



Optimizing Large Language Models for OpenAPI Code Completion

Bohdan Petryshyn  and Mantas Lukoševičius 

Abstract—Recent advancements in Large Language Models (LLMs) and their utilization in code generation tasks have significantly reshaped the field of software development. Despite the remarkable efficacy of code completion solutions in mainstream programming languages, their performance lags when applied to less ubiquitous formats such as OpenAPI definitions. This study evaluates the OpenAPI completion performance of GitHub Copilot, a prevalent commercial code completion tool, and proposes a set of task-specific optimizations leveraging Meta’s open-source model Code Llama. A semantics-aware OpenAPI completion benchmark proposed in this research is used to perform a series of experiments through which the impact of various prompt-engineering and fine-tuning techniques on the Code Llama model’s performance is analyzed. The fine-tuned Code Llama model reaches a peak correctness improvement of 55.2% over GitHub Copilot despite utilizing 25 times fewer parameters than the commercial solution’s underlying Codex model. Additionally, this research proposes an enhancement to a widely used code infilling training technique, addressing the issue of underperformance when the model is prompted with context sizes smaller than those used during training. The dataset, the benchmark, and the model fine-tuning code are made publicly available.

Index Terms—Large Language Models, Code Llama, GitHub Copilot, OpenAPI, code completion, code generation, fine-tuning, prompt-engineering, benchmarking.

I. INTRODUCTION

IN the world of constantly growing digitalization and automation, Application Programming Interfaces (APIs) play a central role in enabling communication between software systems. The OpenAPI specification is a widely adopted standard for defining APIs in both machine and human-readable format [1]. An OpenAPI definition describes all the requests, their parameters, and responses supported by the API. Since the format is machine-readable, a variety of applications were developed to leverage OpenAPI definitions for generating API documentation and client libraries, performing automated testing, or even automatically integrating with the API using the recent advancements in Artificial Intelligence (AI) and Large Language Models (LLMs).

Although OpenAPI definitions can be fully generated based on the existing API implementation, writing a definition manually before implementing the API is a common practice that enables the development of higher-quality APIs. Writing OpenAPI definitions by hand is often a time-consuming and error-prone process due to the verbosity of the format. Modern polyglot code completion solutions like GitHub Copilot [2]

B. Petryshyn and M. Lukoševičius are with Faculty of Informatics, Kaunas University of Technology, 44249 Kaunas, Lithuania (e-mail: {bohdan.petryshyn, mantas.lukosevicius}@ktu.edu).

and Amazon Web Services (AWS) Code Whisperer [3] have shown that LLMs can be used to increase developer productivity by providing code suggestions based on the surrounding context. These solutions are integrated into ubiquitous Integrated Development Environments (IDEs) and are capable of suggesting high-quality code completions in popular programming languages such as Python, JavaScript, and Java. However, the performance of these solutions in less popular formats like OpenAPI leaves much to be desired.

This research evaluates the OpenAPI completion performance of one of the most widely used commercial code completion solutions, GitHub Copilot, and investigates the feasibility of developing a better-performing solution based on a modern open-source LLM Code Llama [4]. To achieve this, a semantics-aware OpenAPI completion benchmark is proposed and used to evaluate the performance of both solutions. An extensive set of experiments is performed to analyze the impact of a variety of variables, prompt engineering, and fine-tuning techniques on the performance of the Code Llama model. The results of the experiments reveal that the original version of the model, when used with optimal prompt format is already capable of outperforming GitHub Copilot in OpenAPI completion task. The fine-tuned version of the model further widens the performance gap with the commercial solution. To the best of the authors’ knowledge, this is the first work investigating the Code Llama fine-tuning for code infilling in detail. The OpenAPI completion benchmark as well as the fine-tuning pipeline were made publicly available at <https://github.com/BohdanPetryshyn/openapi-completion-benchmark> and <https://github.com/BohdanPetryshyn/code-llama-fim-fine-tuning>.

II. STATE OF THE ART IN CODE GENERATION

As technology advances, the landscape of code generation continues to evolve, with transformer-based models and integrated solutions revolutionizing the field. This section delves into the state-of-the-art solutions for code generation, explores integrated solutions that streamline the development process, assesses methodologies for evaluating the quality of generated code, and highlights the challenges in OpenAPI autocompletion as a specific case of code generation.

A. Transformer Models for Code Generation

Transformer-based models are the most widely used type of models for Natural Language Processing (NLP) and code generation in particular [5]. The original transformer model was introduced in the famous “Attention Is All You Need” article in 2017 [6]. The model was based on the trainable

neural attention mechanism proposed by Bahdanau Dzmitry et al. [7] earlier in 2014. This mechanism presented a solution to several problems of Recurrent Neural Network (RNN)-based models that were one of the primary methods for NLP at the time [8].

One of the main shortcomings of RNN-based models is the difficulty of capturing long-term dependencies in the input due to vanishing and exploding gradient problems and limited hidden state size [9]. The attention mechanism allows the model to capture dependencies between any two tokens in the input sequence, making the likelihood of a long-term dependency being captured the same as the likelihood of any short-term dependency.

Another major drawback of RNN-based models is the sequential nature of processing which makes it impossible to efficiently parallelize the computation. A significant part of the attention computation can be done independently for each token in the input sequence which dramatically speeds up the processing on modern highly parallel hardware like Graphics Processing Units (GPUs) and Tensor Processing Units (TPUs).

Transformer model architecture can consist of an encoder, decoder, or combination of both. Each of these parts leverages the attention mechanism in different ways to achieve its objectives. The role of an encoder is to analyze the input sequence and encode it into an internal representation that captures dependencies between the input tokens. This is done using a mechanism called self-attention where the attention score is calculated between each pair of the input tokens and a new value is calculated for each token as a sum of all the values in the input sequence weighted by the attention score. The role of a decoder is to autoregressively generate the output sequence leveraging a mechanism called encoder-decoder attention. In this process, each output token is generated as an attention-weighted sum of the encoded input tokens and the previously generated output tokens.

The encoder-only models like BERT from Google [10], RoBERTa [11], AIBERT [12], and DeBERTa [13] only consist of an encoder and do not produce sequences as the output. Instead, these models only produce encoded representations of the input sequence that can be used for tasks like input classification, sentiment analysis, or feature extraction. The objective used during training is often not well-aligned with the inference objective, so these models usually require fine-tuning on the specific target task.

The encoder-decoder models like the original transformer suggested by Ashish Vaswani et al. [6], T5 [14], and BART from Meta [15] incorporate both encoder and decoder in their architecture. These models are capable of sequence-to-sequence (seq-to-seq) generation and are commonly used for machine translation and summarization where the input and the expected output sequences can be clearly defined.

The decoder-only models like the Generative Pre-trained Transformer (GPT) family of models from OpenAI [16], PaLM from Google [17], Chinchilla from Deepmind [18], and Llama from Meta [19], [20] only consist of a decoder. This preserves the ability to analyze the input sequence and generate a sequence as the output but also eliminates the boundary between the input and the output. The latter characteristic

makes this type of models more generic and often abolishes the need for fine-tuning after the initial training, as most of the tasks can be represented as a natural language seq-to-seq function [21]. Compared to the encoder-decoder architectures, which are also capable of seq-to-seq generation, decoder-only architectures further streamline the training process by removing the need to define the input and output formats for the model to operate on. These models are often used for language modeling and text generation but are also capable of input classification, sentiment analysis, and machine translation if invoked with a corresponding prompt. The capabilities of the large decoder-only models can be further expanded with prompting techniques like few-shot prompting [22].

