

# EXECUTION-EVAL: CAN LANGUAGE MODELS EXECUTE REAL-WORLD CODE?

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

As language models advance, traditional benchmarks face challenges of dataset saturation and disconnection from real-world performance, limiting our understanding of true model capabilities. We introduce EXecution-Eval (EXE), a benchmark designed to assess LLMs’ ability to execute code and predict program states. EXE attempts to address key limitations in existing evaluations: difficulty scaling, task diversity, training data contamination, and cost-effective scalability. Comprising over 30,000 tasks derived from 1,000 popular Python repositories on GitHub, EXE spans a wide range of lengths and algorithmic complexities. Tasks require models to execute code, necessitating various operations including mathematical reasoning, logical inference, bit manipulation, string operations, loop execution, and maintaining multiple internal variable states during computation. Our methodology involves: (a) selecting and preprocessing GitHub repositories, (b) generating diverse inputs for functions, (c) executing code to obtain ground truth outputs, and (d) formulating tasks that require models to reason about code execution. This approach allows for continuous new task generation for as few as 1,123 tokens, significantly reducing the risk of models “training on the test set.” We evaluate several state-of-the-art LLMs on EXE, revealing insights into their code comprehension and execution capabilities. Our results show that even the best-performing models struggle with complex, multi-step execution tasks, highlighting specific computational concepts that pose the greatest challenges for today’s LLMs. Furthermore, we review EXE’s potential for finding and predicting errors to aid in assessing a model’s cybersecurity capabilities. We propose EXE as a sustainable and challenging testbed for evaluating frontier models, offering insights into their internal mechanistic advancement.

## 1 INTRODUCTION

Language model benchmarks are facing challenges of rapid saturation (Ott et al., 2022) and an increasing disconnect from real-world performance perceived by end-users (Zheng et al., 2023). Due to this, benchmarks are being continually created to address failure modes; e.g. SuperGLUE targeting GLUE’s low problem difficulty (Wang et al., 2019), BIG-bench targeting general low evaluation diversity (Srivastava et al., 2022) and Auto-Arena-Hard targeting training-set contamination and data diversity in Chatbot-Arena (Li et al., 2024)(Chiang et al., 2024). These failure modes all demonstrate the challenge in linking the mechanistic improvements within language models to human understandable tasks.

Hence, to maximise an evaluation’s utility we aim to minimise the common failure modes of; a) difficulty, not ensuring an unbound scale of small trivial problems to complex multi-step problems, b) diversity, not ensuring a representative distribution across a large space of problems, c) novelty, not ensuring continually fresh, out-of-training data samples can be generated and, d) scalability, not ensuring tasks are cost-effective to generate in the thousands and beyond.

Motivated by these challenges we introduce EXecutionEval (EXE), an evaluation replicating one of the primary tasks humans perform while coding; predicting and comparing a final program state for a given set of inputs - seen in Figure 1. EXE is designed to avoid the aforementioned failure modes; emphasising difficulty (smooth scale from trivial 1-step, one-line functions to difficult 100s-of-step, multi-layer functions), diversity (unbound number of test cases generatable for tasks from

054  
055  
056  
057  
058  
059  
060  
061  
062  
063  
064  
065  
066  
067  
068  
069  
070  
071  
072  
073  
074  
075  
076  
077  
078  
079  
080  
081  
082  
083  
084  
085  
086  
087  
088  
089  
090  
091  
092  
093  
094  
095  
096  
097  
098  
099  
100  
101  
102  
103  
104  
105  
106  
107

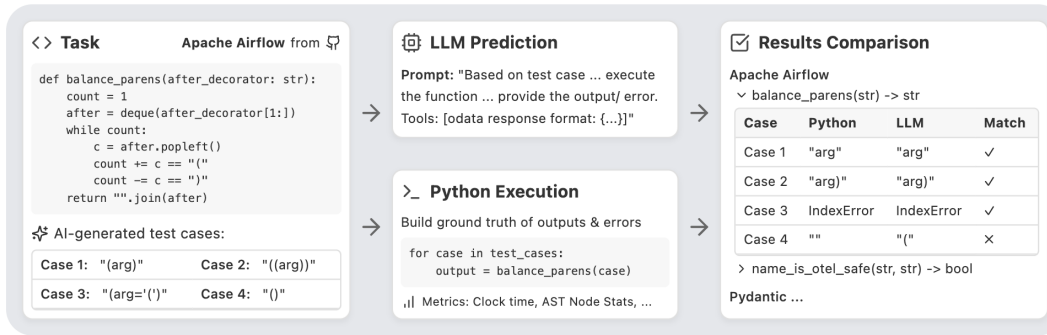


Figure 1: An example task from Apache Airflow’s Github repository (code simplified to fit within diagram). EXE sources tasks from 1,000 Python repositories, generates test cases for them, and compares the LLM’s ability to execute code against python’s interpreter.

1,000 GitHub Repos), novelty (program inputs can be continually generated) and scalability (initial release containing 30,000+ problems at a cost of \$33).

EXE also holds theoretical inspiration. (Fowler et al., 2022) et al have replicated positive pedagogical correlations found by (Lopez et al., 2008) between the abilities of CS1 students to “trace” programs (i.e. manually predict outputs and write the internal state out line by line) and their abilities to pass code writing and explanation exams. This is mirrored in CRUX-Eval’s (Gu et al., 2024) findings, where they observe a moderate correlation between a model’s ability to execute a block of code and a model’s HumanEval (Chen et al., 2021) code writing Pass@1 rate.

## 2 EVALUATION FRAMEWORK

As seen in Figure 1, an EXE task is to predict a function’s return value or error from: a) a code snippet and b) a set of input arguments. Code snippets are extracted from PyPi’s most popular 1,000 python projects hosted on GitHub, we select our snippets to be pure (i.e. deterministic, no side effects), language model generatable (i.e. arg types of ints, lists, ...) and to only require builtins (local imports and external libraries are inlined for the snippet). To realise this we follow the following three stage pipeline 2:




Figure 2: Three stage EXE task generation pipeline. Detailed example tasks and generated inputs can be found in Appendix A.1.

**1. Repo Selection and Code Scraping.** We first select the top 1,000 most popular pypi packages and collate the corresponding github repos where possible, similar to (Jimenez et al., 2023). Repositories are filtered to include only those with permissive licences that allow derivative works with attribution. These repos are then pulled down locally and filtered based on a static Abstract Syntax Tree (AST) analysis determining which repositories contain type-annotated code.

**2. Function Selection and Dependency Collation.** We perform a static AST analysis to filter to functions with LLM generatable argument and return type annotations. Further AST analysis

then recursively identifies dependent elements (modules, functions, classes, variables, ...) across files, builds a dependency graph, and inlines them into a base task. Finally, base tasks containing side effects or non-deterministic code such as environment variables, process calls, randomness or network requests are filtered out. See Appendix A.3 for step-by-step methodology and A.5 for detail on acceptable type annotations and filtering.

**3. Test Case Generation.** Using the argument type annotations we construct a LLM function calling schema that generates a diverse set of inputs. The base task code is then executed with each generated input and the result with runtime statistics are logged. This forms the test case (base task code + generated input), output (returned result or error from executed code) and statistics (runtime statistics + static AST analysis statistics). See Appendix A.2 for step-by-step methodology and Appendix A.6 for details on statistics.




Figure 3: We observe task counts per repository to have a near logarithmic falloff. Note: Based on manual observations, several repositories are removed from EXE due to thousands of similar functions with only single modifications, for example changing a url address.

Through these stages of filtering, the original top 1,000 repositories are filtered down to the 33,875 task instances which comprise EXE. A high level breakdown of these task instances across repositories is presented in Figure 3. We note some repositories are overrepresented primarily due to being more modern (using type annotations) and the style of code (shorter deterministic pieces).

## 2.1 TASK FORMATION

**Model input.** The model is given a complete snippet of code alongside the input state to be executed. The model is then tasked to predict the resulting return value, or in the case that an exception is raised the model is instructed to generate an exception type and value. In practice, we prompt models with an odata json representation and use a parser to ensure valid generations. We do append one additional user reply with the parsing error if the model’s response fails to parse. Examples of input instances can be found in Appendix A.1.

**Evaluation metrics.** To evaluate a proposed solution, we use the pass@k metric (Chen et al., 2021), comparing the ground truth and the generated prediction as json objects (set and frozenset are sorted before conversion to json lists). If the original code produced an exception, we compare the type and message (excluding stacktrace) using a language model comparison. See detailed methodology in Appendix A.7 and see examples of generated outputs in Appendix A.1.

## 2.2 FEATURES OF EXE

**Diversity of inputs and outputs.** Unlike many benchmarks focused on a particular subject matter area, a task in this eval may require a model to perform mathematical reasoning, logical inference, bit manipulation, string operations, loop execution, or to maintain multiple internal variables during computation. Furthermore, these may only form part of an algorithm that the model has to execute. Our random human inspection has uncovered algorithmic time complexities spanning from  $O(1)$  to  $O(x^n)$  and structured analysis has found tasks with code context lengths ranging from 440 to 311,000 tokens. Ensuring this broad diversity reduces the risk of hitting a local maxima and increases our opportunity to measure internal capabilities across a range of difficulties.

**Continually updatable.** Both our code collection and task input generation processes can create new tasks with minimal human oversight. Simply re-running our code collection to pull the latest

commits or directing it towards an uncollected Python GitHub repository will create new task instances. Furthermore we can continue to generate new test cases for existing tasks, our test case generator automatically avoids generating seen inputs. Hence, EXE can be extended continually with new task instances, ensuring answers were not included in training corpuses of models for evaluation.

**Cost effective scalability.** With generation of new tasks requiring an average of 1,112 input tokens (batch of 15) and evaluation of tasks typically requiring 1,123 tokens, ExecEval can be generated, tested and continually updated at a fraction of the cost of human-curated benchmarks. Our initial dataset of 33,875 cases has only incurred an approximate costing of \$33 to produce and \$95 to test on.

**Long multi-step problems with smooth difficulty scaling.** We provide a continuous spectrum of task difficulties, ranging from 1-step, one-line functions to multi-file, multi-class, multi-100-step tasks. Our most complex tasks include function call depths (non-recursive) of up to 13 levels (median: 2), separate identifier counts (i.e. variable names, function names, ...) of up to 823 (median: 16) and up to 63 if statements (median: 1). This smooth scaling of difficulty allows for a more detailed measurement of model coherence along multi-step problems than what is typically seen in traditional evaluations. However, as language models continue to advance rapidly, even this wide range of difficulties may eventually face saturation.

To address this, we observe a mechanism inspired by the SKILL-MIX evaluation (Yu et al., 2023) that leverages the typed nature of our function selection process. This approach allows us to create even more complex tasks by chaining functions where the output type of one matches the input type of another, or by combining multiple outputs into a composite input. The number of potential new tasks can be upper bounded by  $n^2 \cdot (T_{\max})^k \cdot C$ , where  $n$  is the total number of types,  $T_{\max} = \max_{i,j} T_{i,j}$  is the maximum number of existing tasks between any two types,  $k$  is the number of functions to chain, and  $C$  is the average number of test cases per task. While this is an upper bound and the actual number of valid composite tasks would be lower due to specific type compatibility constraints, it still represents a significant expansion of our task space. We view this as an opportunity to trade some of the 'realism' of using 100% real-world code for the ability to probe the upper bounds of model capabilities. For constant compute models, this approach allows us to test their internal mechanistic capabilities in handling increasingly complex, multi-step problems. And for chain-of-thought models, it provides a test of increasingly long-term agentic coherency.

**Error prediction.** To test the full spectrum of code execution we further generate test cases designed to trigger exceptions. Many of these require in-depth analysis to see ahead of time, for example predicting an invalid array index through multiple functions. While debugging exceptions is one of the more challenging software engineering tasks, we are yet to see it commonly evaluated in benchmarks.

### 3 RESULTS

We report our evaluation results across different SOTA models alongside our findings across different task statistics below.

Table 1: EXE Pass@1 results

Model	EXE dataset (Pass@1)	Errors (Pass@1)
GPT-4o	72.4	49.5
GPT-4o-mini	60.9	32.0
Llama3.1-8B	37.4	2.1
Llama3.1-405B	71.4	34.3
Claude3.5-Sonnet	76.1	45.8
Mistral-Large-2407	71.5	33.7

**LLMs can execute real-world code, achieving results in-line with code generation benchmarks.** We find EXE shows similar relative model performance between models as seen in coding benchmarks such as HumanEval (Chen et al., 2021) and as seen in benchmarks requiring logical inference

such as (Lu et al., 2023). Furthermore we find a similar diversity of performance across packages as seen in agentic benchmarks such as (Jimenez et al., 2023). We show our findings in Figure 4.

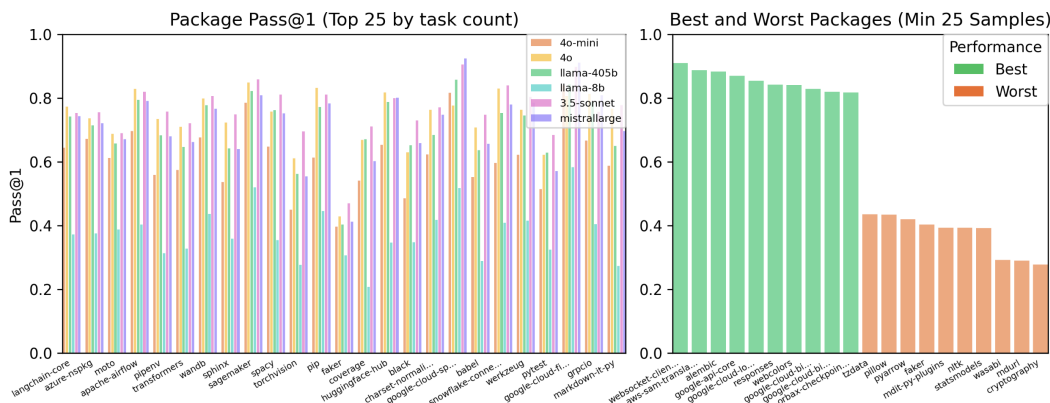


Figure 4: Left - We show the relative accuracy of different models across the top 20 packages by task count. Both the relative differences between models and the relative differences between packages are within expectations from other coding benchmarks (Jimenez et al., 2023). Right - We show the magnitude of diversity across packages (mean performance across all models).

Prior works such as Learning To Execute (Zaremba & Sutskever, 2014) and CRUX-Eval (Gu et al., 2024) have placed justifiable limitations on code complexity; removing mathematical operations, limiting line count, disallowing custom classes and only having one singular function to name a few. We hypothesised that these are no longer necessary, and to understand the true internal capabilities of a constant compute model (i.e. no Chain of Thought) we must test on real-world code, only applying limitations where forced (i.e. no arbitrary object inputs, as LLMs can’t generate them). Our results as seen in table 1 provide initial evidence towards our hypothesis.

**ExecEval provides a smooth curve of task difficulties.** We set out to ensure a) our evaluation does not induce saturation from a bounded distribution of task difficulties, b) our evaluation does not induce an “AI overhang” by not having a smooth transition between difficulties and, c) the correlated factors affecting difficulty are human interpretable.

As shown in Figure 5 several task statistics such as “lines of code”, “processing time” and “number of function calls” all correlate log-linearly with a model’s achieved pass@1 score. These correlations provide preliminary evidence towards c) as they align with simplistic human intuition, i.e. more lines of code, more compute cycles, higher difficulty. Furthermore, we view the log-linear relationships as evidence towards b), i.e. EXE provides a smooth transition between difficulties. And finally, we view the relationships as a demonstration of difficulty being affected by factors within our control, i.e. number of function calls - providing empirical evidence towards a).

Beyond evaluation-wide difficulty scaling, EXE also demonstrates diversity and varying difficulty levels within individual task sets. Each function has up to 15 generated test cases, allowing us to analyse variance per task set. To measure execution path diversity, we collect runtime statistics (detailed in Appendix A.6) and find a mean Coefficient of Variation (CV) of 0.61 for “Count of conditionals executed”, indicating substantial variation in code paths taken. Furthermore we find a CV of 0.20 for “lines executed”, showing significant diversity in the number of steps required to answer. Finally, we measure diversity in generated task difficulty through model performance - GPT-4o achieves a mean pass rate of 0.742 ( $\sigma = 0.293$ ) per function, providing empirical evidence test cases present a difficulty scale.

