

DOMAINEVAL: An Auto-Constructed Benchmark for Multi-Domain Code Generation

Qiming Zhu^{1,2*}, Jialun Cao^{3*}, Yaojie Lu^{1†}, Hongyu Lin¹,
Xianpei Han¹, Le Sun¹, Shing-Chi Cheung³

¹Chinese Information Processing Laboratory, Institute of Software, Chinese Academy of Sciences, Beijing, China

²University of Chinese Academy of Sciences, Beijing, China

³The Hong Kong University of Science and Technology, Hong Kong, China

{zhuqiming2022, luyaojie, hongyu, xianpei, sunle}@iscas.ac.cn
{jcaoap, scc}@cse.ust.hk

Abstract

Code benchmarks such as HumanEval are widely adopted to evaluate capabilities of Large Language Models (LLMs), providing insights into their strengths and weaknesses. However, current benchmarks primarily exercise LLMs' capability on *common coding tasks* (e.g., bubble sort, greatest common divisor), leaving *domain-specific coding tasks* (e.g., computation, system, cryptography) unexplored. To fill this gap, we propose a multi-domain code benchmark, DOMAINEVAL, designed to evaluate LLMs' coding capabilities thoroughly. Our pipeline works in a fully automated manner, enabling a push-button pipeline from code repositories into formatted subjects under study. Interesting findings are observed by evaluating 12 representative LLMs against DOMAINEVAL. We notice that LLMs are generally good at *computation* tasks while falling short on *cryptography and system* coding tasks. The performance gap can be as much as 68.94% (80.94% - 12.0%) in some LLMs. We also observe that generating more samples can increase the overall performance of LLMs, while the domain bias may even increase. The contributions of this study include a code generation benchmark dataset DOMAINEVAL, encompassing six popular domains, a fully automated pipeline for constructing code benchmarks, and an identification of the limitations of LLMs in code generation tasks based on their performance on DOMAINEVAL, providing directions for future research improvements.

Code — <https://github.com/domaineval>

Introduction

Large Language Models (LLMs) have revolutionized various areas such as question answering (Rogers, Gardner, and Augenstein 2023), math reasoning (Imani, Du, and Shrivastava 2023), and especially *software development* (Lozhkov et al. 2024). Stakeholders are eager to know whether and to what extent LLMs can improve development efficiency.

To this end, a variety of code generation benchmarks such as HumanEval (Chen et al. 2021) and MBPP (Austin et al.

*These authors contributed equally.

†Corresponding authors.

```

Network
import ipaddress
def is_valid_ipv6_address(ip_str):
    try:
        ipaddress.IPv6Address(ip_str)
    except ValueError:
        return False
    return True

Computation
def calculate_mean(lst):
    total = sum(lst)
    mean = total / len(lst)
    return mean

Basic
import datetime as dt
def get_weekday_occurrence(
    d: dt.date):
    xthday = 1+(d.day-1)//7
    return d.weekday(),xthday

Cryptography
def derive_key(
    master: bytes
) -> bytes:
    derived=HKDF(
        master,32,b'',SHA256
    )
    return derived

System
def free_physmem():
    out = sh(
        ['free', '-b'],
        env={'LANG': 'C.UTF-8'})
    lines = out.split('\n')
    for line in lines:
        if line.startswith('Mem'):
            total,used,free,shared=(int(x) for x in line.split()[1:5])
            nt = collections.namedtuple(
                'free','total used free shared output'
            )
            return nt(total, used, free, shared, out)
    raise ValueError("can't find 'Mem'")
  
```

Figure 1: Domain-specific code from DOMAINEVAL, illustrating distinct functional characteristics across domains.

2021) have been introduced and intensively used to evaluate LLMs' coding capability. They primarily consist of *common coding tasks* such as sorting an array or computing the greatest common divisor. Furthermore, to meet more *realistic* needs, more benchmarks are emerging, expanding mainly along with two perspectives. First, *linguistic diversity*. This line of work (Wang et al. 2024; Cassano et al. 2022; Chai et al. 2024) mitigates linguistic bias in both natural languages (e.g., English, Chinese) and programming languages (e.g., Python, Java, C/C++). Second, *code scale diversity*. This line of work aims to scale up the code granularity from