Code generation is a broad term that is used to describe a wide range of activities. One such activity is code solution generation for a natural language problem definition. This task is usually approached as a seq-to-seq generation problem where the task definition is the input, and the code solution is the output of the model. One of the successful solutions in this area is AlphaCode published by DeepMind [23]. The model leverages the encoder-decoder architecture. It is pre-trained on a variety of public GitHub repositories in different programming languages and fine-tuned on a curated dataset of competitive programming tasks. As stated by the authors, the model achieved a ranking of top 54.3% among more than 5 000 participants on the Codeforces platform.

Code generation from a natural language task definition might be useful for no-code or low-code engineers but code completion, where the model generates continuation for an incomplete code input, represents everyday programmers' work much closer. Models like Codex from OpenAI [24], CodeGen from Salesforce [25], [26], and Code Llama from Meta [4] leverage decoder-only architectures for autoregressive code modeling. The fill-in-the-middle (FIM) also referred to as infilling transformation and training objective allows decoder-only models to include the context from both sides of the generated part, removing the limitation of strict left-to-right generation while maintaining the ability to output long sequences of tokens [27]. This approach not only enables context-aware code completion but also supports natural language task definition to some degree. These models can be prompted with a natural language code comment to generate an implementation. This can be explained by a weak natural language signal that might be captured during training on large corpora of code with comments and code-related discussions [25].

B. Transformer Model Performance Optimization

Prompt engineering and model fine-tuning are two primary methods for large language model performance optimization. A common workflow for improving the performance of a transformer-based model in a specific domain is to use prompt engineering methods to achieve the best possible performance with the base model and then fine-tune the model on a task-specific dataset if further performance improvement is desirable. This research will utilize both prompt engineering and fine-tuning to improve the performance of the Code Llama model on the OpenAPI definitions generation task.

Prompt engineering is based on the idea that the model can be guided to generate more accurate output by optimizing the prompt in a series of human-insight-driven experiments [28]. Common prompt engineering techniques for instruction-tuned models include few-shot prompting [22], chain-of-thought prompting [29], and generated knowledge prompting [30]. For autoregressive code generation, prompt engineering techniques include embedding additional context to the prompt in the form of comments or other language-specific constructs [25] and optimizing the prompt structure and size to maximize the model’s performance and resource efficiency.

Fine-tuning is the second primary method for improving the performance of large language models. The idea behind this method is that the model that has already been trained for a similar or broader task can be further fine-tuned for a specific task. The fine-tuning process is similar to training, but it is performed on a pre-trained model and requires significantly less data and computational resources. Two fine-tuning approaches are commonly used: full fine-tuning and Parameter-Efficient Fine-Tuning (PEFT) [31].

Full fine-tuning is performed by continuing the training process of an already pre-trained model. During this process, all the model parameters are updated just like during the original training. This means that the hardware requirements for the full fine-tuning are also similar. This approach demonstrates the potential for achieving the best results, but it also requires the most computational resources and data. The ability to change all the model parameters during full fine-tuning can also lead to catastrophic forgetting, a training failure mode where the model loses the ability to perform the original task or stops generating adequate outputs altogether.

Parameter-efficient fine-tuning, as an alternative to full fine-tuning, is performed by freezing most of the pre-trained model parameters and only training a relatively small number of additional parameters. This approach requires much less computational resources and tends to yield comparable results to the full-finetuning [31]. The PEFT approach is also considered to be less prone to catastrophic forgetting as the majority of the pre-trained model parameters are left unchanged during the fine-tuning process. Notable PEFT techniques include adapter tuning [32] which inserts additional adapter layers between the original model layers and only updates them during training, BitFit [33] which only trains the bias-terms of the original model, and Low-Rank Adaptation (LoRA) [34] which decomposes the original parameters into low-rank matrices and only trains them. The LoRA method claims to reduce the number of trainable parameters by 10 000 times and memory requirements by 3 times while maintaining or surpassing the performance of full fine-tuning when applied to models like GPT-3 and RoBERTa.

C. Integrated Code Generation Solutions

Integrated code generation and completion solutions are end-to-end systems that combine code generation models, prompt engineering techniques, developer experience optimizations, and IDE integrations to solve a real-world problem of improving developer productivity. These solutions are often

proprietary, so the details of their implementation are not available to the public. However, some solutions are open-source and can be analyzed in detail.

The most notable commercial solution is GitHub Copilot. Developed by GitHub, OpenAI, and Microsoft, this code assistant is based on the Codex [24] model. The solution is integrated into several IDEs like Visual Studio Code, NeoVim, and JetBrains IDEs. As discovered by a reverse engineering effort [35], Copilot’s prompt-building approach is based on several language-specific heuristics. For example, for programming languages like Python, a list of imported libraries and declared functions within the file is extracted and always included in the prompt. The code assistant underwent manual evaluation in several studies [36]–[39] demonstrating decent performance in benchmarks like HumanEval [24] and solving moderate complexity tasks at platforms like LeetCode. No studies were found that would evaluate the solution performance in OpenAPI completion. In this research, GitHub Copilot is evaluated as the baseline solution for the OpenAPI definitions completion task.

Other proprietary solutions worth mentioning are Tabnine [40] which emphasizes the security of the user’s data and the code base, AWS Code Whisperer [3] which demonstrates additional expertise in AWS services and client libraries and performs additional checks for license compliance and common security vulnerabilities in the generated code, and Sourcegraph Cody [41] which demonstrates awareness of the whole user’s code base.

Continue [42] is an example of an open-source integrated code completion solution that aims to be model-agnostic and supports multiple models with the most commonly used being Code Llama Instruct running in a local or a cloud-based LLM server. Other open-source models supported by the solution include WizardCoder [43], Phind Code Llama [44], Mistral [45], CodeUp [46], and Zephyr [47]. Commercial models like GPT-3, GPT-4, and PaLM-2 [48] can be integrated as well. During code assessment, it was discovered that Continue prompts the model in an instructional manner by providing it with the surrounding code and asking to generate the missing part. This explains the fact that the solution only works with models that support instruction execution. Another open-source solution, LLM-vscode [49], is a Visual Studio Code extension developed by Hugging Face for demonstrational purposes. The extension can be used with many models but in contrast to Continue, it supports both instruction execution and code infilling models. LLM-vscode is usually used with models like Code Llama, Phind Code Llama, and WizardCoder.

Another research on applying LLMs to a real-world code completion use case was published by Microsoft in 2020 [50]. Apart from evaluating transformer optimization techniques for code like applying tokenization methods that can overcome the closed vocabulary problem, this research describes multiple developer experience optimization techniques like client-side tree-based caching, parallel beam search decoder, and compute graph optimizations.

D. Code Generation Evaluation

Evaluation of LLMs has always been a challenging task in the field of NLP. The evaluation methods have been evolving in step with the abilities of the models and the complexity of the tasks they perform.

As the first applications of language models started appearing in the field of natural language modeling and machine translation, the evaluation methods were designed to measure the quality of the generated text by comparing it to the expected output. The most widely used metrics for this purpose are BLEU [51], ROUGE [52], and METEOR [53]. These metrics are based on calculating the similarity by counting the overlapping n-grams or measuring the longest common sequences.

Code generation can not be reliably evaluated with machine translation metrics due to the fundamental differences between natural language and code. The code is structured and possesses a more rigid syntax than natural language, making it more sensitive to small character-level changes. Another source of complexity in code evaluation is the fact that a single code problem can have multiple equally correct solutions. This makes it impossible to define an exhaustive set of expected outputs for a given input.

One of the current state-of-the-art evaluation methods for code generation is the HumanEval benchmark [24]. Originally proposed by OpenAI, this benchmark was designed to evaluate Python code generation by providing a set of problems each with a description and a set of automated test cases. Later, the method was extended to support 18 additional programming languages by the MultiPL-E benchmark [54]. These methods take language semantics into account but are only applicable to a limited set of programming languages and do not support specific code semantics like OpenAPI definitions.