**ExecEval’s test case generation scales.** While EXE today includes up to 15 test cases per task, our analysis demonstrates EXE’s generation pipeline can scale significantly further without plateauing. As shown in Figure 6, generation of novel test case continues well beyond 300 cases per task while maintaining all quality controls (detailed in Appendix A.2) - implying a potential dataset scale-up lower bound of 20x. Growth rates vary across specific functions - for example, langchain-core’s image formatting function, which requests a base64 encoded image string, shows the lowest growth

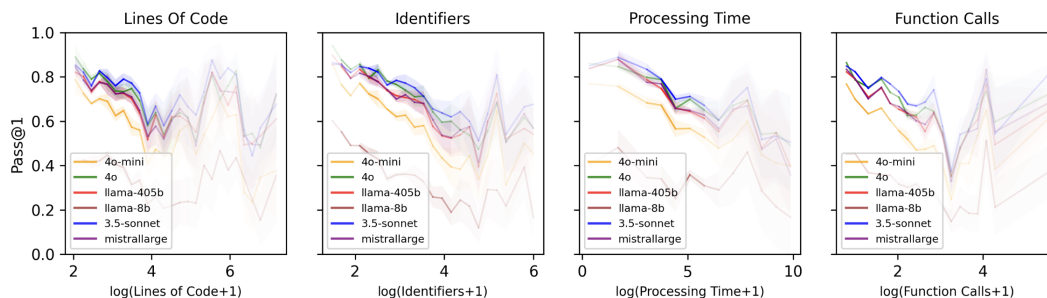


Figure 5: Pass@1 for all tasks across four of our code metrics. The shaded area represents variance, and the opacity is scaled with count of samples. Processing time is measured in microseconds.

rate. This aligns with intuition - generating novel, base64 images poses significantly more difficulty than generating diverse string or numeric inputs.

Importantly, our token efficiency analysis (right plot) reveals that significant scaling is possible without proportional prompt growth. By randomly selecting and injecting just 60 prior cases into the generation prompt, we can effectively generate over 1,000 novel cases. This sublinear token growth suggests the potential for substantial dataset expansion without incurring prohibitive costs. Detailed examples of tasks and their generated test cases are provided in Appendix A.8.

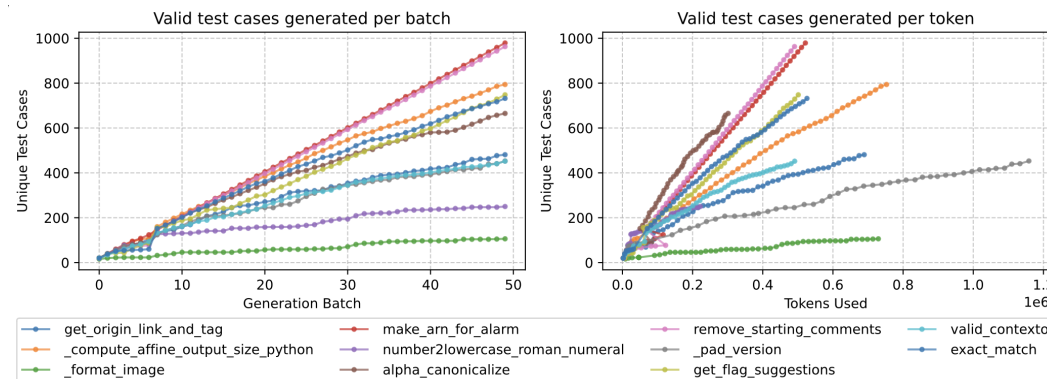


Figure 6: Test case generation analysis across eleven diverse Python functions sourced from popular libraries including Azure, PyTorch, Langchain, and NLTK. Functions range from geometric computations (torchvision) to SQL regex (snowflake-python-connector). Left: Cumulative unique validated test cases per generation batch. Right: Same data plotted against token usage, showing generation cost is largely constant per batch (primary factor is initial task code length). Further methodology and source code for tested functions are provided in Appendix A.8.

**Stylistic coding patterns shape the metrics.** As can be seen in Figure 5 the pass@1 rate of function calls hits an elbow and then surprisingly improves as the call count increases. During our investigation we found several of these occurrences, and not only with call count. These were found to be largely driven by specific coding patterns and complex tasks that LLMs excel at. We show in Figure 7 below three example tasks, and more specifically coding patterns driving this anomaly.

**LLMs struggle with certain coding features.** As EXE contains a diverse set of tasks, we are able to observe model performance differing greatly based on coding features used in any task. To illustrate: floating point math operations such as multiplications (GPT-4o: 43 mean Pass@1) significantly increase task difficulty, however bit manipulation and boolean operations only showed a minor negative impact. Iterative operations such as compound assignment operations i.e. "i += 1" (56 Pass@1), list slicing (65 Pass@1) and list comprehensions (68 Pass@1) all increased difficulty, however for loops on (73 Pass@1) on average did not have a significant impact.

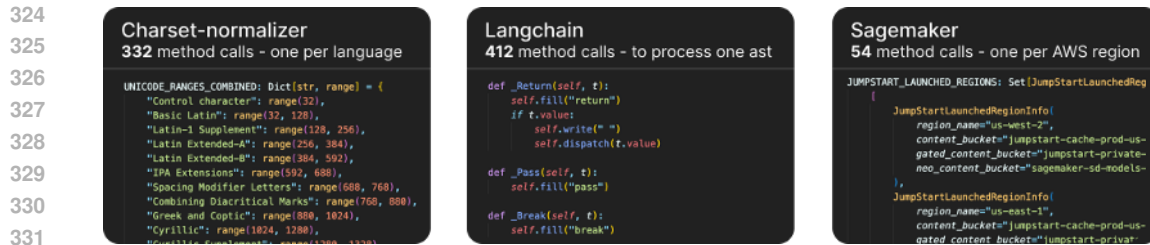


Figure 7: Three examples of high pass@1 rate tasks that contain large amounts of function calls. Left - Charset-normaliser performs 300+ function calls to define ranges of unicode characters upon initialisation; this constant has little effect on task difficulty but is used frequently and hence appears in many tasks. Middle - Langchain’s Unparser class traverses an AST and regenerates source code. The calling method in our dataset is ”add\_last\_line\_print(str) → str” which takes in code, parses it and then uses Unparse(...) to unparse it; this is a prime example of a ”directly predictable task”, i.e. one not requiring line by line code execution to predict a result. Right - Similar to Charset-normaliser, AWS’s Sagemaker has a module level constant with 10s of calls; not creating a large impact on task difficulty but frequent in its use.

With the above metrics, and those seen in Figure 7, their mean Pass@k decreases as their count increases. To reduce the risk of our metrics being a proxy for longer problems we show the effects can still be seen below in Figure 8 after normalisation by lines of code (only lines with executable syntax tokens are counted).




Figure 8: Pass@1 for all tasks across four of our code metrics normalised by line of code count (limited to GPT models for readability). All four of the above metrics previously showed a negative impact as they increased, interestingly we now observe branching statements having little to no impact and return statements surprisingly driving an increase in Pass@1 score. Our strong negative factors i.e. function calls and identifiers created, still are seen increasing task difficulty as they take up ever greater percentages of the task.

## 4 RELATED WORK

There is a rich history of work on evaluating language models’ abilities in reasoning, execution, and multi-step problem-solving across various domains. These efforts span from natural language processing to mathematical reasoning, and from code generation to program execution. Our work, EXExecution-Eval (EXE), builds upon this foundation while addressing key challenges in benchmark design and evaluation.

Code generation benchmarks have been the foundation of evaluating the coding abilities of language models. Works like HumanEval (Chen et al., 2021) and MBPP (Austin et al., 2021) established standardised datasets for assessing code synthesis from natural language descriptions. These efforts have expanded to cover multiple programming languages (Cassano et al., 2022; Khan et al., 2023) and more complex domains such as algorithmic problem solving (Huang et al., 2023). While these benchmarks focus primarily on the task of code generation, we believe additional focus on the tasks of code execution and error prediction have been overlooked and may offer additional insight into the internal capabilities of frontier models.

378 The concept of "learning to execute" itself has a long history, Zaremba & Sutskever (2014) explored  
379 neural networks' ability to learn and execute simple programs. Graves et al. (2014) constructed the  
380 first Neural Turing Machines with (Kaiser & Sutskever, 2015; Reed & de Freitas, 2015; Dehghani  
381 et al., 2018) all building further into this domain. This line of research has evolved, with recent  
382 works like Bieber et al. (2020); Nye et al. (2021) and Gu et al. (2024) applying graph and language  
383 models to execute synthetic or simplistic Python programs. EXE builds upon these foundations by  
384 evaluating execution capabilities on complex, messy, real-world code from diverse GitHub reposi-  
385 tories, providing a more challenging, scaleable and realistic test bed.

386 Recent trends in benchmark design have emphasised the importance of diverse, multi-step problems  
387 and agentic capabilities. Works like Jimenez et al. (2023) have introduced benchmarks that require  
388 solving real world software engineering problems while Zhou et al. (2023) has enabled evaluation  
389 of complex instruction following and performing multi-step reasoning. In the mathematical domain,  
390 benchmarks like those by Hendrycks et al. (2021) and Lu et al. (2023) have pushed models to solve  
391 intricate, multi-step problems.

392 The challenge of benchmark saturation and the need for continually updated evaluations has been  
393 recognized in recent works (Ott et al., 2022). Live benchmarks such as those proposed by Li et al.  
394 (2024), (Chiang et al., 2024) and Kiela et al. (2021) aim to address this issue. Skill-Mix (Yu et al.,  
395 2023) takes a novel approach, combining separate skills required to solve a problem they are able  
396 to increase task difficulty non-linearly with  $k$  skills. EXE has been inspired by both these concepts,  
397 hence the focus on enabling continual generation of new coding tasks and test cases, as well as the  
398 potential extension into chaining functions.

399 While many existing benchmarks use curated or synthetic datasets, EXE leverages real-world code  
400 from popular Python repositories. This approach is inspired by works like CodeNet (Puri et al.,  
401 2021) and The Stack (Kocetkov et al., 2022) which demonstrated the value of diverse, real-world  
402 data in training and evaluating language models.

## 403 404 405 5 EXTENSIONS

406  
407 **Expanding the scope and diversity** We believe scaling EXE to include more repositories by as  
408 much as 100x would significantly reduce the noise seen in our coding metrics and provide a more  
409 resilient baseline for future frontier models. By incorporating additional Python functions — po-  
410 tentially using language models to predict missing type annotations — and including a diversity  
411 of other programming languages such as C++, Go and JavaScript, we believe there is even further  
412 opportunity to scale. This would offer further insights into the generalisability of a model's code  
413 understanding, pose new challenges for analysis such as pointers, macros and type-free codebases.

414 **Probing code execution mechanisms with simple functions** We believe there is an opportunity  
415 to align code execution with mechanistic interpretability, to gain an understanding of how constant  
416 compute language models can execute complex multi-step instructions. To illustrate, if we select  
417 the simplest function that a language model can not directly predict the outcome of, a hash function  
418 for example (one that doesn't use floating point math in this case), one requiring compute at each  
419 iteration. This would force the network to perform the computation step by step, and for a constant  
420 compute feed-forward network, layer by layer. Hence, performing a single iteration that may not  
421 lead to anything interesting, however as we increase the iteration count one by one, the model now  
422 must find a repeated circuit to perform the same computation in the later layers. For every increase it  
423 must find another circuit or a more optimal way of performing its work until it fails. We believe this  
424 would present an interesting approach alongside standard mechanistic interpretability techniques for  
425 circuit discovery and understanding of control flow, variable tracking and computational logic at the  
mechanistic level.

426 **Breakpoint analysis for validating code execution granularly** Rather than evaluating the final  
427 return value, including multiple evaluation points within code execution may assist verification of if  
428 models are performing the step-by-step computations to reach a return value. Furthermore by insert-  
429 ing 'breakpoints' throughout the execution process, we can transform a single return state prediction  
430 task into numerous intermediate state prediction tasks. To illustrate, given a code snippet with a  
431 breakpoint at a specific line, a model would be tasked to determine the values of the local variables  
when the breakpoint is triggered. This mirrors common human debugging practices and may reveal



432 discrepancies between final output accuracy and intermediate state understanding, offering further  
 433 resistance against tasks where their final outcome can be directly predicted.

434 **Connection to cybersecurity threat model.** Software vulnerability research techniques are largely  
 435 <sup>1</sup> enabled by the ability to predict and reason about expected program outcomes. For example,  
 436 code injection, path resolution and memory buffer attacks are often found through manual human  
 437 analysis; tracing inputs through the control flow, predicting output states and reasoning if there  
 438 are opportunities to exploit. As EXE contains parsers such as seen in Appendix A.1 we see an  
 439 opportunity to select a subset of EXE where prediction of error would imply language models have  
 440 the internal capability to comprehend and aid humans with crafting vulnerabilities.

## 442 6 CONCLUSIONS

443  
 444 In this paper, we introduced EXecution-Eval (EXE), a benchmark designed to evaluate whether lan-  
 445 guage models can execute real-world code. By collecting over 30,000 tasks from 1,000 popular  
 446 Python repositories, EXE presents a diverse range of problems requiring computational operations  
 447 such as mathematical reasoning, logical inference, and state maintenance. Our evaluations suggest  
 448 that while language models demonstrate some capability in executing code, they often struggle with  
 449 complex, multi-step tasks—particularly those involving many identifiers, function calls and iterative  
 450 operations. Our findings indicate that although current models have limitations in accurately rea-  
 451 soning about and executing real-world code, they perform surprisingly well on average, prompting  
 452 several opportunities extending this investigation.

453 EXE aims to address limitations of existing benchmarks by providing a scalable, diverse, and con-  
 454 tinually updatable framework. Its design targets a smooth difficulty scale and easy generation of  
 455 new tasks with minimal human oversight with the goal to reduce the risk of models "training on the  
 456 test set."

## 458 REFERENCES

- 459  
 460 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,  
 461 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large  
 462 language models. August 2021.
- 463 David Bieber, Charles Sutton, Hugo Larochelle, and Daniel Tarlow. Learning to execute programs  
 464 with instruction pointer attention graph neural networks. October 2020.
- 465 Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald  
 466 Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, Arjun Guha,  
 467 Michael Greenberg, and Abhinav Jangda. MultiPL-E: A scalable and extensible approach to  
 468 benchmarking neural code generation. August 2022.
- 469 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared  
 470 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri,  
 471 Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan,  
 472 Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian,  
 473 Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fo-  
 474 tios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex  
 475 Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saun-  
 476 ders, Christopher Hesse, Andrew N Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa,  
 477 Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob  
 478 McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating  
 479 large language models trained on code. July 2021.
- 480 Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li,  
 481 Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E Gonzalez, and Ion Stoica.  
 482 Chatbot arena: An open platform for evaluating LLMs by human preference. March 2024.

484 <sup>1</sup>Some techniques such as random fuzzing may not rely on any internal program knowledge. However,  
 485 to find actionable results within realistic computational bounds, fuzzers are often augmented based on this  
 knowledge to limit their generatable space.