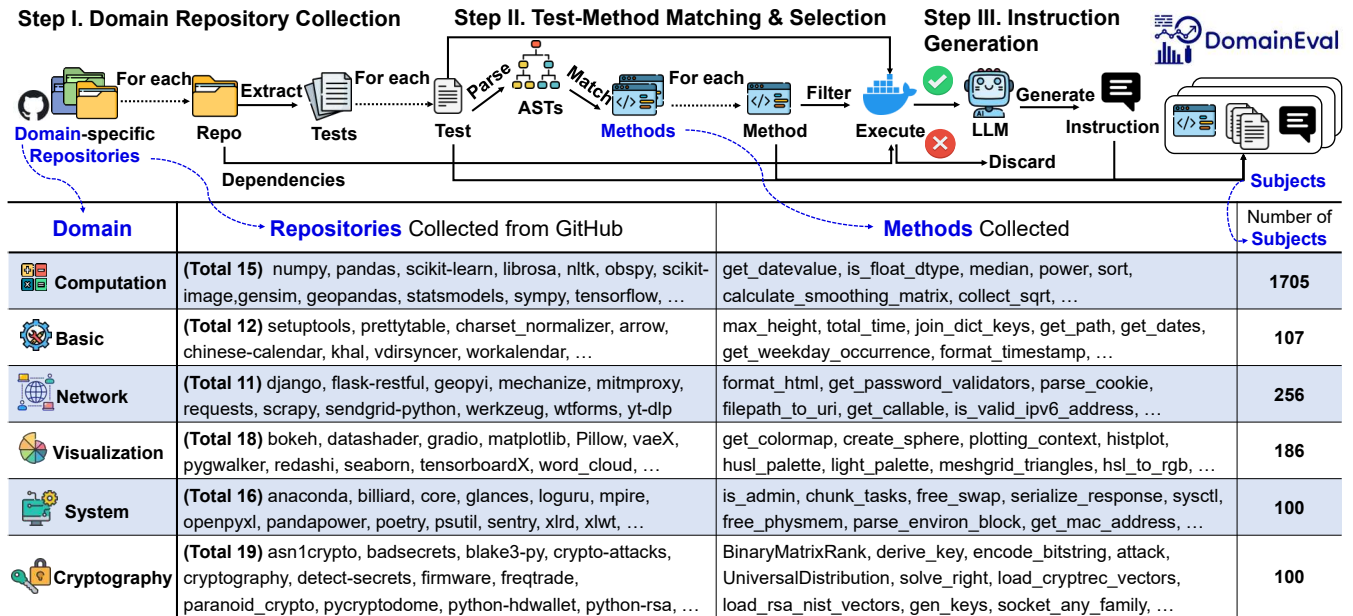


Figure 2: DOMAINEVAL construction pipeline. Figure details the domains, repositories and statistics in DOMAINEVAL.

function-level (Chen et al. 2021; Yu et al. 2024; Austin et al. 2021) to class-level (Du et al. 2023) and repo-level (Zhang et al. 2023; Cao et al. 2024a; Li et al. 2024b).

However, if LLMs are to be applied in real-world industrial scenarios where different product lines prioritize different domains of code, it becomes essential to further understand these LLMs’ coding capabilities across various domains. In other words, while existing benchmarks make significant efforts in varying natural/programming languages and code granularity, *the LLMs’ ability to generate domain-specific code remains under-explored*.

To fill this gap, we introduce DOMAINEVAL, a multi-domain code generation benchmark. It consists of 2454 subjects (*i.e.*, instruction, reference solution, test cases), along with 5892 test cases, covering six domains of code, *i.e.*, computation, network, basic operation, system, visualization, and cryptography. Figure 1 shows representative examples from each domain. We can see clear functional distinctions among codes from different domains. For example, code in *computation* involves computational tasks such as mean calculation. Code in *network* handles network request-communications and remote connections. Code in *cryptography* incorporates cryptographic algorithms, encrypting plain-text with a key or conducting key recovery attacks.

It is noteworthy that in order to facilitate the benchmark construction and welcome future contributions to DOMAINEVAL, we provide *a fully automated test-guided construction pipeline*. The pipeline works in a push-button manner, *i.e.*, given a code repository, it outputs a set of formatted subjects (*i.e.*, instruction, reference solution, test cases). Empowered by this pipeline, DOMAINEVAL is with exceptional *scalability*, capable of incorporating the ever-evolving code corpus into it. Moreover, the continuous in-

flux of updated code through this pipeline fortifies DOMAINEVAL against the data contamination threat (Cao et al. 2024b), thereby maintaining its integrity and novelty.

Our extensive experiments on 10+ LLMs indicate that LLMs are generally good at *computation*, with an average of 82.44% Pass@1, while falling short on *cryptography* and *system* domains, with an average of 33.08% Pass@1 and 37.50% Pass@1, respectively. The performance gap among domains can be as much as 68%+ from *Llama-2-13B-Chat* model, which is observed to have the largest performance variance compared with other LLMs. We also observe that generating more samples can increase the overall performance of LLMs, while the domain bias may even increase.

The contributions can be summarized as follows.

- We introduce DOMAINEVAL, a multi-domain code generation benchmark that consists of 2k+ subjects (*i.e.*, description, reference code and tests) covering six domains.
- We provide a fully automated test-guided construction pipeline to facilitate the benchmark construction.

Benchmark Construction

In this section, we introduce the whole construction pipeline of DOMAINEVAL, as shown in Figure 2. We provide a fully automated, test-guided construction algorithm that transforms code repositories into a collection of formatted subjects for LLM evaluation. Each subject consists of three components: instruction for LLM evaluation, reference solution, and a series of test cases, as illustrated in Figure 3.

The pipeline begins with the collection of raw code snippets from specific domains, which are then systematically transformed into suitable benchmark data. This process involves three key steps: (1) **Domain Repository Collection**: The first step involves collecting raw code data (R) from

