

Can LLM Coding Assistants Support Emerging Programming Languages? An Empirical Study on Cangjie

Junhang Cheng ✉ 

Beihang University, Beijing, China

Fang Liu¹ ✉ 

Beihang University, Beijing, China

Jia Li ✉ 

Wuhan University, Wuhan, China

Chengru Wu ✉ 

Beihang University, Beijing, China

Nanxiang Jiang ✉ 

Beihang University, Beijing, China

Li Zhang ✉ 

Beihang University, Beijing, China

Abstract

Background. LLM coding assistants are widely used in development, but evidence about their reliability mostly comes from mainstream languages such as Python and Java. Their behaviour on emerging languages, where public code and documentation are scarce, is far less clear.

Aims. We ask how capable current LLMs are on Cangjie, an emerging language in the HarmonyOS ecosystem, and which assistive technique is worth its token cost.

Method. We build CangjieBench, 248 manually translated tasks derived from HumanEval and ClassEval, run under Text-to-Code and Code-to-Code modalities. We evaluate six LLMs under five methods: direct generation, syntax-constrained prompting, retrieval over documentation, retrieval over code, and agent-based workflows. We report Pass@1, compile rate, token cost, and a seven-category failure taxonomy validated by two LLM annotators ($\kappa \geq 0.98$).

Results. Direct generation is unusable on most models, and over four-fifths of failed samples carry foreign syntax that the Cangjie compiler rejects. A 2,146-token syntax reference gives the best accuracy-to-cost trade-off among prompt-based methods. The best agent configuration reaches the highest Pass@1 at 10–120x the token cost of a single prompt, while weaker agents cannot match a syntax-constrained call on the same backbone. A Python reference does not always help: on Code-to-Code it can pull models toward Python idioms and drop compile rates below Text-to-Code.

Conclusions. No single method dominates both accuracy and token cost. A practical assistant for an emerging language should ship a concise syntax cheat sheet first, escalate to curated retrieval for API-heavy tasks, and reserve agent loops for tasks that need multi-turn repair. The replication package is at <https://github.com/cjhCoder7/CangjieBench>.

2012 ACM Subject Classification Software and its engineering → Software verification and validation; Computing methodologies → Natural language processing; Software and its engineering → General programming languages

Keywords and phrases Large language models, Code generation, Code translation, Emerging programming languages, Empirical software engineering, Benchmark, Cangjie

¹ Corresponding author.



© Junhang Cheng, Fang Liu, Jia Li, Chengru Wu, Nanxiang Jiang, and Li Zhang; licensed under Creative Commons License CC-BY 4.0

20th International Symposium on Empirical Software Engineering and Measurement (ESEM 2026).

Editors: Robert Feldt, Maria Paasivaara, Daniel Mendez, Stefan Wagner, and Marvin Muñoz Barón;

Article No. 44; pp. 44:1–44:21



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Digital Object Identifier 10.4230/LIPIcs.ESEM.2026.44

Category Technical Track Paper

Related Version *Preprint*: <https://arxiv.org/abs/2603.14501>

Supplementary Material

Software (Software, Dataset): <https://github.com/cjhCoder7/CangjieBench> [6]

Funding This research is supported by the National Natural Science Foundation of China (Grant No. 62302021), the State Key Laboratory of Complex & Critical Software Environment (Grant No. CCSE2025ZX-09), and the Fundamental Research Funds for the Central Universities.

1 Introduction

Large language models (LLMs) now power IDE completions and agent CLIs [2, 23, 10] that generate, translate, and repair code in real projects, but most of the evidence that they do this well rests on a narrow set of languages. HumanEval [5] and MBPP [3] evaluate Python, the benchmarks most often reported for new models [13] still centre on Python, Java, and JavaScript, and MultiPL-E [4] shows that once evaluation moves to less represented languages the gap between a model’s Python score and its score elsewhere can be large. For languages outside this core, far less is known about whether current models are usable, which assistive technique is worth its token cost, and what the typical failure modes look like.

Emerging languages make the problem worse. A young language has little public code, its documentation is still being written, and its idioms have not stabilised, so the base model, the retrieval index, and the agent scaffolding all have less material to draw on than on Python or Java, and claims established on Python, such as “retrieval adds a few points of Pass@1” or “agent loops recover most compile errors”, do not automatically carry over and have to be re-measured.

Cangjie is such a language. Released in 2024 as the native language of the HarmonyOS application ecosystem [12], Cangjie is statically typed and memory-safe, with syntax and type rules that differ from both Python and Java. HarmonyOS now ships on hundreds of millions of devices, so its developers are the intended audience for a Cangjie coding assistant, yet the volume of Cangjie seen at training time is small and the assistant has less prior knowledge of the language than developers expect. For a team adopting Cangjie today, three questions are open: how capable current models are without help, whether assistive techniques that work for Python transfer, and how the choice between prompt, retrieval, and agent affects accuracy and token cost.

Existing evaluations do not close this gap. Prior work on low-resource languages has mostly targeted domain-specific languages such as Verilog [20] or Solidity [27], where any performance drop mixes missing domain knowledge with missing language knowledge. Cross-language benchmarks such as MultiPL-E and HumanEval-XL [26] translate HumanEval into many targets but do not compare assistive methods or report token cost, and empirical studies of prompt engineering, retrieval-augmented generation [39], and agent-based workflows [35] are almost entirely set in Python and rarely put these methods on the same accuracy-against-cost axis. A study of a modern, statically-typed emerging general-purpose language that brings these factors together is missing.

We address that gap with an empirical study organised around four research questions:

- **RQ1 (accuracy)**. How do direct generation, syntax-constrained prompting, documentation retrieval, code-example retrieval, and agent-based workflows compare on Cangjie pass rate?

- **RQ2 (token cost).** How much does each method cost in input and output tokens, and how does the cost ranking compare with the pass-rate ranking?
- **RQ3 (transfer).** Is Code-to-Code translation from Python easier than Text-to-Code generation, and does a Python reference induce negative transfer toward Python idioms?
- **RQ4 (failure modes).** What is the distribution of foreign-syntax drift, API misuse, type and mutability errors, duplicate-definition or harness conflicts, runtime exceptions, and logic errors, and how does it shift across methods, models, and tasks?

To answer these questions, we build **CangjieBench**, a 248-task benchmark with 164 function-level problems manually translated from HumanEval and 84 class-level problems from ClassEval [8], each shipping with a Cangjie reference solution and a harness that compiles the generated code, links it against the reference tests, and runs the result. On this benchmark we evaluate six recent LLMs (DeepSeek-V3 [19], ERNIE-4.5 [31], Kimi-K2 [32], Qwen3-235B-2507 and Qwen3-Coder-480B [38], and GPT-5 [24]) under five methods: direct generation, syntax-constrained prompting with a concise rule set, retrieval over Cangjie documentation, retrieval over curated code examples, and an agent-based workflow that reads documentation, runs the compiler, and revises across turns. Every task runs under two modalities, Text-to-Code and Code-to-Code, and for each run we record Pass@1, compile rate, token usage, and the category of the first observed failure, seeded by manual labelling of 788 failed GPT-5 samples and applied to the full benchmark by two LLM annotators whose agreement we cross-check ($\kappa \geq 0.98$).

Four observations recur across RQ1–RQ4. Direct generation is unreliable on most models, as code typically fails before the tests run and over four-fifths of failed samples carry foreign syntax the compiler will not accept. Among prompt-based methods a concise syntax rule set gives the best accuracy-to-cost trade-off and retrieval does not match it once its token cost is counted, while the best agent configuration reaches the highest absolute pass rate but spends an order of magnitude more tokens than any single-shot method. A Python reference does not always help, and on several cells it pulls the model toward Python idioms so that Code-to-Code compile rate drops below Text-to-Code, a negative-transfer pattern rather than a boost. Finally, the failure mix shifts with method, as prompt-only runs are dominated by syntax and API errors while agent runs are the only C2C configuration where logic errors reach a non-trivial share, since their failed samples regularly clear the compiler.