- 486 Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal  
487 transformers. July 2018.  
488
- 489 Max Fowler, David IV, Mohammed Hassan, Seth Poulsen, Matthew West, and Craig Zilles. Reeval-  
490 uating the relationship between explaining, tracing, and writing skills in cs1 in a replication study.  
491 *Computer Science Education*, 32:1–29, 06 2022. doi: 10.1080/08993408.2022.2079866.  
492
- 493 Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. October 2014.
- 494 Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I  
495 Wang. CRUXEval: A benchmark for code reasoning, understanding and execution. January 2024.  
496
- 497 Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song,  
498 and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. March  
499 2021.
- 500 Yiming Huang, Zhenghao Lin, Xiao Liu, Yeyun Gong, Shuai Lu, Fangyu Lei, Yaobo Liang, Yelong  
501 Shen, Chen Lin, Nan Duan, and Weizhu Chen. Competition-Level problems are effective LLM  
502 evaluators. December 2023.  
503
- 504 Carlos E Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik  
505 Narasimhan. SWE-bench: Can language models resolve real-world GitHub issues? October  
506 2023.
- 507 Łukasz Kaiser and Ilya Sutskever. Neural GPUs learn algorithms. November 2015.  
508
- 509 Mohammad Abdullah Matin Khan, M Saiful Bari, Xuan Long Do, Weishi Wang, Md Rizwan  
510 Parvez, and Shafiq Joty. XCodeEval: A large scale multilingual multitask benchmark for code  
511 understanding, generation, translation and retrieval. March 2023.  
512
- 513 Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie  
514 Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian  
515 Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina  
516 Williams. Dynabench: Rethinking benchmarking in NLP. April 2021.
- 517 Denis Kocetkov, Raymond Li, Loubna Ben Allal, Jia Li, Chenghao Mou, Carlos Muñoz Ferrandis,  
518 Yacine Jernite, Margaret Mitchell, Sean Hughes, Thomas Wolf, Dzmitry Bahdanau, Leandro von  
519 Werra, and Harm de Vries. The stack: 3 TB of permissively licensed source code. November  
520 2022.  
521
- 522 Tianle Li, Wei-Lin Chiang, Evan Frick, Lisa Dunlap, Tianhao Wu, Banghua Zhu, Joseph E Gon-  
523 zalez, and Ion Stoica. From crowdsourced data to high-quality benchmarks: Arena-Hard and  
524 BenchBuilder pipeline. June 2024.
- 525 Mike Lopez, Jacqueline Whalley, Phil Robbins, and Raymond Lister. Relationships between  
526 reading, tracing and writing skills in introductory programming. In *Proceedings of the Fourth  
527 International Workshop on Computing Education Research, ICER '08*, pp. 101–112, New  
528 York, NY, USA, 2008. Association for Computing Machinery. ISBN 9781605582160. doi:  
529 10.1145/1404520.1404531. URL <https://doi.org/10.1145/1404520.1404531>.  
530
- 531 Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-  
532 Wei Chang, Michel Galley, and Jianfeng Gao. MathVista: Evaluating mathematical reasoning of  
533 foundation models in visual contexts. October 2023.
- 534 Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David  
535 Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Au-  
536 gustus Odena. Show your work: Scratchpads for intermediate computation with language models.  
537 November 2021.  
538
- 539 Simon Ott, Adriano Barbosa-Silva, Kathrin Blagec, Jan Brauner, and Matthias Samwald. Mapping  
global dynamics of benchmark creation and saturation in artificial intelligence. March 2022.

Ruchir Puri, David S Kung, Geert Janssen, Wei Zhang, Giacomo Domeniconi, Vladimir Zolotov, Julian Dolby, Jie Chen, Mihir Choudhury, Lindsey Decker, Veronika Thost, Luca Buratti, Saurabh Pujar, Shyam Ramji, Ulrich Finkler, Susan Malaika, and Frederick Reiss. CodeNet: A large-scale AI for code dataset for learning a diversity of coding tasks. May 2021.

Scott Reed and Nando de Freitas. Neural Programmer-Interpreters. November 2015.

Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W Kocurek, Ali Safaya, Ali Tazarv, Alice Xiang, Alicia Parrish, Allen Nie, Aman Hussain, Amanda Askell, Amanda Dsouza, Ambrose Slone, Ameet Rahane, Anantharaman S Iyer, Anders Andreassen, Andrea Madotto, Andrea Santilli, Andreas Stuhlmüller, Andrew Dai, Andrew La, Andrew Lampinen, Andy Zou, Angela Jiang, Angelica Chen, Anh Vuong, Animesh Gupta, Anna Gottardi, Antonio Norelli, Anu Venkatesh, Arash Gholamidavoodi, Arfa Tabassum, Arul Menezes, Arun Kirubarajan, Asher Mullokandov, Ashish Sabharwal, Austin Herrick, Avia Efrat, Aykut Erdem, Ayla Karakaş, B Ryan Roberts, Bao Sheng Loe, Barret Zoph, Bartłomiej Bojanowski, Batuhan Özyurt, Behnam Hedayatnia, Behnam Neyshabur, Benjamin Inden, Benno Stein, Berk Ekmekci, Bill Yuchen Lin, Blake Howald, Bryan Orinion, Cameron Diao, Cameron Dour, Catherine Stinson, Cedrick Argueta, and Ramírez. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. June 2022.

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. May 2019.

Dingli Yu, Simran Kaur, Arushi Gupta, Jonah Brown-Cohen, Anirudh Goyal, and Sanjeev Arora. Skill-Mix: a flexible and expandable family of evaluations for AI models. October 2023.

Wojciech Zaremba and Ilya Sutskever. Learning to execute. October 2014.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P Xing, Hao Zhang, Joseph E Gonzalez, and Ion Stoica. Judging LLM-as-a-judge with MT-bench and chatbot arena. June 2023.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models. November 2023.

## A APPENDIX

You may include other additional sections here.

### A.1 EXAMPLE INPUT & OUTPUT

Below is an example from the evaluation set. It is split into three components:

**1. Code Task.** The function `split_email` was found to pass the type requirements, and as such all modules, classes, functions and attributes required to execute it have been recursively inlined.

**2. Test Case Inputs.** Based on the type definition (used for setting the function calling schema) inputs/ output pairs have been generated with the goal of maximising diversity of control flow paths within the function.

**3. Outputs.** Based on the type definition (used for setting the function calling schema) inputs/ output pairs have been generated with the goal of maximising diversity of control flow paths within the function.

### Examples

## 594 A.1.1 EXAMPLE A.

595 **Code**

596 Note: The top 1,000 PyPI repos are used to form EXE, this function is from celery, rank 594

```

599 def abbr(S: str, max: int, ellipsis: str | bool = '...') -> str:
600     """Abbreviate word."""
601     if S is None:
602         return '???'
603     if len(S) > max:
604         return isinstance(ellipsis, str) and (
605             S[: max - len(ellipsis)] + ellipsis) or S[: max]
606     return S
607
608 def abbrtask(S: str, max: int) -> str:
609     """Abbreviate task name."""
610     if S is None:
611         return '???'
612     if len(S) > max:
613         module, _, cls = S.rpartition('.')
614         module = abbr(module, max - len(cls) - 3, False)
615         return module + '[]' + cls
616     return S

```

616 **Test Case Inputs**617 Note: For quick groking, only three inputs are shown for this example. Standard tasks contain 15  
618 generated inputs.

```

619 [
620     {
621         "input": [{"module.ClassName", 15}, {}],
622         "output": "mod[]ClassName",
623     },
624     {
625         "input": [{"long.module.name.with.many.parts.ClassName", 25}, {}],
626         "output": "long.module.n[]ClassName",
627     },
628     {
629         "input": [{"module.ClassName", 3}, {}],
630         "output": "[]ClassName",
631     }
632 ]

```

633 **Generated Outputs**

```

634 [
635     {
636         "input": [{"module.ClassName", 15}, {}],
637         "output": "mod[]ClassName",
638         "prediction": "module[]ClassName",
639         "result": false,
640         "answer_tokens": {"completion": 15, "prompt": 781, "total": 796}
641     },
642     {
643         "input": [{"long.module.name.with.many.parts.ClassName", 25}, {}],
644         "output": "long.module.n[]ClassName",
645         "prediction": "long.module.name[]ClassName",
646         "result": false,
647         "answer_tokens": {"completion": 17, "prompt": 787, "total": 804}

```

```

648 {
649     "input": [{"module.ClassName", 3}, {}],
650     "output": "[.]ClassName",
651     "prediction": "[.]ClassName",
652     "result": true,
653     "answer_tokens": {"completion": 14, "prompt": 781, "total": 795}
654 },
655 ]

```

### A.1.2 EXAMPLE B.

#### Code

This function is from email-validator, rank 345.

```

663 from typing import Optional, Tuple
664 import re
665 import unicodedata
666
667 class EmailNotValidError(ValueError):
668     """Parent class of all exceptions raised by this module."""
669     pass
670
671 class EmailSyntaxError(EmailNotValidError):
672     """Exception raised when an email address fails validation because
673     ↪ of its form."""
674     pass
675
676 ATEXT = r'a-zA-Z0-9_!#$%&'\*\+\-/\=\\^`{|}\~'
677
678
679 def safe_character_display(c: str) -> str:
680     # Return safely displayable characters in quotes.
681     if c == '\\':
682         return f"\"{c}\"" # can't use repr because it escapes it
683     if unicodedata.category(c)[0] in ("L", "N", "P", "S"):
684         return repr(c)
685
686     # Construct a hex string in case the unicode name doesn't exist.
687     if ord(c) < 0xFFFF:
688         h = f"U+{ord(c):04x}".upper()
689     else:
690         h = f"U+{ord(c):08x}".upper()
691
692     # Return the character name or, if it has no name, the hex string.
693     return unicodedata.name(c, h)
694
695 ATEXT_RE = re.compile('[.]' + ATEXT + '.') # ATEXT plus dots
696
697 def check_unsafe_chars(s: str, allow_space: bool = False) -> None:
698     # Check for unsafe characters or characters that would make the
699     ↪ string
700     # invalid or non-sensible Unicode.
701     bad_chars = set()
702     for i, c in enumerate(s):

```

```

702     category = unicodedata.category(c)
703     if category[0] in ("L", "N", "P", "S"):
704         # Letters, numbers, punctuation, and symbols are permitted.
705         pass
706     elif category[0] == "M":
707         # Combining character in first position would combine with
708         # ↪ something
709         # outside of the email address if concatenated, so they are
710         # ↪ not safe.
711         # We also check if this occurs after the @-sign, which would
712         # ↪ not be
713         # sensible because it would modify the @-sign.
714         if i == 0:
715             bad_chars.add(c)
716     elif category == "Zs":
717         # Spaces outside of the ASCII range are not specifically
718         # ↪ disallowed in
719         # internationalized addresses as far as I can tell, but they
720         # ↪ violate
721         # the spirit of the non-internationalized specification that
722         # ↪ email
723         # addresses do not contain ASCII spaces when not quoted.
724         # ↪ Excluding
725         # ASCII spaces when not quoted is handled directly by the
726         # ↪ atom regex.
727         #
728         # In quoted-string local parts, spaces are explicitly
729         # ↪ permitted, and
730         # the ASCII space has category Zs, so we must allow it here,
731         # ↪ and we'll
732         # allow all Unicode spaces to be consistent.
733         if not allow_space:
734             bad_chars.add(c)
735     elif category[0] == "Z":
736         # The two line and paragraph separator characters (in
737         # ↪ categories Zl and Zp)
738         # are not specifically disallowed in internationalized
739         # ↪ addresses
740         # as far as I can tell, but they violate the spirit of the
741         # ↪ non-internationalized
742         # specification that email addresses do not contain line
743         # ↪ breaks when not quoted.
744         bad_chars.add(c)
745     elif category[0] == "C":
746         # Control, format, surrogate, private use, and unassigned
747         # ↪ code points (C)
748         # are all unsafe in various ways. Control and format
749         # ↪ characters can affect
750         # text rendering if the email address is concatenated with
751         # ↪ other text.
752         # Bidirectional format characters are unsafe, even if used
753         # ↪ properly, because
754         # they cause an email address to render as a different email
755         # ↪ address.
756         # Private use characters do not make sense for publicly
757         # ↪ deliverable
758         # email addresses.
759         bad_chars.add(c)
760     else:
761         # All categories should be handled above, but in case there
762         # ↪ is something new
763         # to the Unicode specification in the future, reject all
764         # ↪ other categories.
765         bad_chars.add(c)

```

```

756
757 if bad_chars:
758     raise EmailSyntaxError("The email address contains unsafe
759     ↪ characters: "
760     ↪         + ", ".join(safe_character_display(c) for
761     ↪         ↪ c in sorted(bad_chars)) + ".")
762
763 def split_email(email: str) -> Tuple[Optional[str], str, str, bool]:
764     # Return the display name, unescaped local part, and domain part
765     # of the address, and whether the local part was quoted. If no
766     # display name was present and angle brackets do not surround
767     # the address, display name will be None; otherwise, it will be
768     # set to the display name or the empty string if there were
769     # angle brackets but no display name.
770
771     # Typical email addresses have a single @-sign and no quote
772     # characters, but the awkward "quoted string" local part form
773     # (RFC 5321 4.1.2) allows @-signs and escaped quotes to appear
774     # in the local part if the local part is quoted.
775
776     # A `display name <addr>` format is also present in MIME messages
777     # (RFC 5322 3.4) and this format is also often recognized in
778     # mail UIs. It's not allowed in SMTP commands or in typical web
779     # login forms, but parsing it has been requested, so it's done
780     # here as a convenience. It's implemented in the spirit but not
781     # the letter of RFC 5322 3.4 because MIME messages allow newlines
782     # and comments as a part of the CFWS rule, but this is typically
783     ↪ not
784     # allowed in mail UIs (although comment syntax was requested once
785     ↪ too).
786     #
787     # Display names are either basic characters (the same basic
788     ↪ characters
789     # permitted in email addresses, but periods are not allowed and
790     ↪ spaces
791     # are allowed; see RFC 5322 Appendix A.1.2), or or a quoted string
792     ↪ with
793     # the same rules as a quoted local part. (Multiple quoted strings
794     ↪ might
795     # be allowed? Unclear.) Optional space (RFC 5322 3.4 CFWS) and
796     ↪ then the
797     # email address follows in angle brackets.
798     #
799     # We assume the input string is already stripped of leading and
800     ↪ trailing CFWS.
801
802 def split_string_at_unquoted_special(text: str, specials:
803     ↪ Tuple[str, ...]) -> Tuple[str, str]:
804     # Split the string at the first character in specials (an
805     ↪ @-sign
806     # or left angle bracket) that does not occur within quotes and
807     # is not followed by a Unicode combining character.
808     # If no special character is found, raise an error.
809     inside_quote, escaped, left_part = False, False, ""
810     for i, c in enumerate(text):
811         # < plus U+0338 (Combining Long Solidus Overlay) normalizes
812         ↪ to
813         # U+226E (Not Less-Than), and it would be confusing to
814         ↪ treat
815         # the < as the start of "<email>" syntax in that case.
816         ↪ Likewise,

```

```

810
811     # if anything combines with an @ or ", we should probably
812     ↪ not
813     # treat it as a special character.
814     if unicodedata.normalize("NFC", text[i:])[0] != c:
815         left_part += c
816
817     elif inside_quote:
818         left_part += c
819         if c == '\\': and not escaped:
820             escaped = True
821         elif c == '"': and not escaped:
822             # The only way to exit the quote is an unescaped
823             ↪ quote.
824             inside_quote = False
825             escaped = False
826         else:
827             escaped = False
828     elif c == "'":
829         left_part += c
830         inside_quote = True
831     elif c in specials:
832         # When unquoted, stop before a special character.
833         break
834     else:
835         left_part += c
836
837
838 if len(left_part) == len(text):
839     raise EmailSyntaxError("An email address must have an
840     ↪ @-sign.")
841
842 right_part = text[len(left_part):] # The right part is whatever
843     ↪ is left.
844
845 return left_part, right_part
846
847 def unquote_quoted_string(text: str) -> Tuple[str, bool]:
848     # Remove surrounding quotes and unescape escaped backslashes
849     # and quotes. Escapes are parsed liberally. I think only
850     ↪ backslashes
851     # and quotes can be escaped but we'll allow anything to be.
852     quoted, escaped, value = False, False, ""
853     for i, c in enumerate(text):
854         if quoted:
855             if escaped:
856                 value += c
857                 escaped = False
858             elif c == '\\':
859                 escaped = True
860             elif c == '"':
861                 if i != len(text) - 1:
862                     raise EmailSyntaxError("Extra character(s) found
863                     ↪ after close quote: "
864                                             + ",
865                                             ↪ ".join(safe_character_display(c)
866                                             ↪ for c in text[i + 1:]))
867                 break
868             else:
869                 value += c
870         elif i == 0 and c == "'":