📌 Instruction	✅ Reference Solution	📄 Test Cases
<p>[instruction]</p> <p>Functionality: The Commutator function is designed to calculate the commutator of ...</p> <p>Inputs:</p> <ul style="list-style-type: none"> - a (Expr): The first argument for ... - b (Expr): The second argument for ... <p>Outputs:</p> <ul style="list-style-type: none"> - Expr: The result of the commutator ... 	<p>[method_name] eval</p> <p>[full_method_name] Commutator.eval</p> <p>[method_path] sympy/sympy/physics/secondquant.py</p> <p>[method_code]</p> <pre>from sympy.core.singleton import S ... class Commutator(Function): is_commutative = False @classmethod def eval(cls, a, b): """built-in docstring""" if not (a and b): return S.Zero </pre>	<p>Test 0</p> <p>[test_path] .../sympy/physics/tests/test_secondquant.py</p> <p>[test_code]</p> <pre>from sympy.core.symbol import symbols ... def test_commutation(): ... a, b, c, d = symbols('a,b,c,d', above_fermi=True) ... assert Commutator.eval(c1, c1)==0 ... test_commutation() </pre>
<p>[method_code_mask]</p> <pre>from sympy.core.singleton import S ... class Commutator(Function): """built-in comments""" is_commutative = False @classmethod def eval(cls, a, b): [MASK] </pre>	<pre>class Commutator(Function): is_commutative = False @classmethod def eval(cls, a, b): """built-in docstring""" if not (a and b): return S.Zero </pre>	<p>Test N</p>

Figure 3: Examples of subjects in DOMAIN-EVAL. Each subject consists of three components: an instruction for LLM evaluation, a reference solution, and a set of test cases.

selected domains. These snippets are categorized into two groups: C_0 , containing function code, and T_0 , including tests designed to validate the correctness of the code in C_0 .

(2) **Test-Method Matching & Selection:** This step aims to create a set of candidate subjects $S' = (c, \hat{T})$ from the collected domain code, where each pair consists of a function code c and its corresponding test cases \hat{T} . To achieve this, we first match each function code c with its corresponding set of test cases \hat{T} , ensuring they are compatible with the execution environments E . We then filter C_0 to retain only those code snippets that are executable and accompanied by valid test functions, resulting in a refined set C_1 and its associated test cases, which together form the candidate subject set S' . (3) **Instruction Generation:** For each code snippet c in C_1 , we employ a LLM to generate corresponding instruction $i = \text{LLM}(c)$. The instruction i , along with the associated function code c and test cases \hat{T} , are combined to create the final subjects in benchmark dataset, with each entry represented as (i, c, \hat{T}) .

Domain Repository Collection

Referring to the work (Zhuo et al. 2024), we select six domains for code generation: *computation*, *network*, *basic operation*, *system*, *visualization*, and *cryptology*, to serve as the domain divisions for DOMAIN-EVAL.

To ensure that DOMAIN-EVAL closely aligns with the real-world code requirements of human engineers, we chose GitHub repositories as the source of raw data. We cloned a selection of representative code repositories from GitHub, particularly those with at least 100 stars, as these are considered high-quality code data and reflect the actual needs of engineers. The code repositories used in each domain are

shown in Figure 2. Given that each GitHub repository represents a specific real-world application scenario, we classify the code snippets based on the repository’s domain. In practice, we use the repository’s topic labels and README files to accurately assign the code to the appropriate domain.

Test-Method Matching & Selection

This section describes how to obtain reference solutions and their corresponding tests from code repositories to construct candidate subjects for LLM evaluation and the automated construction pipeline.

Given that code snippets from GitHub repositories are written by human engineers in real-world production environments, they tend to be more complex, with more dependencies and higher levels of encapsulation than standalone functions in some code datasets. This complexity makes it challenging to select function code that can serve as benchmark data while simultaneously acquiring corresponding test cases to form complete candidate subjects.

To address this challenge, we propose a test-method tracing strategy aimed at constructing candidate subjects for benchmarking. Our approach has two main steps: (1) Search the code repository for tests related to the reference code and perform Test-Method matching, where each function code c is automatically paired with its corresponding tests \hat{T} . (2) To ensure the final candidates are suitable for LLM evaluation and can smoothly pass through the automated construction pipeline, we filter the matched results using three criteria: *executable*, *significant*, and *appropriate difficulty*.

Test-Method Matching To match the function code with corresponding tests and package them into candidate subjects, we start with the test code and trace it back to the asso-

ciated function code. Specifically, we search for Python code files in repositories and use the library `ast` to parse Python code into an abstract syntax tree. Then, we sequentially traverse the nodes of the syntax tree, extracting all functions and their class context (if present). Next, we select test code snippets based on two heuristic rules: First, the function or class name of a test should contain `test` or `Test`. Second, the test code snippet should contain an `assert` statement.

To match the selected test code with the corresponding function code, we identify `ast.Call` nodes as function calls. To retrieve the implementation of the called function, we consider two scenarios: If the function and test are in the same file, we traverse the abstract syntax tree to locate and unpars the function node. If the function call spans files, we use recursive path matching by analyzing Python’s `import` behaviors until the specific file is found. Specifically, we continuously retrieve and replace the name of the next level path based on the content of `__init__.py` and the `import` statement in `from... import... as...` format, until the specific Python file is located. Once located, the process is the same as in the first scenario, using `ast` to parse and identify the function node.

In real-world code repositories, the correlation between tests and function code is often not one-to-one. To address this, we identify all functions within the test code and pair them with the corresponding function code. The function code is then used as the reference solution, and we group all related tests into a test suite. This process allows us to package function code snippets from GitHub repositories with their corresponding test cases, creating candidate subjects.

Test-Method Selection After packaging the candidate subjects, we continue to construct our benchmark dataset. We notice that not every function can directly convert to benchmark data suitable for LLM evaluation. Therefore, to facilitate the automatic construction of code benchmarks, we impose three criteria on the candidate subject.

