

Structured Chain-of-Thought Prompting for Code Generation

Jia Li 
lijia@stu.pku.edu.cn
Peking University
Beijing, China

Yongmin Li
Peking University
Beijing, China
liyongmin@pku.edu.cn

Ge Li
Peking University
Beijing, China
lige@pku.edu.cn


Zhi Jin
Peking University
Beijing, China
zhijin@pku.edu.cn

ABSTRACT

Large Language Models (LLMs) (e.g., ChatGPT) have shown impressive performance in code generation. LLMs take prompts as inputs, and Chain-of-Thought (CoT) prompting is the state-of-the-art prompting technique. CoT prompting asks LLMs first to generate CoTs (i.e., intermediate natural language reasoning steps) and then output the code. However, CoT prompting is designed for natural language generation and has low accuracy in code generation.

In this paper, we propose Structured CoTs (SCoTs) and present a novel prompting technique for code generation, named SCoT prompting. Our motivation is source code contains rich structural information and any code can be composed of three program structures (i.e., sequence, branch, and loop structures) [3]. Intuitively, structured intermediate reasoning steps make for structured source code. Thus, we ask LLMs to use program structures to build CoTs, obtaining SCoTs. Then, LLMs generate the final code based on SCoTs. Compared to CoT prompting, SCoT prompting explicitly constrains LLMs to think about how to solve requirements from the view of source code and further the performance of LLMs in code generation. We apply SCoT prompting to two LLMs (i.e., ChatGPT and Codex) and evaluate it on three benchmarks (i.e., HumanEval, MBPP, and MBCPP). (1) **SCoT prompting outperforms the state-of-the-art baseline - CoT prompting by up to 13.79% in Pass@1.** (2) Human evaluation shows human developers prefer programs from SCoT prompting. (3) SCoT prompting is robust to examples and achieves substantial improvements.

ACM Reference Format:

Jia Li , Ge Li, Yongmin Li, and Zhi Jin. 2023. Structured Chain-of-Thought Prompting for Code Generation. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Code generation aims to automatically generate a program that satisfies a given natural language requirement [13, 14, 38]. Large

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

Conference'17, July 2017, Washington, DC, USA

© 2023 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

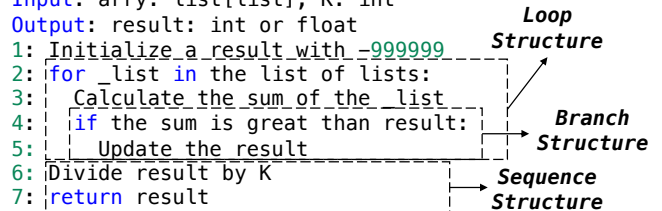
1. Initialize a result with -999999
2. Iterate through the list of lists
3. Initialize a sum with 0
4. Iterate through the list
5. Add the element to the sum
6. Update result with the maximum of sum and result
7. Divide the result by K
8. Return the result

(a) Chain-of-Thought

Input: array: list[list], K: int

Output: result: int or float

```
1: Initialize a result with -999999
2: for _list in the list of lists:
3:   Calculate the sum of the list
4:   if the sum is great than result:
5:     Update the result
6: Divide result by K
7: return result
```



(b) Structured Chain-of-Thought

Figure 1: The comparison of a Chain-of-Thoughts (CoT) and our Structured Chain-of-Thought (SCoT).

Language Models (LLMs) have recently shown impressive performance in code generation, such as ChatGPT [18], and CodeGen [17]. During the inference, LLMs take a prompt as input that consists of several examples (e.g., <requirement, code> pairs) and a new requirement. LLMs learn code generation from examples and analogously generate a new program. The performance of LLMs heavily relies on the prompt [39]. Nowadays, how to make an effective prompt (i.e., *Prompting technique*) for code generation is still an open question.

Chain-of-Thought (CoT) prompting [35] is the state-of-the-art (SOTA) prompting technique. CoT Prompting asks LLMs first to generate a CoT and then output the code. A CoT is several intermediate natural language reasoning steps that describe how to write code step by step. Figure 1 (a) shows a CoT on code generation. However, CoT prompting brings slight improvements in code generation. For example, it only improves ChatGPT by 0.82 points in Pass@1 upon a real-world benchmark [7].

In this paper, we propose a Structured CoT for code generation. Our motivation is that code generation aims to convert a natural language requirement to source code. Different from natural languages, source code contains rich structural information [22, 30, 37]. For example, source code contains three basic structures

[3], including sequence, branch, and loop structures. Intuitively, intermediate reasoning steps leading to the structured code should also be structured. Consider a human developer's thought process when solving a requirement (e.g., find the maximum number in a list). It is typical to come up with a solving process with program structures: "Initialize a result with -inf; for each number in the list; if the number is greater than result: Update result with the number ...". Our idea is to enable LLMs to generate similar structured CoTs - a coherent series of intermediate reasoning steps constructed by program structures. Besides, LLMs' training data contains lots of code data, so they have the ability to generate program structures. However, standard CoT ignores the program structures and has low accuracy in code generation. Thus, it is necessary to design a structured CoT to unlock the reasoning ability of LLMs in code generation.

Figure 1 (b) shows a SCoT. The design of our SCoT has two inspirations. First, existing work [3] proved that any simple or complex program can be composed of three basic structures, i.e., sequence structure, branch structure, and loop structure. Thus, we introduce three basic structures and constrain LLMs to use them to generate CoTs. As shown in Figure 1 (b), the SCoT uses a loop structure to clearly describe an iteration in line 2. While in the CoT, the scopes of two iterations in lines 2 and 4 are ambiguous. It shows the superiority of SCoT in code generation. Besides, every program contains a required input-output structure, which includes the input-output parameters and their types (e.g., Input: array: list[list]; Output: result in Figure 1 (b)). By generating the input-output structure, LLMs are asked to analyze requirements and determine the entry and exit of the code, which benefits the following solving process.

Based on the SCoT, we present a new prompting technique named SCoT prompting. It asks LLMs first to generate a SCoT using program structures and then implement the code. Compared to CoT prompting, SCoT prompting explicitly introduces program structures into intermediate reasoning steps and constraints LLMs to think about how to solve requirements from the view of programming languages. It further unlocks the reasoning ability of LLMs in code generation, thus achieving higher accuracy.

We apply SCoT prompting to two popular LLMs (i.e., ChatGPT [18] and Codex [7]) and evaluate it on three representative benchmarks (i.e., HumanEval [7], MBPP [2], and MBCPP [1]). We use unit tests to measure the correctness of generated programs and report the Pass@ k ($k \in [1, 3, 5]$) [7]. Based on experimental results, we obtain four findings. (1) SCoT prompting significantly improves the accuracy of LLMs on code generation. **Compared to the SOTA baseline - Chain-of-Thought prompting, in terms of Pass@1, SCoT prompting outperforms it by up to 13.79% in HumanEval, 12.31% in MBPP, and 6.63% in MBCPP.** (2) Human evaluation shows that human developers prefer programs generated by SCoT prompting. (3) SCoT prompting is effective for different LLMs and different programming languages. In terms of Pass@1, it improves ChatGPT by up to 13.79% and Codex by up to 13.77%. Besides, SCoT prompting is language-agnostic and effective in multiple languages (e.g., Python and C++). (4) We explore the robustness of SCoT prompting to examples. Results show that SCoT prompting does not depend on specific examples or writing styles.