```



```

864         quoted = True
865     else:
866         value += c
867
868
869     return value, quoted
870
871     # Split the string at the first unquoted @-sign or left angle
872     ↪ bracket.
873     left_part, right_part = split_string_at_unquoted_special(email,
874     ↪ ("@", "<"))
875
876     # If the right part starts with an angle bracket, then the left
877     ↪ part
878     # is a display name and the rest of the right part up to the
879     # final right angle bracket is the email address, .
880     if right_part.startswith("<"):
881         # Remove space between the display name and angle bracket.
882         left_part = left_part.rstrip()
883
884     # Unquote and unescape the display name.
885     display_name, display_name_quoted =
886     ↪ unquote_quoted_string(left_part)
887
888     # Check that only basic characters are present in a non-quoted
889     ↪ display name.
890     if not display_name_quoted:
891         bad_chars = {
892             safe_character_display(c)
893             for c in display_name
894             if (not ATEXT_RE.match(c) and c != ' ') or c == '.'
895         }
896         if bad_chars:
897             raise EmailSyntaxError("The display name contains
898             ↪ invalid characters when not quoted: " + ",
899             ↪ ".join(sorted(bad_chars)) + ".")
900
901     check_unsafe_chars(display_name, allow_space=True) # Check for
902     ↪ other unsafe characters.
903
904     # Check that the right part ends with an angle bracket
905     # but allow spaces after it, I guess.
906     if ">" not in right_part:
907         raise EmailSyntaxError("An open angle bracket at the start
908         ↪ of the email address has to be followed by a close angle
909         ↪ bracket at the end.")
910     right_part = right_part.rstrip(" ")
911     if right_part[-1] != ">":
912         raise EmailSyntaxError("There can't be anything after the
913         ↪ email address.")
914
915     # Remove the initial and trailing angle brackets.
916     addr_spec = right_part[1:].rstrip(">")
917
918     # Split the email address at the first unquoted @-sign.
919     local_part, domain_part =
920     ↪ split_string_at_unquoted_special(addr_spec, ("@",))

```

```

918
919
920     # Otherwise there is no display name. The left part is the local
921     # part and the right part is the domain.
922     else:
923         display_name = None
924         local_part, domain_part = left_part, right_part
925
926     if domain_part.startswith("@"):
927         domain_part = domain_part[1:]
928
929     # Unquote the local part if it is quoted.
930     local_part, is_quoted_local_part =
931     ↪ unquote_quoted_string(local_part)
932
933     return display_name, local_part, domain_part, is_quoted_local_part
934
935

```

### Test Case Inputs

```

938 [
939     {
940         "input": [{"simple@example.com"}, {}],
941         "output": [null, "simple", "example.com", false],
942     },
943     {
944         "input": [{"user+name@sub.domain.com"}, {}],
945         "output": [null, "user+name", "sub.domain.com", false],
946     },
947     {
948         "input": [{"user.name@domain.co.uk"}, {}],
949         "output": [null, "user.name", "domain.co.uk", false],
950     },
951     {
952         "input": [{"\"quoted@local\"@example.com"}, {}],
953         "output": [null, "quoted@local", "example.com", true],
954     },
955     {
956         "input": [{"display name <user@domain.com>"}, {}],
957         "output": ["display name", "user", "domain.com", false],
958     },
959     {
960         "input": [{"user@localhost"}, {}],
961         "output": [null, "user", "localhost", false],
962     },
963     {
964         "input": [{"user@[IPv6:2001:db8::1]"}, {}],
965         "output": [null, "user", "[IPv6:2001:db8::1]", false],
966     },
967     {
968         "input": [{"\"escaped\\\"quote\"@example.com"}, {}],
969         "output": [null, "escaped\"quote", "example.com", true],
970     },
971     {
972         "input": [{"user.name@longsubdomain.example.com"}, {}],
973         "output": [null, "user.name", "longsubdomain.example.com", false],
974     },
975     {
976         "input": [{"very.common@example.com"}, {}],
977         "output": [null, "very.common", "example.com", false],
978     },
979 ]

```

```

972 {
973   "input": ["user@domain-with-dash.com"], {},
974   "output": [null, "user", "domain-with-dash.com", false],
975 },
976 {
977   "input": ["user@123.123.123.123"], {},
978   "output": [null, "user", "123.123.123.123", false],
979 },
980 {
981   "input": ["\"much.more unusual\"@example.com"], {},
982   "output": [null, "much.more unusual", "example.com", true],
983 },
984 {
985   "input": ["user@xn--exmple-cua.com"], {},
986   "output": [null, "user", "xn--exmple-cua.com", false],
987 },
988 {
989   "input": ["user@domain_with_underscore.com"], {},
990   "output": [null, "user", "domain_with_underscore.com", false],
991 }
992 ]

```

## Generated Outputs

```

994 [
995   {
996     "input": ["simple@example.com"], {},
997     "output": [null, "simple", "example.com", false],
998     "prediction": [null, "simple", "example.com", false],
999     "result": true,
1000     "answer_tokens": {"completion": 18, "prompt": 4610, "total": 4628}
1001   },
1002   {
1003     "input": ["user+name@sub.domain.com"], {},
1004     "output": [null, "user+name", "sub.domain.com", false],
1005     "prediction": [null, "user+name", "sub.domain.com", false],
1006     "result": true,
1007     "answer_tokens": {"completion": 21, "prompt": 4614, "total": 4635}
1008   },
1009   {
1010     "input": ["user.name@domain.co.uk"], {},
1011     "output": [null, "user.name", "domain.co.uk", false],
1012     "prediction": [null, "user.name", "domain.co.uk", false],
1013     "result": true,
1014     "answer_tokens": {"completion": 20, "prompt": 4613, "total": 4633}
1015   },
1016   {
1017     "input": ["\"quoted@local\"@example.com"], {},
1018     "output": [null, "quoted@local", "example.com", true],
1019     "prediction": [null, "quoted@local", "example.com", true],
1020     "result": false,
1021     "answer_tokens": {"completion": 20, "prompt": 4615, "total": 4635}
1022   },
1023   {
1024     "input": ["display name <user@domain.com>"], {},
1025     "output": [null, "display name", "user", "domain.com", false],
1026     "prediction": [null, "display name", "user", "domain.com", false],
1027     "result": true,
1028     "answer_tokens": {"completion": 19, "prompt": 4615, "total": 4634}
1029   },
1030   {
1031     "input": ["user@localhost"], {},
1032     "output": [null, "user", "localhost", false],

```

```

1026     "prediction": [null, "user", "localhost", false],
1027     "result": true,
1028     "answer_tokens": {"completion": 17, "prompt": 4610, "total": 4627}
1029   },
1030   {
1031     "input": [{"user@[IPv6:2001:db8::1]"}, {}],
1032     "output": [null, "user", "[IPv6:2001:db8::1]", false],
1033     "prediction": "EmailSyntaxError: An email address must have an
1034     ↪ @-sign.",
1035     "result": false,
1036     "answer_tokens": {"completion": 24, "prompt": 4620, "total": 4644}
1037   },
1038   {
1039     "input": [{"\"escaped\\\"quote\"@example.com"}, {}],
1040     "output": [null, "escaped\"quote", "example.com", true],
1041     "prediction": [null, "escaped\"quote", "example.com", true],
1042     "result": false,
1043     "answer_tokens": {"completion": 20, "prompt": 4615, "total": 4635}
1044   },
1045   {
1046     "input": [{"user.name@longsubdomain.example.com"}, {}],
1047     "output": [null, "user.name", "longsubdomain.example.com", false],
1048     "prediction":
1049     ↪ [null, "user.name", "longsubdomain.example.com", false],
1050     "result": true,
1051     "answer_tokens": {"completion": 22, "prompt": 4615, "total": 4637}
1052   },
1053   {
1054     "input": [{"very.common@example.com"}, {}],
1055     "output": [null, "very.common", "example.com", false],
1056     "prediction": [null, "very.common", "example.com", false],
1057     "result": true,
1058     "answer_tokens": {"completion": 19, "prompt": 4611, "total": 4630}
1059   },
1060   {
1061     "input": [{"user@domain-with-dash.com"}, {}],
1062     "output": [null, "user", "domain-with-dash.com", false],
1063     "prediction": [null, "user", "domain-with-dash.com", false],
1064     "result": true,
1065     "answer_tokens": {"completion": 21, "prompt": 4614, "total": 4635}
1066   },
1067   {
1068     "input": [{"user@123.123.123.123"}, {}],
1069     "output": [null, "user", "123.123.123.123", false],
1070     "prediction": [null, "user", "123.123.123.123", false],
1071     "result": true,
1072     "answer_tokens": {"completion": 23, "prompt": 4616, "total": 4639}
1073   },
1074   {
1075     "input": [{"\"much.more unusual\"@example.com"}, {}],
1076     "output": [null, "much.more unusual", "example.com", true],
1077     "prediction": [null, "much.more unusual", "example.com", true],
1078     "result": true,
1079     "answer_tokens": {"completion": 20, "prompt": 4615, "total": 4635}
1080   },
1081   {
1082     "input": [{"user@xn--exmple-cua.com"}, {}],
1083     "output": [null, "user", "xn--exmple-cua.com", false],
1084     "prediction": [null, "user", "xn--exmple-cua.com", false],
1085     "result": true,
1086     "answer_tokens": {"completion": 24, "prompt": 4617, "total": 4641}
1087   },
1088   }

```

```

1080     "input": [{"user@domain_with_underscore.com"}, {}],
1081     "output": [null, "user", "domain_with_underscore.com", false],
1082     "prediction": "EmailSyntaxError: The email address contains unsafe
1083     ↪ characters: 'U+005F'.",
1084     "result": false,
1085     "answer_tokens": {"completion": 28, "prompt": 4614, "total": 4642}
1086   }
1087 ]

```

## 1090 A.2 INPUT GENERATION

1092 Test case generation is performed through a three-stage pipeline: schema construction, test genera-  
 1093 tion, and validation.

### 1095 A.2.1 SCHEMA CONSTRUCTION

1097 Using our AST analysis’s findings (see Section A.5), we construct OpenAPI-compatible JSON  
 1098 schemas from identified argument and return types. Consider a type-annotated function from our  
 1099 dataset:

```

1101 from typing import Dict, List, Optional, Tuple, Union
1102
1103 def get_tree_starting_at(module: str, edges: List[Tuple[str, str]]) ->
1104 ↪ List[Union[str, List[str]]]:
1105     """
1106     Returns the tree starting at a given module following all edges.
1107
1108     Args:
1109     module (`str`): The module that will be the root of the subtree
1110     ↪ we want.
1111     eges (`List[Tuple[str, str]]`): The list of all edges of the
1112     ↪ tree.
1113
1114     Returns:
1115     `List[Union[str, List[str]]`: The tree to print in the
1116     ↪ following format: [module, [list of edges
1117     starting at module], [list of edges starting at the preceding
1118     ↪ level], ...]
1119     """
1120     vertices_seen = [module]
1121     new_edges = [edge for edge in edges if edge[0] == module and edge[1]
1122     ↪ != module and "__init__.py" not in edge[1]]
1123     tree = [module]
1124     while len(new_edges) > 0:
1125         tree.append(new_edges)
1126         final_vertices = list({edge[1] for edge in new_edges})
1127         vertices_seen.extend(final_vertices)
1128         new_edges = [
1129             edge
1130             for edge in edges
1131             if edge[0] in final_vertices and edge[1] not in
1132             ↪ vertices_seen and "__init__.py" not in edge[1]
1133         ]
1134     return tree

```

1131 This generates the following schema for language model function calling (note: the case below  
 1132 shows a json schema further wrapped in OpenAI’s specific "tool" schema):  
 1133

```

1134 {"tools": [{
1135   "type": "function",
1136   "function": {
1137     "name": "FunctionTestCaseModel",
1138     "description": "Correctly extracted `FunctionTestCaseModel` with
1139     ↪ all the required parameters with correct types",
1140     "parameters": {
1141       "$defs": {
1142         "ArgsModel": {
1143           "properties": {
1144             "module": {
1145               "description": "Positional argument 'module' with
1146               ↪ type '<class 'str'>'",
1147               "title": "Module",
1148               "type": "string"
1149             },
1150             "edges": {
1151               "description": "Positional argument 'edges' with
1152               ↪ type 'typing.List[typing.Tuple[str, str]]'",
1153               "items": {
1154                 "items": {"type": "string"},
1155                 "type": "array"
1156               },
1157               "title": "Edges",
1158               "type": "array"
1159             }
1160           },
1161           "required": ["module", "edges"],
1162           "title": "ArgsModel",
1163           "type": "object"
1164         },
1165         "KwargsModel": {
1166           "properties": {},
1167           "title": "KwargsModel",
1168           "type": "object"
1169         },
1170         "TestCase": {
1171           "properties": {
1172             "args": {
1173               "allOf": [{"$ref": "#/$defs/ArgsModel"}],
1174               "description": "Positional args."
1175             },
1176             "kwargs": {
1177               "allOf": [{"$ref": "#/$defs/KwargsModel"}],
1178               "description": "Keyword args."
1179             }
1180           },
1181           "required": ["args", "kwargs"],
1182           "title": "TestCase",
1183           "type": "object"
1184         }
1185       },
1186       "required": ["test_cases"],
1187       "type": "object"
1188     }
1189   }
1190 ]}

```

```

1188     }
1189   }}
1190 }}

```

This schema is then embedded within our test case generation prompt:

```

1194 You are an expert software tester tasked with generating diverse test
1195 ↪ cases for a given function. Your goal is to create a comprehensive
1196 ↪ set of tests that cover various scenarios and edge cases.
1197
1198 First, let's review the previously generated test cases to ensure we
1199 ↪ explore new scenarios:
1200 <previously_generated_test_cases>
1201 {seen or "No test cases have been generated yet."}
1202 </previously_generated_test_cases>
1203
1204 Now, let's examine the context and function details:
1205
1206 <module_code>
1207 {module_code}
1208 </module_code>
1209
1210 Now, let's look at the specific function we need to test:
1211
1212 <function_signature>
1213 {func.signature}
1214 </function_signature>
1215
1216 <function_docstring>
1217 {func.docstring}
1218 </function_docstring>
1219
1220 <function_implementation>
1221 {func.code}
1222 </function_implementation>
1223
1224 Before generating the test cases, let's think through the process:
1225
1226 <test_case_analysis>
1227 1. Analyze the function signature, docstring, and implementation to
1228 ↪ understand its purpose and expected behavior.
1229 2. Identify the input parameters and their types.
1230 3. Determine the function's return type and expected output format.
1231 4. Consider the following categories of test cases:
1232 a. Simple and straightforward cases
1233 b. Complex cases with multiple inputs
1234 c. Edge cases with large values or sizes
1235 d. Edge cases with small values or sizes
1236 e. Cases that may require significant processing time
1237 f. Cases that might cause errors or exceptions
1238 g. Cases with invalid inputs that should raise specific exceptions
1239 5. For numerical arguments:
1240 - Include positive and negative integers/floats
1241 - Include zero
1242 - Include prime numbers
1243 - Include maximum and minimum possible integer values
1244 - Include very large floats and very small floats (close to zero)
1245 6. For string arguments:
1246 - Include empty strings
1247 - Include strings with special characters
1248 - Include very long strings
1249 - Include strings in different languages or with Unicode characters
1250 7. For boolean arguments:

```

```

1242     - Include both True and False cases
1243 8. For dynamic containers (e.g., lists, dictionaries):
1244     - Include cases with many elements
1245     - Include cases with no elements
1246     - Include cases with deeply nested objects
1247     - Include cases with mixed data types
1248 9. For each test case, predict the expected output or exception.
1249 10. Ensure that each test case is unique and covers a different
1250     ↪ scenario.
1251 11. Consider any specific constraints or requirements mentioned in the
1252     ↪ docstring.
1253 </test_case_analysis>
1254
1255 Now, generate 15 diverse test cases based on this analysis. Present each
1256 ↪ test case as a Python dictionary with 'args' and 'kwargs' keys, even
1257 ↪ if one of them is empty. Do not include any additional text or
1258 ↪ formatting.

```

### 1259 A.2.2 TEST GENERATION AND EXECUTION

1260 After generation, each test case is executed against the original function in a controlled environment.  
 1261 We capture:

- 1262 • Return values or raised exceptions
- 1263 • Runtime statistics (see Section A.6)

### 1264 A.2.3 VALIDATION PIPELINE

1265 Generated test cases are tested against seven validators for quality control. Each validator, upon  
 1266 failure, appends specific feedback as part of a reply to the conversation with the language model:

- 1267 1. **Schema Conformance:** Test cases must parse as valid function inputs
- 1268 2. **Duplication:** Each test case input must be unique
- 1269 3. **Coverage:** Minimum 10 test cases per function
- 1270 4. **Non-triviality:** Less than 50% of cases can return unmodified input
- 1271 5. **Output Diversity:** No single output as 66% of cases
- 1272 6. **Error Balance:** Exception cases limited to 50% of total
- 1273 7. **Runtime Bounds:** CPU time under 10 seconds per case

1274 We provide examples of validation feedback messages in Section A.4.

### 1275 A.2.4 REGENERATION STRATEGY

1276 The system allows two full generation attempts, each permitting three validation/reply/regeneration  
 1277 cycles. To maximise task breadth while maintaining quality, we may still preserve some test cases  
 1278 from a task that fails to pass all validators. We do this by relaxing some validator requirements:

- 1279 • The minimum test case count requirement (criterion 3) is waived for the final generation  
 1280 attempt
- 1281 • Test cases that contain duplicates or exceed runtime bounds are individually filtered out  
 1282 (criteria 2 and 7)
- 1283 • The task's remaining test cases must still meet our core quality requirements: non-triviality,  
 1284 output diversity, and a balanced error rate (criteria 4, 5, and 6)

1285 This approach using GPT-4o-latest (generation spanned multiple versions) yields our current dataset  
 1286 of 33,875 test cases across 1,000 repositories, with an average generation cost of 1,123 tokens per



1296 test case. Failed generations primarily occur due to schema conformance (criterion 1 - schema con-  
 1297 formance poses an outsized challenge to smaller models i.e. llama3.1-8b; mirroring the execution  
 1298 prediction task), duplication and output diversity (criterion 2 and 5 - both commonly observed in  
 1299 functions with a limited input/output domains, i.e. single boolean args/ returns).

