In addition, as shown in Figure 3, to make APICoder more robust, we shuffle the APIs in each set  $\mathcal{A}_i$  and also add noise APIs, since APIRetriever does not know the order of APIs in the target code and often finds incorrect APIs.

During the training phase of CODEGENAPI, unlike the previous settings that force all files to have the same priority, we design a resampling strategy to enable high-quality python files to appear more frequently and vice versa. The strategy considers the star number of the repository, the unit test function rate of the code file, and the API rate of the code file. More details can be found in Appendix B.

## 4 Benchmark Construction

Private libraries are commonly used in practice, but few attempts have been made to evaluate the performance of generating code invoking private APIs.



To fill this gap, we craft three benchmarks, called TorchDataEval, MonkeyEval, and BeatNumEval. Each programming problem consists of context, target code, and the corresponding test cases.

To create a realistic benchmark for evaluating code generation for private library, we use TorchData, a Python library released just recently<sup>4</sup>. We carefully learnt the official API documentation of TorchData and made sure we were proficient in all APIs. Then, we manually created 50 programming problems based on the API usage examples in the documentation. Two volunteers with extensive experience in Python were invited to check the correctness of each problem. We control the difficulty of the programming problems by the number of APIs in the target code. The percentage of programming problems containing 1 API, 2 APIs, and more APIs is set to 6:3:1.

We also construct two benchmarks using pseudo private libraries, named MonkeyEval and BeatNumEval, each containing 101 programming problems. They are modified from PandasEval and NumpyEval, which were proposed for the public libraries Pandas and Numpy (Zan et al., 2022). In detail, we manually modified all library-related keywords in PandasEval and NumpyEval. For example, as shown in Figure 1, pandas is converted to monkey, dataframe is converted to knowledgeframe, and the API name isin is converted to iscontain. For more details on keyword conversion, please refer to Appendix C. To craft the API documentation for Monkey and BeatNum, we manually paraphrased the descriptions of all the new APIs to ensure that the pre-trained language models have never seen them.

## 5 Experiments

In this section, we conduct experiments to illustrate the superiority of our proposed framework.

### 5.1 Experimental Setup

**API Information.** As shown in the second column (APIs) in Table 1, there are four settings for prompting API information before the context:

- **No API:** there is nothing to be prompted;
- **Perfect:** the information of golden APIs in the target code is prompted;

<sup>4</sup>Our base model, CODEGEN, is pre-trained with GitHub data before 2021-10. TorchData was released after this time point and no code files using it are available on GitHub so far; hence we can consider it as a private library.

- **Top- $N$ :** the information of top  $N$  APIs retrieved by APIRetriever is prompted, where  $N \in \{1, 2, 3, 5\}$ ;
- **Human:** the information of the APIs chosen by users is prompted. In our experiments, we invited three volunteers who are programmers familiar with Python but without any background in our three benchmarks. As in Figure 4, they interacted with the APIRetriever and provided their choices for all programming problems. The final APIs are determined by voting on their choices.

**Baselines.** Our contributions can be reviewed in terms of both APIRetriever and APICoder. For APIRetriever, all models in the No API setting are our baseline, while we propose the Perfect, Top- $N$ , and Human settings. For APICoder, the main baseline is our base model, CODEGEN-MONO 350M (Nijkamp et al., 2022), in the same API information setting. We use CODEGEN for short in the following. In addition, we include advanced pre-trained code generation models that are comparable in parameter size: CodeT5 (Wang et al., 2021), CodeGPT (Lu et al., 2021), CodeClippy<sup>5</sup> and CodeParrot<sup>6</sup>. Codex 12B (Chen et al., 2021) is also used to show the performance of giant models.

**Evaluation Metrics.** Followed by Chen et al. (2021), we regard  $\text{pass}@k$  as our metric. For each programming problem, we sample  $n \geq k$  code snippets, and then count the number of correct ones  $c$ , where passing all test cases is considered as correct. If  $n - c < k$ , then  $\text{pass}@k$  equals 1; otherwise, equals  $1 - \prod_{i=n-c+1}^n (1 - \frac{k}{i})$ . In our experiments,  $k$  is set to one of  $[1, 10, 100]$  and  $n$  is set to 200.

**Implementation Details.** We implement our approach based on PyTorch (Paszke et al., 2019) and Huggingface’s transformers (Wolf et al., 2019). We use a dense retrieval toolkit<sup>7</sup> to train APIRetriever by setting the batch size to 10 per device, the learning rate to  $1e-5$ , the ratio of positive vs. negative samples to 1:8, and the vector dimensions  $z$  of  $\mathbf{p}$  and  $\mathbf{a}$  to 768. The model uses cross-entropy as the loss function and Adam (Kingma and Ba, 2014) as the parameters optimizer. It is trained for 100K steps on a cluster of 8 NVIDIA V100 GPUs with 32GB memory. The training time is about 3 days.