The OpenAPI specification is not a programming language but a data format specification. This means that the existing code generation evaluation methods can not be easily extended to support OpenAPI completion evaluation. To enable the development of OpenAPI completion solutions, a new benchmark has to be proposed that supports the specific semantics of OpenAPI definitions.

III. METHODS AND DATASETS

As outlined in the previous section, OpenAPI completion requires specialized evaluation and optimization methods. This section describes the proposed OpenAPI completion benchmark, the foundational model selected for the study as well as the prompt engineering and fine-tuning techniques used to optimize the model’s performance.

A. OpenAPI Completion Evaluation Criteria and Benchmarking

As discussed in the previous section, existing text or code generation evaluation methods can not be considered optimal for OpenAPI generation evaluation. These methods are not semantics-aware as in the case with machine language translation metrics or do not support specific OpenAPI semantics

as in the case with the widely used HumanEval and MultiPL-E benchmarks.

The proposed benchmark aims to closely resemble the real-world usage of the solution in terms of input modeling and output evaluation. Figure 1 outlines the steps for a single evaluation iteration. The masking step, where part of the input definition is replaced with a special mask token, is responsible for modeling the input. To equalize the evaluation conditions for all the solutions, the same set of masked test cases is used for all the evaluations. At the completion step, the evaluated model is used to replace the mask token with the generated OpenAPI definition. The evaluation step, where the completed definition is compared to the original, is responsible for the completion result evaluation.




Fig. 1. Proposed OpenAPI completion benchmark pipeline.

The masking step can be seen as splitting the original OpenAPI definition into two parts (prefix and suffix) and removing some amount of the definition from in between. This way, the input modeling is defined by the position of the prefix-suffix split and the size of the masked part between them. When editing an OpenAPI definition, the user does not always add new endpoints or parameters and can expect to get a completion suggestion even when they have not finished typing the previous word or syntactic construct. This means that the prefix and suffix split should be randomly selected and not tied to a newline or a keyword. Conversely, the end of the masked sequence should be on a newline, as it is reasonable to assume that users are more likely to edit a document at the end of a line rather than in the middle of it when describing new API operations or adding parameters to existing ones. In this study, the masked part is always 10 lines long (till the end of the 10th line), which is a reasonable size for a single parameter or a response definition.

The following metrics are measured at the evaluation step:

- 1) Correctness – how often the generated OpenAPI definition is semantically identical to the original one;
- 2) Validity – how often the generated OpenAPI definition is syntactically valid;
- 3) Speed – how fast the solution generates the OpenAPI definition.

To be semantically identical, two snippets of OpenAPI definition do not have to be identical at the character level. The YAML and JSON specifications allow for different ways of representing the same data. For example, keys can be optionally quoted, and the order of keys is not significant. The OpenAPI specification allows for even more flexibility with the support of references and other advanced syntactic features. Additionally, OpenAPI definitions contain subjective information such as descriptions, titles, summaries, and examples of data. These fields are not required to be identical for the

generated OpenAPI definition to be considered semantically identical to the expected one.

To evaluate correctness, the benchmark uses the Oasdiff tool [55] which is capable of calculating the semantic difference between two OpenAPI definitions. The calculated difference is not affected by the YAML, JSON, or OpenAPI specification flexibility described above. To tolerate subjectiveness in the definition, the following set of heuristics is applied to the calculated difference:

- 1) The definitions are semantically identical if the calculated difference is empty;
- 2) The definitions are semantically identical if the calculated difference only contains insignificant changes;
- 3) Changes in the following subjective fields are considered insignificant: description, summary, title;
- 4) Changes in the examples of data are considered insignificant if they do not change the structure and type of the data.

1) **Evaluation dataset:** The dataset for the benchmark was collected from the APIs-guru definitions directory [56] which is a weekly-updated collection of public OpenAPI definitions. The directory contains more than 4 000 definitions in YAML format.

Ideally, the definitions have to be unseen by the model during training or fine-tuning. However, as the Code Llama model authors do not reveal the details of the training dataset, it is not possible to guarantee that the definitions were not used during training. Only the newest definitions in the directory were selected for the evaluation dataset to minimize the chance of being seen by the model. The definitions were also checked for presence in The Stack [57] dataset which is one of the biggest open code datasets collected from GitHub. The definitions that were present in The Stack dataset were excluded from the evaluation dataset.

Another requirement for the definitions in the evaluation dataset is that they have to represent different domains and styles. The APIs-guru collection contains definitions for APIs from a variety of domains such as payment systems, cloud services, and electronic commerce. The directory also contains definitions from many producers such as Amazon, Meta, and Google. The definitions were selected from different domains and producers to ensure that the benchmark is not biased toward a specific field or style of OpenAPI definitions. Automatically generated definitions were detected and excluded from the evaluation dataset as they do not represent the future real-world use of the solution.

Only OpenAPI definitions larger than 3 000 lines were selected for the evaluation dataset with the preference given to definitions larger than 20 000 lines when all the other criteria are equal. Large definitions are preferred since they are more likely to contain complex syntactic constructs and semantic dependencies. Small definitions that fit into the context size of the model are often too trivial to complete correctly and do not present any practical value for the benchmark.

Ten OpenAPI definitions satisfying the above criteria were selected for the evaluation dataset (Table I). Considering the size of the selected definitions, multiple test cases could be extracted from each definition. The test cases were extracted

by randomly selecting a position in the definition and masking the rest of the line and ten following lines. The evaluation objective was to generate the infilling for the masked part so that the resulting definition is semantically identical to the original one.

TABLE I
OPENAPI DEFINITIONS INCLUDED IN THE EVALUATION DATASET

API Title	Producer	Doman	Size, lines
Adyen API	Adyen	Payments	20 359
Google Video API	Google	Video	16 008
Mux API	Mux	Video	7 128
Rubrik API	Rubrik	Security	43 540
Severa Public API	Visma	Project mngmt.	36 871
Sinao API	Sinao	Accounting	12 215
SpaceTraders API	SpaceTraders	Gaming	3 121
The Racing API	The Racing API	Sports	9 650
VPC Lattice API	AWS	Cloud	6 709
Visma e-conomic	Visma	Accounting	10 162

The benchmark was implemented with the idea of being fully automated and reproducible. This made it possible to easily evaluate the performance of many solutions and configurations. However, the fact that commercial solutions like GitHub Copilot are only available as IDE plugins and require manual interaction with the IDE imposed a limitation on the size of the evaluation dataset. To keep the manual evaluation time reasonable while still providing a representative evaluation, the dataset size was constrained to 100 test cases in total (10 test cases from each of the 10 definitions) which is comparable to the size of the HumanEval benchmark with 164 code problems.

B. Code Llama Model

Training LLMs from scratch is a highly computationally expensive process [18]. For practical reasons, it is often more efficient to incorporate a generic pre-trained model that can be fine-tuned for a specific task or used in a zero-shot manner. For this research, the Code Llama family of models from Meta [4] was selected as the foundational model for the OpenAPI completion task.

The Code Llama model is trained on an unbound set of programming languages which not only ensures its ability to generate OpenAPI definitions but also allows it to generate code samples that are commonly found in OpenAPI definitions in a variety of programming languages. The model supports large context sizes (up to 100 000 tokens) which is important for the OpenAPI completion task as the definitions can reach tens of thousands of lines in size.

This code generation model is based on another widely adopted natural language model from the same company – Llama 2. The model was trained for code infilling on an additional 500 billion tokens of code. Figure 2 outlines the training pipeline of the Code Llama model family. The model was further trained for Python specialization and fine-tuned for instruction execution, but for the sake of this research, the base Code Llama model is used. The 70-billion-parameter version of the model was trained and released months after

the original Code Llama family and might follow a slightly different training pipeline.