1300

### 1301 A.3 FUNCTION SELECTION AND DEPENDENCY RESOLUTION

1302

1303 The function selection and dependency collation process comprises three main stages: type annota-  
 1304 tion analysis, dependency graph construction, and code inlining, followed by a final filtering stage.  
 1305 Here we detail each stage:

1306

#### 1307 A.3.1 TYPE ANNOTATION ANALYSIS

1308

1309 Function selection begins with a recursive AST analysis of type annotations. Each candidate func-  
 1310 tion must have both its arguments and return type validated as "LLM-generatable" - meaning they  
 1311 can be reliably produced by a language model. As detailed in Section A.5, we recursively validate  
 1312 against a predefined set of acceptable types.

1313

1314 For example, when processing complex nested types like 'List[Tuple[str, int]]', the analyzer first  
 1315 validates 'List', then 'Tuple', and finally 'str' and 'int'. Functions with arguments or return types  
 1316 containing non-LLM-generatable elements (e.g., file handles, sockets, custom objects) are filtered  
 out during this stage.

1317

#### 1318 A.3.2 DEPENDENCY GRAPH CONSTRUCTION

1319

1320 Once a function passes type validation, we construct a dependency graph to identify all code ele-  
 1321 ments required for the function's execution. This process involves:

1322

1. **Symbol Analysis:** For each function, we perform an AST walk to identify:

1323

- Local variables: We track symbols defined within the current scope including but not limited to assignments, function arguments, loop variables, comprehension variables, and lambda parameters. These are excluded from dependency tracking as they are part of the function's internal logic.

1324

- Used symbols: We collect all variable references, function calls, type annotations (e.g., in 'x: List[Prompt]', both 'List' and 'Prompt' need resolving), and attribute accesses (e.g., in 'library.varname', both 'library' and 'varname' need resolving). By comparing against the local variables, we identify which symbols must be resolved externally. For each symbol, we walk the AST to find its definition.

1325

- Nested scopes: We handle nested functions and classes by treating their names as local variables in the outer scope while tracking their internal symbol usage separately.

1326

2. **Import Resolution:** For each identified external dependency, we:

1327

- Resolve relative imports based on the file's location in the package and module imports based on the package structure

1328

- Track aliases and renamed imports, mapping against accessed attributes (e.g. for 'lib.var' where we 'import x as lib', we must find 'var' in 'x')

1329

- Ignore builtin imports, treating them as standard code blocks

1330

- Recursively process imported modules, classes, functions and variables through Step 1. Symbol Analysis

1331

- Handle special cases such as '\_init\_.py' files, complex imports 'from x import \*' and more

1332

3. **Graph Construction:** We build a directed graph where nodes represent code blocks (functions, classes, assignments) and edges represent dependencies between these blocks. The graph maintains the minimal set of dependencies required for each function while preserving their original relationships.

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

1345

1346

1347

1348

1349

1350 **4. Symbol Resolution Validation:** Before a function is accepted, we verify that every used symbol  
 1351 has been successfully resolved to a definition. This validation is crucial as it ensures we can create  
 1352 a complete, self-contained version of the function. Functions using runtime code generation (e.g.,  
 1353 ‘exec’, ‘eval’), dynamic attribute access (e.g., ‘getattr’ with variable names), or other patterns that  
 1354 prevent static resolution are largely filtered out at this stage.

1355 To illustrate this process with a simple example, consider the following from the Azure SDK Ta-  
 1356 bles package. The original code was spread across two files in ‘azure-nspkg/sdk/tables/azure-data-  
 1357 tables/azure/data/tables/'. The extracted minimal dependency chain (debug output preserved to show  
 1358 file origins and dependency types) is shown below:  
 1359

```

1360 # _common_conversion.py | resolved_import_from/defaultlib -> from
1361   ↳ datetime import timezone
1362 from datetime import timezone
1363
1364 # _common_conversion.py | function -> _to_utc_datetime
1365 def _to_utc_datetime(value):
1366     try:
1367         value = value.astimezone(timezone.utc)
1368     except ValueError:
1369         # Before Python 3.8, this raised for a naive datetime.
1370         pass
1371     try:
1372         return value.strftime("%Y-%m-%dT%H:%M:%S.%fZ")
1373     except ValueError:
1374         return value.strftime("%Y-%m-%dT%H:%M:%SZ")
1375
1376 # _serialize.py | resolved_import_from/defaultlib -> from datetime
1377   ↳ import datetime
1378 from datetime import datetime
1379 # _serialize.py | resolved_import_from/defaultlib -> from uuid import
1380   ↳ UUID
1381 from uuid import UUID
1382 # _serialize.py | resolved_import_from/defaultlib -> from typing import
1383   ↳ Dict, Optional, Union
1384 from typing import Dict, Optional
1385 # _serialize.py | resolved_import_from/defaultlib -> from binascii
1386   ↳ import hexlify
1387 from binascii import hexlify
1388
1389 # _serialize.py | function -> _parameter_filter_substitution
1390 def _parameter_filter_substitution(parameters: Optional[Dict[str, str]],
1391   ↳ query_filter: str) -> str:
1392     """Replace user defined parameters in filter.
1393     :param parameters: User defined parameters
1394     :type parameters: dict[str, str]
1395     :param str query_filter: Filter for querying
1396     :return: A query filter replaced by user defined parameters.
1397     :rtype: str
1398     """
1399     if parameters:
1400         filter_strings = query_filter.split(" ")
1401         for index, word in enumerate(filter_strings):
1402             if word[0] == "@":
1403                 val = parameters[word[1:]]
1404                 if val in [True, False]:
1405                     filter_strings[index] = str(val).lower()
1406                 elif isinstance(val, (float)):
1407                     filter_strings[index] = str(val)
1408                 elif isinstance(val, int):
1409                     if val.bit_length() <= 32:
1410                         filter_strings[index] = str(val)
1411                 else:
1412                     filter_strings[index] = f"{str(val)}L"

```

```

1404
1405         elif isinstance(val, datetime):
1406             filter_strings[index] =
1407                 ↪ f"datetime'{_to_utc_datetime(val)}'"
1408         elif isinstance(val, UUID):
1409             filter_strings[index] = f"guid'{str(val)}'"
1410         elif isinstance(val, bytes):
1411             v = str(hexlify(val))
1412             v = v[2:-1] # Python 3 adds a 'b' and quotations
1413             filter_strings[index] = f"X'{v}'"
1414         else:
1415             val = val.replace('"', "'")
1416             filter_strings[index] = f"'{val}'"
1417     return " ".join(filter_strings)
1418     return query_filter

```

Note that these functions have been extracted from much larger source files (indicated by the commented file names) - we only collect the minimal code required for execution.

### A.3.3 CODE INLINING

The final stage involves generating a self-contained version of the function with all dependencies inlined. Rather than attempting to strip back the original files to their minimal form, we are motivated to inline as it ensures the language model executes exactly the same code as our interpreter.

The inlining process:

1. Performs a topological sort of the dependency graph to determine the correct order of declarations.
2. Inlines code based on its original structure:
  - Most code, including functions, classes, and variables, is inlined directly at the appropriate scope.
  - When an entire module has been imported (e.g., ‘import random’), we create a namespace class to maintain proper module-level scoping.
3. Generates the final code by maintaining the original code structure and ensuring all dependencies are declared before use.

After code inlining, we perform a final filtering pass to remove functions with side effects or non-deterministic behaviour. This filtering must occur after inlining as many problematic patterns only become apparent once we have the complete code context. For example, network requests might be hidden behind multiple layers of function calls, or environment variables might be accessed through utility functions in separate modules. Functions that use system calls, file I/O, network operations, random number generation, or environment variables are filtered out at this stage.

While our dependency resolution system handles many common Python patterns, including dynamic imports and aliased imports, there remain some challenges. Functions with circular dependencies between modules cannot currently be processed, and certain package initialization patterns that rely on import-time side effects are not supported. These limitations primarily affect a small percentage of candidate functions.

## A.4 VALIDATOR EXAMPLES

Each validator appends specific feedback to guide the model in correcting errors. Below are the prompts used for each of these feedback messages:

### A.4.1 SCHEMA CONFORMANCE VALIDATOR

```

1456 Validation Error found while parsing test case JSON:
1457 <exception>

```

```

1458 {exception}
1459 </exception>
1460 Recall the function correctly, fix these errors and generate a valid
1461 ↪ test case following the schema.
1462
1463
1464

```

#### 1465 A.4.2 DUPLICATION VALIDATOR

```

1467 Validation Error: Duplicate test case inputs detected
1468 The following test cases have identical inputs:
1469 <duplicate_cases>
1470 {json.dumps(duplicate_cases)}
1471 </duplicate_cases>
1472 Recall the function correctly and generate test cases with unique
1473 ↪ input combinations.
1474
1475

```

#### 1476 A.4.3 COVERAGE VALIDATOR

```

1478 Validation Error: Insufficient test coverage ({len(cases)}/10 required
1479 ↪ minimum cases).
1480 Generate additional unique test cases to cover these scenarios.
1481
1482

```

#### 1483 A.4.4 NON-TRIVIALITY VALIDATOR

```

1486 Validation Error: Test cases too simple. Greater than 50% of test cases
1487 ↪ are returning their inputs as outputs. Inputs must undergo some
1488 ↪ transformation during processing.
1489 <test_cases_with_results>
1490 {json.dumps(cases)}
1491 </test_cases_with_results>
1492 Fix these errors by generating test cases that:
1493 1. Explore different code paths within the function
1494 2. Trigger transformation of the inputs so that they differ from the
1495 ↪ outputs
1496

```

#### 1497 A.4.5 OUTPUT DIVERSITY VALIDATOR

```

1499 Validation Error: Insufficient output diversity in test cases. One
1500 ↪ output is returned by more than 2/3s of all cases.
1501 <test_cases_counted_outputs>
1502 {json.dumps(output_counter)}
1503 </test_cases_counted_outputs>
1504 <test_cases_with_results>
1505 {json.dumps(cases)}
1506 </test_cases_with_results>
1507 Generate additional test cases that contain differing outputs to the
1508 ↪ most popular above.
1509

```

#### 1510 A.4.6 ERROR BALANCE VALIDATOR

1511

```

1512
1513 Validation Error: Too many error-inducing test cases
1514 ↪ ({{len(error_cases)}}/{{len(cases)}})
1515 <test_cases_with_results>
1516 {json.dumps(cases)}
1517 </test_cases_with_results>

```

1518  
1519

## 1520 A.5 ACCEPTABLE TYPES & FILTERING CRITERIA

1521  
1522  
1523  
1524  
1525  
1526

**Acceptable types.** To find functions where the inputs and outputs are LLM generatable, we recursively parse both arguments and return types as AST objects i.e. for `list[tuple[str, False]]` we first check `list` is an acceptable type, then recurse down into `tuple`, following that we then check `str` and finally we check `False`. `False` isn't an acceptable type but it is an acceptable constant and hence accepted. Note: certain acceptable types and constants are not allowed as return values, i.e. `None` is not an accepted return constant

1527  
1528  
1529  
1530

```

acceptable_types = { 'int', 'str', 'float', 'bool', 'none', 'list', 'dict',
'tuple', 'set', 'datetime.date', 'date', 'literal', 'optional', 'union',
'sequence', 'iterable', 'frozenset', 'mapping' }

```

1531

```

acceptable_constants = { 'ellipsis', True, False, None }

```

1532  
1533  
1534  
1535  
1536  
1537

**Filtering functions.** When filtering functions we maintain four separate block lists, 1) a list of banned imports (including direct and aliases), 2) a list of banned functions (some common libraries have a limited set of non-deterministic methods, we don't want to fully exclude them), 3) a list of banned variables (some variables such as `__version__` are likely to be environment based), 4) a list of banned repos (some repos from cloud providers provide thousands of near identical methods with different urls, we remove these as they are not a valuable contribution to the evaluation).

1538

## 1539 A.6 STATIC AND RUNTIME CODE STATISTICS

1540  
1541

Given a task from the evaluation set we perform the following static and runtime analyses:

1542

### 1542 Static Analysis:

1543  
1544  
1545  
1546  
1547  
1548  
1549  
1550  
1551

1. **Lines of Code Count.** Total number of lines, excluding blanks and comments.
2. **AST Node Types Count.** Count of all Python Abstract Syntax Tree (AST) node types present in the code, e.g. `FunctionDef()`, `AsyncFunctionDef()`, `Assign()`, `For()`, ...
3. **Cyclomatic Complexity.** An estimate of the number of linearly independent paths through a program's source code. Note: There are several limitations in the implementation of this metric as we only parse python source code, and some modern python features such as pattern matching statements are yet to be supported.
4. **Maintainability Index.** A estimate of code maintainability and quality incorporating several other estimated measures (e.g. Halstead Volume, Cyclomatic Complexity, and lines of code). Note: Faces the same aforementioned limitations.

1552  
1553  
1554

### 1555 Runtime Analysis:

1556  
1557  
1558  
1559  
1560  
1561  
1562  
1563  
1564  
1565

1. **CPU Time.**
2. **Loop Iterations.** Including for loops, while loops and list comprehensions.
3. **Arithmetic Operations.** Including addition, subtraction, multiplication, division and power operations.
4. **Execution Metrics.** Including lines executed, library lines executed and conditional statements executed.
5. **Function Calls.** Including builtin function calls, user-defined function calls and total function calls.
6. **Variable Usage.** Including variables declared and variables used