First, **Executable**. To use Pass@k (Chen et al. 2021) for evaluation, reference solution code must be executable to verify semantic consistency with generated code against test cases. To ensure the security of our benchmark data, we utilize a sandbox to isolate code execution and maintain a list of banned keywords (Xie et al. 2024). If any of these prohibited keywords are detected, we consider the execution and evaluation of such code to be potentially risky and consequently discard it. The test environment first consists of a basic Python environment and necessary packages. To execute a piece of code, it needs adequate dependencies, so we concatenate the required context (e.g., `import` statements, class context, static variables) for the function. After preparing the context, we run the function code along with its tests.

Second, **Significant**. The code used for LLM evaluation should be important and meaningful, crucial in real-world production scenarios. For example, `__init__` functions primarily involve repetitive variable assignments, which are mechanical and lack significance, failing to reflect the capabilities of LLMs. In contrast, human engineers tend to write tests for code that implements critical functionality. Therefore, we extract the code with tests as candidates.

```
## System Message
You are preparing an interview for software engineers. The interviewees are going to complete the ${FUNCTION_NAME} function. Write a clear instruction describing this function in detail, which includes the functionality, the input arguments (if any), and the outputs (if any). Do not reveal test cases. Generate the instruction with the following format:
'''
Functionality: ...
Inputs: ...
Outputs: ...
'''

## Instruction
${CODE}
```

Figure 4: Dialogue prompt template for guiding the LLM to generate detailed *instruction* fields.

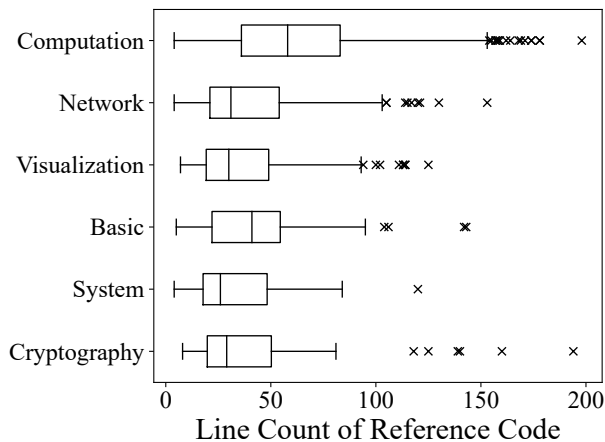


Figure 5: Distribution of line counts for the reference code (context included) across domains in `DOMAIN EVAL`. Box plots provide a visualization of data distribution, showing the interquartile range, median, normal range, and outliers.

Third, **Appropriate Difficulty**. The number of lines of reference code is one of the direct indicators of the complexity of the code generation task. We set a limit on the function implementation of reference code, *i.e.*, the standard answer, used in the task. We restrict them to between 3 and 100 lines. On the one hand, functions with fewer than three lines typically have overly simple logic (Yu et al. 2024) and do not effectively reveal the shortcomings of LLMs. On the other hand, functions exceeding 100 lines may contain overly complex logic, which can present significant challenges for evaluation. These challenges include exceeding LLMs’ context limitations or overwhelming their information capacity, making it difficult to generate precise instructions in subsequent steps and thus hindering the automatic construction pipeline.

Instruction Generation

After filtering the candidate subjects, our subjects still lack the *instruction* field. To ensure the reproducibility of the dataset construction process, we employ open-source LLM, *i.e.* *Qwen2-72B-Instruct-GPTQ-Int4*, for data generation. We utilize a *dialogue prompt template*, as depicted in Figure 4, to guide the LLM in generating detailed *instruction* fields for each validated subject. These *instruction* fields comprehensively describe the function, including its purpose, input arguments, and expected outputs, which serve as the input for evaluating the code generation capabilities of LLMs. By leveraging LLMs, we can generate natural language descriptions that outline the desired functionality, inputs, and outputs of the code (Xie et al. 2024; Li et al. 2024a).

Finally, to construct the template for code generation, we mask the function code, following prior work (Xie et al. 2024), by [MASK] as shown in the *method.code_mask* field in Figure 3, and instruct LLMs to complete the masked segments during the evaluation.

Benchmark Statistics

The above pipeline generates instructions, reference code, and a series of tests shown in Figure 3, which combine to form complete subjects. As a result, *DOMAINEVAL* consists of 2454 code subjects. Figure 2 shows the number of subjects constructed from repositories across different domains. Overall, the number of subjects from each domain is at least 100. The *computation* domain encompasses the greatest number of subjects, with 1705 subjects, compared to 100 ~ 256 subjects in other domains. This is because the computation-related repositories offer more function code and also include a larger number of test cases to ensure the accuracy of each computational operation.

Figure 5 illustrates the distribution of lines of reference code within *DOMAINEVAL*. Overall, the lines of reference code (context included) across six domains are similar, ranging from 4 to 198, with an average of 55.69. In particular, code in *computation* has slightly more lines of code than that in other domains, with an average of 63.20 lines of code compared with 33.95 ~ 42.03 in other domains.

Experiments

Experiment Setup

Studied LLMs. We assess 12 representative instruction-tuned LLMs against *DOMAINEVAL*, including GPT-3.5-turbo, GPT-4o-mini (Brown et al. 2020; Achiam et al. 2023), Qwen2 (Yang et al. 2024), Phi-3 (Abdin et al. 2024), DeepSeek-Coder series (Zhu et al. 2024; Guo et al. 2024), Llama2 (Touvron et al. 2023) and CodeLlama series (Roziere et al. 2023), CodeQwen1.5 (Bai et al. 2023) with size of open-source models from 6.7B to 72B. These models exhibit proficiency in following instructions and delivering appropriately formatted responses.