We summarize our contributions in this paper as follows.

- We propose a Structured Chain-of-Thought (SCoT), which utilizes program structures to build the intermediate reasoning steps.
- We propose a novel prompting technique for code generation, named SCoT Prompting. It prompts large language models first to generate a SCoT and then implement the code.
- We conduct extensive experiments on three benchmarks. Qualitative and quantitative experiments show that SCoT prompting significantly outperforms SOTA baselines (e.g., Chain-of-Thought prompting).
- We discuss the contributions of different program structures and the robustness of SCoT prompting.

Data Availability. We open source our replication package [?], including the datasets and the source code of SCoT prompting, to facilitate other researchers and practitioners to repeat our work and verify their studies.

2 METHODOLOGY

In this section, we propose a Structured Chain-of-Thought (SCoT). A SCoT denotes several intermediate reasoning steps constructed by program structures. Then, we present a novel prompting technique for code generation named SCoT prompting. SCoT prompting asks LLMs first to generate a SCoT and then output the final code. In the subsections, we first describe the design of our SCoT and further show the details of SCoT prompting.

2.1 Structured Chain-of-Thought

Standard Chain-of-Thought (CoT) is several intermediate natural language reasoning steps that lead to the final answer [35]. The CoT is initially designed for natural language generation (e.g., common-sense reasoning [26]). Thus, the CoT only uses natural languages to sequentially describe how to solve a problem step by step. Figure 1 (a) shows a CoT on code generation. A limitation is that CoT brings slight improvements in code generation. For example, adding the CoT only improves ChatGPT by 0.82 points in Pass@1 upon a real-world benchmark - HumanEval [7].

In this paper, we propose a Structured CoT. Our motivation is that, unlike natural language generation, the goal of code generation is highly structured code. Source code solves a problem through special structures, including sequence structures, branch structures, and loop structures. For example, given a requirement - reading text from a given file, imagine a human developer's thought process. The developer will use program structures to design an initial idea: "if the given file exists: read text from the file; else: raise an error;". The program structures clearly show the solving process and benefit the following code implementation. Thus, intermediate reasoning steps leading to the code should also be structured.

Figure 2 shows some examples of SCoT. Compared to the CoT, our SCoT explicitly introduces program structures. Existing work [3] proved that any simple or complex program can be composed of three basic structures, i.e., sequence structure, branch structure, and loop structure. Thus, we introduce three basic structures, whose details are shown as follows.

- **Sequence Structure.** The intermediate steps are sequentially placed and all steps are at the same level.
- **Branch Structure.** It starts with a condition and places different intermediate steps for different results of the condition. In this

```

Input: paren_string: str
Output: list_of_int: List[int]
1: Initialize list_of_int to an empty list
2: for each string in paren_string do
3:   Initialize depth to 0
4:   for each character in string do
5:     if character is '(' then
6:       depth += 1
7:     elif character is ')' then
8:       depth -= 1
9:   append depth to list_of_int
10: return list_of_int
    
```

(a)

```

Input: string: str, substring: str
Output: count: int
1: Initialize count to 0
2: while substring is not found in string do
3:   if string is empty then
4:     return 0
5:   increment count
6:   remove the first character of string
7: return count
    
```

(b)

Figure 2: Examples of SCoT in code generation.

paper, branch structures contain three formats, *i.e.*, if ..., if ... else, and if ... elif ... else.

- **Loop Structure.** A set of intermediate steps are repeatedly conducted until given conditions are not met. In this paper, loop structures contain two basic formats, including the for loop and the while loop.

We allow the nesting between different program structures. It allows LLMs to design more complex SCoT for some difficult requirements. As shown in Figure 2, the SCoT flexibly uses various program structures to build a solving process.

Besides three basic structures, we add the input-output structure, which contains input-output parameters and their types. Our motivation is that an input-output structure is required for a program, which indicates the entry and exit. Generating the input-output structure is beneficial to clarify requirements and generate the following solving process.

2.2 SCoT prompting

Based on the SCoT, we propose a new prompting technique for code generation, named SCoT prompting. It asks LLMs first to generate a SCoT and then output the final code.

To implement SCoT prompting, we design two special prompts. The first prompt is used to generate a SCoT, and an example of the prompt is shown in Figure 3. The prompt starts with several examples (*i.e.*, <requirement, SCoTs>). These examples cover three basic program structures and the input-output structure. Next, the *italic sentences* are instructions for LLMs, which indicate the goal of LLMs and related constraints. Finally, the prompt ends with a new requirement and is fed into LLMs. We expect LLMs to learn from examples and generate a SCoT for the new requirement.

After generating a SCoT, we design the second prompt for generating the final code. An example of the prompt is shown in Figure 4.

```

def first_Repeated_Char(str):
    """
    Write a python function to find the first repeated
    character in a given string.
    """
    Pass
    
```

Please understand the requirement and write a rough solving process. It starts with a input-output structure. You should use three basic structures to build the solving process, including sequences, branches, and loops. The necessary details should be written in natural languages.

```

Input: str: a string
Output: ch: a repeated character in str
1: for each character ch in str:
2:   if ch appears more than once in str:
3:     return ch
4: return None
    
```

(more examples...)

```

def text_lowercase_underscore(text):
    """
    Write a function to find sequences of lowercase letters
    joined with an underscore.
    """
    pass
    
```

Please understand the requirement and write a rough solving process. It starts with a input-output structure. You should use three basic structures to build the solving process, including sequences, branches, and loops. The necessary details should be written in natural languages.

An Example

A new input

Figure 3: A prompt for generating a SCoT.

```

def first_Repeated_Char(str):
    """
    Write a python function to find the first repeated
    character in a given string.
    Input: str: a string
    Output: ch: a repeated character in str
    1: for each character ch in str:
    2:   if ch appears more than once in str:
    3:     return ch
    4: return None
    """
    # Please check the above solving process and write a
    code based on it. Note that the solving process may contain
    errors.
    h = {}
    for ch in str:
        if ch in h:
            return ch;
        else:
            h[ch] = 0
    return None
    
```

(more examples...)

```

def text_lowercase_underscore(text):
    """
    Write a function to find sequences of lowercase letters
    joined with an underscore.
    Input: text: str
    Output: seq: list of str
    1: Initialize seq as an empty list
    2: for each word in text split by space do
    3:   if word matches the pattern of lowercase
    letters joined with an underscore then
    4:     append the word to seq
    5: return seq
    """
    
```

Please check the above solving process and write a code based on it. Note that the solving process may contain errors.

An Example

A new input

Figure 4: A prompt for generating the code.

The prompt starts with several examples (*i.e.*, <requirement, SCoT, code>). The *italic sentences* are instructions. We consider the SCoT as a soft template and ask LLMs to implement a program. Finally,

the prompt ends with a new requirement and its SCoT, and is input into LLMs. By learning from examples, LLMs generate a new program based on the requirement and SCoT.