<sup>5</sup><https://github.com/CodedotAI/gpt-code-clippy>

<sup>6</sup><https://huggingface.co/lvwerra/codeparrot>

<sup>7</sup><https://github.com/luyug/Dense>

APICoder	APIs	TorchDataEval			MonkeyEval			BeatNumEval		
		pass@1	pass@10	pass@100	pass@1	pass@10	pass@100	pass@1	pass@10	pass@100
CodeT5 220M	Top-2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
CodeGPT 124M	Top-2	0.67	2.78	7.72	0.82	0.99	1.73	0.52	1.88	4.70
CodeClippy 125M	Top-2	0.04	0.39	2.75	0.10	0.76	1.86	0.03	0.33	2.11
CodeParrot 110M	No API	4.04	7.11	13.26	0.54	2.04	7.38	2.67	7.66	18.86
	Perfect	4.86 <sup>+0.82</sup>	8.88 <sup>+1.77</sup>	17.25 <sup>+3.99</sup>	2.39 <sup>+1.85</sup>	3.33 <sup>+1.29</sup>	9.99 <sup>+2.61</sup>	5.01 <sup>+2.34</sup>	11.30 <sup>+3.64</sup>	26.36 <sup>+7.5</sup>
	Top-1	4.02 <sup>-0.02</sup>	8.35 <sup>+1.24</sup>	18.17 <sup>+4.91</sup>	2.54 <sup>+2.00</sup>	3.43 <sup>+1.39</sup>	11.39 <sup>+4.01</sup>	4.32 <sup>+1.65</sup>	9.39 <sup>+1.73</sup>	19.91 <sup>+1.05</sup>
	Top-2	4.64 <sup>+0.60</sup>	8.96 <sup>+1.85</sup>	17.48 <sup>+4.22</sup>	1.52 <sup>+0.98</sup>	2.96 <sup>+0.92</sup>	9.32 <sup>+1.94</sup>	2.77 <sup>+0.10</sup>	8.57 <sup>+0.91</sup>	19.74 <sup>+0.88</sup>
	Top-3	4.00 <sup>-0.04</sup>	7.51 <sup>+0.40</sup>	15.13 <sup>+1.87</sup>	1.32 <sup>+0.78</sup>	3.16 <sup>+1.12</sup>	10.59 <sup>+3.21</sup>	1.69 <sup>-0.98</sup>	9.01 <sup>+1.35</sup>	19.90 <sup>+1.04</sup>
	Top-5	4.22 <sup>+0.18</sup>	7.51 <sup>+0.40</sup>	15.43 <sup>+2.17</sup>	0.99 <sup>+0.45</sup>	2.78 <sup>+0.74</sup>	11.76 <sup>+4.38</sup>	1.74 <sup>-0.93</sup>	8.11 <sup>+0.45</sup>	17.54 <sup>-1.32</sup>
	Human	4.01 <sup>-0.03</sup>	7.60 <sup>+0.49</sup>	14.47 <sup>+1.21</sup>	2.44 <sup>+1.90</sup>	3.62 <sup>+1.58</sup>	9.83 <sup>+2.45</sup>	5.23 <sup>+2.56</sup>	11.78 <sup>+4.12</sup>	22.81 <sup>+3.95</sup>
CODEGEN 350M	No API	6.72	15.71	22.00	0.95	4.90	8.89	5.15	11.96	18.79
	Perfect	9.84 <sup>+3.12</sup>	22.62 <sup>+6.91</sup>	34.00 <sup>+12.00</sup>	2.14 <sup>+1.19</sup>	6.41 <sup>+1.51</sup>	11.86 <sup>+2.97</sup>	9.47 <sup>+4.32</sup>	17.05 <sup>+5.09</sup>	28.67 <sup>+9.88</sup>
	Top-1	8.72 <sup>+2.00</sup>	19.22 <sup>+3.51</sup>	27.97 <sup>+5.97</sup>	2.22 <sup>+1.27</sup>	7.20 <sup>+2.30</sup>	12.85 <sup>+3.96</sup>	7.52 <sup>+2.37</sup>	15.25 <sup>+3.29</sup>	24.71 <sup>+5.92</sup>
	Top-2	7.52 <sup>+0.80</sup>	16.36 <sup>+0.65</sup>	26.00 <sup>+4.00</sup>	2.46 <sup>+1.51</sup>	6.35 <sup>+1.45</sup>	9.89 <sup>+1.00</sup>	6.65 <sup>+1.50</sup>	13.68 <sup>+1.72</sup>	22.74 <sup>+3.95</sup>