Evaluation Metrics. Our evaluation uses the unbiased version of Pass@k (Chen et al. 2021) to accurately assess the functional correctness of code snippets generated by LLMs. Following prior work (Zhuo et al. 2024), we report Pass@1 and Pass@5 for the experiment in zero-shot setting and use

macro-average as scores. For Pass@1 metric, we use greedy decoding, *i.e.* set temperature to 0.0. For Pass@5 metric, we opt for the minimum sample size $N = 5$ and maintain temperature at 0.2 and top-p at 0.95. For code generation tasks, we use *torch.bfloat16* when loading LLMs.

Evaluation Process. During the evaluation, to prevent LLMs from failing execution due to omitted import statements, which is a tolerable flaw but could potentially distort assessment results, we implement a corrective measure, *i.e.*, completing the missing dependencies based on the import scenario mentioned in the instruction.

Overall Result

Figure 6 shows LLMs' Pass@1 and Pass@5 against *DOMAINEVAL*. Columns plotted in blue show the Pass@1/5 values; the bluer, the larger. The columns in orange and green highlight the average ("Mean") and standard deviation ("Std") of the corresponding rows, respectively. Overall, the average performance across studied LLMs is similar, ranging from 49.89% ~ 67.13%. At the same time, the performance across domains varies, *i.e.*, the performance in *Computation* reaches the top among all LLMs, with an average of 82.44% Pass@1 and 88.57% Pass@5, while the worst scores are observed in *Cryptography* domain, with 33.08% Pass@1 and 40.25% Pass@5. In other words, the performance gaps between different domains cannot be ignored.

Domain Biases

From Figure 6, we can see significant gaps across six domains. Horizontally, LLMs are generally *good at computation tasks* while *falling short on cryptography and system* coding tasks. In particular, LLMs excel in *computation* domain, where Pass@k metrics all exceed 75%, with some reaching over 90%. When it comes to *cryptography* and *system* domains, LLMs exhibit significantly lower performance, with average Pass@1 of 33.08% and 37.50%, respectively. The performance gap can be as much as 68.94% (80.94% - 12.0%) Pass@1 in Llama-2-13B-Chat. Vertically, all LLMs exhibit similar domain biases, *i.e.*, LLMs universally show consistent performance gaps with *a shared trend of strengths and weaknesses*.

In addition, a recent work (Zhuo et al. 2024) also explored the coding capability (*i.e.*, using APIs to implement domain-specific code) across various domains and concluded that LLMs are good at *cryptography* domain. Our conclusion does not conflict with theirs because we generate the domain-specific code while they invoke the domain-specific APIs. In other words, *being good at calling APIs in a domain does not mean being good at implementing code in the domain*. Therefore, our finding serves as a supplement to previous work (Zhuo et al. 2024).

LLMs Biases

Among 12 studied LLMs, the closed-source model *GPT-4o-mini* exhibits the average highest performance, with a 67.13% Pass@5. *Qwen2-72B-Instruct-GPTQ-Int4* has the best overall performance among open-source models, with a 64.25% Pass@5. Moreover, considering the variation across

Pass@1 (Greedy Search N=1)	Size	Comp	Network	Visual	Basic	System	Crypt	Mean	Std
GPT-4o-mini	\	90.38	70.31	59.68	69.16	51.00	43.00	63.92	16.68
GPT-3.5-turbo	\	83.40	58.98	48.92	56.07	32.00	31.00	51.73	19.50
Qwen2-72B-Instruct-GPTQ-Int4	72B	86.86	66.80	49.46	69.16	41.00	36.00	58.21	19.39
DeepSeek-Coder-33b-instruct	33B	83.93	64.45	50.54	59.81	46.00	35.00	56.62	16.94
DeepSeek-Coder-V2-Lite-Instruct	16B	86.04	62.11	50.00	65.42	41.00	38.00	57.10	17.92
DeepSeek-Coder-6.7b-instruct	6.7B	83.52	58.98	45.70	57.94	36.00	40.00	53.69	17.32
CodeLlama-34b-Instruct	34B	76.07	60.16	41.94	55.14	35.00	31.00	49.89	17.09
CodeLlama-13b-Instruct	13B	80.29	62.11	42.47	58.88	34.00	27.00	50.79	19.90
CodeLlama-7b-Instruct	7B	77.13	60.55	43.55	52.34	36.00	32.00	50.26	16.82
CodeQwen1.5-7B-Chat	7B	85.16	60.94	47.85	60.75	37.00	37.00	54.78	18.31
Phi-3-medium-4k-instruct	14B	75.54	60.16	45.16	61.68	42.00	35.00	53.26	15.10
Llama-2-13b-chat	13B	80.94	53.12	34.95	44.86	19.00	12.00	40.81	24.97
Average	\	82.44	61.56	46.69	59.27	37.50	33.08	53.42	18.33