Related work [25] has found that generative models may be negatively affected by error accumulation. Similarly, in SCoT prompting, the generated SCoT may contain noises (*e.g.*, errors or missing steps). These noises will further negatively affect code implementation. In this paper, we utilize two approaches to alleviating error accumulation. First, as shown in Figure 4, we ask LLMs to double-check the SCoT and fix possible noises. It allows LLMs to adaptively refer to the SCoT and filter out noises. Second, SCoT prompting utilizes a two-step generation pipeline. It provides a window of opportunity to debug where the SCoT goes wrong. In practice, human developers can first check the generated SCoT and fix possible errors. Then, the SCoT is used to generate code.

2.3 Implementation Details

SCoT prompting is a prompting technique for code generation, which does not rely on specific LLMs. In this paper, we consider ChatGPT as the default LLM. We select a few (*e.g.*, three) <requirement, code> pairs from real-world benchmarks (*i.e.*, training data) as example seeds. Then, we manually write the SCoT for seeds and obtain examples - <requirement, SCoT, code> triples, which are used to make prompts in Figure 3 and 4. A prompt contains three examples by default. The examples and prompt templates are available in our replication package. In the future, users can flexibly apply our approach to more powerful LLMs in a plug-and-play fashion.

3 STUDY DESIGN

To assess SCoT prompting, we conduct a large-scale study to answer four research questions. In this section, we present the details of our study, including datasets, evaluation metrics, comparison baselines, and implementation details.

3.1 Research Questions

Our study aims to answer the following research questions (RQ).

RQ1: How does SCoT prompting perform in terms of accuracy compared to baselines? This RQ aims to verify that SCoT prompting has a higher accuracy than existing prompting techniques on code generation. We apply three existing prompting techniques and SCoT prompting to two LLMs. Then, we use unit tests to measure the correctness of programs generated by different approaches and report the Pass@ k .

RQ2: Do developers prefer programs generated by SCoT prompting? The ultimate goal of code generation is to assist human developers in writing code. In this RQ, we hire 10 developers (including industry employees and academic researchers) to manually review the programs generated by SCoT prompting and baselines. We measure the quality of programs in three aspects, including correctness, code smell, and maintainability.

RQ3: Is SCoT prompting robust to examples? Prompting techniques may be sensitive to example [39]. In this RQ, we measure the robustness of SCoT prompting to examples. Specifically, we measure the performance of SCoT prompting with different example seeds and different example writing styles.

Table 1: Statistics of the datasets in our experiments.

Statistics	HumanEval	MBPP	MBCPP
Language	Python	Python	C++
# Train	–	474	413
# Test	164	500	435
Avg. tests per sample	7.7	3	3

RQ4: What are the contributions of different program structures in SCoT prompting? As stated in Section 2.1, SCoT prompting introduces three basic structures and the input-output structure. This RQ is designed to analyze the contributions of different structures. We select an LLM as the base model. Then, we individually remove a program structure and report the fluctuations in performance.

3.2 Benchmarks

Following previous studies [6, 7, 17, 40], we conduct experiments on three representative code generation benchmarks, including the HumanEval in Python, MBPP in Python, and MBCPP in C++. The details of the benchmarks are described as follows.

- **HumanEval** [7] is a Python function-level code generation benchmark, which contains 164 hand-written programming problems. Each programming problem consists of an English requirement, a function signature, and several test cases, with an average of 7.7 test cases per problem.
- **MBPP** [2] is a Python function-level code generation benchmark. It contains 974 programming problems that involve simple numeric manipulations or basic usage of standard libraries. Each problem contains an English requirement, a function signature, and three manually written test cases for checking functions.
- **MBCPP** [1] is a C++ function-level code generation benchmark. It consists of 848 programming problems that are collected by crowd-sourcing. Each problem contains an English description, a function signature, and three test cases for checking the correctness of functions.

We follow the original splits of three datasets. The statistics of the benchmarks are shown in Table 1. We randomly pick several samples from training data to make examples in prompts (Section 2.3). Then, we measure the performance of different approaches on test data. Because HumanEval does not contain train data, we reuse examples from MBPP in HumanEval.

3.3 Evaluation Metrics

Following previous code generation studies [6, 7, 17, 40], we use Pass@ k as our evaluation metrics. Specifically, given a requirement, a code generation model is allowed to generate k programs. The requirement is solved if any generated programs pass all test cases. We compute the percentage of solved requirements in total requirements as Pass@ k . For Pass@ k , a higher value is better. In our experiments, k is set to 1, 3, and 5, because we think that developers mainly use Top-5 outputs in real-world scenarios.

Previous work [1, 6, 7] found that standard Pass@ k has high variance and proposed an unbiased Pass@ k . We follow previous

work and employ the unbiased Pass@ k . Specifically, we generate $n \geq k$ programs per requirement (in this paper, we use $n = 20$, $k \in [1, 3, 5]$), count the number of solved requirements c , and calculate the unbiased Pass@ k :

$$\text{Pass@}k := \mathbb{E}_{\text{Problems}} \left[1 - \frac{\binom{n-c}{k}}{\binom{n}{k}} \right] \quad (1)$$

We also notice that previous code generation studies use text-similarity-based metrics (e.g., BLEU [21]). These metrics are initially designed for natural language generation and are poor in measuring the correctness of programs [7]. Thus, we omit these metrics in our experiments.

3.4 Comparison Baselines

This paper proposes a new prompting technique for code generation. To assess the effectiveness of our approach, we select three mainstream prompting techniques as baselines.

- **Zero-shot prompting** [7] directly feeds the requirement into LLMs without examples. Then, it extracts a generated program from LLMs' outputs.
- **Few-shot prompting** [7] randomly selects several < requirement, code > pairs as examples. Given a requirement, it concatenates examples and the requirement together, making a prompt. Then, the prompt is fed into LLMs, and LLMs predict a new program.
- **Chain-of-Thought (CoT) prompting** [35] is a variant of few-shot prompting. CoT prompting produces a special prompt consisting of < requirement, CoT, code > triples as examples. A CoT is several intermediate natural language reasoning steps.

To ensure the fairness of comparison, all baselines and SCoT prompting have the same number of examples (i.e., three examples) and example seeds.

The criteria for selecting baselines are three-fold. (1) SCoT prompting is a prompting technique for code generation. Thus, we directly compare it to existing prompting techniques for code generation. We also notice some emerging prompting techniques in other fields, such as Self-Consistency [31] and Least-to-Most [41]. But these approaches are designed for specific tasks (e.g., Arithmetic reasoning) and can not be directly applied to code generation. Thus, we omit them in this paper. (2) Our approach is to augment LLMs and can be flexibly applied to different LLMs. Thus, we do not directly compare LLMs to our approach. (3) We also omit some rank techniques for code generation [6]. They first use LLMs to generate many candidates and then leverage test cases or neural networks to rerank candidates. We think our work and these rank techniques are complementary. Users can use our approach to generate programs and then use post-processing techniques to select the final output. We further discuss the complementarity through experiments in Section 5.2.

3.5 Base Large Language Models

There are many available LLMs for source code. Our motivation is that existing LLMs can be divided into two categories: standard

language models and instruction-tuned models. For each category, we pick a representative model as the base model.

(1) Standard language models are pre-trained on a large-scale corpus with the next-token prediction objective. They are mainly used to continually complete the given content, such as code completion. Thus, we pick the state-of-the-art completion model for code - Codex [7] as a base model.