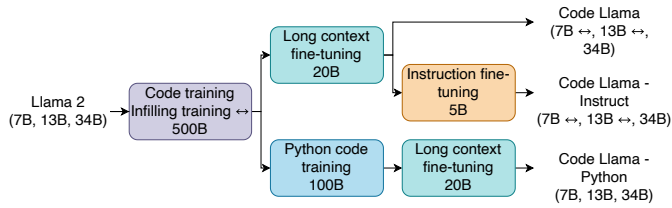


Fig. 2. Code Llama model family training pipeline, adapted from [4].

After the recent release of the biggest version, the model is now available in four sizes: 7, 13, 34, and 70 billion parameters. This study is focused on the smaller 7 and 13-billion-parameter versions of the model because the bigger 34-billion-parameter version does not support the infilling mode of operation (“↔” in Figure 2, discussed below) and the computational requirements of the 70-billion-parameter version exceed the capabilities of the evaluation platform.

The Code Llama model can operate in two modes: autoregressive code generation, where the model takes a code snippet and generates a continuation for it token by token, and code infilling, where the model takes two snippets of code and generates the completion that would fill the gap between them. The mode of operation is determined by the format of the prompt. For code generation, the model should be prompted with the code snippet as is. For code infilling, the two snippets (prefix and suffix) should be formatted in a special way that was defined during the model training.

Following the training methods originally described in the work of Bavarian et al. [27], the Code Llama model supports two prompt formats for code infilling: Prefix Suffix Middle (PSM) and Suffix Prefix Middle (SPM). The study has shown that the SPM format produces slightly better results even when the model is trained on the PSM format. The SPM format is also expected to result in a higher key-value cache hit rate as the suffix is less likely to change in a real-world code completion scenario. Both formats are evaluated in this research.

The Code Llama model introduces four special tokens to the existing Llama 2 tokenizer: <PRE>, <SUF>, <MID>, <EOT>. Using these tokens, the PSM prompt format is defined in [27] as

$$\begin{aligned} <PRE> \circ \text{Enc}(\text{prefix}) \circ <SUF> \circ \\ & \text{Enc}(\text{suffix}) \circ <MID> \circ \text{Enc}(\text{middle}), \end{aligned} \quad (1)$$

where the $\text{Enc}(\cdot)$ is the encoding operation and \circ is a token-level concatenation. The SPM format is defined in [27] as

$$\begin{aligned} <PRE> \circ <SUF> \circ \text{Enc}(\text{suffix}) \circ \\ & <MID> \circ \text{Enc}(\text{prefix}) \circ \text{Enc}(\text{middle}). \end{aligned} \quad (2)$$

The model autoregressively generates tokens until the <EOT> token is generated indicating a successful connection of the middle part with the suffix. The fact that in the SPM format, the middle part immediately follows the prefix makes the format even more similar to the simple regression case.

C. Prompt Engineering Techniques for OpenAPI Completion

As the expected OpenAPI definition size is significantly larger than the maximum context size supported by Code Llama, it is not possible to fit the whole definition in the prompt. In this research, the impact of multiple variables on the completion performance is evaluated to find the optimal prompt format for a real-world OpenAPI completion task.

Code Llama’s infilling mode is more suitable for the OpenAPI completion task than the autoregressive generation mode as it allows the model to consider the context both before and after the cursor position. The part of the definition that goes after the cursor should allow the model to take into account components such as schemas, requests, and responses that are commonly placed at the end of the document. The suffix is also important for the model to generate syntactically correct definitions that are integral to the rest of the document. Editing at the end of a document can be considered an infilling task where the suffix is empty. In the real world, API designers are more likely to edit existing definitions by adding new operations or parameters in the middle of the document rather than at the end of it.

One of the factors affecting completion performance is the prefix-to-suffix ratio. Theoretically, prefixes and suffixes should contain the same amount of useful information for the model to generate a semantically correct completion. However, the model can have a preference for one of the parts inherited from the training data. The optimal ratio can be found by evaluating the completion performance for a range of ratios.

Another factor that affects both the completion accuracy and speed is the context size. The context size is the number of tokens available to the model at the time of completion. A larger context size allows the model to take into consideration more information from the surrounding context but also requires more computational time and resources to process. The optimal context size can be found by evaluating the completion performance for a range of context sizes while keeping the prefix-to-suffix ratio constant.

OpenAPI specification allows including metadata about the API in the definition. This metadata can include information like the API’s title, description, version, contact information, or license type. The metadata is not required for the definition to be considered valid and can be neglected during OpenAPI validation or comparison. However, the metadata may contain useful information for the model to generate a more semantically correct definition.

The components section of the OpenAPI definition is crucial for the model to generate valid definitions since it defines the set of available schemas, requests, responses, and other components that can be referenced in the definition. To test this hypothesis, the prompt builder is modified to extract the components section if available, format it as a comment with the list of component names, and include it in the prompt prefix.

D. Code Llama Model Fine-tuning for OpenAPI Completion

The OpenAPI completion performance of the model can be further improved by fine-tuning it on task-specific data. This

section describes the dataset used for fine-tuning the 7-billion-parameter version of the Code Llama model as well as the full details of the fine-tuning process.

The dataset for fine-tuning was collected from the APIs-guru OpenAPI definitions directory [56]. The directory contains more than 4000 definitions in YAML format. Analysis of the repository revealed that about 75% of the definitions in the directory are produced by a handful of major companies like Amazon, Google, and Microsoft. To avoid the dataset bias towards a specific producer, the maximum number of definitions from a single producer was limited to 20. Multiple versions of the same API were also excluded from the dataset as they are likely to contain similar definitions. Definitions from producers used in the evaluation dataset were also excluded from the fine-tuning dataset to avoid data leakage. The resulting dataset contains 990 definitions and is accessible at the HuggingFace Hub at <https://huggingface.co/datasets/BohdanPetryshyn/openapi-completion-refined>.

To avoid the objective mismatch between the training and fine-tuning processes, the infilling prompt format was reconstructed to mirror the format described by the Code Llama model authors. Nevertheless, the specifics of the OpenAPI format had to be taken into account. Similarly to the Code Llama training, the prefix-middle-suffix split was performed at the character level with the split positions selected randomly. The same tokens (<PRE>, <SUF>, <MID>, <EOT>) were used to format the prompt in the PSM (1) and SPM (2) formats and the leading spaces implicitly generated by the SentencePiece tokenizer were suppressed. However, unlike the Code Llama training, the split was performed at the context level and was not limited to documents that fit into the context size as most of the OpenAPI definitions are much larger than the context size used during fine-tuning. In the SPM format, the prefix and middle parts were concatenated after the tokenization to improve the model performance with split tokens which are expected to be encountered often in the real-world OpenAPI completion scenarios.

Another consequence of the large size of the OpenAPI definitions is the bias towards training-time context size infilling that might be developed when fine-tuning the model using the conventional FIM training objective as described in the work of Bavarian et. al. [27]. This bias can cause the model to show optimal performance only when prompted with the same context size as during fine-tuning. As can be seen in Figure 3, the distribution of the training sample lengths is skewed toward the maximum context size of 5120 tokens when the conventional context-level FIM transformation is used. To mitigate the bias, the input documents were split into smaller pieces of random length between 4096 and 6144 tokens (0.8-1.2 of the 5120 tokens context size) before proceeding to the context-level FIM transformation. This made sure the context chunks included multiple documents of various lengths more often and evened the distribution of the training sample lengths. This technique is further referred to as document splitting. When used before the FIM transformation, document splitting should not be distinguishable from using full documents of uniform length distribution as the truncated prefixes and suffixes correspond to a normal real-world code completion scenario when the

context size is limited and does not cover the entire document.