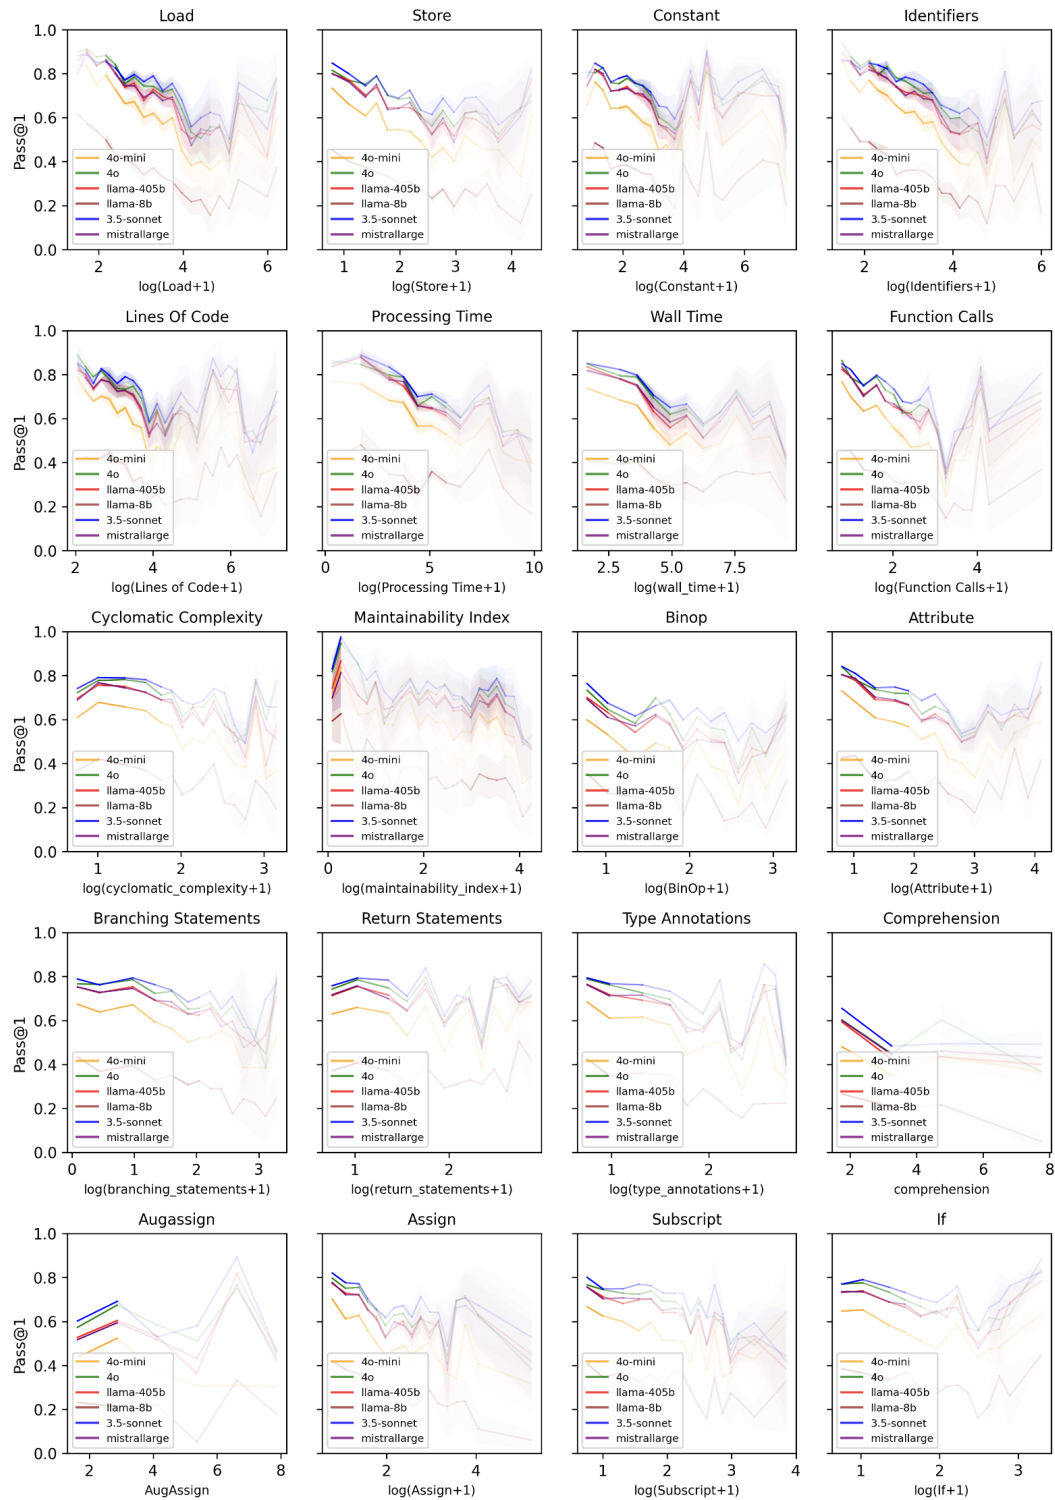


Figure 9: Top static statistics visualised against Pass@1 rate for all models tested

## A.7 OUTPUT COMPARISON AND VALIDATION

When evaluating model outputs against ground truth values, we employ two distinct comparison strategies depending on whether the output represents a successful execution result or an error case.

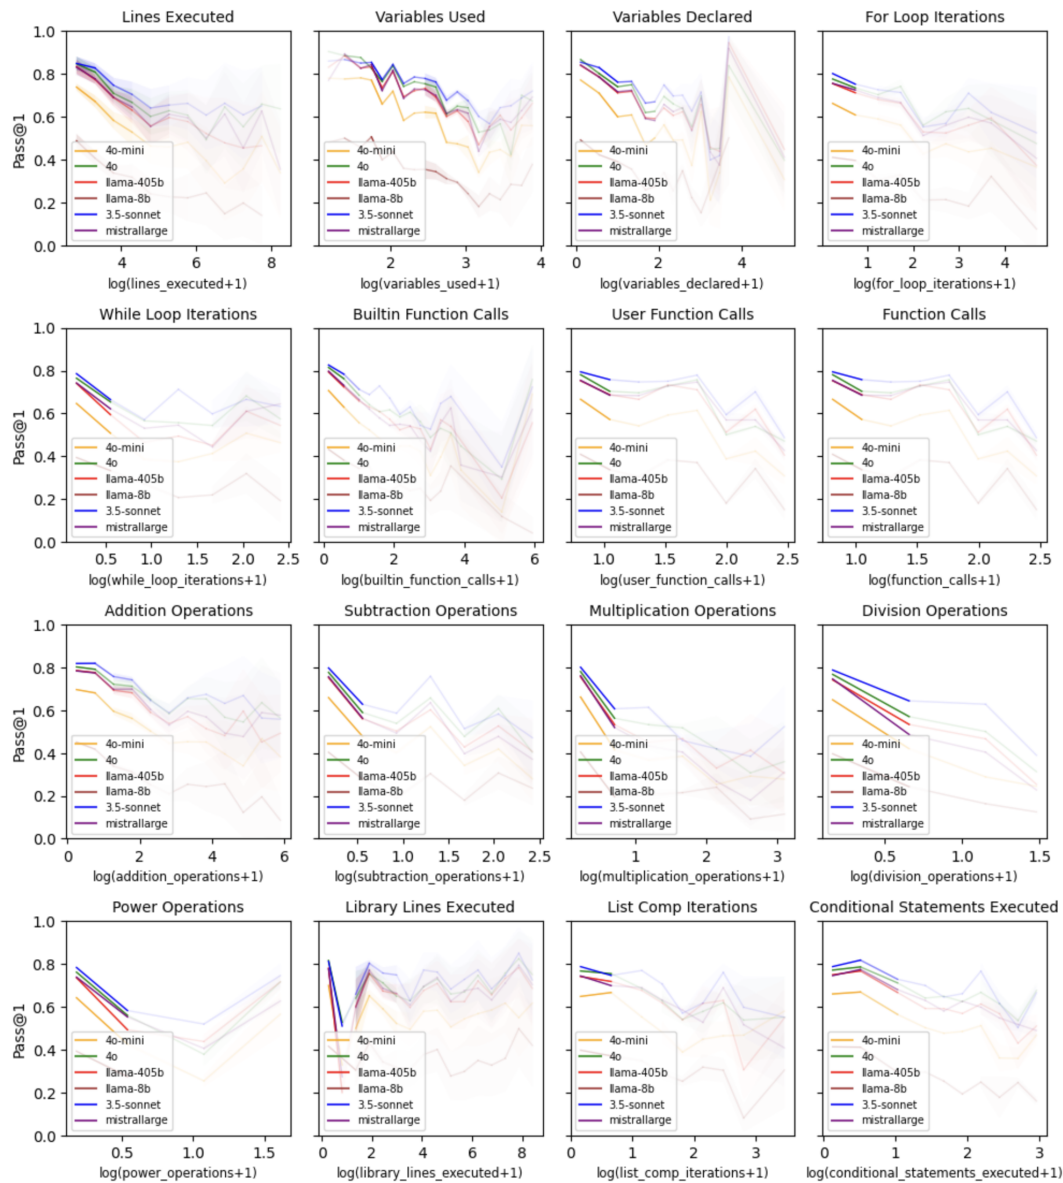


Figure 10: Runtime statistics visualised against Pass@1 rate for all models tested

This dual approach is necessary because error messages often contain version or implementation-specific details while maintaining semantic equivalence.

#### A.7.1 DIRECT VALUE COMPARISON

For successful execution results, we perform limited preprocessing (unsorted container objects e.g. `set` and `frozenset` are sorted before conversion to json lists, iterable types i.e. tuples are converted to lists, numbers are consistently formatted), then make a direct comparison between the model output and ground truth as json objects.

#### A.7.2 ERROR MESSAGE COMPARISON

For error cases, we use a language model-based comparison approach that focuses on specific error patterns and known version differences. This structured approach is necessary as error messages have evolved across Python versions while maintaining the same underlying causes.

1674 **Stacktrace Handling.** We explicitly exclude stacktraces from comparison as they contain  
 1675 execution-specific information like file paths, and external details that the model is not privy to.  
 1676

1677 **Version-Specific Error Messages.** Python has evolved to provide more helpful error messages  
 1678 in recent versions, with significant changes between major releases. Our comparison system must  
 1679 handle these variations appropriately. Examples of version-specific differences:  
 1680

```

1681 # Python 3.9
1682 my_list = [1, 2 3]
1683 SyntaxError: invalid syntax

1684 # Python 3.10
1685 my_list = [1, 2 3]
1686 SyntaxError: invalid syntax. Perhaps you forgot a comma?

1687
1688 # Both indicate the same missing comma issue

1689
1690
1691 # Python 3.11
1692 my_string = f"{x z y}" + f"{1 + 1}"
1693 SyntaxError: f-string: invalid syntax. Perhaps you forgot a comma?

1694 # Python 3.12
1695 my_string = f"{x z y}" + f"{1 + 1}"
1696 SyntaxError: invalid syntax. Perhaps you forgot a comma?

1697
1698 # While the messages differ, they point to the same syntactic error
  
```

1699  
 1700 To handle these variations, our error comparison system uses a prompt that encourages human-like  
 1701 reasoning about error equivalence:  
 1702

```

1703 You are an expert Python developer looking at two error messages.
1704 ↪ Determine if they are describing the same underlying issue, even if
1705 ↪ expressed differently. Consider:

1706 - Different Python versions might provide different levels of detail for
1707 ↪ the same error
1708 - The core issue might be described in more or less helpful ways
1709 - Extra hints or suggestions don't change the fundamental error
1710 - Line numbers and file paths are irrelevant

1711
1712 Message 1: {error1}
1713 Message 2: {error2}

1714 Would a Python developer consider these to be the same error? Answer
1715 ↪ only 'True' or 'False'.
  
```

1716  
 1717  
 1718 This structured approach to error comparison improves consistency in evaluation across different  
 1719 Python versions and implementation variations while maintaining the ability to identify truly distinct  
 1720 error cases.  
 1721

## 1722 A.8 PER FUNCTION TASK SET DIVERSITY

1723  
 1724 To measure EXE's potential to scale in the future, we analyse a model's ability to continually gen-  
 1725 erate new test cases given a single function. This is performed by:

- 1726 1. Sampling functions from EXE's dataset (samples detailed below).
- 1727 2. Generating a batch of test cases in accordance with A.2, recording token usage.



- 1728 3. Running validators in accordance with A.2, removing cases that are duplicates, fail to exe-  
 1729 cute, fail to be parsed, or that trigger any validator.  
 1730  
 1731 4. Continue generating new batches of test cases, injecting a random selection of (up to 60)  
 1732 previously generated cases into the prompt (detailed samples of test cases generated can be  
 1733 seen at the end of this appendix).

#### 1734 A.8.1 INLINED CODE TASKS FOR GENERATION

1735 Example 1. `get_origin_link_and_tag` from `utils.py` in `azure-nspkg`:

```
1736
1737
1738 from typing import List
1739
1740 def get_origin_link_and_tag(issue_body_list: List[str]) -> (str, str):
1741     link, readme_tag = '', ''
1742     for row in issue_body_list:
1743         if 'link' in row.lower() and 'release request' not in
1744             ↪ row.lower() and link == '':
1745             link = row.split(":", 1)[-1].strip()
1746         if 'readme tag' in row.lower() and readme_tag == '':
1747             readme_tag = row.split(":", 1)[-1].strip()
1748         if link and readme_tag:
1749             break
1750
1751     if link.count('https') > 1:
1752         link = link.split(']')[0]
1753         link = link.replace('[', "").replace(']', "").replace('(',
1754             ↪ "").replace(')', "")
1755     return link, readme_tag
```

1755 Example 2. `_compute_affine_output_size_python.py` from `geometry.py` in  
 1756 `torchvision`:

```
1757
1758 from typing import List, Tuple
1759
1760 import math
1761
1762 def _compute_affine_output_size_python(matrix: List[float], w: int, h:
1763     ↪ int) -> Tuple[int, int]:
1764     # Mostly copied from PIL implementation:
1765     # The only difference is with transformed points as input matrix has
1766     ↪ zero translation part here and
1767     # PIL has a centered translation part.
1768     # https://github.com/python-pillow/Pillow/blob/11de3318867e43980573_
1769     ↪ 73ee9f12dcb33db7335c/src/PIL/Image.py#L2054
1770
1771     a, b, c, d, e, f = matrix
1772     xx = []
1773     yy = []
1774
1775     half_w = 0.5 * w
1776     half_h = 0.5 * h
1777     for x, y in ((-half_w, -half_h), (half_w, -half_h), (half_w,
1778         ↪ half_h), (-half_w, half_h)):
1779         nx = a * x + b * y + c
1780         ny = d * x + e * y + f
1781         xx.append(nx + half_w)
1782         yy.append(ny + half_h)
1783
1784     nw = math.ceil(max(xx)) - math.floor(min(xx))
1785     nh = math.ceil(max(yy)) - math.floor(min(yy))
1786     return int(nw), int(nh) # w, h
```

1782 Example 3. `_format_image.py` from `_chat_models.py` in `langchain-core`:

```
1783
1784 from typing import Dict
1785
1786 import re
1787
1788 def _format_image(image_url: str) -> Dict:
1789     """
1790     Formats an image of format data:image/jpeg;base64,{b64_string}
1791     to a dict for anthropic api
1792
1793     {
1794         "type": "base64",
1795         "media_type": "image/jpeg",
1796         "data": "/9j/4AAQSkZJRg...",
1797     }
1798
1799     And throws an error if it's not a b64 image
1800     """
1801     regex = r"^data:(?P<media_type>image/.+);base64,(?P<data>.+)$"
1802     match = re.match(regex, image_url)
1803     if match is None:
1804         raise ValueError(
1805             "Anthropic only supports base64-encoded images currently."
1806             " Example: data:image/png;base64,'/9j/4AAQSk'..."
1807         )
1808     return {
1809         "type": "base64",
1810         "media_type": match.group("media_type"),
1811         "data": match.group("data"),
1812     }
```

1811 Example 4. `make_arn_for_alarm.py` from `utils.py` in `moto`:

```
1812
1813 REGION_PREFIX_TO_PARTITION = {
1814     # (region prefix, aws partition)
1815     "cn-": "aws-cn",
1816     "us-gov-": "aws-us-gov",
1817     "us-iso-": "aws-iso",
1818     "us-isob-": "aws-iso-b",
1819 }
1820
1821 DEFAULT_PARTITION = "aws"
1822
1823 PARTITION_NAMES = list(REGION_PREFIX_TO_PARTITION.values()) +
1824 ↪ [DEFAULT_PARTITION]
1825
1826 def get_partition(region: str) -> str:
1827     if not region:
1828         return DEFAULT_PARTITION
1829     if region in PARTITION_NAMES:
1830         return region
1831     for prefix in REGION_PREFIX_TO_PARTITION:
1832         if region.startswith(prefix):
1833             return REGION_PREFIX_TO_PARTITION[prefix]
1834     return DEFAULT_PARTITION
1835
1836 def make_arn_for_alarm(region: str, account_id: str, alarm_name: str) ->
1837 ↪ str:
1838     return
1839     ↪ f"arn:{get_partition(region)}:cloudwatch:{region}:{account_id}:alarm:{alarm_name}"
```

1836 Example 5. number2lowercase\_roman\_numeral.py from page\_labels.py in pypdf2:  
 1837  
 1838

```

1839 from typing import Iterator
1840
1841 def number2uppercase_roman_numeral(num: int) -> str:
1842     roman = [
1843         (1000, "M"),
1844         (900, "CM"),
1845         (500, "D"),
1846         (400, "CD"),
1847         (100, "C"),
1848         (90, "XC"),
1849         (50, "L"),
1850         (40, "XL"),
1851         (10, "X"),
1852         (9, "IX"),
1853         (5, "V"),
1854         (4, "IV"),
1855         (1, "I"),
1856     ]
1857
1858     def roman_num(num: int) -> Iterator[str]:
1859         for decimal, roman_repr in roman:
1860             x, _ = divmod(num, decimal)
1861             yield roman_repr * x
1862             num -= decimal * x
1863             if num <= 0:
1864                 break
1865
1866     return "".join(list(roman_num(num)))
1867
1868 def number2lowercase_roman_numeral(number: int) -> str:
1869     return number2uppercase_roman_numeral(number).lower()
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889