Codex [7] is a powerful language model for code generation, which supports a commercial application - GitHub Copilot [9]. Codex's training data contains both natural language and billions of lines of code. We use OpenAI's APIs to access the latest version of Codex with 175 billion parameters, i.e., code-davinci-002 [19].

(2) Instruction-tuned models refer to LLMs after instruction tuning. Instruction tuning trains LLMs to understand human users' instructions and perform tasks based on the instructions. We select the state-of-the-art instruction-tuned model - ChatGPT [18] as a base model.

ChatGPT [18] is the state-of-the-art LLM for code generation. ChatGPT is trained with extensive natural language text and code files. Then, it is trained with reinforcement learning and learns to follow human instructions. We use OpenAI's APIs to access the ChatGPT, i.e., gpt-3.5-turbo-0301 [18].

Our approach does not rely on specific LLMs and can be applied to different LLMs in a plus-and-play fashion. In the future, we will explore its usage on more powerful LLMs.

3.6 Sampling Settings

Following previous studies [7, 17, 40], we use nucleus sampling [11] to decode programs from LLMs. To ensure the fairness of experiments, all baselines and SCoT prompting generate 20 programs per requirement. The details of sampling settings are shown as follows.

Baselines. The temperature is 0.8 and the top- p is 0.95. For zero-shot and few-shot prompting, the maximum generated length is 300 tokens. The maximum generated length of CoT prompting is 600 tokens. Our motivation is that CoT prompting needs to generate intermediate reasoning steps and code. Thus, it requires a larger generation length.

SCoT prompting. In the first step, we sample 20 individual SCoTs from LLMs per requirement. The temperature is 0.8 and the top- p is 0.95. The maximum generated length is 300 tokens. Then, for each SCoT, we use LLMs to generate a corresponding program. The temperature is 0 and the maximum generated length is 300 tokens. Finally, we obtain 20 programs for each requirement. The total generation length of two steps is the same as CoT prompting.

4 RESULTS AND ANALYSIS

4.1 RQ1: How does SCoT prompting perform in terms of accuracy compared to baselines?

In the first research question, we apply SCoT prompting and baselines to three benchmarks and use unit tests to measure the correctness of generated programs.

Setup. We apply baselines and SCoT prompting to two LLMs (Section 3.5). Then, we measure the performance of different approaches on three code generation benchmarks (Section 3.2) using the Pass@ k (Section 3.3).

Table 2: The Pass@k (%) of SCoT prompting and baselines on three code generation benchmarks. The numbers in red denote SCoT prompting’s relative improvements compared to the SOTA baseline - CoT prompting.

Base Model	Prompting Technique	HumanEval			MBPP			MBCPP		
		Pass@1	Pass@3	Pass@5	Pass@1	Pass@3	Pass@5	Pass@1	Pass@3	Pass@5
ChatGPT	Zero-shot prompting	49.73	66.07	71.54	37.07	43.54	48.58	47.53	60.09	64.22
	Few-shot prompting	52.47	69.32	74.10	40.00	49.82	53.13	52.58	63.03	66.11
	CoT prompting	53.29	69.76	75.52	41.83	51.04	54.57	53.51	63.84	67.03
	SCoT Prompting	60.64	73.53	77.32	46.98	55.31	58.36	57.06	65.70	68.70
Relative Improvement		13.79%	5.40%	2.38%	12.31%	8.37%	6.95%	6.63%	2.91%	2.49%
Codex	Zero-shot prompting	40.20	61.78	68.11	27.07	43.81	47.93	40.25	54.17	60.65
	Few-shot prompting	42.93	62.96	70.10	33.17	45.72	49.62	44.12	57.65	62.45
	CoT prompting	43.79	63.41	71.56	35.66	46.57	50.11	45.79	58.92	62.56
	SCoT Prompting	49.82	66.56	75.14	38.29	50.74	53.16	48.34	60.77	64.19
Relative Improvement		13.77%	4.97%	5.00%	7.38%	8.95%	6.09%	5.57%	3.14%	2.61%

Results. The Pass@k ($k \in [1, 3, 5]$) of different approaches are shown in Table 2. The numbers in red denote SCoT prompting’s relative improvements compared to the SOTA baseline - CoT prompting.

Analyses. (1) SCoT prompting achieves the best results among all baselines. Table 2 shows that SCoT prompting can generate more correct programs than baselines on three benchmarks. Compared to the SOTA baseline - CoT prompting, in terms of Pass@1, SCoT prompting outperforms it by up to 13.79% in HumanEval, 12.31% in MBPP, and 6.63% in MBCPP. Pass@1 is a strict metric and it is difficult to improve. The results show that SCoT prompting can significantly improve the accuracy of LLMs on code generation and is more promising than existing prompting techniques. (2) SCoT prompting is effective in different LLMs and programming languages. SCoT prompting is effective in different LLMs. Compared to CoT prompting, in terms of Pass@1, SCoT prompting further improves ChatGPT by up to 13.79% and Codex by up to 13.77%. Besides, SCoT prompting is language-agnostic and can be applied to different programming languages. As shown in Table 2, SCoT prompting brings substantial improvements in Python (*i.e.*, HumanEval and MBPP) and C++ (*i.e.*, MBCPP). (3) SCoT prompting unlocks the reasoning ability of LLMs on code generation. LLMs can benefit from generating intermediate reasoning steps. The baseline - CoT prompting utilizes natural language steps but only brings slight improvements. In terms of Pass@1, CoT prompting improves few-shot prompting by up to 2% in HumanEval, 7.51% in MBPP, and 3.79% in MBCPP. In this paper, we introduce program structures into intermediate reasoning steps and propose a Structured Chain-of-Thought (SCoT). The SCoT constrains LLMs to use program structures to generate intermediate steps, moving in the direction of code. In terms of Pass@1, SCoT prompting improves few-shot prompting by up to 16.05% in HumanEval, 17.45% in MBPP, and 9.56% in MBCPP. The improvements show that SCoT prompting further unlocks the reasoning ability of LLMs in code generation.

Table 3: The results of human evaluation in three aspects. The numbers in red denote SCoT prompting’s relative improvements compared to the SOTA baseline - CoT prompting. All the p -values are substantially smaller than 0.05.

Approach	Correctness	Code Smell	Maintainability
Zero-shot prompting	1.012	1.523	1.372
Few-shot prompting	1.119	1.653	1.552
CoT prompting	1.225	1.689	1.616
SCoT prompting	1.412	1.869	1.873
Relative Improvement	15.27%	10.66%	15.90%

Answer to RQ1: SCoT prompting achieves higher accuracy than baselines on three benchmarks. In terms of Pass@1, SCoT prompting outperforms the SOTA baseline by up to 13.79% in HumanEval, 12.31% in MBPP, and 6.63% in MBCPP. The significant improvements show the effectiveness of SCoT prompting in code generation.

4.2 RQ2: Do developers prefer programs generated by SCoT prompting?

The ultimate goal of code generation is to assist developers in writing programs. In this RQ, we hire 10 developers (including industry employees and academic researchers) to manually review the programs generated by SCoT prompting and baselines.