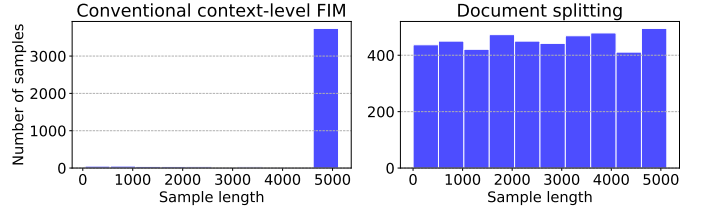


Fig. 3. Training sample length distribution with and without the document splitting technique.

Parameter efficient fine-tuning was performed with the following hyperparameters: LoRA rank of 32, LoRA alpha of 64, LoRA dropout of 0.1, AdamW optimizer with the learning rate of $2e-4$, cosine scheduler with the warmup ratio of 0.1, the batch size of 4 with the gradient accumulation steps of 8 and the maximum number of optimization steps of 1000.

IV. RESULTS

Following the prompt engineering and fine-tuning methods described in the previous section, the Code Llama model was optimized for the OpenAPI completion task. This section presents the results of the evaluation of the optimized model and compares it to the performance of GitHub Copilot. For all the evaluations described in this section, the same set of masked OpenAPI definitions was used ensuring consistent evaluation conditions.

A. GitHub Copilot Performance in OpenAPI Completion

GitHub Copilot was selected as the baseline solution for OpenAPI generation as at the time of writing it is the most widely used coding assistant on the market. The solution is a commercial product available with a monthly subscription plan.

As GitHub Copilot does not provide an API for programmatic access, the completion cannot be automated and must be performed manually. For the evaluation results to be comparable with the automatic completion results, the manual evaluation process resembles the automatic completion procedure as closely as possible and uses the same set of masked OpenAPI definitions. Manual completion was performed for each masked definition following these steps:

- 1) The <<MASK>> marker is removed, and the cursor is set in place of the marker;
- 2) If Copilot refuses to generate a suggestion, the last character is removed and added again to trigger the completion;
- 3) Copilot's suggestions are accepted without modifications one by one until one of the following conditions is reached:
 - a) Suggestions are no longer generated (Copilot considers the infilling complete);
 - b) Maximum number of completion lines is reached (15 lines).

The speed of suggestion generation using GitHub Copilot was measured by tracking the time between triggering completion with a key press and seeing the generated suggestion

on the screen. The number of characters generated was divided by the generation time. The experiment was repeated ten times to take human error as well as random service latency factors into account. The average code generation speed was measured to be 19.3 characters per second with a standard deviation of 1.9.

Following the completion procedure described, GitHub Copilot correctly completed 29% of the OpenAPI definitions. 68% of the completed definitions were valid.

B. Original Code Llama Model for OpenAPI Completion

Using the benchmark proposed in the previous chapter, an extensive set of experiments was performed to analyze the impact of a variety of variables and prompt-engineering techniques on the performance of the Code Llama model in the OpenAPI completion task. The experiments were designed to analyze the impact of a single variable at a time. When not stated otherwise, the 7-billion-parameter version of the Code Llama model, the PSM prompt format (1), the 50/50 prefix-to-suffix ratio, and the 4096 tokens context size were used.

The Code Llama models were evaluated at the Hugging Face Inference Endpoints platform [58] using the simple greedy search decoding strategy. Without quantization, a single Nvidia A10G GPU was used for the 7-billion-parameter model. For the 13-billion-parameter model, four Nvidia A10G GPUs were used.

1) **Optimal prefix-to-suffix ratio:** A range of prefix-to-suffix ratios from 30/70 up to 90/10 was evaluated to find the optimal proportion. The evaluation was performed on the 7-billion-parameter version of the Code Llama model with a constant context size of 4096 tokens. The optimal prefix-to-suffix ratio can depend on the above-mentioned factors but for parameter space reduction, the dependencies are intentionally disregarded in this research.

Figure 4 depicts the dependency of correctness and validity of the OpenAPI completion on the prefix-to-suffix ratio. As can be seen from the graph, the highest percentage of correctly infilled definitions of 36% corresponds to 50/50 and 60/40 prefix-to-suffix ratios. The 50/50 ratio demonstrates a higher validity percentage of 63% compared to 61% for the 60/40 ratio. Even though the dip at the 60/40 ratio can be attributed to the factor of randomness, the 50/50 ratio can still be considered optimal as it does not discriminate the importance of either prefix or suffix without sound justification.

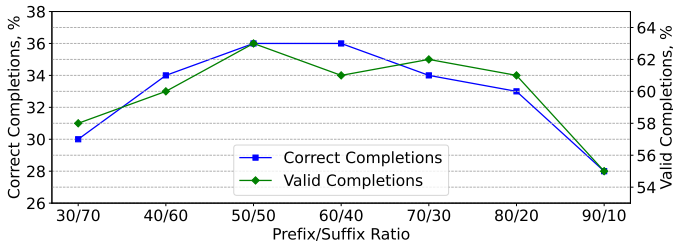


Fig. 4. OpenAPI completion performance of the original Code Llama 7B model for different prefix-to-suffix ratios with a fixed context size of 4096 tokens.

2) **Optimal context size:** Context size is another factor that affects not only the solution performance but also the speed of code generation and, consequently, throughput and resource efficiency of a real-world implementation. A range of context sizes from 1024 to 7168 tokens was evaluated with a constant prefix-to-suffix ratio of 50/50 and the 7-billion-parameter version of the Code Llama model.

As can be seen in Figure 5, OpenAPI completion correctness and validity increase with the expanding context size. The correctness rate, which is the primary target metric in this research, demonstrates a much smoother trend with a significant increase up to the context size of 4096 tokens. After this point, the correctness level stabilizes at around 36% and even decreases by 1% after 6144 tokens. The speed, as expected, follows the downward trend with a dramatic decline after the 3072 tokens point. This can be explained by the quadratic complexity of the input processing in decoder-only transformers.

With a correctness rate of 36%, validity rate of 63%, and a generation speed of 79 characters per second, the context size of 4096 tokens can be considered optimal for real-world scenarios since after this point, neglectable correctness and validity increases come at a cost of significant speed and resource-efficiency deterioration.

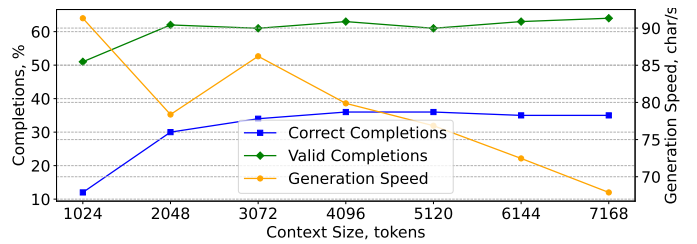


Fig. 5. OpenAPI completion performance of the original Code Llama 7B model for different context sizes with a fixed prefix-to-suffix ratio of 50/50.

3) **PSM and SPM infilling formats:** SPM format (2) is an alternative prompt format supported by the Code Llama model in the infilling mode. Contrary to the fact that the SPM format demonstrated higher performance in the original work of Bavarian et al. [27], the authors of Code Llama designed PSM (1) to be the primary prompting format for the model stating that the SPM implementation might require additional token healing on the tokenization step at the inference time. In this experiment, the default Code Llama tokenizer is used. As can be seen in Figure 6, the SPM prompting format underperforms at all context sizes compared to PSM. For example, the model correctly completed 24% of the definitions and generated valid completions 53% of the time at 4096 tokens, which is a 33.3% correctness decrease compared to the same configuration with the PSM prompt format.

4) **OpenAPI metadata in prompt:** This experiment evaluates the impact of including the components section of the OpenAPI definition in the prompt prefix as described in Section III-C. The experiment was performed on the 7-billion-parameter Code Llama model with a 60/40 prefix-to-suffix ratio and a range of context sizes from 4096 to 7168 tokens.