Taking the benchmark, the evaluation, and the failure analysis together, this paper makes three contributions.

- **CangjieBench**, a 248-task function- and class-level benchmark for Cangjie, with reference solutions, compile-and-test harnesses, and a public replication package.
- A head-to-head evaluation of six LLMs across five methods and two modalities that reports pass rate, compile rate, and token cost separately, making the accuracy-cost trade-off explicit rather than collapsing it into one score.
- A failure taxonomy for Cangjie code generation, with per-method distributions and an analysis of where current models break on an emerging language.

2 Background and Related Work

Low-Resource Language Benchmarks. LLMs perform well on high-resource languages [4] but drop sharply on low-resource ones, largely due to training-data scarcity. MultiPL-E extended HumanEval and MBPP to eighteen languages, and more recent benchmarks focus on individual targets such as Verilog [20, 29], Solidity [27], TeX [16], VHDL [33], shell scripting [34], GPU kernels [25], and SQL [18]. On the translation side, ClassEval-T [37]

extends class-level evaluation to translation tasks, and MERA Code [7] unifies tasks across languages. A recent survey [15] organises this space. Two gaps remain for our study. First, **entanglement with domain knowledge**: domain-specific language (DSL) benchmarks mix language novelty with domain constraints, so an observed drop mixes unfamiliar syntax with unfamiliar hardware or contract semantics. Second, **data-leakage risk**: languages traditionally labelled low-resource, such as Lua [22] and R [43], have long histories and still appear in modern pre-training corpora, weakening contamination controls. Cangjie addresses both: as a general-purpose language it keeps the domain constant, and its 2024 release keeps pre-training exposure small.


Methods of Code Generation. Modern code generation builds on foundation models such as GPT-5, Qwen3, and Gemini [9], together with three broad families of assistive methods: *Prompt Engineering* [30, 28], *Retrieval-Augmented Generation* [21], and *Agentic Frameworks* [35]. At the project level, recent work also studies multi-agent collaboration over larger codebases [42]. For Python-to-target-language translation specifically, prior work spans parallel corpora such as AVATAR [1], function-level decomposition [41], execution-guided repair [11], interpretable error correction [36], and language-pair benchmarks such as CRUST-Bench for C-to-Rust [17]. Previous studies on low-resource code, however, usually limit evaluation to direct zero-shot or few-shot prompting and rarely report token cost alongside accuracy. We place all three families on the same footing, with the same tasks, models, and target language, so that differences in accuracy and token cost can be attributed to the method rather than to incomparable setups.

3 CangjieBench Construction

Multilingual suites such as MultiPL-E do not include Cangjie, and DSL benchmarks entangle language novelty with domain knowledge. We therefore build **CangjieBench**, which targets Cangjie directly, runs both Text-to-Code generation and Python-to-Cangjie translation on the same problem set, and grades each sample by compiling and executing it on the Cangjie toolchain. As Figure 1 shows, it has three components: a manually translated dataset of 248 Cangjie problems, two evaluation tasks under different input modalities, and a framework that compiles and runs each sample inside a Docker sandbox (Section 9).

3.1 Dataset Construction

Publicly available Cangjie code is scarce and dominated by official samples a pre-trained model has likely seen, so the usual crawl-and-deduplicate pipeline is not viable. We instead build CangjieBench by manually translating two widely used Python benchmarks, *HumanEval* and *ClassEval*, covering function-level and class-level problems. *HumanEval* [5] consists of 164 hand-written Python problems, each with a function signature, docstring, and unit tests, and is the de facto standard for function-level generation despite a known bias toward algorithmic tasks and limited API diversity. *ClassEval* [8] extends this to the class level with 100 problems that require complete Python classes, testing the multi-method cohesion and cross-method state management absent from function-level benchmarks. Both datasets are self-contained, one task per file with no cross-file dependencies, unlike repository-level benchmarks such as *CoderEval* [40] and *SWE-bench* [14], which keeps the ground truth easy to verify by hand and makes each task small enough that any failure can be attributed to the model’s handling of Cangjie rather than to ambiguity about which file to edit.



■ **Figure 1** Overview of the CangjieBench framework: dataset construction, evaluation tasks, and evaluation sandbox.

To keep the translation faithful, we follow a fixed set of construction rules throughout, derived from the authors’ experience building Cangjie programs since its 2024 release. No existing benchmark methodology addresses Cangjie, so rather than adopt one we developed the rules during a pilot translation of roughly 30 tasks and applied them to all 248.

- **Type Adaptation.** We map Python primitives onto their strict Cangjie equivalents (`int`→`Int64`, `float`→`Float64`, `str`→`String`, `bool`→`Bool`) and Python lists and dictionaries onto `ArrayList` and `HashMap`. Where a Python signature relies on dynamic typing, we use Cangjie’s `Any` interface and add explicit casts only where the harness needs a concrete value.
- **Naming Convention.** We keep the original *snake_case* identifiers rather than camel case, which preserves the link to the source task and rules out a trivial failure where the model looks for the wrong identifier.
- **Algorithm Utilization.** Cangjie reference solutions reproduce the control flow of the Python solutions step by step, including branching, loops, and recursion, without substituting a more idiomatic algorithm, since the goal is to grade how well a model expresses a known algorithm in Cangjie rather than the algorithm itself.
- **Prompt Transformation.** HumanEval and ClassEval prompts embed Python fragments such as signatures, docstrings with type hints, and few-shot examples, and we rewrite every fragment into Cangjie form so that the description and any embedded code both refer to the target language, since a mixed-language prompt would leave us unable to tell whether the model was responding to the prompt or patching over the inconsistency.
- **Dependency Management.** We substitute the matching Cangjie standard library for standard Python modules. Tasks depending on third-party Python libraries with no Cangjie equivalent, such as `sqlite3`, `PIL`, and `sklearn`, are excluded rather than reimplemented, since a from-scratch replacement would measure library quality rather than model capability, and Table 1 lists all 16, but we implement very small utilities such as `hashlib.md5` so that tasks needing only a hash function stay in.

The construction was carried out in three steps over roughly 1.5 months. First, the first author translated every prompt and reference solution and wrote a matching Cangjie test suite. Second, two other authors with substantial Cangjie experience independently reviewed each translated problem, solution, and test case, resolving disagreements by discussion and,

■ **Table 1** ClassEval tasks excluded from CangjieBench and the Python-specific library responsible for the exclusion.

Problem ID	Problem Name	Exclusion Reason
ClassEval_14	BookManagementDB	Relies on <code>sqlite3</code> for database operations
ClassEval_25	CookiesUtil	Relies on <code>json</code> for data serialisation
ClassEval_28	DatabaseProcessor	Relies on <code>sqlite3</code> for database operations
ClassEval_34	DocFileHandler	Relies on <code>docx</code> for Word document processing
ClassEval_38	ExcelProcessor	Relies on <code>openpyxl</code> for spreadsheet manipulation
ClassEval_44	HtmlUtil	Relies on <code>bs4</code> and <code>gensim</code> for HTML/NLP processing
ClassEval_45	ImageProcessor	Relies on <code>PIL</code> for image processing
ClassEval_50	JSONProcessor	Relies on <code>json</code> for data serialisation
ClassEval_52	Lemmatization	Relies on <code>nltk</code> for NLP processing
ClassEval_60	MovieTicketDB	Relies on <code>sqlite3</code> for database operations
ClassEval_69	PDFHandler	Relies on <code>PyPDF2</code> for PDF manipulation
ClassEval_83	StudentDatabaseProcessor	Relies on <code>sqlite3</code> for database operations
ClassEval_84	TextFileProcessor	Relies on <code>json</code> for data serialisation
ClassEval_92	UserLoginDB	Relies on <code>sqlite3</code> for database operations
ClassEval_98	XMLProcessor	Relies on <code>xml</code> for XML parsing
ClassEval_99	ZipFileProcessor	Relies on <code>zipfile</code> for archive management

where needed, by rerunning the tests against alternative references. Third, we ran every suite against its reference solution and against deliberate perturbations such as wrong return types, off-by-one boundaries, and known-buggy substitutes, confirming that the tests both accept the correct answer and reject common incorrect ones.