Setup. To ensure the fairness of evaluation, we follow settings of human evaluation in previous studies [10, 14]. We have carefully checked the evaluation settings and think our settings are reliable. Specifically, we manually evaluate generated programs in the following aspects:

- **Correctness (whether the program satisfies the requirement).** 0 point: the program is totally inconsistent with the requirement. 1 point: the program is implemented, but misses some details. 2 points: the program is correctly implemented.
- **Code Smell (whether the program contains bad code smell).** 0 point: There is a serious code smell. 1 point: some details are not

```

Requirement:
# Return True is list elements are monotonically increasing or
decreasing.
-----
Few-shot prompting:
def monotonic(l: list):
    if all(l[i] <= l[i+1] for i in range(len(l)-1)) or all(l[i] >=
l[i+1] for i in range(len(l)-1)):
        return True
    else:
        return False
-----
SCoT prompting (SCoT & Source Code):
Input: l: list
Output: True or False
1: for each element in l do
2:   if the element is greater than its
   previous element then
3:     l increases
4:   if the element is less than its
   previous element then
5:     l decreases
6: if both increase and decrease then
7:   return False
8: else
9:   return True
-----
def monotonic(l: list):
    increasing = False
    decreasing = False
    for i in range(1, len(l)):
        if l[i] > l[i-1]:
            increasing = True
        if l[i] < l[i-1]:
            decreasing = True
    if increasing and decreasing:
        return False
    else:
        return True

```

Figure 5: Two programs generated by few-shot prompting and SCoT prompting, respectively.

in place. There is code smell of low severity. 2 points: the details are in place. No obviously better code in terms of performance exists. If possible, resources are released accordingly. No obvious code smell.

- **Maintainability (whether the implementation is standardized and has good readability).** 0 point: the program does not follow a consistent specification, or there are many meaningless names in variable naming, or there are certain repetitions and redundant code. 1 point: the program implementation meets certain specifications. But some variable names can be further refined. 2 points: the program implementation is relatively standardized. The variable naming is basically semantically straightforward, and the readability is good.

We explain the above aspects to evaluators through some examples. We also discuss with evaluators and set the score of each aspect to an integer, ranging from 0 to 2 (from bad to good). For SCoT prompting and baselines, we select a fixed LLM as the base model (*i.e.*, ChatGPT) and collect 200 generated programs per approach. Finally, we obtain 800 programs for evaluation. We invite 10 developers with 3-5 years of development experience to evaluate the programs in the form of a questionnaire. The evaluators include industry employees and academic researchers that are not co-authors of this paper. The 800 programs are divided into 5 groups, with each questionnaire containing one group. The programs are randomly shuffled and anonymously reviewed by evaluators. Each group is evaluated by two evaluators, and the final score is the average of two evaluators' scores. Evaluators are allowed to search the Internet for unfamiliar concepts.

Results. The results of the human evaluation are shown in Table 3. The numbers in red denote SCoT prompting's relative improvements compared to the SOTA baseline - CoT prompting. All the p -values are substantially smaller than 0.05.

Analyses. SCoT prompting achieves the highest scores in all three aspects among all baselines. Specifically, SCoT prompting outperforms the SOTA baseline - CoT prompting by 15.27% in correctness, 10.66% in code smell, and 15.90% in maintainability.

We attribute the improvements to our proposed SCoT. The SCoT constrains LLMs to use program structures to generate intermediate reasoning steps. It allows LLMs to explore diverse solutions with three basic structures, improving the correctness of the code. Then, based on the SCoT, LLMs focus on implementing a program in a standardized way. Thus, the generated programs contain fewer code smells than ones from baselines.

Figure 5 shows two programs generated by SCoT prompting and few-shot prompting, respectively. Both programs pass unit tests. But the program from few-shot prompting contains a very complex statement highlighted in Figure 5). Developers have to spend lots of effort to understand and maintain this program. In contrast, the program from SCoT prompting has good readability, and the SCoT clearly explains the behavior of the code. Developers can further use the SCoT as comments of the program for future maintenance.

Answer to RQ2: Human developers prefer programs generated by SCoT prompting. Specifically, SCoT prompting outperforms the SOTA baseline by 19.93% in correctness, 11.25% in code smell, and 16.17% in maintainability. A case study also shows the program from SCoT prompting is easy to read and maintain.

4.3 RQ3: Is SCoT prompting robust to examples?

As stated in Section 2.3, SCoT prompting requires manually written examples to make prompts. In practice, people may write different examples, which makes the performance of SCoT prompting varies. Thus, in this RQ, we explore the robustness of SCoT prompting to examples.

Setup. As stated in Section 2.3, we select some <requirement, code> pairs as example seeds and manually write SCoTs for them, obtaining examples in prompts. In this RQ, we measure the robustness of SCoT prompting to examples in two aspects, *i.e.*, seed selection and writing style.

(1) Seed Selection. It aims to validate SCoT prompting does not rely on specific seeds. We select three groups of <requirement, code> pairs as seeds and ask an annotator to write SCoTs for them. Then, we obtain three groups of examples. We measure the performance of SCoT prompting with different groups of examples. **(2) Writing Style.** People have different writing styles. It aims to validate that SCoT prompting does not rely on specific writing styles. We hire three annotators to independently write SCoTs for the same example seed, and obtain three groups of examples. Annotator A is a Ph.D. student in software engineering. Annotator B is a product manager from the industry. Annotator C is a developer from the industry. Then, we measure the performance of SCoT prompting with different annotators.

For comparison, we also measure the robustness of CoT prompting in the above settings. We select ChatGPT as the base model and conduct evaluations in HumanEval.

Results. The results are shown in Table 5 and 6, respectively.

Analyses. SCoT prompting is robust to examples. As shown in Table 5 and 6, SCoT prompting substantially outperforms CoT prompting when using different example seeds or annotators. It validates that SCoT prompting does not depend on specific seeds or writing styles. It also shows that the improvements of SCoT

Table 4: The results of ablation study. The base model is ChatGPT.

Prompting Technique	HumanEval			MBPP			MBCPP		
	Pass@1	Pass@3	Pass@5	Pass@1	Pass@3	Pass@5	Pass@1	Pass@3	Pass@5
CoT prompting	53.29	69.76	75.52	41.83	51.04	54.57	53.51	63.84	67.03
SCoT prompting	60.64	73.53	77.32	46.98	55.31	58.36	57.06	65.70	68.70
w/o Basic structures	55.67	70.94	76.13	43.36	53.64	56.57	54.79	64.32	67.77
w/o IO structure	59.65	72.79	77.12	46.13	54.76	57.88	56.61	65.01	68.42

Table 5: The Pass@k of CoT prompting and SCoT prompting with different example seeds.

Seed	CoT prompting			SCoT prompting		
	Pass@1	Pass@3	Pass@5	Pass@1	Pass@3	Pass@5
Seed A	53.29	69.76	75.52	60.64	73.53	77.32
Seed B	52.81	68.97	74.55	60.27	73.11	77.16
Seed C	51.36	67.44	73.62	59.36	72.88	76.79

Table 6: The Pass@k of CoT prompting and SCoT prompting with different annotators.

Annotator	CoT prompting			SCoT prompting		
	Pass@1	Pass@3	Pass@5	Pass@1	Pass@3	Pass@5
Annotator A	53.29	69.76	75.52	60.64	73.53	77.32
Annotator B	51.43	67.92	73.44	59.48	72.16	76.44
Annotator C	52.18	68.45	74.71	60.02	73.15	77.24

prompting come from the program structures instead of specific details in examples.