	Top-3	7.92 <sup>+1.20</sup>	18.65 <sup>+2.94</sup>	28.00 <sup>+6.00</sup>	2.02 <sup>+1.07</sup>	5.26 <sup>+0.36</sup>	8.89 <sup>+0.00</sup>	6.26 <sup>+1.11</sup>	16.12 <sup>+4.16</sup>	24.72 <sup>+5.93</sup>
	Top-5	6.08 <sup>-0.64</sup>	17.48 <sup>+1.77</sup>	25.95 <sup>+3.95</sup>	1.58 <sup>+0.63</sup>	5.45 <sup>+0.55</sup>	9.88 <sup>+0.99</sup>	6.34 <sup>+1.19</sup>	15.05 <sup>+3.09</sup>	21.76 <sup>+2.97</sup>
	Human	8.08 <sup>+1.36</sup>	19.85 <sup>+4.14</sup>	31.95 <sup>+9.95</sup>	2.14 <sup>+1.19</sup>	6.14 <sup>+1.24</sup>	11.86 <sup>+2.97</sup>	9.47 <sup>+4.32</sup>	17.12 <sup>+5.06</sup>	28.67 <sup>+9.88</sup>
CODEGENAPI 350M	No API	7.19	16.93	23.97	1.19	4.68	7.91	4.44	8.24	13.83
	Perfect	20.23 <sup>+13.04</sup>	33.37 <sup>+16.44</sup>	41.97 <sup>+18.00</sup>	4.59 <sup>+3.40</sup>	9.14 <sup>+4.46</sup>	13.85 <sup>+5.94</sup>	9.62 <sup>+5.18</sup>	16.51 <sup>+8.27</sup>	22.75 <sup>+8.92</sup>
	Top-1	12.89 <sup>+5.70</sup>	24.26 <sup>+7.33</sup>	31.97 <sup>+8.00</sup>	2.89 <sup>+1.70</sup>	8.28 <sup>+3.60</sup>	12.86 <sup>+4.94</sup>	6.61 <sup>+2.17</sup>	12.62 <sup>+4.38</sup>	17.80 <sup>+3.97</sup>
	Top-2	10.41 <sup>+3.22</sup>	23.50 <sup>+6.57</sup>	31.98 <sup>+8.01</sup>	3.41 <sup>+2.22</sup>	8.33 <sup>+3.65</sup>	11.87 <sup>+8.90</sup>	5.90 <sup>+1.46</sup>	11.79 <sup>+3.55</sup>	15.83 <sup>+2.00</sup>
	Top-3	10.49 <sup>+3.30</sup>	25.45 <sup>+8.52</sup>	35.98 <sup>+12.01</sup>	3.17 <sup>+1.98</sup>	7.51 <sup>+2.83</sup>	10.88 <sup>+2.97</sup>	5.11 <sup>+0.67</sup>	11.40 <sup>+3.16</sup>	15.82 <sup>+1.99</sup>
	Top-5	10.34 <sup>+3.15</sup>	23.04 <sup>+6.11</sup>	27.99 <sup>+4.02</sup>	1.94 <sup>+0.75</sup>	4.75 <sup>+0.07</sup>	7.91 <sup>+0.00</sup>	5.07 <sup>+0.63</sup>	9.64 <sup>+1.40</sup>	13.84 <sup>+0.01</sup>
	Human	15.57 <sup>+8.38</sup>	27.76 <sup>+10.83</sup>	33.97 <sup>+10.00</sup>	3.76 <sup>+2.57</sup>	8.32 <sup>+3.64</sup>	12.86 <sup>+4.95</sup>	9.39 <sup>+4.95</sup>	16.40 <sup>+8.16</sup>	23.74 <sup>+9.91</sup>
Codex 12B	No API	7.16	14.46	23.75	1.47	3.53	7.31	6.95	17.54	25.57
	Perfect	25.03 <sup>+17.87</sup>	51.26 <sup>+36.80</sup>	56.75 <sup>+33.00</sup>	3.58 <sup>+2.11</sup>	7.48 <sup>+3.95</sup>	12.61 <sup>+5.30</sup>	8.59 <sup>+1.64</sup>	23.75 <sup>+6.21</sup>	36.99 <sup>+11.42</sup>
	Top-2	17.98 <sup>+10.82</sup>	32.75 <sup>+18.29</sup>	41.51 <sup>+17.76</sup>	1.92 <sup>+0.45</sup>	5.91 <sup>+2.38</sup>	11.08 <sup>+3.77</sup>	9.54 <sup>+2.59</sup>	21.77 <sup>+4.23</sup>	32.45 <sup>+6.88</sup>