After review, CangjieBench contains **248** tasks: all **164** HumanEval problems, which map cleanly onto Cangjie’s standard library, plus **84** of the original 100 ClassEval problems, with Table 1 listing the 16 excluded for Python-specific dependencies. The two subsets give two granularities: HumanEval probes whether a model can write a single correctly typed routine, while ClassEval asks it to lay out a multi-method class and respect the Cangjie object model across method boundaries.


3.2 Evaluation Tasks

Every task runs under two modalities that share the same reference solution and test harness. In *Text-to-Code*, the model receives only a natural-language description and produces a Cangjie implementation from scratch, isolating its ability to generate Cangjie without a per-task anchor. In *Code-to-Code*, the model additionally receives the corresponding Python reference and must produce a behaviourally equivalent Cangjie implementation, mirroring a team porting an existing Python codebase and letting us measure whether a known-good Python solution helps or instead anchors the model on Python idioms.

We grade correctness by running the per-task Cangjie test suite. A HumanEval-derived sample passes if the generated function compiles and clears all test cases, and a ClassEval sample must pass every method-level test plus a class-level integration test, so that a class whose methods pass in isolation but fail when composed is still flagged incorrect.

3.3 Evaluation Framework

The evaluation framework is a Docker-based sandbox that pins the Cangjie toolchain, standard library, and test runner to a single image and executes each sample in a fresh container. Pinning lets us compare timing and error messages across models, methods, and machines on



■ **Figure 2** The four method families evaluated on CangjieBench. Direct generation is the baseline. The other three families add Cangjie-specific context in different ways.

equal terms, and isolating each run protects the host from agent-generated code the model has only partially understood. For every sample the harness records whether compilation succeeded, whether all tests passed, the compiler and runtime stderr, and the wall-clock time, the raw signals on which the failure analysis in Section 5 is built.

4 Methods

We exclude training methods, since the goal is not to engineer a Cangjie-specialised model but to ask how far an off-the-shelf foundation model can be pushed on an emerging language using techniques a team would normally have at hand. We evaluate four method families common in practice, summarised in Figure 2: *Direct Generation*, *Syntax-Constrained Prompting*, *Retrieval-Augmented Generation* with two variants (RAG-Docs and RAG-Code), and *Agent-based workflows*.

4.1 Method Families

Direct Generation. The model receives the natural-language problem description, plus the Python reference solution under Code-to-Code, wrapped in a fixed prompt template with no grammar reference, no retrieved snippets, and no tools.

Syntax-Constrained Prompting. Since most of the LLMs we evaluate have seen little Cangjie at training time, a common engineering response is to ship a small “language pack”, a syntax cheat-sheet injected into every prompt, and we prepend such a reference to the direct-generation template.

Retrieval-Augmented Generation. A more involved response gives the assistant access to a Cangjie corpus at inference time and pastes a few retrieved snippets into context. We evaluate two variants matching the two corpora a typical assistant would index, official documentation (RAG-Docs) and crawled code (RAG-Code), and for an emerging language both corpus and retrieval signal are weaker than for Python, so whether this pays off is an open question.

Agent-based Workflow. A more open-ended response, similar to how a new developer learns an unfamiliar language, lets the assistant read documentation, write code, run the compiler, and revise its output across several turns, and we want to know what such a loop costs in tokens and buys in pass rate on a language with thin pre-training coverage.

4.2 Configuration

Syntax reference. The grammar reference used by Syntax-Constrained Prompting covers 20 categories: program structure, variables, numeric types, strings and characters, the boolean type, arrays and dynamic arrays, hash maps, hash sets, sorting, tuples, conditional statements, loop statements, the option type, match expressions, lambda expressions, input and output, functions, classes, interfaces, and reserved keywords. It totals 2,146 GPT-5 tokens, well below every model’s context limit.

RAG corpora and retriever. Both RAG variants use Okapi BM25 with $\text{top-}k = 3$. The *Docs* corpus is the Cangjie Developer Manual and standard-library API reference from the official repository, with HTML stripped and text split on section headers, and because problem statements and API reference rarely share surface vocabulary, we add a query-transformation step in which the model first extracts technical keywords such as library and API method names, which we issue to BM25 in place of the raw description. The *Code* corpus is Cangjie source crawled from public GitCode repositories before 10 November 2025, filtered to files at least 80% Cangjie, parsed with the Cangjie AST tool, and split into 30,341 functions and 31,483 class definitions, and here we issue the prompt directly to BM25 because it already contains code-shaped material such as signatures and type hints.

Agent harnesses. We use three CLI agent harnesses without modification: **Codex CLI** version 0.63.0, **Qwen Code CLI** version 0.2.3, and **iFlow CLI** version 0.4.8. All three can read files, execute shell commands, and inspect shell output, and their built-in retrieval differs, itself a method-level variable. Each agent receives the problem description, the direct-generation template, and read access to the same Cangjie documentation tree used by RAG-Docs, kept under Git and rolled back after every interaction step, and terminates when its harness reports completion or hits the step budget.

5 Results

5.1 Experimental Setup

Subjects. The open-weights subjects are DeepSeek-V3 (671B parameters), ERNIE-4.5 (300B), Kimi-K2-0905 (1T), Qwen3-235B-A22B-Instruct-2507 (235B), and Qwen3-Coder-480B, and the closed-source subject is GPT-5. We restrict the selection to models above 200B parameters because smaller models bottomed out at near-zero Pass@1 in pilot runs, leaving no headroom to separate the methods, and the six we keep span two families of general-purpose models and two code-specialised ones, so any method effect we report is not tied to a single training recipe. The four agent configurations are Codex CLI / GPT-5, Qwen Code CLI / Qwen3-Coder, iFlow CLI / Qwen3-Coder, and iFlow CLI / Kimi-K2-0905, using the three CLI harnesses of Section 4.2, chosen because they are the harnesses a Cangjie developer can install and run today.

■ **Table 2** Pass@1 (%) and Compile rate (%) on CangjieBench. Light gray columns indicate Text-to-Code (T2C), and dark gray columns indicate Code-to-Code (C2C). The best result per method is in **bold**, and the overall best across all methods is **red and underlined**. Avg. averages over HumanEval (164 tasks) and ClassEval (84 tasks) weighted by task count.