We also notice that there are slight variances in the performance of SCoT prompting under different settings. It is expected for prompting techniques using examples. Similar variances can be found in SCoT prompting, and SCoT prompting still outperforms CoT prompting in different settings.

Answer to RQ3: SCoT prompting is robust to examples. Under different example seeds or writing styles, SCoT prompting substantially outperforms the SOTA baseline - CoT prompting.

4.4 RQ4: What are the contributions of different program structures in SCoT prompting?

As stated in Section 2.1, SCoT prompting introduces basic structures (i.e., sequence, branch, and loop) and the input-output structure. This RQ is designed to analyze the contributions of different program structures.

Setup. We select ChatGPT as the base model. Then, we conduct an ablation study by independently removing basic structures and the input-output (IO) structure. When removing basic structures, we use a CoT with an IO structure as the intermediate steps. When removing the IO structure, the SCoT only contains a solving process with basic structures. We select ChatGPT as the base model.

Results. The results are shown in Table 4. “w/o” is the abbreviation of without.

Analyses. (1) Three basic structures are beneficial to design a feasible solving process. In Table 4, after removing basic structures,

SCoT prompting without basic structures:

```
Input: array: list[list]
Output: result: int or float
1. Initialize a result with -999999
2. Iterate through the list of lists ?
3. Calculate the sum of the list
4. Update the result with the maximum of sum and result
5. Return the result
```

SCoT prompting:

```
Input: array: list[list]
Output: result: int or float
1: Initialize a result with -999999
2: for _list in the list of lists: ✓
3: Calculate the sum of the _list
4: Update the result with the maximum of sum and result
5: return the result
```

Figure 6: The comparison of SCoT prompting and SCoT prompting without basic structures.

the performance of SCoT prompting drops obviously. We carefully inspect failed cases and find that LLMs benefit from using basic structures to clearly write a solving process. Figure 6 shows the intermediate steps of SCoT prompting and SCoT prompting without basic structures. SCoT prompting without basic structures uses CoTs, which sequentially describe how to write the code line by line and contain many ambiguities. For example, the scopes of two iterations on lines 2 and 4 are unclear. LLMs are likely to misunderstand the CoT and generate incorrect code. In contrast, SCoT prompting uses three basic structures to describe the solving process. The SCoT is clear and is similar to code, benefiting the following code implementation.

(2) The IO structure benefits the requirement understanding. In Table 4, after deleting the IO structure, the performance of SCoT prompting has a slight decrease. We analyze failed cases and think the IO structure benefits the requirement understanding. Figure 7 shows two programs from SCoT prompting and SCoT prompting without the IO structure. We can see that SCoT prompting without the IO structure wrongly understands the output format and generates an incorrect program. After adding the IO structure, LLMs first reason about the input-output format and correctly return a boolean value.

Table 7: The comparison of SCoT-P prompting and SCoT prompting. The numbers in red denote SCoT prompting’s relative improvements compared to SCoT-P prompting.

Approach	HumanEval			MBPP			MBCPP		
	Pass@1	Pass@3	Pass@5	Pass@1	Pass@3	Pass@5	Pass@1	Pass@3	Pass@5
CoT prompting	53.29	69.76	75.52	41.83	51.04	54.57	53.51	63.84	67.03
SCoT-P prompting	55.23	70.33	75.94	43.28	52.16	55.77	54.25	64.09	67.78
SCoT prompting	60.64	73.53	77.32	46.98	55.31	58.36	57.06	65.70	68.70
Relative Improvement	9.80%	4.55%	1.82%	8.55%	6.04%	4.64%	5.18%	2.51%	1.36%

SCoT prompting without IO structure:
<pre>def test_duplicate(arraynums): num_set = set(arraynums) if len(num_set) < len(arraynums): print('Find duplicate elements') else: print('No duplicate elements')</pre>

SCoT prompting:
<pre>def test_duplicate(arraynums): # Input: arraynums, a list of integers # Output: True if exist duplicate element, # False otherwise num_set = set(arraynums) if len(num_set) < len(arraynums): return True else: return False</pre>

Figure 7: The comparison of SCoT prompting and SCoT prompting without the IO structure.

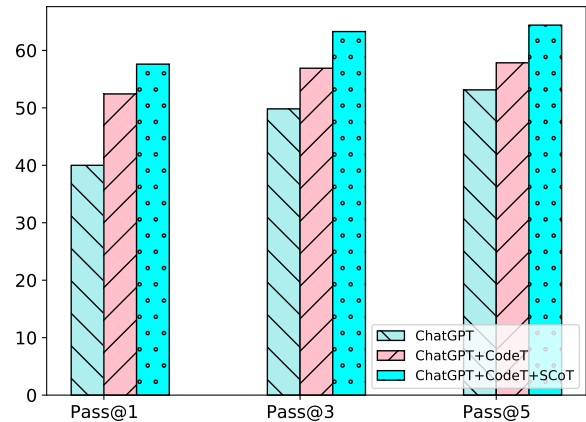
Answer to RQ3: The input-output structure helps LLMs understand requirements and improves ChatGPT by up to 6.37% in Pass@1. Three basic structures are beneficial to clearly describe a solving process and improve ChatGPT by up to 12.73% in Pass@1.

5 DISCUSSION

5.1 SCoT vs. Pseudocode

We notice that the SCoT is similar to the pseudocode. The SCoT and pseudocode both contain an input-output structure and a solving process. We randomly select 100 generated SCoTs and manually review them. We find that 26% of SCoTs are very close to the pseudocode. On one hand, we think the similarity enhances the usability of our approach. For example, users can quickly know the behavior of a program based on its SCoT. The SCoT also can be inserted into the comment and benefits future maintenance. On the other hand, the majority of SCoTs (74%) are different from the pseudocode because they are more abstract. Specifically, SCoTs tend to use natural languages to summarize an operation, e.g., calculate the sum of list1. But the pseudocode contains more implementation details, e.g., `sum ← 0; for i in list1: sum ← sum + i;`.

Compared to the pseudocode, we think the SCoT is a better choice for intermediate steps. Because a SCoT naturally decomposes code generation into two steps. LLMs first focus on exploring diverse solutions and then implement a program in a standardized way. To validate this point, we design a variant of SCoT prompting,

**Figure 8: The complementarity between CodeT and SCoT prompting.**

named SCoT-P Prompting. It is similar to SCoT prompting, but considers the pseudocode as intermediate steps. We apply SCoT-P Prompting and SCoT prompting to ChatGPT and measure their accuracy. The results are shown in Table 7. SCoT prompting substantially outperforms SCoT-P Prompting on three benchmarks. The improvements show the superiority of our SCoT.

5.2 SCoT prompting vs. Rank Techniques

Some recent studies [6, 12] propose *rank techniques* to improve the performance of LLMs on code generation. Given a requirement, they first sample many programs from LLMs and then use test cases or neural networks to rerank sampled programs. For example, CodeT [6] is a popular rank technique. CodeT does large-scale sampling and executes sampled programs on auto-generated test cases. Based on execution results, the programs are reranked. In this paper, we do not directly compare our approach to rank techniques due to two reasons.