Pass@5 (Sampling Search N=5)	Size	Comp	Network	Visual	Basic	System	Crypt	Mean	Std
GPT-4o-mini	\	91.26	72.66	61.83	71.03	57.00	49.00	67.13	14.75
GPT-3.5-turbo	\	87.33	62.89	52.15	60.75	36.00	34.00	55.52	19.74
Qwen2-72B-Instruct-GPTQ-Int4	72B	90.15	70.70	54.84	73.83	50.00	46.00	64.25	16.90
DeepSeek-Coder-33b-instruct	33B	89.79	70.70	55.38	68.22	57.00	42.00	63.85	16.34
DeepSeek-Coder-V2-Lite-Instruct	16B	88.91	65.62	53.76	68.22	49.00	44.00	61.59	16.35
DeepSeek-Coder-6.7b-instruct	6.7B	89.79	63.67	55.38	67.29	49.00	44.00	61.52	16.36
CodeLlama-34b-Instruct	34B	85.10	63.28	48.39	62.62	41.00	42.00	57.07	16.83
CodeLlama-13b-Instruct	13B	89.85	65.62	51.61	66.36	38.00	35.00	57.74	20.55
CodeLlama-7b-Instruct	7B	86.80	63.67	51.61	64.49	43.00	40.00	58.26	17.28
CodeQwen1.5-7B-Chat	7B	91.03	64.06	55.38	68.22	45.00	42.00	60.95	17.95
Phi-3-medium-4k-instruct	14B	85.10	67.58	54.30	67.29	47.00	44.00	60.88	15.45
Llama-2-13b-chat	13B	87.68	55.86	39.78	48.60	26.00	21.00	46.49	24.10
Average	\	88.57	65.53	52.87	65.58	44.83	40.25	59.60	17.72

Figure 6: Pass@1 and Pass@5 of LLMs against DOMAINEVAL. We use *Comp* for *computation*, *Visual* for *visualization*, and *Crypt* for *cryptology*. Mean represents the macro-average of Pass@k across different domains, which is used to reflect the overall performance of LLMs. Std indicates the standard deviation of Pass@k across different domains, which is used to reflect the different degrees of performance of LLM across various domains. To highlight the differences, we use color scales.

domains, *GPT-4o-mini* exhibits the most stable performance, with a 14.75 standard deviation in Pass@5, compared with 15.45 ~ 24.10 of other LLMs.

Notably, *CodeLlama-13B*, which fine-tuned from *Llama-2-13B*, achieves an 11.25% (57.74% - 46.49%) average improvement, while the deviation across domains still remains. It indicates that *although fine-tuning can bring about overall improvement, while the domain gaps still exist*.

Impact of Generated Samples

Finally, we analyze the impact of generated samples (*i.e.*, (*Pass@1 Greedy Search N=1*) with sub-table (*Pass@5 Sampling Search N=5*)). From Figure 6, we can see that after increasing the number of samples from 1 to 5, the average performance increases from 53.42% to 59.60%, with consistent improvements on all six domains. Yet, in terms of standard deviation (the smaller, the less bias, the better), there is little improvement, from an average 18.33 to 17.72. What is worse, *CodeLlama-13B-instruct* even observes an increased deviation, from 19.90 to 20.55, indicating a more bias as generation goes on. In other words, *generating more sam-*

ples can increase the overall performance, while the domain bias may even increase.

Case Study

Despite impressive performance exhibited by LLMs in *computation* domain, their shortcomings in other domains cannot be overlooked. This claim is supported by two indicative cases from *cryptology* and *system* domains, which highlight the challenges faced by LLMs. LLMs need to acquire more background knowledge to enhance their code generation capabilities in specific vertical domains.

The context of this subject is a classical *cryptology* scenario, the attack method targeting the RSA encryption algorithm. This involves the recovery of the two prime factors, p and q , of the RSA modulus N , given the public key (consisting of the modulus N and the public exponent e) and the private key (the private exponent d). When encrypting, in order to enhance the security and impede decryption attempts, the selected p and q by human are both extremely large prime numbers. Although the instruction mentioned keywords such as RSA, the model does not realize that in

```

import math
from random import randrange
def attack(N, e, d):
    ...
    if discriminant >= 0:
- sqrt_discriminant = int(discriminant**0.5)
+ sqrt_discriminant = math.isqrt(discriminant))
    ...

```

Figure 7: A flawed response from *GPT-4o-mini* in *cryptography* domain. The red segment represents incorrect code, while the green segment represents the correct reference.

```

import collections
import os
import re
from psutil.tests import sh
def free_physmem():
- output = os.popen('free -b').read()
+ output = sh(['free', '-b'], env={'LANG': 'C.UTF-8'})
    ...
    mem_values = re.split(r'\s+', mem_line.strip())
    ...
- shared = int(mem_values[5])
+ shared = int(mem_values[4])
    ...

```

Figure 8: A flawed response from *GPT-4o-mini* in *system* domain. The red segment represents incorrect code, while the green segment represents the correct reference.

the context of *cryptography*, rounding the square root of a large number cannot be directly converted to `*0.5`, and instead `math.isqrt` should be used to avoid *OverflowError*.

The response presented in Figure 7 is from *GPT-4o-mini*, which is considered as the premier model within *cryptography* domain. The error in its response is *failed: int too large to convert to float*. Similar errors are also observed in some responses from *DeepSeek* and *Qwen* series models.