Method	Subject	HumanEval				ClassEval				Avg.			
		Pass@1		Compile		Pass@1		Compile		Pass@1		Compile	
		T2C	C2C	T2C	C2C	T2C	C2C	T2C	C2C	T2C	C2C	T2C	C2C
Direct	DeepSeek-V3	2.4	2.4	3.0	3.0	1.2	1.2	1.2	1.2	2.0	2.0	2.4	2.4
	ERNIE-4.5	3.0	4.3	4.3	4.9	0.0	1.2	0.0	1.2	2.0	3.2	2.8	3.6
	Kimi-K2	20.7	21.3	23.8	23.8	4.8	8.3	7.1	8.3	15.3	16.9	18.1	18.5
	Qwen3	3.7	3.7	4.3	4.3	1.2	2.4	1.2	2.4	2.8	3.2	3.2	3.6
	Qwen3-Coder	3.7	7.3	4.3	7.9	1.2	1.2	1.2	1.2	2.8	5.2	3.2	5.6
	GPT-5	6.7	7.9	7.3	8.5	1.2	3.6	1.2	3.6	4.8	6.5	5.2	6.9
Syntax-Constrained	DeepSeek-V3	40.9	42.7	47.6	44.5	9.5	6.0	16.7	6.0	30.2	30.2	37.1	31.5
	ERNIE-4.5	36.6	32.3	39.0	35.4	1.2	4.8	2.4	6.0	24.6	23.0	26.6	25.4
	Kimi-K2	55.5	53.0	62.2	56.1	11.9	14.3	22.6	15.5	40.7	39.9	48.8	42.3
	Qwen3	53.0	45.1	57.3	47.6	13.1	10.7	22.6	14.3	39.5	33.5	45.6	36.3
	Qwen3-Coder	42.7	48.8	47.6	51.8	16.7	11.9	22.6	13.1	33.9	36.3	39.1	38.7
	GPT-5	64.0	44.5	67.1	45.1	23.8	25.0	40.5	31.0	50.4	37.9	58.1	40.3
RAG (Code)	DeepSeek-V3	15.2	14.0	16.5	15.9	2.4	4.8	3.6	4.8	10.9	10.9	12.1	12.1
	ERNIE-4.5	13.4	17.7	15.9	18.9	4.8	3.6	6.0	3.6	10.5	12.9	12.5	13.7
	Kimi-K2	30.5	26.8	33.5	31.7	8.3	9.5	13.1	9.5	23.0	21.0	26.6	24.2
	Qwen3	12.8	7.3	13.4	7.9	3.6	7.1	3.6	7.1	9.7	7.3	10.1	7.7
	Qwen3-Coder	23.2	17.1	25.0	18.3	3.6	6.0	7.1	7.1	16.5	13.3	19.0	14.5
	GPT-5	47.0	45.1	49.4	47.0	8.3	10.7	13.1	13.1	33.9	33.5	37.1	35.5
RAG (Docs)	DeepSeek-V3	27.4	21.3	33.5	22.6	6.0	8.3	9.5	9.5	20.2	16.9	25.4	18.1
	ERNIE-4.5	11.6	11.6	13.4	12.8	2.4	3.6	3.6	3.6	8.5	8.9	10.1	9.7
	Kimi-K2	32.3	30.5	34.8	34.1	7.1	16.7	11.9	16.7	23.8	25.8	27.0	28.2
	Qwen3	11.6	10.4	12.8	12.2	1.2	4.8	1.2	4.8	8.1	8.5	8.9	9.7
	Qwen3-Coder	19.5	17.7	22.6	18.9	3.6	8.3	6.0	9.5	14.1	14.5	16.9	15.7
	GPT-5	36.0	31.1	37.2	32.3	7.1	10.7	15.5	16.7	26.2	24.2	29.8	27.0
Agent	Kimi-K2 (iFlow CLI)	39.0	48.2	44.5	51.8	17.9	22.6	26.2	27.4	31.9	39.5	38.3	43.5
	Qwen3-Coder (iFlow CLI)	29.9	29.9	32.3	33.5	4.8	14.3	9.5	17.9	21.4	24.6	24.6	28.2
	Qwen3-Coder (Qwen Code CLI)	25.0	23.8	30.5	27.4	7.1	7.1	8.3	8.3	19.0	18.1	23.0	21.0
	GPT-5 (Codex CLI)	84.8	86.6	87.2	87.8	32.1	58.3	67.9	65.5	66.9	77.0	80.6	80.2

Decoding protocol. All cells use greedy decoding with `temperature = 0` and `max_tokens = 8192`, drawing a single sample per task, with open-source models served through the SiliconFlow API and GPT-5 through the OpenAI API. This removes sampling variance from the comparison but forgoes the smoothing of Pass@ k with $k > 1$, so we add a Pass@10 robustness check under syntax-constrained prompting in Section 6. For agent runs we cap each task at the harness default step budget and count a run as failed if it exhausts the budget without a compiling submission.

Metrics. *Pass@1* is the fraction of tasks whose single sample compiles and passes every test, and *Compile rate* is the fraction that compile regardless of tests, so their gap separates logic-only failures from cases where no compilable Cangjie program was produced. *Token usage* sums input and output tokens over the 248 tasks, aggregating every turn for agent runs, and *Failure category distribution* reports the share of unsuccessful samples in each of the seven failure classes of Section 5.5. We report Pass@1 and Compile rate as percentages weighted by task count across the two subsets, and the per-cell records behind every number are in the replication package (Section 9).

5.2 RQ1: How do the four method families compare on pass rate?

In Table 2 the four method families do not order themselves on a single accuracy axis, as the three findings below make precise.



■ **Figure 3** Per-task token usage of Qwen3-Coder across the five methods on (a) Text-to-Code and (b) Code-to-Code. Total tokens are on a log scale, and each bar is split into Input (Prompt) and Output (Completion) percentages, with the absolute total printed above.

Direct generation fails at the compile step, not the logic step. Without Cangjie-specific scaffolding, five of the six base models score under 7 % Pass@1 in either modality, only Kimi-K2 breaks 10 % at 15.3 % T2C and 16.9 % C2C on average, and GPT-5, the strongest model elsewhere, reaches only 4.8 % on Direct T2C. Compile rate sits within about 3 pp of Pass@1 on every Direct cell, so a failing model usually fails before the tests run, unable to produce a program the toolchain accepts.

Syntax-Constrained Prompting dominates the prompt-based methods. On T2C average, every base model gains 22–46 pp over Direct once given the grammar reference: GPT-5 climbs from 4.8 % to 50.4 %, Kimi-K2 from 15.3 % to 40.7 %, and Qwen3 from 2.8 % to 39.5 %. Both RAG variants help less, as averaged over T2C and C2C GPT-5 reaches 33.7 % under RAG-Code, 25.2 % under RAG-Docs, and 44.2 % under Syntax-Constrained, with the same ordering on the open-weights subjects. That a 2,146-token grammar reference outperforms BM25 retrieval over megabytes of Cangjie source shows that **fixing the syntax model matters more than supplying example code**, since retrieved examples cannot fix syntax the model has not yet seen.

Agent reaches the highest pass rate, but only with a strong backbone. Codex CLI on GPT-5 averages 66.9 % T2C and 77.0 % C2C, the best on every subset and modality and the only Agent configuration that dominates Syntax-Constrained on its own backbone, while the other three sit in the 18–40 % range, comparable to or below Syntax-Constrained on the same backbones. The agentic loop amplifies what the backbone can already do, but does not lift a weak backbone past a stronger one running with the syntax reference.

5.3 RQ2: How much does each method cost in tokens?

Figure 3 reports Qwen3-Coder’s per-task token usage across all five methods, spanning two orders of magnitude. Direct generation consumes roughly 1k tokens per task, Syntax-Constrained about 3.3k once the 2,146-token grammar header is added, and the two RAG variants between about 1.7k and 4.6k, while Agent under iFlow CLI jumps to 34–36k and Agent under Qwen Code CLI to 221–248k, roughly 220–250x Direct. Output share falls from about half of Direct’s tokens to below 3 % for both agent runs, so the assistant spends almost all of its budget reading documentation and prior turns rather than writing code.





Figure 4 Pass@1 per 1,000 input tokens for Kimi-K2, Qwen3-Coder, and GPT-5 on (a) Text-to-Code and (b) Code-to-Code. Pass@1 and input tokens are task-count-weighted over HumanEval (164) and ClassEval (84), matching Table 2. Higher is better.

Cost reverses the accuracy ranking. Figure 4 plots Pass@1 per 1,000 input tokens, on which a method that doubles accuracy at tenfold the tokens loses. On Kimi-K2, Direct is the most efficient cell at around 40 Pass@1 per 1k tokens on both modalities despite scoring only 15–17 % Pass@1, since a single 500-token prompt yields a little correct code very cheaply. On Qwen3-Coder and GPT-5, where Direct yields almost no passing samples, Syntax-Constrained and RAG-Code lead, and GPT-5 RAG-Code on T2C reaches 26.3 Pass@1 per 1k tokens, more than 4x RAG-Docs at 6.2 on the same backbone.