Table 1: Pass@ $k$ (%) results on the three benchmarks. The blue background means no API as extra prompt; the yellow background means perfect APIs as extra prompt; the write background means top-1, 2, 3, or 5 APIs retrieved by APIRetriever as extra prompt; and the purple background means the APIs chosen by human from top-5 of APIRetriever as extra prompt. Numbers in red and green indicate the absolute changes over no API setting.

For pre-training CODEGENAPI, we set the code block size to 1,024, the batch size to 4, the learning rate to  $5e-4$ , the gradient accumulation steps to 4, the weight decay to 0.1, and the warm up steps to 1,000. Noise APIs are added at a rate of 0.05. It is trained for 100K steps about 1.6 days on 8 32GB NVIDIA V100 GPUs. In all of our training phases, we use mixed precision FP16 to speed up. When generating code snippets using pre-trained models, we conduct various temperatures ranging from 0.1 to 1.0 with the interval of 0.1. All results are reported with the best values across these hyper-parameters.

## 5.2 Main Results

Table 1 summarizes the performance of our framework and all baselines on TorchDataEval, MonkeyEval, and BeatNumEval. Based on numerous experimental results, we derived plausible observations and valuable insights to answer the following research questions.

*“Is API information useful for private library oriented code generation?”* As we can see in Table 1, all models without prompting any APIs (the No API setting) achieve relatively poor per-

formance on all benchmarks. Especially, Codex 12B, a powerful code generation model with large parameters, can only achieve similar performance to CODEGEN and CODEGENAPI 350M in the No API setting. This indicates that even with gigantic models, the task of code generation with private libraries is extremely challenging. Encouragingly, with prompted API information (the Perfect, Top- $N$ , Human settings), both the off-the-shelf models (e.g., CodeParrot, CODEGEN, and Codex) and our continually pre-trained CodeGenAPI achieve consistent performance gains compared to those in the No API setting. Moreover, the more powerful the model itself in code generation (i.e., Codex 12B > CODEGEN 350M > CodeParrot 110M), the more benefits that API information can bring. For example, on TorchDataEval in the Perfect setting, Codex 12B brings pass@10 an absolute improvement of 36.89%, while CodeParrot 110M only improves 1.77%. This observation also suggests that prompting API information can unleash the potential of gigantic models towards invoking private APIs. All the above results prove the usefulness of API information for code generation for private libraries.

*“Is the APIRetriever effective in finding useful*

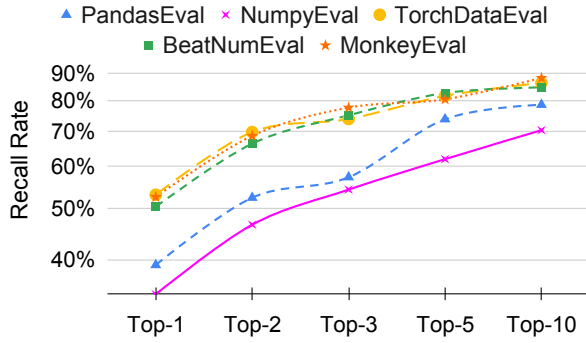


Figure 5: The recall rates of retrieved APIs.

API information?” All models in the Top- $N$  setting outperform the same models in the No API setting, suggesting that APIRetriever is able to find useful APIs. For a certain model, we observe that the Top-1/Top-2 settings usually perform better than the Top-3/Top-5 settings due to the fact that the latter introduces more noise APIs to the APICoder. In addition, involving humans (the Human setting) in the selection of APIs can further improve performance, suggesting the effectiveness of the human interaction we designed. Note that the Top- $N$  and Human settings are occasionally superior to the Perfect setting, which is reasonable because the noise APIs exist when training the model.

“Is the APICoder effective in invoking private APIs?” As shown in Table 1, off-the-shelf models like CODEGEN are capable of handling private library invocations. To seek more extraordinary performance, we continually pre-train CODEGEN and obtain a new model CODEGENAPI. We can observe that CODEGENAPI consistently outperforms its base model CODEGEN on TorchDataEval and MonkeyEval, which proves the effectiveness of CODEGENAPI. However, on BeatNumEval, CODEGENAPI is inferior to CODEGEN. After careful troubleshooting, we reveal that the process of continual pre-training aims to essentially learn how to invoke the correct APIs with maximum likelihood, while the key obstacle to using BeatNum modified from Numpy lies in the numerical calculation like ‘a[: ,None]+b\*2’, instead of invoking the correct APIs. Therefore, CODEGENAPI fails to yield benefits for BeatNumEval. Overall, APICoder has the capability to invoke private APIs.

### 5.3 Closer Analysis

We have demonstrated the effectiveness of our framework. In this subsection, we provide several closer analyses to inspire future work in this

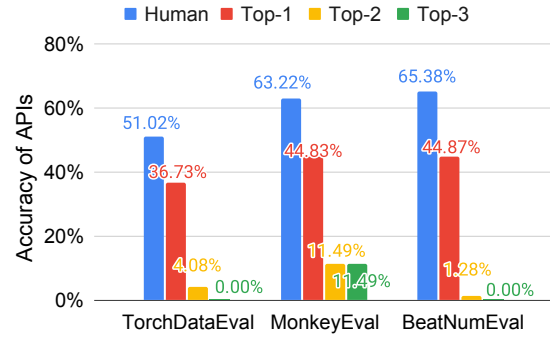


Figure 6: Accuracy of retrieved APIs.

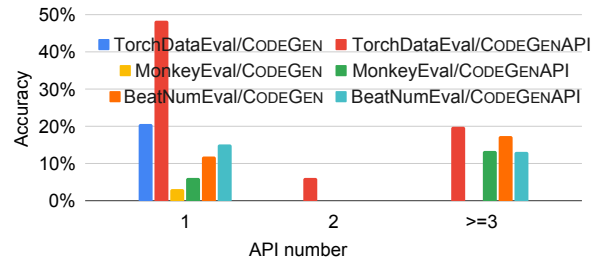


Figure 7: Accuracy of CODEGENAPI and CODEGEN with respect to the number of APIs. The problem is solved if one of 200 samples passes all test cases.

direction.