(1) SCoT prompting and rank techniques have different focuses, and they are complementary. Our work aims to design a new prompting technique and improve the accuracy of LLMs in code generation. Rank techniques do not care about LLMs and aim to pick the best one from LLMs’ multiple outputs. In practice, users can use SCoT prompting to generate many programs and then use rank techniques to pick a final output.

To verify the complementarity between SCoT prompting and rank techniques, we conduct an exploratory experiment. We select ChatGPT as a base model and progressively introduce CodeT and SCoT prompting. The results on MBPP are shown in Figure 8. We

can see that the performance of ChatGPT is continually improved by adding CodeT and SCoT prompting.

(2) Rank techniques approaches rely on execution environments.

Rank techniques require executing programs on test cases and using execution results to rerank programs. In many realistic programming scenarios, users want to get code suggestions for an unfinished project. It is infeasible to execute auto-generated programs. Thus, we think rank techniques have limited application scenarios and make additional use of the execution results. Our approach works in a general scenario and does not use execution results. Thus, it is unfair to directly compare SCoT prompting to rank techniques.

5.3 Threats to Validity

There are three main threats to the validity of our work.

(1) The generalizability of experimental results. To mitigate this threat, we carefully select the benchmarks, metrics, and baselines. Following previous studies [1, 2, 7], we pick three representative code generation benchmarks. They are hand-written or collected from real-world programming communities, and cover two popular languages (*i.e.*, Python and C++). For evaluation metrics, we select a widely used metric - Pass@*k*, which utilizes test cases to check the correctness of programs. We use the unbiased Pass@*k* which is more reliable [7]. For comparison baselines, we select the SOTA prompting techniques and conduct a comprehensive comparison in Section 4. SCoT prompting and baselines have the same example seeds and maximum generation lengths.

(2) The impact of the two-step pipeline. CoT prompting generates a CoT and the code in one step. Our SCoT prompting generates the code in two steps. It first generates SCoTs and then generates the code. It is possible that the improvements come from the two-step pipeline. To solve this threat, we have two considerations. First, LLMs in our experiments are auto-regressive language models. For an auto-regressive language model, a one-step pipeline and a two-step pipeline are theoretically equivalent. Second, we conduct an ablation study in Section 4.4. We keep the two-step pipeline unchanged and remove program structures. The results in Table 4 show that SCoT prompting without prompt structures has a significant drop in the Pass@*k*. It shows that the improvements of SCoT prompting are brought by program structures instead of the two-step pipeline.

(3) The data leakage. Existing LLMs are trained with extensive code files from open-source communities. It is possible that their training data contains the experimental benchmarks, leading to data leakage. But we think that it does not affect the fairness of our experiments. In this paper, we select a specific LLM (*e.g.*, ChatGPT) as the base model and apply different prompting techniques to it. Thus, the reported relative improvements between baselines and our approach are credible. In the future, we will add the latest benchmarks to alleviate this threat.

6 RELATED WORK

Large language models (LLMs) for Source Code are large-scale neural networks that are pre-trained with a large corpus consisting of natural language text and source code. Nowadays, LLMs for source code have been expanding and can be divided into two categories: standard language models and instruction-tuned models.

Standard Language models are pre-trained on a large-scale corpus with the next-token prediction objective. They are mainly used to continually complete the given context, such as code completion. After the success of GPT series [4, 23, 24] in NLP, OpenAI fine-tunes GPT models on code to produce closed-source Codex [7]. There follow many open-source replication attempts, *e.g.*, CodeParrot [29], CodeGen [17], CodeGeeX [40], InCoder [8], StarCoder [15] and CodeT5+ [33].

Instruction-tuned models are models after instruction tuning [34]. Instruction tuning trains models to understand human users' instructions and perform tasks by following instructions. ChatGPT [18] is trained with human feedback [20], powerful on both natural language tasks and programming tasks. Many attempt to train an "open-source ChatGPT". Alpaca [27] is LLaMA [28] tuned using self-instruct [32] and ChatGPT feedback. Code Alpaca [5] is LLaMA tuned using self-instruct and ChatGPT feedback, with instructions focusing on programming tasks. WizardCoder [16] is StarCoder [15] tuned using Evol-Instruct [36] and ChatGPT feedback with Code Alpaca's dataset as seed dataset. InstructCodeT5+ [33] is CodeT5+ [33] tuned on Code Alpaca's dataset.

Prompting Techniques. With the enormous number of parameters (*e.g.*, Codex: 175 billion parameters), it is hard to directly fine-tune LLMs on code generation. *Prompting techniques* are a popular approach, which leverages LLMs to generate code by inputting a special prompt.

Early, researchers proposed zero-shot prompting and few-shot prompting. Zero-shot prompting concatenates a task instruction (*e.g.*, please generate a program based on the requirement) and a requirement together, making a prompt. Based on the zero-shot prompting, few-shot prompting further adds several `< requirement, code >` pairs to the prompts, so that LLMs can learn code generation from given examples.

The Chain-of-Thought (CoT) prompting [35] is a recently proposed prompting technique. CoT prompting asks LLMs first to generate CoTs (*i.e.*, intermediate natural language reasoning steps) and then output the final code. It allows LLMs to first design a solving process that leads to the code. CoT prompting has achieved the SOTA results in natural language generation and sparked lots of follow-up research, such as self-consistency prompting [31], least-to-most prompting [41]. But these prompting techniques are designed for natural language generation and bring slight improvements in code generation.

In this paper, we propose a novel prompting technique named Structured Chain-of-Thought (SCoT) prompting. Different from standard CoT prompting, SCoT prompting explicitly introduces program structures and asks LLMs to generate intermediate reasoning steps with program structures. We compare CoT prompting and SCoT prompting in Section 4. The results show that SCoT prompting significantly outperforms CoT prompting in three benchmarks.

7 CONCLUSION AND FUTURE WORK

Large Language Models (LLMs) with Chain-of-Thought (CoT) prompting is the state-of-the-art (SOTA) approach to generating code. It first generates a CoT and then outputs the code. A CoT is several intermediate natural language reasoning steps. However, CoT prompting still has low accuracy in code generation. This paper

proposes a Structured CoT (SCoT) and presents a new prompting technique for code generation, named SCoT prompting. SCoT prompting asks LLMs to generate a SCoT using program structures (i.e., sequence, branch, and loop structures). Then, LLMs generate the code based on the SCoT. A large-scale study on three benchmarks shows that SCoT prompting significantly outperforms CoT prompting in Pass@k and human evaluation. Besides, SCoT prompting is robust to examples and obtains stable improvements.

In the future, we will explore new prompting techniques for code generation. For example, source code can be represented by a tree (e.g., abstract syntax tree). We can design a tree-based prompting technique, which uses LLMs to generate a tree.