The Agent ceiling collapses on the per-token axis. Codex CLI on GPT-5 sets the absolute Pass@1 ceiling, but Figure 4 shows it consumes those gains at 0.14 Pass@1 per 1k tokens on T2C and 0.20 on C2C, tens to nearly two hundred times below the prompt-based methods on the same backbone, and the two iFlow configurations sit at 0.6–1.8. Method choice therefore becomes a question of token budget rather than absolute accuracy: when tokens are unconstrained Codex CLI on GPT-5 is the right choice, but under any practical budget **Syntax-Constrained Prompting is the dominant strategy on every backbone we evaluate**, reaching on Qwen3-Coder 33.9 % T2C and 36.3 % C2C at 10–75x lower cost than the same backbone in either Agent harness.

5.4 RQ3: Does a Python reference help or hurt?

Code-to-Code is not uniformly easier than Text-to-Code. Averaged across all six base models and all prompt-based methods, C2C and T2C are within about 1 pp on HumanEval and T2C is slightly ahead on ClassEval. The Python reference neither reliably helps nor



■ **Figure 5** $\Delta \text{Pass@1}$ (C2C - T2C, percentage points), task-count-weighted over HumanEval and ClassEval. Negative (red) means a Python reference hurts on Code-to-Code, and positive (blue) means it helps. Top: prompt-based methods $\times$ base models. Bottom: each of the four agent configurations.

hurts, since its direction depends on the method and its effect on the average is small next to the method-level variance in Figure 5, which shows $\Delta \text{Pass@1}$ in percentage points, C2C minus T2C, for every cell, with the Cangjie-anchored prompt-based methods on the red side where the reference hurts and Direct, RAG-Docs, and the Agent configurations on the blue side where it helps.

The most striking case is GPT-5 under Syntax-Constrained on HumanEval, where Pass@1 drops from 64.0 % T2C to 44.5 % C2C, a 19.5 pp regression on the same tasks with the Python solution available, and the average column drops 12.5 pp. Compile rate moves with it, from 67.1 % to 45.1 %, so the regression is a syntax issue rather than a logic one, with the reference pulling GPT-5 toward Python idioms that do not survive the Cangjie compiler.

The pattern splits cleanly by method type: under *Syntax-Constrained* and *RAG-Code*, which put Cangjie syntax or code in the prompt, T2C beats C2C on 4 of 6 base models on the average column, so negative transfer concentrates where the model already has Cangjie-shaped material to anchor on. Two apparent exceptions confirm the mechanism rather than break it. Syntax-Constrained on Qwen3-Coder shows a slight positive Δ of about +2.4 pp, the same direction it takes under Direct, because this code-specialised model carries structure over from the Python reference faster than it takes on the Python idioms the compiler rejects. RAG-Docs on GPT-5 shows a negative Δ of about -2.0 pp despite no Cangjie code in the prompt, because retrieved documentation fragments introduce inconsistent Cangjie API calls that the reference compounds. Under *Direct* and *RAG-Docs*, by contrast, the reference helps on most subjects, since with no Cangjie scaffold to anchor on the Python solution at least gives a structural skeleton.

The Agent rows show a sharply different pattern: three of the four configurations score higher under C2C than T2C on the average column, by 10.1 pp for Codex on GPT-5, 7.7 pp for iFlow on Kimi-K2, and 3.2 pp for iFlow on Qwen3-Coder, with only Qwen Code on Qwen3-Coder flat at -0.8 pp, and the gap is largest on ClassEval, where Codex on GPT-5 jumps from 32.1 % to 58.3 %. The multi-turn loop absorbs the negative-transfer pressure, since the extra turns let the agent revise Python-specific syntax that would land verbatim in a single-shot prompt, so **negative transfer is a property of the consuming method, not of the Python reference itself.**

■ **Table 3** Failure distribution per method and modality, primary annotator glm-5.1. Cells show the share of failed samples in each category. Avg. aggregates failed samples across methods within each modality. The last row gives the total failure count for each column.

Failure category	Direct		Syntax-Constrained		RAG (Code)		RAG (Docs)		Agent		Avg.	
	T2C	C2C	T2C	C2C	T2C	C2C	T2C	C2C	T2C	C2C	T2C	C2C
Syntax (foreign)	83.9	82.0	30.0	27.2	56.6	59.1	60.9	61.0	46.1	50.4	58.8	58.7
API/library	12.4	14.0	34.6	42.4	32.5	30.4	26.4	26.6	30.0	30.0	26.0	27.5
Type/mutability	2.6	3.2	16.3	16.8	8.2	8.7	4.6	6.8	13.3	13.7	7.9	8.9
Duplicate def.	0.0	0.1	8.9	10.2	0.0	0.0	4.5	3.7	0.0	0.0	2.6	2.7
Runtime exc.	0.4	0.5	0.7	0.8	0.6	0.7	0.4	0.5	1.2	1.0	0.6	0.7
Logic error	0.6	0.1	9.2	2.6	2.0	1.0	3.1	1.4	9.4	4.7	4.0	1.6
Timeout	0.0	0.0	0.2	0.0	0.1	0.0	0.1	0.0	0.0	0.2	0.1	0.0
<i>Total failures</i>	1,414	1,396	944	990	1,229	1,243	1,238	1,243	647	597	5,472	5,469

5.5 RQ4: Where do the failures come from?

We labelled every failed sample with one of seven error categories.

Taxonomy. The taxonomy was derived through open coding of 788 failed GPT-5 samples on T2C across all five methods. We selected GPT-5 because it reaches the widest range of failure modes, whereas weaker models stall almost entirely on foreign-syntax errors, and we seeded on T2C because the model cannot attribute failures to the Python reference, so each error traces to a gap in its Cangjie knowledge. The first author grouped samples by recurring surface patterns and merged or split categories until no large cluster remained unlabelled, producing six categories, with a seventh, *Timeout*, added for runs that do not finish within the evaluation budget.

Annotation procedure. For each failed sample we supplied the LLM annotator with the compiler or runtime error output, the first failing error line, and the category descriptions with one short example each, and instructed it to assign exactly one label, report a confidence level, give a one-sentence rationale, and quote the primary evidence.

Annotator selection. We chose glm-5.1 and kimi-k2.6 as annotators because both have strong Chinese-language understanding, which matters as Cangjie documentation and error messages often mix English identifiers with Chinese descriptions, and they come from different organisations, Zhipu AI and Moonshot AI, so agreement across two independently trained models is stronger evidence of label reliability than two runs of one model. Treating glm-5.1 as primary and kimi-k2.6 as a cross-check, applying both to the 788 manually annotated samples reproduced the manual labels on almost every sample, so we extended LLM labelling to the full benchmark.

The seven categories are *Syntax (foreign)* for invalid syntax mirroring another language, *API/library* for nonexistent identifiers or wrong call shapes, *Type/mutability* for type or mutability mismatches, *Duplicate def.* for redeclarations or harness conflicts, *Runtime exception* for uncaught exceptions, *Logic error* for wrong test outputs, and *Timeout* for runs that did not finish. Table 3 reports the failure mix per method and modality.




Figure 6 Inter-annotator confusion matrix between glm-5.1 (rows) and kimi-k2.6 (columns) for (a) T2C and (b) C2C. Each cell shows the row-normalised percentage and the absolute count.