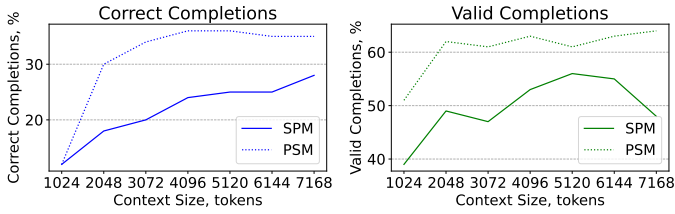


Fig. 6. OpenAPI completion performance of the original Code Llama 7B model with the PSM and SPM prompt formats.

The extended prefix size was expected to compensate for the space taken by the components section in the prompt.

As can be seen in Figure 7, adding the components list to the prompt prefix results in moderate validness and correctness improvements at the larger context sizes (11.4% correctness and 1.6% validness improvement at 6144 tokens). The results at a smaller context size of 4096 tokens, on the other hand, demonstrate a significant drop in both correctness (25%) and validity (6.3%).

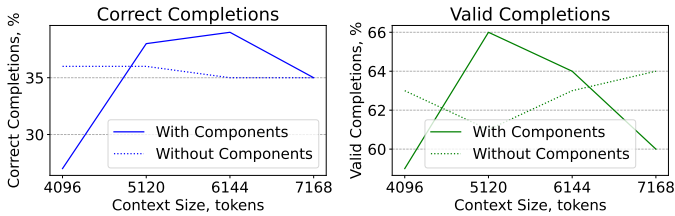


Fig. 7. OpenAPI completion performance of the original Code Llama 7B model with and without OpenAPI components in the prompt prefix.

The obtained results might mean that the components are not necessarily required for correct OpenAPI completion. The model can deduce the correct component names from just the surrounding context. When used with a relatively small context size, the suggested method can even be harmful to the solution performance as it consumes space that could otherwise be taken by more valuable information in the prefix.

5) **Optimal model size:** The Code Llama model is available in four sizes: 7, 13, 34, and 70 billion parameters. The impact of the model size on the solution performance is evaluated by comparing the performance of the 7 and 13-billion-parameter models. Due to the resource constraints and the limitations of the larger models discussed in Section III-B, the 34 and 70-billion-parameter models were not evaluated.

A machine with four Nvidia A10G GPUs and a total of 96GB of GPU memory was used for the evaluation. This machine made it possible to evaluate the model with up to 4096 tokens of context. The evaluation was performed on a range of context sizes from 1024 to 4096 tokens with a constant prefix-to-suffix ratio of 50/50.

As can be seen in Figure 8, the bigger 13-billion-parameter model outperforms the smaller model at all context sizes with the only exception being correctness after 3072 tokens. At the context size of 4096 tokens, Code Llama 13B surpasses the results obtained with Code Llama 7B by only 1% in validity while falling behind in correctness by 5%.

Figure 9 compares the generation speed trends of the models. Due to the GPU computing power redundancy caused

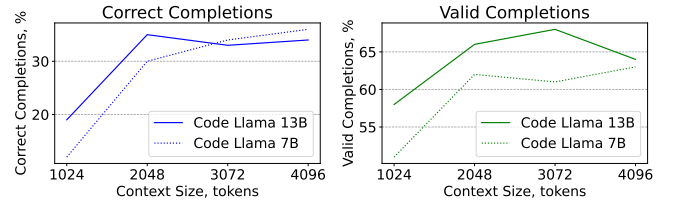


Fig. 8. OpenAPI completion performance of the 7 and 13-billion-parameter versions of the original Code Llama model.

by the high memory requirements of the model, the 13 billion Code Llama outperforms the smaller model at 1024 and 2048 tokens. However, the quadratic complexity of the attention mechanism quickly makes up for the computing power redundancy, and the generation speed drops below the speed of the smaller model at 3072 and 4096 tokens. The speed reduction trend suggests that the difference in the generation speed will only increase with the context size.

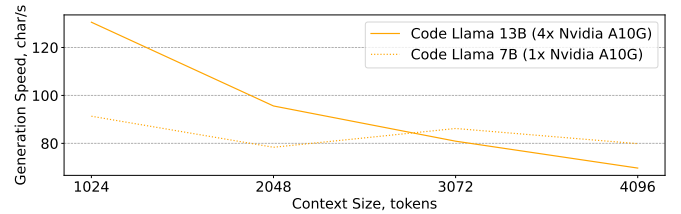


Fig. 9. Generation speed of the 7 and 13-billion-parameter versions of the original Code Llama model.

C. Fine-tuning Code Llama Model for OpenAPI Completion

As demonstrated in the previous section, the original Code Llama model is already capable of OpenAPI completion and even outperforms GitHub Copilot. This section evaluates the fine-tuning approaches that can further improve the model performance.

1) **Benefits of using a mixture of PSM and SPM infilling formats:** The study of Bavarian et al. suggests that using a mix of SPM and PSM formats during training leads to an increase in the infilling performance [27]. At the same time, the original Code Llama model has a known problem with infilling in the SPM format IV-B3. It is worth evaluating the efficiency of including this format in the fine-tuning process. This section compares the performance of the 7-billion-parameter Code Llama model fine-tuned solely with the PSM format and with a mixture of PSM and SPM formats (joint fine-tuning). When used, the SPM format was applied to half of the training samples. Context size of 4096 tokens and a 50/50 prefix-suffix split was used during fine-tuning.

As can be seen in Figure 10, both fine-tuning approaches demonstrate similar performance in the PSM format with the joint fine-tuning showing a slight improvement in correctness and validity. The joint fine-tuning also results in significantly higher performance when the SPM format is used for inference. These results suggest that it always makes sense to




Fig. 10. OpenAPI completion performance of the Code Llama 7B model fine-tuned solely with the PSM format and with a mixture of PSM and SPM formats.

include the SPM format when fine-tuning the Code Llama model for code infilling tasks.

Compared to the original Code Llama model, the model fine-tuned with the joint format demonstrates a peak improvement of 28.6% in both correctness and validity at the context size of 6144 tokens when prompted with the PSM format. Performance in the SPM format is still inferior to PSM but the gap is reduced from 28.6% to 11.1% in correctness and from 12.7% to 4.9% in validity at the same context size. This shows that the selected fine-tuning approach can also lead to solving the original model’s problem of split-token handling in the SPM format.

2) **Training-time context size infilling bias and document splitting technique:** As revealed in the previous section, the fine-tuned model underperforms when prompted with context sizes smaller than used during fine-tuning (Figure 10). This can be caused by the bias towards training-time context size infilling discussed in Section III-D. This section further evaluates the impact of the fine-tuning context size as well as the suggested document splitting technique on the performance of the fine-tuned model.

As can be seen in Figure 11, the fine-tuned models with both 4096 tokens and 5120 tokens underperform compared to the original Code Llama model at context sizes smaller than the corresponding fine-tuning context sizes. After this point, the performance of the fine-tuned models stabilizes. Detailed inspection of the infilling results shows that the fine-tuned models tend to generate completions that are too long and contain irrelevant information when prompted with smaller context sizes.

A tempting idea to mitigate the problem is to use a much smaller context size for fine-tuning. As demonstrated in Figure 11, using the context size of 2048 tokens for fine-tuning indeed decreases the performance decline at smaller context sizes. However, using smaller context also reduces the




Fig. 11. OpenAPI completion performance of the Code Llama 7B model fine-tuned at different context sizes and with the document splitting technique.

performance gains at larger context sizes to the point where the fine-tuned model underperforms compared to the original Code Llama model. Nevertheless, the model still demonstrates the best validity rates at all context sizes.