**Quality of Retrieved APIs.** Retrieving the correct APIs as prompts can enhance the code generation performance for private libraries, so we would like to evaluate the effectiveness of APIRetriever. Figure 5 shows the recall rate of APIRetriever on five benchmarks. We can see that the recall rates of top-5 are already high, demonstrating that it is reasonable to provide 5 API candidates for users to choose from. Furthermore, as shown in Figure 6, we analyze the accuracy of APIs chosen by users. We observe that it dramatically exceeds the accuracy of top 1, 2 or 3 APIs retrieved by APIRetriever. This suggests that it is feasible to involve humans in the retrieval of APIs.

**Different Difficulty.** We would like to explore the performance of CODEGENAPI on varying difficulty problems. So we calculate its accuracy across various numbers of APIs in target code  $y$ . Each benchmark is divided into 3 parts, according to the number of APIs. Figure 7 shows that CODEGENAPI outperforms CODEGEN by a large margin on the problems containing only one API. The trend still holds as the number of APIs increases. It demonstrates CODEGENAPI can boost the performance of generating code snippets using private library on varying difficulty.

APICoder	TorchDataEval			MonkeyEval			BeatNumEval		
	pass@1	pass@10	pass@100	pass@1	pass@10	pass@100	pass@1	pass@10	pass@100
CODEGENAPI	<b>10.41</b>	<b>23.50</b>	<b>31.98</b>	<b>3.41</b>	<b>8.33</b>	<b>11.87</b>	<b>5.90</b>	<b>11.79</b>	<b>15.83</b>
-w/ noise rate 0%	9.41	22.88	31.08	2.69	8.03	11.18	5.77	11.01	14.52
-w/ noise rate 10%	9.19	22.87	30.98	3.04	7.67	11.10	4.99	10.80	15.18
-w/ noise rate 20%	8.92	23.04	30.57	2.00	7.39	10.64	4.48	10.97	13.41
-w/o resampling	8.65	21.00	29.71	2.47	7.96	10.13	5.21	8.68	14.75

Table 2: Ablation studies for CODEGENAPI in the Top-2 setting (top 2 APIs provided by APIRetriever are prompted). The default setting of CODEGENAPI is to use the resampling strategy and a noise rate of 5%.

**Noise Rate.** A well-chosen noise rate can improve the robustness of CODEGENAPI against a variety of APIs. If we set the noise rate too large, it may change the original distribution of the code corpus, while too small will lose the capability to deal with noise APIs. The default noise rate is 5%, and we also try 0%, 10%, and 20%. As shown in Table 2, both too large and too small noise rates can degrade the performance.

**Resampling Strategy.** Making high-quality python files high-priority, and vice versa, is in line with our intuition. To demonstrate it, we remove the resampling strategy as mentioned in Section 3.3. As shown in Table 2, we observe a steady decline in performance on the three benchmarks. Such an observation demonstrates the effectiveness of the sampling strategy.

**CODEGENAPI for Public Library.** Technically speaking, CODEGENAPI also can be employed for generating code for public libraries. So, we do experiments on PandasEval and NumpyEval and show the results in Table 3. We find that the performance improvement of CODEGENAPI over the base model on public libraries is not as significant as on private libraries. One major reason is that the models have seen the public libraries during pre-training, so prompting API information yields limited benefit. We can see CODEGENAPI excels over CODEGEN when prompting perfect APIs. But when prompting top-2 APIs, the advantages of CODEGENAPI are not exhibited. This means that CODEGENAPI can also work on third-party public libraries, but it depends heavily on the performance of APIRetriever.

## 6 Related Work

### 6.1 Code Generation

Thanks to the recent development of pre-training techniques, a lot of pre-trained language models have been proposed for code-related tasks.

APICoder	APIs	PandasEval		
		pass@1	pass@10	pass@100
CODEGEN	No API	14.24	30.71	46.04
	Perfect	11.21	33.59	48.47
	Top-2	9.54	29.02	40.56
CODEGENAPI	No API	13.58	34.95	46.51
	Perfect	19.96	42.36	53.43
	Top-2	11.25	28.61	39.48
NumpyEval				
CODEGEN	No API	19.31	40.89	60.58
	Perfect	21.41	41.08	56.38
	Top-2	18.30	35.12	48.46
CODEGENAPI	No API	16.55	29.48	42.52
	Perfect	24.83	41.47	54.41
	Top-2	12.67	27.32	35.62

Table 3: Results of CODEGEN and CODEGENAPI on PandasEval and NumpyEval.