REFERENCES

- [1] Ben Athiwaratkun, Sanjay Krishna Gouda, Zijian Wang, Xiaopeng Li, Yuchen Tian, Ming Tan, Wasi Uddin Ahmad, Shiqi Wang, Qing Sun, Mingyue Shang, et al. 2022. Multi-lingual Evaluation of Code Generation Models. *arXiv preprint arXiv:2210.14868* (2022).
- [2] 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* (2021).
- [3] Corrado Böhm and Giuseppe Jacopini. 1966. Flow diagrams, turing machines and languages with only two formation rules. *Commun. ACM* 9, 5 (1966), 366–371. <https://doi.org/10.1145/355592.365646>
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [5] Sahil Chaudhary. 2023. Code Alpaca: An Instruction-following LLaMA model for code generation. <https://github.com/sahil280114/codealpaca>.
- [6] Bei Chen, Fengji Zhang, Anh Nguyen, Daoguang Zan, Zeqi Lin, Jian-Guang Lou, and Weizhu Chen. 2022. CodeT5: Code Generation with Generated Tests. <https://doi.org/10.48550/ARXIV.2207.10397>
- [7] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolos Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. (2021). [arXiv:2107.03374](https://arxiv.org/abs/2107.03374) [cs.LG]
- [8] Daniel Fried, Armen Aghajanyan, Jessy Lin, Sida Wang, Eric Wallace, Freda Shi, Ruiqi Zhong, Scott Yih, Luke Zettlemoyer, and Mike Lewis. 2023. InCoder: A Generative Model for Code Infilling and Synthesis. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=hQwblM6EL>
- [9] GitHub. 2022. GitHub Copilot. <https://github.com/features/copilot>.
- [10] Yiyang Hao, Ge Li, Yongqiang Liu, Xiaowei Miao, He Zong, Siyuan Jiang, Yang Liu, and He Wei. 2022. AixBench: A Code Generation Benchmark Dataset. *CoRR* abs/2206.13179 (2022). <https://doi.org/10.48550/arXiv.2206.13179> [arXiv:2206.13179](https://arxiv.org/abs/2206.13179)
- [11] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. The Curious Case of Neural Text Degeneration. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net. <https://openreview.net/forum?id=ryGQYrFvH>
- [12] Jeevana Priya Inala, Chenglong Wang, Mei Yang, Andres Codaş, Mark Encarnación, Shuvendu Lahiri, Madanlal Musuvathi, and Jianfeng Gao. 2022. Fault-aware neural code rankers. *Advances in Neural Information Processing Systems* 35 (2022), 13419–13432.
- [13] Jia Li, Ge Li, Zhuo Li, Zhi Jin, Xing Hu, Kechi Zhang, and Zhiyi Fu. 2023. CodeEditor: Learning to Edit Source Code with Pre-Trained Models. *ACM Trans. Softw. Eng. Methodol.* (may 2023). <https://doi.org/10.1145/3597207> Just Accepted.
- [14] Jia Li, Yongmin Li, Ge Li, Zhi Jin, Yiyang Hao, and Xing Hu. 2023. SkCoder: A Sketch-based Approach for Automatic Code Generation. In *45th IEEE/ACM International Conference on Software Engineering, ICSE 2023, Melbourne, Australia, May 14-20, 2023*. IEEE, 2124–2135. <https://doi.org/10.1109/ICSE48619.2023.00179>
- [15] Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, et al. 2023. StarCoder: may the source be with you! *arXiv preprint arXiv:2305.06161* (2023).
- [16] Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. 2023. WizardCoder: Empowering Code Large Language Models with Evol-Instruct. *arXiv preprint arXiv:2306.08568* (2023).
- [17] Erik Nijkamp, Bo Pang, Hiroaki Hayashi, Lifu Tu, Huan Wang, Yingbo Zhou, Silvio Savarese, and Caiming Xiong. 2022. CodeGen: An Open Large Language Model for Code with Multi-Turn Program Synthesis. *arXiv preprint* (2022).
- [18] OpenAI. 2022. ChatGPT. <https://openai.com/blog/chatgpt>.
- [19] OpenAI. 2022. Codex. <https://beta.openai.com/docs/api-reference/completions>.
- [20] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems* 35 (2022), 27730–27744.
- [21] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*. 311–318.
- [22] Han Peng, Ge Li, Wenhan Wang, Yunfei Zhao, and Zhi Jin. 2021. Integrating Tree Path in Transformer for Code Representation. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, Marc'Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.). 9343–9354. <https://proceedings.neurips.cc/paper/2021/hash/4e0223a87610176ef0d24ef6d2cdc3a-Abstract.html>
- [23] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training. (2018).
- [24] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. (2019).
- [25] Marc'Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. Sequence Level Training with Recurrent Neural Networks. In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). <http://arxiv.org/abs/1511.06732>
- [26] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, 4149–4158. <https://doi.org/10.18653/v1/n19-1421>
- [27] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford Alpaca: An Instruction-following LLaMA model. https://github.com/tatsu-lab/stanford_alpaca.
- [28] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models. *arXiv preprint arXiv:2302.13971* (2023).
- [29] Lewis Tunstall, Leandro Von Werra, and Thomas Wolf. 2022. *Natural language processing with transformers*. " O'Reilly Media, Inc."
- [30] Wenhan Wang, Ge Li, Sijie Shen, Xin Xia, and Zhi Jin. 2020. Modular Tree Network for Source Code Representation Learning. 29, 4, Article 31 (sep 2020), 23 pages. <https://doi.org/10.1145/3409331>
- [31] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. <https://openreview.net/pdf?id=1PL1NIMMrw>
- [32] Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khoshabi, and Hannaneh Hajishirzi. 2023. Self-Instruct: Aligning Language Models with Self-Generated Instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Toronto, Canada, 13484–13508. <https://aclanthology.org/2023.acl-long.754>
- [33] Yue Wang, Hung Le, Akhilesh Deepak Gotmare, Nghi DQ Bui, Junnan Li, and Steven CH Hoi. 2023. Codet5+: Open code large language models for code understanding and generation. *arXiv preprint arXiv:2305.07922* (2023).
- [34] Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V Le. 2022. Finetuned Language Models are Zero-Shot Learners. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=gEzRGCozdqR>
- [35] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. Chain of Thought Prompting Elicits Reasoning in Large Language Models. In *Advances in Neural Information Processing Systems*, Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun

- Cho (Eds.). https://openreview.net/forum?id=_VjQlMeSB_J
- [36] Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. 2023. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244* (2023).
- [37] Jian Zhang, Xu Wang, Hongyu Zhang, Hailong Sun, Kaixuan Wang, and Xudong Liu. 2019. A novel neural source code representation based on abstract syntax tree. In *Proceedings of the 41st International Conference on Software Engineering, ICSE 2019, Montreal, QC, Canada, May 25-31, 2019*, Joanne M. Atlee, Tefvik Bultan, and Jon Whittle (Eds.). IEEE / ACM, 783–794. <https://doi.org/10.1109/ICSE.2019.00086>
- [38] Kechi Zhang, Zhuo Li, Jia Li, Ge Li, and Zhi Jin. 2023. Self-Edit: Fault-Aware Code Editor for Code Generation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, 769–787. <https://doi.org/10.18653/v1/2023.acl-long.45>
- [39] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International Conference on Machine Learning*. PMLR, 12697–12706.
- [40] Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Zihan Wang, Lei Shen, Andi Wang, Yang Li, Teng Su, Zhilin Yang, and Jie Tang. 2023. CodeGeeX: A Pre-Trained Model for Code Generation with Multilingual Evaluations on HumanEval-X. *arXiv:2303.17568* [cs.LG]
- [41] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V. Le, and Ed H. Chi. 2023. Least-to-Most Prompting Enables Complex Reasoning in Large Language Models. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net. <https://openreview.net/pdf?id=WZH7099tgfM>