In the same way, although *GPT-4o-mini* exhibits the best performance in *system* domain, Figure 8 shows a flawed response from it. The function is designed to parse the output of `free` command (run with `-b` option) to determine the physical memory state on a Linux system. However, the model fails to ensure that the output format of `free` command remains consistent across different language and regional settings, resulting in incorrect string matching during subsequent parsing. Similar errors are observed in models such as *CodeQwen1.5-7B-Chat*, *Qwen2-72B-Instruct-GPTQ-Int4*, *Phi-3-medium-4k-instruct*, the *DeepSeek* series, and the *CodeLlama* series. Additionally, *GPT-4o-mini* lacks sufficient understanding of `free` command, leading to misjudgments regarding the position of `shared`. Ultimately, none of LLMs are able to pass this subject.

Related Work

To assess the code generation capabilities of models, numerous code benchmarks were introduced such as APPS (Hendrycks et al. 2021) and CodeContests (Li et al. 2022), which are datasets sourced from algorithmic design compe-

titions that emphasize logical thinking. MBPP (Austin et al. 2021) and HumanEval (Chen et al. 2021) are datasets crafted manually for testing. These datasets are designed for evaluation, featuring highly independent functions and simplistic, idealized problem scenarios. Standalone functions are predominantly focused on by these benchmarks; however, non-standalone functions are commonly encountered in pragmatic code generation scenarios (Yu et al. 2024).

Some other benchmarks attached importance on real-world problems by sourcing data from real scenarios such as StackOverflow, GitHub, and curating data manually to form benchmark datasets. For instance, PandasEval, NumpyEval (Zan et al. 2022), and SecurityEval (Siddiq and Santos 2022) are tailored benchmarks for specific scenarios (Zan et al. 2023). Focusing on specific scenarios is their limitation.

CoderEval (Yu et al. 2024), ClassEval (Du et al. 2023), DevEval (Li et al. 2024b), and ODEX (Wang et al. 2023) were constructed in open domains, but they required significant investment of human labor in the stages of data curating, filtering, or annotation.

RepoEval (Zhang et al. 2023), MultiPL-E (Cassano et al. 2022), and Exec-CSN (Xie et al. 2024) were constructed with little human involvement. However, RepoEval is focused on repository level task, which differs from ours. MultiPL-E obtained its data by translating other code datasets, and thus, it was not oriented towards open domains. Exec-CSN employed LLMs to curate the CodeSearchNet dataset and generate test cases, resulting in a final dataset that has a gap from the real-world code on GitHub.

Moreover, previous benchmarks including (Chen et al. 2024; Li et al. 2024a; Ding et al. 2023; Liu et al. 2024) have not explored the difference of code generation capability of LLMs across multiple domains. Though a recent work (Zhuo et al. 2024) exercised capability of LLMs in using APIs from several domains, but its benchmark need human-LLM collaboration to construct. Our **DOMAINEVAL** benchmark is designed not for tool utilization, but for concrete implementation of APIs and functions. Furthermore, **DOMAINEVAL** can be constructed through a fully automated pipeline across open domains.

Conclusion

We introduce **DOMAINEVAL**, a function code generation benchmark across multiple programming domains. Our research underscores the need for a comprehensive benchmark to evaluate LLMs’ code generation capabilities, both broadly and within specific verticals. The automated pipeline we introduce ensures dataset diversity and real-time updates, while enabling the creation of custom domain benchmarks for others. Preliminary results show that while LLMs excel in computation, their performance in *cryptography* and *system* requires enhancement.

For the future, several research directions merit exploration. First, developing specialized training strategies and data augmentation techniques to improve LLMs’ code generation capabilities for specific domains like *cryptography* and *system*. Second, leveraging our automated pipeline to create benchmarks for a broader range of private and domain-specific code, aiming to assess and enhance LLMs.

Acknowledgements

We sincerely thank the reviewers for their insightful comments and valuable suggestions. This work was supported by Beijing Natural Science Foundation (L243006), the Natural Science Foundation of China (No. 62122077, 62306303, 62106251), the Basic Research Program of IS-CAS (Grant No. ISCAS-ZD-202402) and RGC/GRF grant 16206524.