For example, CuBERT (Kanade et al., 2020), CodeBERT (Feng et al., 2020), GraphCodeBERT (Guo et al., 2020), CodeT5 (Wang et al., 2021), CodeGPT (Lu et al., 2021), PLBART (Ahmad et al., 2021), PyCodeGPT (Zan et al., 2022), CODEGEN (Nijkamp et al., 2022), Codex (Chen et al., 2021), AlphaCode (Li et al., 2022), and InCoder (Fried et al., 2022). Almost all of them focus on standalone code, while JigSaw (Jain et al., 2021) and CERT (Zan et al., 2022) are presented for generating code using public libraries. In this paper, we aim to generate code invoking private APIs, which is a common scenario in practice. It is more challenging because pre-trained language models have never seen any information about private libraries. As for benchmarks, HumanEval (Chen et al., 2021), APPs (Hendrycks et al., 2021), P3 (Schuster et al., 2021), MBPP (Austin et al., 2021), BIG-bench (Srivastava et al., 2022), and CodeContests (Li et al., 2022) were proposed to evaluate the performance of generating standalone code. GSM8K-Python (Cobbe et al., 2021) and MathQA-Python (Austin et al., 2021) were engaged in evaluating the capability of solving mathematical problems. PandasEval and NumpyEval (Zan et al., 2022) were released to evaluate the code generation for public library. We propose



three benchmarks, called TorchDataEval, MonkeyEval, and BeatNumEval, aiming to evaluate the performance of code generation for private library.

## 6.2 Retrieval-Based Generation

In the natural language field, retrieval-based generation is a hot topic. A lot of works (Izacard and Grave, 2020; Karpukhin et al., 2020; Qu et al., 2020; Xiong et al., 2020; Santhanam et al., 2021; Formal et al., 2022) have emerged under this topic. Therefore, we refer to the above methods and design our APIRetriever for private API retrieval. In the programming language field, there are also several attempts to use retrieval techniques, such as DEEPAPI (Gu et al., 2016), REDCODER (Parvez et al., 2021), ReACC (Lu et al., 2022), and DocCoder (Zhou et al., 2022). Our work is fundamentally different from them. They all aim to retrieve public code snippets or other resources on GitHub/StackOverFlow based on the user query, while our goal is to retrieve APIs from the API documentation of private library based on code comments. Besides, we design retrieval because we focus on private APIs, which have not been seen by the pre-trained generative language models.

## 7 Conclusion

In this paper, we propose a novel framework for code generation for private library. There are two modules: for a specific programming problem, APIRetriever first finds out the useful private APIs from API documentation, and then APICoder leverages these APIs to generate the code. We craft three benchmarks, including TorchDataEval, MonkeyEval, and BeatNumEval, for better evaluating private library oriented code generation. The experimental results and thorough analysis demonstrate the reasonableness and effectiveness of our framework. In future work, we would like to explore how to make better use of API documentation for code generation and improve the approach for real use when programming with private libraries.

## Limitations

While our proposed approach exhibits many advantages, it also has a few limitations. (1) As stated in Section 5.2, our approach that prompts APIs for programming problem relies heavily upon the code generation capacity of the language model itself. The more powerful the model itself, the more benefits the prompting APIs bring. Likewise,

we also find that if a model itself shows very poor performance, prompting APIs will not bring any benefit to it or even bring negative effects. (2) As the first navigator to explore code generation with private library, we have built three private libraries, but they all include a relatively small number of APIs (<200). With these APIs, our APIRetriever can exhibit decent performance. But we surmise that it may become more challenging for APIRetriever as the number of APIs increases. (3) It is extremely challenging to find a real private library and craft a benchmark like TorchDataEval. To evaluate our idea quickly and cost-effectively, besides TorchDataEval, we also crafted two pseudo private libraries that are modified from the existing public ones as mentioned in Section 4. Although we have done our best to preserve the two pseudo private libraries in line with the real private library, it may still pose some threats to the fair evaluation of code generation for private library. (4) We can see from Table 1 that most models with the Top-N setting fall behind the same model with the Perfect setting. Such observation demonstrates that APIRetriever we designed has a big room for improvement. (5) Our experiments show that our framework can enhance the quality of private library oriented code generation on Python. Limitations may exist when we generalize it to other programming languages such as Java, C, and C++ since the characteristics of libraries for different programming languages are slightly different.

## References

- Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Unified pre-training for program understanding and generation. In *North American Chapter of the Association for Computational Linguistics*, pages 2655–2668.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Jacob Hilton, Reiichiro Nakano, Christopher

- Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186.
- Zhangyin Feng, Daya Guo, Duyu Tang, Nan Duan, Xiaocheng Feng, Ming Gong, Linjun Shou, Bing Qin, Ting Liu, Daxin Jiang, et al. 2020. CodeBERT: A pre-trained model for programming and natural languages. In *Findings of the Conference on Empirical Methods in Natural Language Processing*, pages 1536–1547.
- Thibault Formal, Carlos Lassance, Benjamin Piwowarski, and Stéphane Clinchant. 2022. From distillation to hard negative sampling: Making sparse neural ir models more effective. *arXiv preprint arXiv:2205.04733*.
- Daniel Fried, Armen Aghajanyan, Jessy Lin, Sida Wang, Eric Wallace, Freda Shi, Ruiqi Zhong, Wen-tau Yih, Luke Zettlemoyer, and Mike Lewis. 2022. InCoder: A generative model for code infilling and synthesis. *arXiv preprint arXiv:2204.05999*.
- Xiaodong Gu, Hongyu Zhang, Dongmei Zhang, and Sunghun Kim. 2016. Deep api learning. In *ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 631–642.
- Daya Guo, Shuo Ren, Shuai Lu, Zhangyin Feng, Duyu Tang, et al. 2020. GraphCodeBERT: Pre-training code representations with data flow. In *International Conference on Learning Representations*.
- Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, et al. 2021. Measuring coding challenge competence with apps. In *Neural Information Processing Systems Datasets and Benchmarks Track*.
- Gautier Izacard and Edouard Grave. 2020. Leveraging passage retrieval with generative models for open domain question answering. *arXiv preprint arXiv:2007.01282*.
- Naman Jain, Skanda Vaidyanath, Arun Iyer, Nagarajan Natarajan, Suresh Parthasarathy, Sriram Rajamani, and Rahul Sharma. 2021. Jigsaw: Large language models meet program synthesis. *arXiv preprint arXiv:2112.02969*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547.
- Aditya Kanade, Petros Maniatis, Gogul Balakrishnan, and Kensen Shi. 2020. Learning and evaluating contextual embedding of source code. In *International Conference on Machine Learning*, pages 5110–5121.
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2004.04906*.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Yujia Li, David Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, et al. 2022. Competition-level code generation with alphacode. *arXiv preprint arXiv:2203.07814*.
- Shuai Lu, Nan Duan, Hojae Han, Daya Guo, Seung-won Hwang, and Alexey Svyatkovskiy. 2022. ReACC: A retrieval-augmented code completion framework. In *Association for Computational Linguistics*, pages 6227–6240.
- Shuai Lu, Daya Guo, Shuo Ren, Junjie Huang, Alexey Svyatkovskiy, Ambrosio Blanco, Colin Clement, et al. 2021. CodeXGLUE: A machine learning benchmark dataset for code understanding and generation. *arXiv preprint arXiv:2102.04664*.
- Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, et al. 2022. A conversational paradigm for program synthesis. *arXiv preprint arXiv:2203.13474*.
- Md Rizwan Parvez, Wasi Ahmad, Saikat Chakraborty, Baishakhi Ray, and Kai-Wei Chang. 2021. Retrieval augmented code generation and summarization. In *Findings of EMNLP*, pages 2719–2734.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, et al. 2019. *PyTorch: An imperative style, high-performance deep learning library*. *Neural Information Processing Systems*.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2020. RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering. *arXiv preprint arXiv:2010.08191*.
- Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2021. Colbertv2: Effective and efficient retrieval via lightweight late interaction. *arXiv preprint arXiv:2112.01488*.
- Tal Schuster, Ashwin Kalyan, Oleksandr Polozov, and Adam Tauman Kalai. 2021. Programming puzzles. *arXiv preprint arXiv:2106.05784*.
- Eric Snodgrass and Soon Winnie. 2019. Api practices and paradigms: Exploring the protocological parameters of apis as key facilitators of sociotechnical forms of exchange. *First Monday*, 24(2).

Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, et al. 2022. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint arXiv:2206.04615*.

Priyan Vaithilingam, Tianyi Zhang, and Elena L Glassman. 2022. Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–7.

Yue Wang, Weishi Wang, Shafiq Joty, and Steven CH Hoi. 2021. CodeT5: Identifier-aware unified pre-trained encoder-decoder models for code understanding and generation. In *Empirical Methods in Natural Language Processing*, pages 8696–8708.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul Bennett, Junaid Ahmed, and Arnold Overwijk. 2020. Approximate nearest neighbor negative contrastive learning for dense text retrieval. *arXiv preprint arXiv:2007.00808*.

Daoguang Zan, Bei Chen, Dejian Yang, Zeqi Lin, Minsu Kim, Bei Guan, Yongji Wang, Weizhu Chen, and Jian-Guang Lou. 2022. CERT: Continual pre-training on sketches for library-oriented code generation. In *The 2022 International Joint Conference on Artificial Intelligence*, pages 3653–3660.

Shuyan Zhou, Uri Alon, Frank F Xu, Zhengbao Jiang, and Graham Neubig. 2022. DocCoder: Generating code by retrieving and reading docs. *arXiv preprint arXiv:2207.05987*.

## A Collection of API Documentation

We aim to use data from public libraries for training and generalize the models to private libraries. Thus, we crawled the API documentation of the 31 most popular public libraries in Python. Table 4 summarizes the number of APIs we extracted for each library.