```

1866 Example 6. alpha\_canonicalize.py from parser.py in opt-einsum:  
 1867

```

1868 _einsum_symbols_base =
1869     ↪ "abcdefghijklmnopqrstuvwxyzABCDEFGHIJKLMNOPQRSTUVWXYZ"
1870
1871 from typing import Dict
1872
1873 def get_symbol(i: int) -> str:
1874     """Get the symbol corresponding to int `i` - runs through the
1875     ↪ usual 52
1876     letters before resorting to unicode characters, starting at
1877     ↪ `chr(192)` and skipping surrogates.
1878
1879     **Examples:**
1880
1881     ``python
1882     get_symbol(2)
1883     #> 'c'
1884
1885     get_symbol(200)
1886     #> 'Ĥ'
1887
1888     get_symbol(20000)
1889     #> ''
1890     ...
1891     """
1892     if i < 52:
1893         return _einsum_symbols_base[i]
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000

```

```

1890
1891     elif i >= 55296:
1892         # Skip chr(57343) - chr(55296) as surrogates
1893         return chr(i + 2048)
1894     else:
1895         return chr(i + 140)
1896
1897 def alpha_canonicalize(equation: str) -> str:
1898     """Alpha convert an equation in an order-independent canonical way.
1899
1900     Examples
1901     -----
1902     >>> oe.parser.alpha_canonicalize("dcba")
1903     'abcd'
1904
1905     >>> oe.parser.alpha_canonicalize("â€ï¿½ö")
1906     'abccd'
1907     """
1908     rename: Dict[str, str] = {}
1909     for name in equation:
1910         if name in ".,->":
1911             continue
1912         if name not in rename:
1913             rename[name] = get_symbol(len(rename))
1914     return "".join(rename.get(x, x) for x in equation)
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932

```

Example 7. `remove_starting_comments.py` from `sql_util.py` in `snowflake-connector-python`:

```

1915 import re
1916
1917 COMMENT_START_SQL_RE = re.compile(
1918     r"""
1919         ^\s*(?:
1920             /\*[\w\W]*?\*/
1921         ) """,
1922     re.VERBOSE,
1923 )
1924
1925 def remove_starting_comments(sql: str) -> str:
1926     """Remove all comments from the start of a SQL statement."""
1927     commentless_sql = sql
1928     while True:
1929         start_comment = COMMENT_START_SQL_RE.match(commentless_sql)
1930         if start_comment is None:
1931             break
1932         commentless_sql = commentless_sql[start_comment.end() :]
1933     return commentless_sql
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943

```

Example 8. `_pad_version.py` from `specifiers.py` in `poetry-core`:

```

1935 import itertools
1936
1937 from typing import List, Tuple
1938
1939 def _pad_version(left: List[str], right: List[str]) -> Tuple[List[str],
1940     ↪ List[str]]:
1941     left_split, right_split = [], []
1942
1943     # Get the release segment of our versions
1944     left_split.append(list(itertools.takewhile(lambda x: x.isdigit(),
1945     ↪ left)))

```

```

1944     right_split.append(list(itertools.takewhile(lambda x: x.isdigit(),
1945     ↪ right)))
1946
1947     # Get the rest of our versions
1948     left_split.append(left[len(left_split[0]) :])
1949     right_split.append(right[len(right_split[0]) :])
1950
1951     # Insert our padding
1952     left_split.insert(1, ["0" * max(0, len(right_split[0]) -
1953     ↪ len(left_split[0]))])
1954     right_split.insert(1, ["0" * max(0, len(left_split[0]) -
1955     ↪ len(right_split[0]))])
1956
1957     return (list(itertools.chain(*left_split)),
1958     ↪ list(itertools.chain(*right_split)))
1959
1960
1961

```

Example 9. `get_flag_suggestions.py` from `_helpers.py` in `absl-py`:

```

1961     _SUGGESTION_ERROR_RATE_THRESHOLD = 0.50
1962
1963     from typing import List, Sequence
1964
1965     def _damerau_levenshtein(a, b):
1966         """Returns Damerau-Levenshtein edit distance from a to b."""
1967         memo = {}
1968
1969         def distance(x, y):
1970             """Recursively defined string distance with memoization."""
1971             if (x, y) in memo:
1972                 return memo[x, y]
1973             if not x:
1974                 d = len(y)
1975             elif not y:
1976                 d = len(x)
1977             else:
1978                 d = min(
1979                     distance(x[1:], y) + 1, # correct an insertion error
1980                     distance(x, y[1:]) + 1, # correct a deletion error
1981                     distance(x[1:], y[1:]) + (x[0] != y[0]) # correct a wrong
1982                     ↪ character
1983                 )
1984                 if len(x) >= 2 and len(y) >= 2 and x[0] == y[1] and x[1] == y[0]:
1985                     # Correct a transposition.
1986                     t = distance(x[2:], y[2:]) + 1
1987                     if d > t:
1988                         d = t
1989
1990             memo[x, y] = d
1991             return d
1992         return distance(a, b)
1993
1994     def get_flag_suggestions(
1995         attempt: str, longopt_list: Sequence[str]
1996     ) -> List[str]:
1997         """Returns helpful similar matches for an invalid flag."""
1998         # Don't suggest on very short strings, or if no longopts are
1999         ↪ specified.
2000         if len(attempt) <= 2 or not longopt_list:
2001             return []
2002
2003         option_names = [v.split('=')[0] for v in longopt_list]
2004
2005         # Find close approximations in flag prefixes.
2006         # This also handles the case where the flag is spelled right but
2007         ↪ ambiguous.

```

```

1998 distances = [(_damerau_levenshtein(attempt, option[0:len(attempt)]),
1999 ↪ option)
2000                 for option in option_names]
2001 # t[0] is distance, and sorting by t[1] allows us to have stable
2002 ↪ output.
2003 distances.sort()
2004
2005 least_errors, _ = distances[0]
2006 # Don't suggest excessively bad matches.
2007 if least_errors >= _SUGGESTION_ERROR_RATE_THRESHOLD * len(attempt):
2008     return []
2009
2010 suggestions = []
2011 for errors, name in distances:
2012     if errors == least_errors:
2013         suggestions.append(name)
2014     else:
2015         break
2016 return suggestions

```

Example 10. `valid_contexto.py` from `core.py` in `idna`:

```

2018 from types import SimpleNamespace
2019
2020 from typing import Tuple
2021
2022 def _encode_range(start: int, end: int) -> int:
2023     return (start << 32) | end
2024
2025 def _decode_range(r: int) -> Tuple[int, int]:
2026     return (r >> 32), (r & ((1 << 32) - 1))
2027
2028 import bisect
2029
2030 def intranges_contain(int_: int, ranges: Tuple[int, ...]) -> bool:
2031     """Determine if `int_` falls into one of the ranges in `ranges`."""
2032     tuple_ = _encode_range(int_, 0)
2033     pos = bisect.bisect_left(ranges, tuple_)
2034     # we could be immediately ahead of a tuple (start, end)
2035     # with start < int_ <= end
2036     if pos > 0:
2037         left, right = _decode_range(ranges[pos-1])
2038         if left <= int_ < right:
2039             return True
2040     # or we could be immediately behind a tuple (int_, end)
2041     if pos < len(ranges):
2042         left, _ = _decode_range(ranges[pos])
2043         if left == int_:
2044             return True
2045     return False
2046
2047 class idnadataClass(SimpleNamespace):
2048     def __init__(self):
2049         scripts = {
2050             'Greek': (
2051                 0x37000000374,
2052                 0x37500000378,
2053                 0x37a0000037e,
2054                 0x37f00000380,
2055                 0x38400000385,
2056                 0x38600000387,
2057                 0x3880000038b,
2058                 0x38c0000038d,

```

```
2052         0x38e000003a2,  
2053         0x3a3000003e2,  
2054         0x3f000000400,  
2055         0x1d2600001d2b,  
2056         0x1d5d00001d62,  
2057         0x1d6600001d6b,  
2058         0x1dbf00001dc0,  
2059         0x1f0000001f16,  
2060         0x1f1800001f1e,  
2061         0x1f2000001f46,  
2062         0x1f4800001f4e,  
2063         0x1f5000001f58,  
2064         0x1f5900001f5a,  
2065         0x1f5b00001f5c,  
2066         0x1f5d00001f5e,  
2067         0x1f5f00001f7e,  
2068         0x1f8000001fb5,  
2069         0x1fb600001fc5,  
2070         0x1fc600001fd4,  
2071         0x1fd600001fdc,  
2072         0x1fdd00001ff0,  
2073         0x1ff200001ff5,  
2074         0x1ff600001fff,  
2075         0x212600002127,  
2076         0xab650000ab66,  
2077         0x101400001018f,  
2078         0x101a0000101a1,  
2079         0x1d2000001d246,  
2080     ),  
2081     'Han': (  
2082         0x2e8000002e9a,  
2083         0x2e9b00002ef4,  
2084         0x2f0000002fd6,  
2085         0x300500003006,  
2086         0x300700003008,  
2087         0x30210000302a,  
2088         0x30380000303c,  
2089         0x340000004dc0,  
2090         0x4e000000a000,  
2091         0xf9000000fa6e,  
2092         0xfa700000fada,  
2093         0x16fe200016fe4,  
2094         0x16ff000016ff2,  
2095         0x200000002a6e0,  
2096         0x2a7000002b73a,  
2097         0x2b7400002b81e,  
2098         0x2b8200002cea2,  
2099         0x2ceb00002ebe1,  
2100         0x2ebf00002ee5e,  
2101         0x2f8000002fale,  
2102         0x300000003134b,  
2103         0x31350000323b0,  
2104     ),  
2105     'Hebrew': (  
2106         0x591000005c8,  
2107         0x5d0000005eb,  
2108         0x5ef000005f5,  
2109         0xfb1d0000fb37,  
2110         0xfb380000fb3d,  
2111         0xfb3e0000fb3f,  
2112         0xfb400000fb42,  
2113         0xfb430000fb45,  
2114         0xfb460000fb50,  
2115     ),
```

```

2106         'Hiragana': (
2107             0x304100003097,
2108             0x309d000030a0,
2109             0x1b0010001b120,
2110             0x1b1320001b133,
2111             0x1b1500001b153,
2112             0x1f2000001f201,
2113         ),
2114         'Katakana': (
2115             0x30a1000030fb,
2116             0x30fd00003100,
2117             0x31f000003200,
2118             0x32d0000032ff,
2119             0x330000003358,
2120             0xff660000ff70,
2121             0xff710000ff9e,
2122             0xlaff00001aff4,
2123             0xlaff50001affc,
2124             0xlaffd0001afff,
2125             0x1b0000001b001,
2126             0x1b1200001b123,
2127             0x1b1550001b156,
2128             0x1b1640001b168,
2129         ),
2130     }
2131
2132     self.__dict__.update(locals())
2133
2134 idnadata = idnadataClass()
2135
2136 def _is_script(cp: str, script: str) -> bool:
2137     return intranges_contain(ord(cp), idnadata.scripts[script])
2138
2139 def valid_contexto(label: str, pos: int, exception: bool = False) ->
2140     bool:
2141     cp_value = ord(label[pos])
2142
2143     if cp_value == 0x00b7:
2144         if 0 < pos < len(label)-1:
2145             if ord(label[pos - 1]) == 0x006c and ord(label[pos + 1]) ==
2146                 0x006c:
2147                 return True
2148             return False
2149
2150     elif cp_value == 0x0375:
2151         if pos < len(label)-1 and len(label) > 1:
2152             return _is_script(label[pos + 1], 'Greek')
2153             return False
2154
2155     elif cp_value == 0x05f3 or cp_value == 0x05f4:
2156         if pos > 0:
2157             return _is_script(label[pos - 1], 'Hebrew')
2158             return False
2159
2160     elif cp_value == 0x30fb:
2161         for cp in label:
2162             if cp == '\u30fb':
2163                 continue
2164             if _is_script(cp, 'Hiragana') or _is_script(cp, 'Katakana')
2165                 or _is_script(cp, 'Han'):
2166                 return True
2167             return False
2168
2169     elif 0x660 <= cp_value <= 0x669:

```



```

2160
2161     for cp in label:
2162         if 0x6f0 <= ord(cp) <= 0x06f9:
2163             return False
2164     return True
2165
2166     elif 0x6f0 <= cp_value <= 0x6f9:
2167         for cp in label:
2168             if 0x660 <= ord(cp) <= 0x0669:
2169                 return False
2170     return True
2171
2172     return False

```

Example 11. `exact_match.py` from `meteor_score.py` in `nltk`:

```

2174 from typing import Callable, Iterable, List, Tuple
2175
2176 def _match_enums(
2177     enum_hypothesis_list: List[Tuple[int, str]],
2178     enum_reference_list: List[Tuple[int, str]],
2179 ) -> Tuple[List[Tuple[int, int]], List[Tuple[int, str]], List[Tuple[int,
2180     ↪ str]]]:
2181     """
2182     matches exact words in hypothesis and reference and returns
2183     a word mapping between enum_hypothesis_list and enum_reference_list
2184     based on the enumerated word id.
2185
2186     :param enum_hypothesis_list: enumerated hypothesis list
2187     :param enum_reference_list: enumerated reference list
2188     :return: enumerated matched tuples, enumerated unmatched hypothesis
2189     ↪ tuples,
2190             enumerated unmatched reference tuples
2191     """
2192     word_match = []
2193     for i in range(len(enum_hypothesis_list))[:-1]:
2194         for j in range(len(enum_reference_list))[:-1]:
2195             if enum_hypothesis_list[i][1] == enum_reference_list[j][1]:
2196                 word_match.append(
2197                     (enum_hypothesis_list[i][0],
2198                     ↪ enum_reference_list[j][0])
2199                 )
2200                 enum_hypothesis_list.pop(i)
2201                 enum_reference_list.pop(j)
2202                 break
2203     return word_match, enum_hypothesis_list, enum_reference_list
2204
2205 def _generate_enums(
2206     hypothesis: Iterable[str],
2207     reference: Iterable[str],
2208     preprocess: Callable[[str], str] = str.lower,
2209 ) -> Tuple[List[Tuple[int, str]], List[Tuple[int, str]]]:
2210     """
2211     Takes in pre-tokenized inputs for hypothesis and reference and
2212     ↪ returns
2213     enumerated word lists for each of them
2214
2215     :param hypothesis: pre-tokenized hypothesis
2216     :param reference: pre-tokenized reference
2217     :preprocess: preprocessing method (default str.lower)
2218     :return: enumerated words list
2219     """
2220     if isinstance(hypothesis, str):
2221         raise TypeError(

```

```

2214         f'"hypothesis" expects pre-tokenized hypothesis
2215         ↪ (Iterable[str]): {hypothesis}'
2216     )
2217
2218     if isinstance(reference, str):
2219         raise TypeError(
2220             f'"reference" expects pre-tokenized reference
2221             ↪ (Iterable[str]): {reference}'
2222         )
2223
2224     enum_hypothesis_list = list(enumerate(map(preprocess, hypothesis)))
2225     enum_reference_list = list(enumerate(map(preprocess, reference)))
2226     return enum_hypothesis_list, enum_reference_list
2227
2228 def exact_match(
2229     hypothesis: Iterable[str], reference: Iterable[str]
2230 ) -> Tuple[List[Tuple[int, int]], List[Tuple[int, str]], List[Tuple[int,
2231     str]]]:
2232     """
2233     matches exact words in hypothesis and reference
2234     and returns a word mapping based on the enumerated
2235     word id between hypothesis and reference
2236
2237     :param hypothesis: pre-tokenized hypothesis
2238     :param reference: pre-tokenized reference
2239     :return: enumerated matched tuples, enumerated unmatched hypothesis
2240     ↪ tuples,
2241           enumerated unmatched reference tuples
2242     """
2243     enum_hypothesis_list, enum_reference_list =
2244     ↪ _generate_enums(hypothesis, reference)
2245     return _match_enums(enum_hypothesis_list, enum_reference_list)