References

- Abdin, M.; Jacobs, S. A.; Awan, A. A.; Aneja, J.; Awadallah, A.; Awadalla, H.; Bach, N.; Bahree, A.; Bakhtiari, A.; Behl, H.; et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Austin, J.; Odena, A.; Nye, M.; Bosma, M.; and Michalewski, H. 2021. Program Synthesis with Large Language Models.
- Bai, J.; Bai, S.; Chu, Y.; Cui, Z.; Dang, K.; Deng, X.; Fan, Y.; Ge, W.; Han, Y.; Huang, F.; et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Cao, J.; Chen, Z.; Wu, J.; Chi Cheung, S.; and Xu, C. 2024a. Can AI Beat Undergraduates in Entry-level Java Assignments? Benchmarking Large Language Models on JavaBench. *arXiv:2406.12902*.
- Cao, J.; Zhang, W.; Cheung, S.-C.; and on, S. 2024b. Concerned with Data Contamination? Assessing Countermeasures in Code Language Model. *arXiv:2403.16898*.
- Cassano, F.; Gouwar, J.; Nguyen, D.; Nguyen, S.; Phipps-Costin, L.; Pinckney, D.; Yee, M.-H.; Zi, Y.; Anderson, C. J.; Feldman, M. Q.; Guha, A.; Greenberg, M.; and Jangda, A. 2022. MultiPL-E: A Scalable and Extensible Approach to Benchmarking Neural Code Generation. *arXiv:2208.08227*.
- Chai, L.; Liu, S.; Yang, J.; Yin, Y.; Jin, K.; Liu, J.; Sun, T.; Zhang, G.; Ren, C.; Guo, H.; et al. 2024. McEval: Massively Multilingual Code Evaluation. *arXiv preprint arXiv:2406.07436*.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; and Pinto, H. P. d. O. 2021. Evaluating Large Language Models Trained on Code.
- Chen, S.; Feng, X.; Han, X.; Liu, C.; and Yang, W. 2024. PPM: Automated Generation of Diverse Programming Problems for Benchmarking Code Generation Models.
- Ding, Y.; Wang, Z.; Ahmad, W. U.; Ding, H.; and Tan, M. 2023. CrossCodeEval: A Diverse and Multilingual Benchmark for Cross-File Code Completion.
- Du, X.; Liu, M.; Wang, K.; Wang, H.; and Liu, J. 2023. ClassEval: A Manually-Crafted Benchmark for Evaluating LLMs on Class-level Code Generation.
- Guo, D.; Zhu, Q.; Yang, D.; Xie, Z.; Dong, K.; Zhang, W.; Chen, G.; Bi, X.; Wu, Y.; Li, Y.; et al. 2024. DeepSeek-Coder: When the Large Language Model Meets Programming—The Rise of Code Intelligence. *arXiv preprint arXiv:2401.14196*.
- Hendrycks, D.; Basart, S.; Kadavath, S.; Mazeika, M.; and Arora, A. 2021. Measuring Coding Challenge Competence With APPS.
- Imani, S.; Du, L.; and Shrivastava, H. 2023. MathPrompter: Mathematical Reasoning using Large Language Models. *arXiv:2303.05398*.
- Li, J.; Li, G.; Zhang, X.; Dong, Y.; and Jin, Z. 2024a. EvoCodeBench: An Evolving Code Generation Benchmark Aligned with Real-World Code Repositories.
- Li, J.; Li, G.; Zhao, Y.; Li, Y.; and Liu, H. 2024b. DevEval: A Manually-Annotated Code Generation Benchmark Aligned with Real-World Code Repositories.
- Li, Y.; Choi, D.; Chung, J.; Kushman, N.; Schrittwieser, J.; Leblond, R.; Eccles, T.; Keeling, J.; Gimeno, F.; Dal Lago, A.; et al. 2022. Competition-level code generation with alphacode. *Science*, 378(6624): 1092–1097.
- Liu, J.; Xia, C. S.; Wang, Y.; and Zhang, L. 2024. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *Advances in Neural Information Processing Systems*, 36.
- Lozhkov, A.; Li, R.; Allal, L. B.; Cassano, F.; and Lamy-Poirier, J. 2024. StarCoder 2 and The Stack v2: The Next Generation.
- Rogers, A.; Gardner, M.; and Augenstein, I. 2023. QA Dataset Explosion: A Taxonomy of NLP Resources for Question Answering and Reading Comprehension. *ACM Comput. Surv.*, 55(10).
- Roziere, B.; Gehring, J.; Gloeckle, F.; Sootla, S.; Gat, I.; Tan, X. E.; Adi, Y.; Liu, J.; Remez, T.; Rapin, J.; et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.
- Siddiq, M. L.; and Santos, J. C. 2022. SecurityEval dataset: mining vulnerability examples to evaluate machine learning-based code generation techniques. In *Proceedings of the 1st International Workshop on Mining Software Repositories Applications for Privacy and Security*, 29–33.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, C.; Li, Z.; Gao, C.; Wang, W.; Peng, T.; Huang, H.; Deng, Y.; Wang, S.; and Lyu, M. R. 2024. Exploring Multilingual Bias of Large Code Models in Code Generation. *arXiv:2404.19368*.
- Wang, Z.; Zhou, S.; Fried, D.; and Neubig, G. 2023. Execution-Based Evaluation for Open-Domain Code Generation.

Xie, Y.; Xie, A.; Sheth, D.; Liu, P.; and Fried, D. 2024. CodeBenchGen: Creating Scalable Execution-based Code Generation Benchmarks.

Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

Yu, H.; Shen, B.; Ran, D.; Zhang, J.; and Liang, G. 2024. CoderEval: A Benchmark of Pragmatic Code Generation with Generative Pre-trained Models.

Zan, D.; Chen, B.; Yang, D.; Lin, Z.; and Kim, M. 2022. CERT: Continual Pre-Training on Sketches for Library-Oriented Code Generation.

Zan, D.; Chen, B.; Zhang, F.; Lu, D.; and Wu, B. 2023. Large Language Models Meet NL2Code: A Survey. In *ACL 2023*, 7443–7464. Association for Computational Linguistics.

Zhang, F.; Chen, B.; Zhang, Y.; Keung, J.; and Liu, J. 2023. RepoCoder: Repository-Level Code Completion Through Iterative Retrieval and Generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2471–2484. Association for Computational Linguistics.

Zhu, Q.; Guo, D.; Shao, Z.; Yang, D.; Wang, P.; Xu, R.; Wu, Y.; Li, Y.; Gao, H.; Ma, S.; et al. 2024. DeepSeek-Coder-V2: Breaking the Barrier of Closed-Source Models in Code Intelligence. *arXiv preprint arXiv:2406.11931*.

Zhuo, T. Y.; Vu, M. C.; Chim, J.; Hu, H.; and Yu, W. 2024. BigCodeBench: Benchmarking Code Generation with Diverse Function Calls and Complex Instructions.