Foreign-syntax errors dominate, and only the grammar reference cuts them. Direct generation fails almost entirely on syntax, with about 83 % of its failed samples on both T2C and C2C being *Syntax (foreign)*, code that borrows Python, Rust, Java, or Swift constructs the Cangjie compiler rejects, such as `if x > 0:` written without parentheses, `for i in range(n)` loops, and `|acc, x|` closure syntax, matching the roughly 3 pp Pass@1-to-Compile gap on every Direct cell. The pattern is consistent across the weaker models, which almost never emit a construct Cangjie accepts once it diverges from Python, so their low Pass@1 reflects a parser they cannot satisfy rather than an algorithm they cannot find. The other methods address this with mixed success: Syntax-Constrained cuts foreign-syntax failures to about 28–30 %, but RAG-Code and RAG-Docs still leave 56–61 % despite putting Cangjie tokens in the prompt, because retrieval supplies surface text without overriding the model’s syntactic priors, and Agent sits at 46–51 % where the multi-turn loop revises some surface errors.

Where the rest of the failures shift. Once Syntax-Constrained removes most foreign syntax, the mass shifts onto API/library at 34–42 % and Type/mutability at around 16 %, because the grammar reference teaches the parser-level shape of Cangjie but not the standard library or type rules, so models write valid Cangjie that calls methods Cangjie lacks or mixes incompatible numeric types. The C2C column spikes at 42.4 % on API/library versus 34.6 % on T2C, consistent with RQ3, as the Python reference seeds Python-style API names Cangjie rejects. About 9–10 % of Syntax-Constrained failures are Duplicate def., where the model declares its own `main` and conflicts with the harness, and Logic error reaches a non-trivial level only under Agent on C2C, at 4.7 %, the only method where many failed samples reach the test phase with code the toolchain accepts.

Cross-annotator validation. Figure 6 reports the 7×7 confusion matrix between glm-5.1 and kimi-k2.6. Both matrices are almost entirely diagonal, with raw agreement 99.34 % on T2C ($\kappa = 0.989$) and 98.92 % on C2C ($\kappa = 0.981$) and no off-diagonal pair above 2 % of any row, so the percentages in Table 3 do not depend on the choice of annotator.

■ **Table 4** Repeated syntax-guided sampling against single-run agents on Text-to-Code. *Syntax* reports task-level Pass@10 from ten Syntax-Constrained samples per task, and *Agent* reports Pass@1 from one CLI run. $+N.N$ shows the pp gain of Syntax Pass@10 over Agent Pass@1. *Tokens* is the input-plus-output sum per task. Same backbone within each row.

Backbone	Agent harness	HumanEval				ClassEval			
		Syntax		Agent		Syntax		Agent	
		Pass@10	Tokens	Pass@1	Tokens	Pass@10	Tokens	Pass@1	Tokens
Kimi-K2	iFlow CLI	84.1 $+45.1$	27.1k	39.0	24.2k	28.6 $+10.7$	36.6k	17.9	27.5k
Qwen3-Coder	iFlow CLI	64.0 $+34.1$	28.1k	29.9	36.1k	19.0 $+14.2$	42.2k	4.8	36.2k
Qwen3-Coder	Qwen Code CLI	64.0 $+39.0$	28.1k	25.0	224.8k	19.0 $+11.9$	42.2k	7.1	212.9k
GPT-5	Codex CLI	96.3 $+11.5$	46.4k	84.8	391.7k	42.9 $+10.8$	82.8k	32.1	618.2k

6 Discussion

For Cangjie, the bottleneck on every task is parser-level rather than reasoning-level, since direct generation fails almost entirely on syntax and the Compile-to-Pass@1 gap is at most about 3 pp on every Direct cell. The model already knows the algorithm but cannot yet write a program the toolchain accepts, so what separates the four method families is how each closes parser distance to Cangjie and at what token cost.

6.1 Spending the Agent’s Budget on Repeated Syntax Sampling


Section 5.3 shows an agent run costs 10–120x a single Syntax-Constrained call for a single-digit Pass@1 gain, so an obvious alternative is to keep the cheap method and draw several samples instead of one long agent trace. We test this with ten Syntax-Constrained samples per Text-to-Code task, scored under task-level Pass@10, and match it in Table 4 against the Pass@1 of the same backbone wrapped in each agent harness.

Repeated sampling beats the matched agent in every cell, on both subsets. On Kimi-K2 the gap is 45.1 pp on HumanEval at roughly iFlow’s token cost, and on Qwen3-Coder against Qwen Code CLI it is 39.0 pp on HumanEval at one-eighth of the agent’s budget. Even GPT-5 under Codex, the strongest cell in Table 2, gains 11.5 pp on HumanEval and 10.8 pp on ClassEval by switching to parallel syntax-guided samples within roughly one-eighth of Codex’s budget.

This is not an unconditional argument against agents. Pass@10 assumes a test harness that can pick the passing sample, and without tests ten samples are ten guesses. With tests, the budget pays off more on alternative implementations than on long tool-use traces, because most failures the agent repairs are surface syntax errors a fresh draft would not have made, and repeated sampling is parallelisable, so its wall-clock cost can be below an agent run that waits on each compiler call. The edit-and-rerun loop thus earns its tokens only when the failure mix rewards serial repair rather than a fresh draft.

6.2 Why Retrieval Helps Less Than Expected

Both RAG variants trail Syntax-Constrained Prompting on all base models, and failure analysis sharpens the point, since 56–61% of RAG failures are foreign-syntax failures (Table 3) even with Cangjie tokens in the prompt. Retrieved context is plentiful but rarely corrective.



■ **Figure 7** Two representative RAG failures. On `all_prefixes`, RAG-Docs returns `ArrayList` documentation when the task needs Cangjie string slicing, and the model writes `substring` (an API Cangjie does not expose), which the compiler rejects. On `sum_product`, RAG-Code returns syntactically valid but task-irrelevant snippets, and the model still emits `while i < n` without parentheses.

Figure 7 illustrates two cases where retrieval returns documentation or source unrelated to the task and the model falls back to a foreign API name or a Python-style loop the compiler rejects, so retrieval produces context but not the context that would prevent the failure.

For an emerging language this matters more than for Python, because with a thinner corpus a noisy index can no longer compensate by sheer volume, and two failure modes drive the foreign-syntax share even when the prompt is full of Cangjie text. The first is vocabulary mismatch, where a natural-language description and a Cangjie source file share almost no surface tokens, so a BM25 query for “slice the prefix of a string” rarely surfaces the slicing example a developer would open. The second is the small-corpus problem, where among only a few tens of thousands of crawled functions many APIs appear once or not at all, so the snippet the model needs is often absent. A more useful index would be small and compiler-checked, organised by API signature and paired with short examples for the constructs that most often break, loops, string handling, option types, and mutability, so that retrieval aims at the most likely compiler error rather than enlarging the prompt with plausible but inert text. Until such an index exists, the RAG cells in Table 2 are best read as adding tokens that do not move parser distance, not as a real alternative to the syntax reference.

6.3 A Staged Workflow for Emerging-Language Assistants

No single ranking covers both accuracy and token cost, so the practical question is which method to reach for at each stage of a workflow. Direct Generation is best read as a measurement of how much Cangjie a model already knows rather than a deployment choice. Syntax-Constrained Prompting is the appropriate first stage, since it dominates both RAG variants on every base model, addresses the largest failure class, and adds only a 2,146-token header, and when a test harness is available repeated syntax-guided sampling extends it to recover most of the agent’s headline accuracy at a fraction of the budget (Section 6.1).

Curated retrieval, indexed by signatures rather than free text, is worth its tokens on API-heavy tasks where the syntax reference leaves API/library failures unaddressed, and only after a Syntax-Constrained attempt has surfaced a concrete API gap. Agents are best reserved for tasks that genuinely require interaction, such as inspecting files, preserving multi-method class structure, or repairing Python-specific drift across turns, where Section 5.4 shows them earning their cost on Code-to-Code. The staging is therefore an escalation triggered by the failure class the previous step cannot fix.