## B Resampling Strategy

The resampling strategy allows high-quality python files to be more frequently sampled, and vice versa. So the resampling weight ( $w$ ) of each python file is defined in the following aspects: the star number of the corresponding repository ( $N_{\text{star}}$ ), the unit test function rate ( $R_{\text{ut}}$ ) that is the number of unit test functions divided by the number of all functions, the number of API name ( $N_{\text{api}}$ ) in the file, and the

number of APIs ( $M_{\text{api}}$ ) considering one API name may match multiple APIs. Formally, the strategy can be formulated as follows:

$$\begin{aligned} w_{\text{star}} &= 1.0 + \log(N_{\text{star}} + 1) \cdot \text{clip}_{(0)}^{\left(\frac{5}{0}\right)} \times 0.2, \\ w_{\text{ut}} &= (0.5 + (1 - R_{\text{ut}})) \cdot \text{clip}_{(0)}^{\left(\frac{1}{0}\right)}, \\ w_{\text{api}} &= 5.0 - \log\left(\frac{M_{\text{api}}}{N_{\text{api}}}\right) \cdot \text{clip}_{(0)}^{\left(\frac{5}{0}\right)} \times 0.2, \\ w &= w_{\text{star}} \times w_{\text{ut}} \times w_{\text{api}}, \end{aligned} \quad (1)$$

where  $\text{clip}_{(x)}^{\left(\frac{y}{x}\right)}$  limits the value to  $[x, y]$ .

## C Keywords Conversion from Public Library to Private Library.

As mentioned in Section 4, we convert the public library benchmarks (PandasEval and NumpyEval) to the private library benchmarks (MonkeyEval and BeatNumEval) by manually modifying all public library-related keywords. In Table 5, we list all the keywords before and after the conversion.

Pandas	NumPy	sklearn	PyTorch	TensorFlow	Django	selenium	Matplotlib	Flask	SciPy	Seaborn
7,094	12,085	53,166	124,902	32,116	24,375	4,842	439,913	31,867	153,359	161,477
NLTK	BeautifulSoup	pygame	PIL	jieba	Gensim	spaCy	transformers	fairseq	SQLAlchemy	Scrapy
206,816	22,519	70,396	127,212	26,620	37,331	239,945	652,913	158,721	54,765	3,537
AllenNLP	datasets	tokenizers	MXNet	imageio	pytest	MetPy	ansible	requests		
276,088	136,843	195	142,070	175,878	1,047	27,429	40,839	39,333		

Table 4: The number of APIs in the 31 public libraries we crawled.

PandasEval-MonkeyEval						
isnull	mean	pandas	dataframe	df	isin	pd
ifnull	average	monkey	knowledgeframe	kf	incontain	mk
tolist	apply	to_numeric	dropna	append	tail	copy
convert_list	employ	to_num	sipna	adding	last_tail	clone
innull	astype	select_dtypes	iterrows	min	max	map
isnone	totype	choose_dtypes	traversal	get_min	get_max	mapping
last	shift	merge	value_counts	rename_axis	reset_index	sample
final_item	shifting	unioner	counts_value_num	renaming_axis	reseting_index	sample_by_num
concat	to_dict	cumsum	sort_index	to_string	drop_duplicates	duplicated
concating	convert_dict	cumulative_sum	sorting_index	convert_string	remove_duplicates	duplicated_values
round	format	to_pydatetime	div	ceil	assign	intersection
value_round	formatting	convert_pydatetime	division	ceiling	allocate	interst
drop	Series	ravel	any	fillna	all	Pandas
sip	Collections	flat_underlying	whatever	fillnone	total_all	Monkey
reindex	get	std	rename	sum	unique	to_datetime
reindexing	getting	standard	renaming	total_sum	distinctive	convert_datetime
applymap	sort_values	DataFrame	groupby	nlargest	replace	len
conduct_map	sort_the_values	KnowledgeFrame	grouper	nbiggest	replacing	length
head	series	isna				
header_num	collections	ifna				
NumpyEval-BeatNumEval						
to_numpy	ndarray	array	transpose	numpy	Numpy	np
to_beatnum	ndnumset	numset	switching_places	beatnum	Beatnum	bn
column_stack	concatenate	slice	sum	imag	abs	real
stack_col	connect	piece	total_count	imaginary	absolute	reality
fill_diagonal	all	fromstring	in1d	mean	where	std
pad_diagonal	total	come_from_str	intersection1dim	average	filter_condition	standard_op
add	histogram	fromarrays	reshape	filled	stack	cumsum
add_concat	hist_operation	come_from_arrays	change_shape_to	masked_fill	pile_operation	cumulative_sum
astype	arange	setxor1d	compressed		argmin	argmax
convert_type	arr_range	seting_exclusive_or_one_dim	remove_masked_data		get_argmin_value	get_argmax
vstack	squeeze	hstack	asarray	repeat	vectorize	split
vertical_stack	sqz	horizontal_stack	asnumset	duplicate	vectorisation	sep_split
diff	unique	unravel_index	flatten	norm	delete	ones
difference	uniq	convert_index_or_arr	convert_into_one_dim	normlization	remove_operation	create_ones
append	any	logical_and	bincount	isnan	argpartition	ravel
apd	any_condition	logic_and_element_wise	binoccurrence	ifnan	perform_partition	asview
array_split	inv	insert	searchsorted	min	max	full
split_array	inverse	stick	find_sorted	get_min	get_max	full_value_func

Table 5: The keywords of converting PandasEval to MonkeyEval, and NumpyEval to BeatNumEval. The grey background means the original keywords, and the white background means the converted ones.