```

## 2244 A.8.2 GENERATED TEST CASES

2245 Below is a sample of generated test cases (cut down to 3 examples, showing the first 60 cases for  
2246 brevity).

2247 First 60 generated cases for alpha\_canonicalize

```

2248
2249
2250 0:  {"args": ["a1b2c3->d4e5f6"], "kwargs": {}}
2251 1:  {"args": ["    ->    ,    ->    "], "kwargs": {}}
2252 2:  {"args": ["    "], "kwargs": {}}
2253 3:  {"args": ["AAA BBB CCC"], "kwargs": {}}
2254 4:  {"args": ["abcdefghijklmnopqrstuvwxy"], "kwargs": {}}
2255 5:  {"args": ["a\u0000b\u0001c\u0002->d\u0003e\u0004f\u0005"],
2256     ↪ "kwargs": {}}
2257 6:  {"args": ["abcdefghijklmnopqrstuvwxyABCDEFGHJKLMNOPQRSTUVWXYZ"],
2258     ↪ "kwargs": {}}
2259 7:  {"args":
2260     ↪ ["\ud83d\ude00\ud83d\ude03\ud83d\ude04\ud83d\ude01\ud83d\ude06->\ud
2261     ↪ 83d\ude05\ud83d\ude02\ud83e\udd23\ud83d\ude0a\ud83d\ude07"],
2262     ↪ "kwargs": {}}
2263 8:  {"args": ["a->b, c->d"], "kwargs": {}}
2264 9:  {"args": ["\u3053\u3093\u306b\u3061\u306f->\u4e16\u754c,\u4f60\u59
2265     ↪ 7d->\u4e16\u754c,\uc548\u155\ud558\uc138\uc694->\uc138\uacc4"],
2266     ↪ "kwargs": {}}
2267 10: {"args": ["\ud83c\udf1f\ud83c\udf20\u2728\ud83d\udcab\u2b50"],
2268     ↪ "kwargs": {}}
2269 11: {"args": ["ZYXWVUTSRQPONMLKJIHGFEDCBA"], "kwargs": {}}
2270 12: {"args": ["AaBbCcDdEeFfGgHhIiJjKkLlMmNnOoPpQqRrSsTtUuVvWwXxYyZz"],
2271     ↪ "kwargs": {}}

```

```

2268
2269 13: {"args": ["aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
2270 ↪ aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa"], "kwargs":
2271 ↪ {}}
2271 14: {"args": ["abcdefghijklmnopqrstuvwxyABCDEFGHIJKLMNOPQRSTUVWXYZ012_J
2272 ↪ 3456789!@#%^&*()_+"], "kwargs":
2273 ↪ {}}
2274 15: {"args": ["aa"], "kwargs": {}}
2275 16: {"args": [",,,,,->..."], "kwargs": {}}
2276 17: {"args": ["\u03b1\u03b2\u03b3\u03b4\u03b5\u03b6\u03b7\u03b8\u03b9\u03ba\u03bb\u03bc\u03bd\u03be\u03bf\u03c0\u03c1\u03c2\u03c3\u03c4\u03c5\u03c6\u03c7\u03c8\u03c9"], "kwargs":
2277 ↪ {}}
2278 18: {"args": ["1234567890"], "kwargs": {}}
2279 19: {"args": ["a->a,b->b,c->c,d->d"], "kwargs": {}}
2280 20: {"args": ["A1->B2,C3->D4,E5->F6"], "kwargs": {}}
2281 21: {"args": ["123456789"], "kwargs": {}}
2282 22: {"args": ["a->b,c->d,e->f,g->h"], "kwargs": {}}
2283 23: {"args": ["a->a,b->b,c->c"], "kwargs": {}}
2284 24: {"args": ["a\nb\tc\r d->e\nf\tg\r h"], "kwargs": {}}
2285 25: {"args": [""], "kwargs": {}}
2286 26: {"args": ["->->->->->"], "kwargs": {}}
2287 27: {"args": ["a->b->c->d->e->f->g->h->i->j->k->l->m->n->o->p->q->r->s_j
2288 ↪ ->t->u->v->w->x->y->z"], "kwargs":
2289 ↪ {}}
2290 28: {"args": ["\u7532->\u4e59,\u4e19->\u4e01,\u620a->\u5df1"],
2291 ↪ "kwargs": {}}
2292 29: {"args": ["\u6df7\u5408\u5b57\u7b26\u4e32with\u82f1\u6587and\u6570_j
2293 ↪ \u5b57123"], "kwargs":
2294 ↪ {}}
2295 30: {"args": ["\u0124\u011b\u013c\u013c\u00f6"], "kwargs": {}}
2296 31: {"args": ["!@#%^&*()_+"], "kwargs": {}}
2297 32: {"args": ["aaaaabbbbbccccc"], "kwargs": {}}
2298 33: {"args": [".,->.,->.,->.,->.,->."], "kwargs": {}}
2299 34: {"args": ["\u00c4\u00d6\u00dc\u00e4\u00f6\u00fc\u00df"], "kwargs":
2300 ↪ {}}
2301 35: {"args": ["a->b->c->d->e->f->g->h->i->j"], "kwargs": {}}
2302 36: {"args": ["A->1,B->2,C->3,D->4,E->5,F->6,G->7,H->8,I->9,J->0"],
2303 ↪ "kwargs": {}}
2304 37: {"args":
2305 ↪ ["\ud83c\udf1f->\ud83c\udf19,\ud83c\udf1e->\ud83c\udf0d"],
2306 ↪ "kwargs": {}}
2307 38: {"args":
2308 ↪ ["\u03b1\u03b2\u03b3\u03b4\u03b5->\u03b6\u03b7\u03b8\u03b9\u03ba,\u03bb\u03bc\u03bd\u03be\u03bf->\u03c0\u03c1\u03c2\u03c3\u03c4\u03c5"],
2309 ↪ "kwargs": {}}
2310 39: {"args": ["AaAaAa->BbBbBb,CcCcCc->DdDdDd"], "kwargs": {}}
2311 40: {"args": ["\u3053\u3093\u306b\u3061\u306f->\u4e16\u754c"],
2312 ↪ "kwargs": {}}
2313 41: {"args": [".,->"], "kwargs": {}}
2314 42: {"args": ["a->b,c->d,e->f"], "kwargs": {}}
2315 43: {"args": ["!@#%^&*()_+=[{}|;:\'",.<>?/~`"], "kwargs": {}}
2316 44: {"args": ["dcba"], "kwargs": {}}
2317 45: {"args": ["a1->b2,c3->d4,e5->f6"], "kwargs": {}}
2318 46: {"args":
2319 ↪ ["abcdefghijklmnopqrstuvwxyABCDEFGHIJKLMNOPQRSTUVWXYZ0123456789"],
2320 ↪ "kwargs": {}}
2321 47: {"args": ["AaAaAa->BbBbBb,CcCcCc->DdDdDd,EeEeEe->FfFfFf"],
2322 ↪ "kwargs": {}}
2323 48: {"args": ["aaaaaaaaaaaaaaaaaaaaa->bbbbbbbbbbbbbbbbbbbb"], "kwargs":
2324 ↪ {}}
2325 49: {"args": ["ABCDEFGHIJKLMNOPQRSTUVWXYZ"], "kwargs": {}}
2326 50: {"args": ["a\nb\tc\r d"], "kwargs": {}}
2327 51: {"args":
2328 ↪ ["\u0124\u011b\u013c\u013c\u00f6->\u0174\u00f4\u0159\u013c\u010f"],
2329 ↪ "kwargs": {}}

```





2430

```

57: {"args": [{"a", "b", "c"}, ["1", "2", "3"]], "kwargs": {}}
58: {"args": [{"100", "200", "300"}, ["99", "199", "299"]], "kwargs":
  ↪ {}}
59: {"args": [{"1"}, ["1", "0", "0", "0"]], "kwargs": {}}
60: {"args": [{"10", "0", "1"}, ["9", "9", "9"]], "kwargs": {}}

```

2435

2436

2437

First 60 generated cases for get\_flag\_suggestions

2438

2439

2440

2441

2442

2443

2444

2445

2446

2447

2448

2449

2450

2451

2452

2453

2454

2455

2456

2457

2458

2459

2460

2461

2462

2463

2464

2465

2466

2467

2468

2469

2470

2471

2472

2473

2474

2475

2476

2477

2478

2479

2480

2481

2482

2483

```

0: {"args": ["typo", ["typo1", "typo2", "correct"]], "kwargs": {}}
1: {"args": ["ambiguous", ["ambiguous1", "ambiguous2", "ambiguous3",
  ↪ "unambiguous", "ambiguous"]], "kwargs": {}}
2: {"args": ["aaaaaaaaa", ["aaaaaaaaa", "aaaaaaaaa",
  ↪ "aaaaaaaaa"]], "kwargs": {}}
3: {"args": ["\u3053\u3093", ["\u3053\u3093\u306b\u3061\u306f",
  ↪ "\u3053\u3093\u3070\u3093\u306f",
  ↪ "\u3053\u3093\u306a\u306b\u3061\u306f"]], "kwargs": {}}
4: {"args": ["", ["option1", "option2", "option3"]], "kwargs": {}}
5: {"args": ["\u65e5\u672c\u8a9e", ["\u65e5\u672c\u8a9e",
  ↪ "\u4e2d\u6587", "\ud55c\uad6d\uc5b4"]], "kwargs": {}}
6: {"args": ["no_match", ["completely", "different", "options"]],
  ↪ "kwargs": {}}
7: {"args": ["a", ["apple", "banana", "cherry"]], "kwargs": {}}
8: {"args": ["casesensitive", ["CaseSensitive", "casesensitive",
  ↪ "CASESENSITIVE"]], "kwargs": {}}
9: {"args": ["prefix", ["prefix_long_option1", "prefix_long_option2",
  ↪ "different_option"]], "kwargs": {}}
10: {"args": ["typo", ["type", "types", "typescript", "typoo"]],
  ↪ "kwargs": {}}
11: {"args": ["flag123", ["flag123=value", "flag124=value",
  ↪ "flag125=value"]], "kwargs": {}}
12: {"args": ["short", ["s", "sh", "sho", "shor", "short", "shorts"]],
  ↪ "kwargs": {}}
13: {"args": ["very_similar", ["very_similar1", "very_similar2",
  ↪ "very_similar3", "completely_different"]], "kwargs": {}}
14: {"args": ["!@#%^&*", ["!@#%^&*", "special_chars",
  ↪ "normal_option"]], "kwargs": {}}
15: {"args": ["completelydifferent", ["apple", "banana", "cherry",
  ↪ "date"]], "kwargs": {}}
16: {"args": ["flag123", ["flag123=value", "flag124=value",
  ↪ "flag125=value", "flag123"]], "kwargs": {}}
17: {"args": ["abc", ["abcd", "abce", "abcf"]], "kwargs": {}}
18: {"args": ["flag", []], "kwargs": {}}
19: {"args": ["apple", ["apple", "apples", "applesauce"]], "kwargs": {}}
20: {"args": ["verylong", ["verylongoptionname", "anotherlongoption",
  ↪ "yetanotherlongoption"]], "kwargs": {}}
21: {"args": ["verysimilar", ["verysimilar1", "verysimilar2",
  ↪ "verysimilar3"]], "kwargs": {}}
22: {"args": ["aaa", ["aaaa", "aaaaa", "aaaaaa", "bbb"]], "kwargs": {}}
23: {"args": ["exact", ["exact", "exactly", "exacting"]], "kwargs": {}}
24: {"args": ["special!@#", ["special!@#", "special$%^",
  ↪ "special&*()"]], "kwargs": {}}
25: {"args": ["short", ["s", "sh", "sho", "shor", "short"]], "kwargs":
  ↪ {}}
26: {"args": ["hello", [], "kwargs": {}}
27: {"args": ["hel", ["hello", "help", "health"]], "kwargs": {}}
28: {"args": ["longflagname", ["longflagname1", "longflagname2",
  ↪ "longflagname3", "shortflag"]], "kwargs": {}}
29: {"args": ["flag=value", ["flag1=value", "flag2=value",
  ↪ "flag3=value", "flag=othervalue"]], "kwargs": {}}
30: {"args": ["health", ["health", "help", "hello"]], "kwargs": {}}
31: {"args": ["prefix", ["prefix_option1", "prefix_option2",
  ↪ "different_option"]], "kwargs": {}}
32: {"args": ["option", ["option1=value", "option2=value",
  ↪ "option3=value"]], "kwargs": {}}

```

```

2484
2485 33: {"args": ["completelydifferent", ["apple", "banana", "cherry"]],
    ↪ "kwargs": {}}
2486
2487 34: {"args": ["hlp", ["help", "hello", "health"]], "kwargs": {}}
2488
2489 35: {"args": ["num123", ["num1234", "num12345", "num123456"]],
    ↪ "kwargs": {}}
2490
2491 36: {"args": ["verylongflagname", ["verylongflagname1",
    ↪ "verylongflagname2", "verylongflagname3"]], "kwargs": {}}
2492
2493 37: {"args": ["hel", ["help", "hello", "health", "helmet"]], "kwargs":
    ↪ {}}
2494
2495 38: {"args": [" whitespace ", ["whitespace", " whitespace ", "
    ↪ whitespace "]], "kwargs": {}}
2496
2497 39: {"args": ["completelydifferent", ["option1", "option2",
    ↪ "option3"]], "kwargs": {}}
2498
2499 40: {"args": ["mixed_case", ["MIXED_CASE", "mixed_case", "MixedCase"]],
    ↪ "kwargs": {}}
2500
2501 41: {"args": [" whitespace ", ["whitespace", " whitespace ", "
    ↪ whitespace ", "no_whitespace"]], "kwargs": {}}
2502
2503 42: {"args": ["helpp", ["help", "hello", "health"]], "kwargs": {}}
2504
2505 43: {"args": ["option", [], "kwargs": {}}
2506
2507 44: {"args": ["flag=value", ["flag1=value", "flag2=value",
    ↪ "flag3=value"]], "kwargs": {}}
2508
2509 45: {"args": ["multi\nline", ["multi\nline", "multiline", "multi line",
    ↪ "multi\tline"]], "kwargs": {}}
2510
2511 46: {"args": ["test-flag", ["test_flag", "test-flag", "testflag"]],
    ↪ "kwargs": {}}
2512
2513 47: {"args": ["prefix", ["prefix_option1", "prefix_option2",
    ↪ "different_option", "prefixx"]], "kwargs": {}}
2514
2515 48: {"args": ["ambiguous", ["ambiguous1", "ambiguous2", "ambiguous3",
    ↪ "unambiguous"]], "kwargs": {}}
2516
2517 49: {"args": ["flag", ["flag1", "flag2", "flag3", "flag4", "flag5",
    ↪ "flag6", "flag7", "flag8", "flag9", "flag10"]], "kwargs": {}}
2518
2519 50: {"args": ["verylongflagname", ["verylongflagname1",
    ↪ "verylongflagname2", "verylongflagname3", "shortflag"]], "kwargs":
    ↪ {}}
2520
2521 51: {"args": ["mixed_case", ["MIXED_CASE", "mixed_case", "MixedCase",
    ↪ "mixedcase"]], "kwargs": {}}
2522
2523 52: {"args": ["\u3053\u3093\u306b\u3061\u306f",
    ↪ "\u3053\u3093\u306b\u3061\u306f",
2524
2525 ↪ "\u3055\u3088\u3046\u306a\u3089", "\u304a\u306f\u3088\u3046"]],
    ↪ "kwargs": {}}
2526
2527 53: {"args": ["multi\nline", ["multi\nline", "multiline", "multi
    ↪ line"]], "kwargs": {}}
2528
2529 54: {"args": ["aaa", ["aaaa", "aaaaa", "aaaaaa"]], "kwargs": {}}
2530
2531 55: {"args": ["abc", [], "kwargs": {}}
2532
2533 56: {"args": ["verysimilar", ["verysimilar1", "verysimilar2",
    ↪ "completelydifferent"]], "kwargs": {}}
2534
2535 57: {"args": ["verylongflagnamewithmorethanfiftycharacterstotest",
    ↪ ["verylongflagnamewithmorethanfiftycharacterstotest1",
    ↪ "verylongflagnamewithmorethanfiftycharacterstotest2",
    ↪ "shortflag"]], "kwargs": {}}
2536
2537 58: {"args": ["he", ["hello", "help", "health"]], "kwargs": {}}
2538
2539 59: {"args": ["aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
    ↪ aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa",
    ↪ ["aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
    ↪ aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa",
    ↪ "aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
    ↪ aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa",
    ↪ "aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
    ↪ aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa"]], "kwargs":
    ↪ {}}
2540
2541 60: {"args": ["typo", ["type", "types", "typescript"]], "kwargs": {}}

```