7 Threats to Validity

Construct validity. Pass@1 and Compile rate measure whether the toolchain accepts the program and whether it passes the tests we wrote, not whether the code is idiomatic or maintainable in a real codebase. We mitigate this by writing the tests against the same reference solutions we translated and reviewing each suite against deliberate perturbations so that obvious wrong answers are rejected. Token cost counts only API tokens rather than wall-clock time, and agent runs that interleave shell calls may incur additional costs, so agent figures are a lower bound on the true gap.

Internal validity. Every cell uses greedy decoding with a single sample, so the comparison is sensitive to sampling noise, and we report a ten-sample Pass@10 robustness check under Syntax-Constrained Prompting in Section 6.1. The failure categories are produced by an LLM annotator rather than a human, so we cross-checked them against a second LLM and against 788 manually inspected GPT-5 failures before extending the labels to the rest of the benchmark. The translation rules, taxonomy seed set, and category definitions are themselves methodological choices, so we derived the rules from pilot work rather than post-hoc, seeded the taxonomy on the strongest model across all five methods for coverage, and checked that the categories are jointly exhaustive and mutually exclusive on the inspection set. Our Cangjie reference solutions and tests were written by hand and could contain bugs, so two additional authors reviewed every problem, and we verified that each suite accepts the reference solution and rejects deliberately broken variants.

External validity. We evaluate one emerging language, six base models, and three agent harnesses, so any quantitative claim is tied to that subject set. The qualitative pattern, that direct generation fails at the parser before the algorithm and that a small grammar reference closes most of the gap, holds because Cangjie is statically typed with limited public code; we expect it to transfer to other emerging general-purpose languages in the same early-adoption phase, but not to dynamically typed languages or DSLs, where the failure surface is shaped by domain semantics rather than parser distance. Tasks come from HumanEval and ClassEval, which are algorithm-level rather than repository-level, so our results say little about long-context Cangjie work such as multi-file edits or build configuration, where an agent’s ability to navigate a project may matter far more than single-file syntax. Pinning the toolchain and CLI versions means absolute pass rates are a snapshot; the ordering across methods is wide enough that it is unlikely to reverse under minor version changes.

8 Conclusion

We asked whether current LLM coding assistants can support an emerging programming language, and the answer from CangjieBench is that they can, but only when the method is matched to where the model actually fails. For the six models we evaluated, the failure on

Cangjie is not one of reasoning: direct generation fails almost entirely on syntax, a short syntax reference of about two thousand tokens closes most of that gap, and retrieval helps less than its prompt size suggests because a BM25 index over public Cangjie code rarely returns the snippet that would prevent the next compiler error. The best agent configuration reaches the highest absolute Pass@1, but only at ten to one hundred and twenty times the cost of a single prompt, while weaker agents do not match a Syntax-Constrained call on the same backbone. For a team adopting an emerging language today, the implication is to ship a syntax cheat sheet first, add curated retrieval once it is tailored to the language, and reserve an agent loop for tasks that need file inspection or multi-turn repair. The same staging also guards against the negative transfer we observed under Code-to-Code, where a Python reference pulls a single-shot prompt toward idioms the Cangjie compiler rejects and only the multi-turn loop reliably repairs them. We expect the qualitative pattern, that syntax is the barrier before semantics, to hold for other statically typed languages in their early-adoption phase, with the thresholds shifting as public code and documentation accumulate. We release CangjieBench, the harness, and the failure taxonomy, validated by two LLM annotators with $\kappa \geq 0.98$, so that the same comparison can be repeated as models and tooling evolve.

9 Data Availability Statement

All data used in this study are available under the MIT licence at <https://github.com/cjhCoder7/CangjieBench>. The repository contains five components, cross-referenced below to the sections that use them.

1. **CangjieBench dataset** (Section 3.1). The 248-task benchmark ships as two JSONL files, one holding the 164 function-level HumanEval tasks and one the 84 class-level ClassEval tasks, with each entry giving the task identifier, Cangjie prompt, reference solution, and test harness, alongside the per-task Cangjie and Python source files.
2. **Evaluation sandbox** (Section 3.3). A Docker image packaging the Cangjie 1.0 compiler, runtime, standard libraries, and standard extensions (stdx), exposing three REST endpoints (`/evaluate_humaneval`, `/evaluate_classeval`, and `/run_code`) for automated compilation and testing. We used the same image to produce every cell in Table 2.
3. **Per-cell evaluation records** (Sections 5.2–5.4). Per model, method, and task we provide the raw generation, the compile/pass status, compiler and runtime stderr, and the input and output token counts, with per-cell summaries and the Pass@10 data from Section 6.1.
4. **Failure analysis** (Section 5.5). Both annotators, glm-5.1 and kimi-k2.6, labelled every failed sample with one of the seven categories, each record giving the category, annotator confidence, the compiler or runtime evidence, and a short rationale, with per-cell distribution summaries.
5. **Prompt templates and method resources** (Section 4). We provide the four prompt templates for Direct, Syntax-Constrained, RAG-Docs, and RAG-Code together with the agent prompt for both modalities, the 2,146-token Cangjie syntax reference used by Syntax-Constrained Prompting, and the two RAG corpora, namely the official Cangjie documents for RAG-Docs and the crawled GitCode source for RAG-Code.

Generative AI tool usage. We used large language models, glm-5.1 and kimi-k2.6, as automated annotators to label failure categories across the full benchmark, as described in Section 5.5. All model-generated labels were cross-validated against each other and against a manually annotated seed set.

References

- 1 Wasi Ahmad, Md Golam Rahman Tushar, Saikat Chakraborty, and Kai-Wei Chang. Avatar: A parallel corpus for java-python program translation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 2268–2281, 2023.
- 2 Anthropic. Claude code. URL <https://www.claude.com/product/claude-code>, 2025.
- 3 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint*, 2021. [arXiv:2108.07732](https://arxiv.org/abs/2108.07732).
- 4 Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, et al. Multipl-e: A scalable and polyglot approach to benchmarking neural code generation. *IEEE Transactions on Software Engineering*, 49(7):3675–3691, 2023. doi:10.1109/TSE.2023.3267446.
- 5 Mark Chen. Evaluating large language models trained on code. *arXiv preprint*, 2021. [arXiv:2107.03374](https://arxiv.org/abs/2107.03374).
- 6 Junhang Cheng. cjhCoder7/CangjieBench. Software (visited on 2026-09-18). URL: <https://github.com/cjhCoder7/CangjieBench>.
- 7 Artem Chervyakov, Alexander Kharitonov, Pavel Zadorozhny, Adamenko Pavel, Rodion Levichev, Dmitrii Vorobev, Dmitrii Salikhov, Aidar Valeev, Alena Pestova, Maria Dziuba, et al. Mera code: A unified framework for evaluating code generation across tasks. *arXiv preprint*, 2025. [arXiv:2507.12284](https://arxiv.org/abs/2507.12284).
- 8 Xueying Du, Mingwei Liu, Kaixin Wang, Hanlin Wang, Junwei Liu, Yixuan Chen, Jiayi Feng, Chaofeng Sha, Xin Peng, and Yiling Lou. Classeval: A manually-crafted benchmark for evaluating llms on class-level code generation. *arXiv preprint*, 2023. doi:10.48550/arXiv.2308.01861.
- 9 Google. Gemini. URL <https://deepmind.google/models/gemini/>, 2025.
- 10 Google. Gemini cli. URL <https://geminicli.com/>, 2025.
- 11 Minghua He, Yue Chen, Fangkai Yang, Pu Zhao, Wenjie Yin, Yu Kang, Qingwei Lin, Saravan Rajmohan, and Dongmei Zhang. Execoder: Empowering large language models with executability representation for code translation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 7110–7136, 2025.
- 12 Huawei. Cangjie language documentation (v1.0.0). URL <https://cangjie-lang.cn/en/docs?url=%2F1.0.0%2F>, 2025.
- 13 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint*, 2024. doi:10.48550/arXiv.2403.07974.
- 14 Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik Narasimhan. Swe-bench: Can language models resolve real-world github issues? *arXiv preprint*, 2023. doi:10.48550/arXiv.2310.06770.
- 15 Sathvik Joel, Jie Wu, and Fatemeh Fard. A survey on llm-based code generation for low-resource and domain-specific programming languages. *ACM Transactions on Software Engineering and Methodology*, 2024.
- 16 Sahil Kale and Vijaykant Nadadur. Texpert: A multi-level benchmark for evaluating latex code generation by llms. *arXiv preprint*, 2025. doi:10.48550/arXiv.2506.16990.
- 17 Anirudh Khatry, Robert Zhang, Jia Pan, Ziteng Wang, Qiaochu Chen, Greg Durrett, and Isil Dillig. Crust-bench: A comprehensive benchmark for c-to-safe-rust transpilation. *arXiv preprint*, 2025. doi:10.48550/arXiv.2504.15254.
- 18 Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, et al. Can llm already serve as a database interface? a big bench for large-scale database grounded text-to-sqls. *Advances in Neural Information Processing Systems*, 36:42330–42357, 2023.