The model fine-tuned with the document splitting technique demonstrates more equalized performance across the context sizes. The correctness and validity of the completions generated by this model surpass the original Code Llama model at almost all context sizes with a peak increase of 20% at 6144 and 7168 tokens and a marginal decrease of 5.8% at 3072 tokens. The model slightly underperforms compared to the model fine-tuned with the conventional context-level FIM technique at the context size of 4096 tokens when prompted with context sizes larger than 4096 tokens, but it outperforms the same model at smaller context sizes by a significant margin.

V. FINDINGS AND IMPLICATIONS

Through the course of this research, the impact of multiple variables, prompt engineering, and fine-tuning techniques on the OpenAPI completion performance of the Code Llama model was evaluated. The most notable experiment results are summarized and compared in Table II.

1) **Original Code Llama model is capable of outperforming GitHub Copilot:** The original Code Llama 7B model, when used with the optimal prefix-to-suffix ratio of 50/50, the PSM prompt format, and the context size of 4096 tokens, outperforms GitHub Copilot by 24.1% in correctness, generating correct OpenAPI completions 36% of the time. At the same time, the model produces valid completions 63% of the time which is a 7.3% decrease compared to Copilot. It is worth pointing out that the GPT-3 model, which is the foundational model used by GitHub Copilot, uses 175 billion parameters which is 25 times more than in the case of the Code Llama 7B model.

2) **Larger model does not always show significant performance improvement:** The larger 13-billion-parameter version of the Code Llama model demonstrated only a marginal correctness improvement compared to the smaller 7-billion-parameter version. This might mean that the OpenAPI format is not complex enough to leverage the additional model capacity. Considering the quadrupled computational cost, this renders the larger version of the model not practical for use in OpenAPI completion.

TABLE II
OPENAPI COMPLETION PERFORMANCE COMPARISON

Model / Solution	Correctness, max., %	Validity, max., %	Correctness, avg., %	Validity, avg., %
GitHub Copilot	29	68	29.0	68.0
Code Llama 7B	36 (4 096 tokens)	64 (7 168 tokens)	31.1	60.7
Code Llama 13B (evaluated at 1 024-4 096 tokens)	34 (4 096 tokens)	68 (3 072 tokens)	30.2	64.0
Code Llama 7B, fine-tuned at 4 096 tokens	45 (6 144 tokens)	84 (4 096 tokens)	32.0	63.1
Code Llama 7B, fine-tuned at 5 120 tokens	42 (7 168 tokens)	79 (5 120 tokens)	26.8	57.1
Code Llama 7B, fine-tuned with document splitting	42 (6 144 tokens)	76 (7 168 tokens)	34.0	69.1

3) **Task-specific fine-tuning can significantly improve the model performance:** The 7-billion-parameter version of the Code Llama model fine-tuned for OpenAPI completion at the context size of 4 096 tokens following the methods described in Section III-D demonstrates a maximum correctness improvement of 28.6% compared to the original model and further outperforms GitHub Copilot by 55.2%.

4) **Conventional context-level FIM training objective can lead to suboptimal performance at smaller context sizes:** The case of OpenAPI completion revealed that, when used with documents that are significantly larger than the fine-tuning context size, this technique leads to a bias towards training-time context size infilling. This bias can cause the model to generate too long and irrelevant completions when prompted with a context size smaller than the fine-tuning context size. The models trained with this technique still demonstrate the best performance at the context size used during fine-tuning which can be leveraged when the inference context size is known in advance.

5) **Document splitting technique leads to a more unified performance across context sizes:** The document splitting technique proposed in Section III-D can be used to mitigate the problem of the training-time context size infilling bias and lead to a more unified performance across the context sizes. Even though the models fine-tuned without document splitting outperform the model fine-tuned with this technique at the context sizes used during fine-tuning, the latter model demonstrates the best average performance across the context sizes which is crucial for training foundational or open-source models that are expected to be used with a variety of context sizes.

VI. CONCLUSION

In the course of this research, the first semantics-aware OpenAPI completion benchmark was proposed. The development of such a benchmark made it possible to evaluate the task-specific performance of one of the most widely used commercial code completion solutions – GitHub Copilot.

A series of experiments performed in this research demonstrate how modern compact open-source models can be optimized to outperform commercial solutions. The state-of-the-art code generation model Code Llama from Meta is used here for that.

The Code Llama model fine-tuned for OpenAPI completion reaches the peak correctness improvement of 55.2% compared to GitHub Copilot. At the time of writing, this is the first work

investigating the details of fine-tuning the Code Llama model for code infilling.

In the course of the experiments, an improvement to a widely used code infilling training technique was proposed. The optimization improves the stability of the fine-tuned model performance at the entire range of context sizes. This optimization is particularly beneficial for formats like OpenAPI, where the training document size is significantly larger than the fine-tuning context size.

The OpenAPI completion benchmark as well as the fine-tuning pipeline for the Code Llama model were made publicly available at <https://github.com/BohdanPetryshyn/openapi-completion-benchmark> and <https://github.com/BohdanPetryshyn/code-llama-fim-fine-tuning>. The authors hope that these resources will be found helpful by practitioners even outside of the realm of API design tooling.

REFERENCES

- [1] “OpenAPI Specification v3.1.0 — Introduction, Definitions, & More — spec.openapis.org,” <https://spec.openapis.org/oas/v3.1.0>, [Accessed 07-04-2024].
- [2] “GitHub Copilot · Your AI pair programmer — github.com,” <https://github.com/features/copilot>, [Accessed 24-03-2024].
- [3] “AI Code Generator - Amazon CodeWhisperer - AWS — aws.amazon.com,” <https://aws.amazon.com/codewhisperer/>, [Accessed 24-03-2024].
- [4] B. Roziere, J. Gehring, F. Gloeckle, S. Sootla, I. Gat, X. E. Tan, Y. Adi, J. Liu, T. Remez, J. Rapin *et al.*, “Code llama: Open foundation models for code,” *arXiv preprint arXiv:2308.12950*, 2023.
- [5] E. Dehaerne, B. Dey, S. Halder, S. De Gendt, and W. Meert, “Code generation using machine learning: A systematic review,” *Ieee Access*, vol. 10, pp. 82 434–82 455, 2022.
- [6] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [7] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [8] H. Salehinejad, S. Sankar, J. Barfett, E. Colak, and S. Valaee, “Recent advances in recurrent neural networks,” *arXiv preprint arXiv:1801.01078*, 2017.
- [9] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2010, pp. 249–256.
- [10] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [11] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [12] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, “Albert: A lite bert for self-supervised learning of language representations,” *arXiv preprint arXiv:1909.11942*, 2019.