- 19 Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint*, 2024. [arXiv:2412.19437](#).
- 20 Mingjie Liu, Nathaniel Pinckney, Brucek Khailany, and Haoxing Ren. VerilogEval: Evaluating large language models for verilog code generation. In *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, pages 1–8. IEEE, 2023.
- 21 Yang Liu, Li Zhang, Fang Liu, Zhuohang Wang, Donglin Wei, Zhishuo Yang, Kechi Zhang, Jia Li, and Lin Shi. Reposcope: Leveraging call chain-aware multi-view context for repository-level code generation. *arXiv preprint*, 2025. [arXiv:2507.14791](#).
- 22 Enkhbold Nyamsuren. Evaluating quantized large language models for code generation on low-resource language benchmarks. *Journal of Computer Languages*, page 101351, 2025. doi:10.1016/J.COLA.2025.101351.
- 23 OpenAI. Introducing codex. URL <https://openai.com/index/introducing-codex/>, 2025.
- 24 OpenAI. Introducing gpt-5. URL <https://openai.com/index/introducing-gpt-5/>, 2025.
- 25 Anne Ouyang, Simon Guo, Simran Arora, Alex L Zhang, William Hu, Christopher Ré, and Azalia Mirhoseini. Kernelbench: Can llms write efficient gpu kernels? *arXiv preprint*, 2025. doi:10.48550/arXiv.2502.10517.
- 26 Qiwei Peng, Yekun Chai, and Xuhong Li. Humaneval-xl: A multilingual code generation benchmark for cross-lingual natural language generalization. *arXiv preprint*, 2024. doi:10.48550/arXiv.2402.16694.
- 27 Zhiyuan Peng, Xin Yin, Rui Qian, Peiqin Lin, Yongkang Liu, Hao Zhang, Chenhao Ying, and Yuan Luo. Soleval: Benchmarking large language models for repository-level solidity code generation. *arXiv preprint*, 2025. [arXiv:2502.18793](#).
- 28 Sameer Pimparkhede, Mehant Kammakomati, Srikanth G Tamilselvam, Prince Kumar, Ashok Pon Kumar, and Pushpak Bhattacharyya. Doccgen: Document-based controlled code generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 18681–18697, 2024. doi:10.18653/V1/2024.EMNLP-MAIN.1040.
- 29 Nathaniel Pinckney, Christopher Batten, Mingjie Liu, Haoxing Ren, and Brucek Khailany. Revisiting verilogEval: A year of improvements in large-language models for hardware code generation. *ACM Transactions on Design Automation of Electronic Systems*, 2025.
- 30 Haritz Puerto, Martin Tutek, Somak Aditya, Xiaodan Zhu, and Iryna Gurevych. Code prompting elicits conditional reasoning abilities in text+ code llms. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 11234–11258, 2024. doi:10.18653/V1/2024.EMNLP-MAIN.629.
- 31 Baidu ERNIE Team. Ernie 4.5 technical report, 2025. URL: https://ernie.baidu.com/blog/publication/ERNIE_Technical_Report.pdf.
- 32 Kimi Team, Yifan Bai, Yiping Bao, Guanduo Chen, Jiahao Chen, Ningxin Chen, Ruijue Chen, Yanru Chen, Yuankun Chen, Yutian Chen, et al. Kimi k2: Open agentic intelligence. *arXiv preprint*, 2025. [arXiv:2507.20534](#).
- 33 Prashanth Vijayaraghavan, Luyao Shi, Stefano Ambrogio, Charles Mackin, Apoorva Nitsure, David Beymer, and Ehsan Degan. Vhdl-eval: A framework for evaluating large language models in vhdl code generation. In *2024 IEEE LLM Aided Design Workshop (LAD)*, pages 1–6. IEEE, 2024.
- 34 Ngoc Phuoc An Vo, Brent Paulovicks, and Vadim Sheinin. Tackling execution-based evaluation for nl2bash. *CoRR*, 2024. doi:10.48550/arXiv.2405.06807.
- 35 Hanbin Wang, Zhenghao Liu, Shuo Wang, Ganqu Cui, Ning Ding, Zhiyuan Liu, and Ge Yu. Intervenor: Prompting the coding ability of large language models with the interactive chain of repair. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 2081–2107, 2024. doi:10.18653/V1/2024.FINDINGS-ACL.124.
- 36 Min Xue, Artur Andrzejak, and Marla Leuther. An interpretable error correction method for enhancing code-to-code translation. In *The Twelfth International Conference on Learning Representations*, 2024.

- 37 Pengyu Xue, Linhao Wu, Zhen Yang, Chengyi Wang, Xiang Li, Yuxiang Zhang, Jia Li, Ruikai Jin, Yifei Pei, Zhaoyan Shen, et al. Classeval-t: Evaluating large language models in class-level code translation. *Proceedings of the ACM on Software Engineering*, 2(ISSTA):1421–1444, 2025. doi:10.1145/3728940.
- 38 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint*, 2025. arXiv:2505.09388.
- 39 Zezhou Yang, Sirong Chen, Cuiyun Gao, Zhenhao Li, Xing Hu, Kui Liu, and Xin Xia. An empirical study of retrieval-augmented code generation: Challenges and opportunities. *ACM Transactions on Software Engineering and Methodology*, 2025.
- 40 Hao Yu, Bo Shen, Dezhi Ran, Jiaxin Zhang, Qi Zhang, Yuchi Ma, Guangtai Liang, Ying Li, Qianxiang Wang, and Tao Xie. Codereval: A benchmark of pragmatic code generation with generative pre-trained models. In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*, pages 1–12, 2024. doi:10.1145/3597503.3623316.
- 41 Longhui Zhang, Bin Wang, Jiahao Wang, Xiaofeng Zhao, Min Zhang, Hao Yang, Meishan Zhang, Yu Li, Jing Li, and Jun Yu. Function-to-style guidance of llms for code translation. In *Forty-second International Conference on Machine Learning*, 2025.
- 42 Qianhui Zhao, Li Zhang, Fang Liu, Junhang Cheng, Chengru Wu, Junchen Ai, Qiaoyuanhe Meng, Lichen Zhang, Xiaoli Lian, Shubin Song, et al. Towards realistic project-level code generation via multi-agent collaboration and semantic architecture modeling. *arXiv preprint*, 2025. arXiv:2511.03404.
- 43 Zixiao Zhao and Fatemeh Fard. Do current language models support code intelligence for r programming language? *ACM Transactions on Software Engineering and Methodology*, 34(8):1–39, 2025. doi:10.1145/3735635.