- [13] P. He, X. Liu, J. Gao, and W. Chen, “Deberta: Decoding-enhanced bert with disentangled attention,” *arXiv preprint arXiv:2006.03654*, 2020.
- [14] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *Journal of machine learning research*, vol. 21, no. 140, pp. 1–67, 2020.
- [15] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, “Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension,” *arXiv preprint arXiv:1910.13461*, 2019.
- [16] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever *et al.*, “Improving language understanding by generative pre-training,” 2018.
- [17] A. Chowdhery, S. Narang, J. Devlin, M. Bosma, G. Mishra, A. Roberts, P. Barham, H. W. Chung, C. Sutton, S. Gehrmann *et al.*, “Palm: Scaling language modeling with pathways,” *Journal of Machine Learning Research*, vol. 24, no. 240, pp. 1–113, 2023.
- [18] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, T. Rutherford, D. d. L. Casas, L. A. Hendricks, J. Welbl, A. Clark *et al.*, “Training compute-optimal large language models,” *arXiv preprint arXiv:2203.15556*, 2022.
- [19] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, N. Bacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023.
- [20] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, “Llama 2: Open foundation and fine-tuned chat models,” *arXiv preprint arXiv:2307.09288*, 2023.
- [21] T. Wang, A. Roberts, D. Hesslow, T. Le Scao, H. W. Chung, I. Beltagy, J. Launay, and C. Raffel, “What language model architecture and pretraining objective works best for zero-shot generalization?” in *International Conference on Machine Learning*. PMLR, 2022, pp. 22 964–22 984.
- [22] L. Reynolds and K. McDonell, “Prompt programming for large language models: Beyond the few-shot paradigm,” in *Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021, pp. 1–7.
- [23] Y. Li, D. Choi, J. Chung, N. Kushman, J. Schrittwieser, R. Leblond, T. Eccles, J. Keeling, F. Gimeno, A. Dal Lago *et al.*, “Competition-level code generation with alphacode,” *Science*, vol. 378, no. 6624, pp. 1092–1097, 2022.
- [24] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman *et al.*, “Evaluating large language models trained on code,” *arXiv preprint arXiv:2107.03374*, 2021.
- [25] E. Nijkamp, B. Pang, H. Hayashi, L. Tu, H. Wang, Y. Zhou, S. Savarese, and C. Xiong, “Codegen: An open large language model for code with multi-turn program synthesis,” *arXiv preprint arXiv:2203.13474*, 2022.
- [26] E. Nijkamp, H. Hayashi, C. Xiong, S. Savarese, and Y. Zhou, “Codegen2: Lessons for training llms on programming and natural languages,” *arXiv preprint arXiv:2305.02309*, 2023.
- [27] M. Bavarian, H. Jun, N. Tezak, J. Schulman, C. McLeavey, J. Tworek, and M. Chen, “Efficient training of language models to fill in the middle,” *arXiv preprint arXiv:2207.14255*, 2022.
- [28] A. Gao, “Prompt engineering for large language models,” *Available at SSRN 4504303*, 2023.
- [29] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [30] J. Liu, A. Liu, X. Lu, S. Welleck, P. West, R. L. Bras, Y. Choi, and H. Hajishirzi, “Generated knowledge prompting for commonsense reasoning,” *arXiv preprint arXiv:2110.08387*, 2021.
- [31] N. Ding, Y. Qin, G. Yang, F. Wei, Z. Yang, Y. Su, S. Hu, Y. Chen, C.-M. Chan, W. Chen *et al.*, “Parameter-efficient fine-tuning of large-scale pre-trained language models,” *Nature Machine Intelligence*, vol. 5, no. 3, pp. 220–235, 2023.
- [32] N. Houlsby, A. Giurgiu, S. Jastrzebski, B. Morrone, Q. De Laroussilhe, A. Gesmundo, M. Attariyan, and S. Gelly, “Parameter-efficient transfer learning for nlp,” in *International conference on machine learning*. PMLR, 2019, pp. 2790–2799.
- [33] E. B. Zaken, S. Ravfogel, and Y. Goldberg, “Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models,” *arXiv preprint arXiv:2106.10199*, 2021.
- [34] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “Lora: Low-rank adaptation of large language models,” *arXiv preprint arXiv:2106.09685*, 2021.
- [35] “copilot-explorer — thakkarparth007.github.io,” <https://thakkarparth007.github.io/copilot-explorer/posts/copilot-internals.html>, [Accessed 24-03-2024].
- [36] N. Nguyen and S. Nadi, “An empirical evaluation of github copilot’s code suggestions,” in *Proceedings of the 19th International Conference on Mining Software Repositories*, 2022, pp. 1–5.
- [37] D. Sobania, M. Briesch, and F. Rothlauf, “Choose your programming copilot: a comparison of the program synthesis performance of github copilot and genetic programming,” in *Proceedings of the genetic and evolutionary computation conference*, 2022, pp. 1019–1027.
- [38] A. M. Dakhel, V. Majdinasab, A. Nikanjam, F. Khomh, M. C. Desmarais, and Z. M. J. Jiang, “Github copilot ai pair programmer: Asset or liability?” *Journal of Systems and Software*, vol. 203, p. 111734, 2023.
- [39] M. Wermelinger, “Using github copilot to solve simple programming problems.(2023),” 2023.
- [40] “Tabnine is an AI assistant that speeds up delivery and keeps your code safe — Tabnine — tabnine.com,” <https://www.tabnine.com/>, [Accessed 24-03-2024].
- [41] “Cody — AI coding assistant — sourcegraph.com,” <https://sourcegraph.com/cody>, [Accessed 24-03-2024].
- [42] “GitHub - continuedev/continue: The easiest way to code with any LLM—Continue is an open-source autopilot for VS Code and JetBrains — github.com,” <https://github.com/continuedev/continue>, [Accessed 24-03-2024].
- [43] Z. Luo, C. Xu, P. Zhao, Q. Sun, X. Geng, W. Hu, C. Tao, J. Ma, Q. Lin, and D. Jiang, “WizardCoder: empowering code large language models with evol-instruct,” *arXiv preprint arXiv:2306.08568*, 2023.
- [44] “Phind Model,” www.phind.com/blog/phind-model-beats-gpt4-fast, [Accessed 24-03-2024].
- [45] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. d. l. Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier *et al.*, “Mistral 7b,” *arXiv preprint arXiv:2310.06825*, 2023.
- [46] “GitHub - juyongjiang/CodeUp: CodeUp: A Multilingual Code Generation Llama2 Model with Parameter-Efficient Instruction-Tuning on a Single RTX 3090 — github.com,” <https://github.com/juyongjiang/CodeUp>, [Accessed 24-03-2024].
- [47] L. Tunstall, E. Beeching, N. Lambert, N. Rajani, K. Rasul, Y. Belkada, S. Huang, L. von Werra, C. Fourrier, N. Habib *et al.*, “Zephyr: Direct distillation of lm alignment,” *arXiv preprint arXiv:2310.16944*, 2023.
- [48] R. Anil, A. M. Dai, O. Firat, M. Johnson, D. Lepikhin, A. Passos, S. Shakeri, E. Taropa, P. Bailey, Z. Chen *et al.*, “Palm 2 technical report,” *arXiv preprint arXiv:2305.10403*, 2023.
- [49] “GitHub - huggingface/llm-vscode: LLM powered development for VS-Code — github.com,” <https://github.com/huggingface/llm-vscode>, [Accessed 24-03-2024].
- [50] A. Svyatkovskiy, S. K. Deng, S. Fu, and N. Sundaresan, “Intellicode compose: Code generation using transformer,” in *Proceedings of the 28th ACM joint meeting on European software engineering conference and symposium on the foundations of software engineering*, 2020, pp. 1433–1443.
- [51] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “Bleu: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 311–318.
- [52] C.-Y. Lin, “Rouge: A package for automatic evaluation of summaries,” in *Text summarization branches out*, 2004, pp. 74–81.
- [53] S. Banerjee and A. Lavie, “Meteor: An automatic metric for mt evaluation with improved correlation with human judgments,” in *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, 2005, pp. 65–72.
- [54] F. Cassano, J. Gouwar, D. Nguyen, S. Nguyen, L. Phipps-Costin, D. Pinckney, M.-H. Yee, Y. Zi, C. J. Anderson, M. Q. Feldman *et al.*, “Multipl-e: a scalable and polyglot approach to benchmarking neural code generation,” *IEEE Transactions on Software Engineering*, 2023.
- [55] “OpenAPI Diff — oasdiff.com,” <https://www.oasdiff.com/>, [Accessed 24-03-2024].
- [56] “GitHub - APIs-guru/asyncapi-directory: Directory of asynchronous API specifications in AsyncAPI format — github.com,” <https://github.com/APIs-guru/asyncapi-directory>, [Accessed 26-03-2024].
- [57] “bigcode/the-stack · Datasets at Hugging Face — huggingface.co,” <https://huggingface.co/datasets/bigcode/the-stack>, [Accessed 26-03-2024].
- [58] “Inference Endpoints - Hugging Face — huggingface.co,” <https://huggingface.co/inference-endpoints/dedicated>, [Accessed 26-03-2024].