

SWE-BENCH MULTIMODAL: DO AUTONOMOUS PROGRAMMING SYSTEMS GENERALIZE TO NEW SOFTWARE DOMAINS?

Anonymous authors

Paper under double-blind review

ABSTRACT

Autonomous systems for software engineering are now capable of fixing bugs and developing features. These systems are commonly evaluated on SWE-bench (Jimenez et al., 2024a), which assesses their ability to solve software issues from GitHub repositories. However, SWE-bench uses only Python repositories, with problem statements presented predominantly as text and lacking visual elements such as images. This limited coverage motivates our inquiry into how existing systems might perform on unrepresented software engineering domains (e.g., front-end, game development, DevOps), which use different programming languages and paradigms. Therefore, we propose *SWE-bench Multimodal* (SWE-bench M), to evaluate systems on their ability to fix bugs in visual, user-facing JavaScript software. SWE-bench M features 617 task instances collected from 17 JavaScript libraries used for web interface design, diagramming, data visualization, syntax highlighting, and interactive mapping. Each SWE-bench M task instance contains at least one image in its problem statement or unit tests. Our analysis finds that top-performing SWE-bench systems struggle with SWE-bench M, revealing limitations in visual problem-solving and cross-language generalization. Lastly, we show that SWE-agent’s flexible language-agnostic features enable it to substantially outperform alternatives on SWE-bench M, resolving 12% of task instances compared to 6% for the next best system.

1 INTRODUCTION

Language models (LMs) are being increasingly deployed to assist software engineers (Bagalkote, 2024; Yepis & StackOverflow, 2024). As LMs gain in prominence, the research community has been expanding from building LM-based assistants that work at the code line or function level (Chen et al., 2021; Hendrycks et al., 2021) to building *autonomous systems* that can maintain and improve large codebases with hundreds of files and thousands of lines (Wang et al., 2024b; Xia et al., 2024; Yang et al., 2024; Zhang et al., 2024b). These systems provide LMs with tools and environments that let them engage in multi-step interactions to solve complex software development tasks.

SWE-bench (Jimenez et al., 2024a) is the most popular benchmark for evaluating the performance of these systems. Drawn from GitHub issues and pull requests, SWE-bench task instances capture a range of software bugs and verify solution behavior by executing unit tests. Since the introduction of SWE-bench in October 2023, state-of-the-art performance on SWE-bench Lite, the most commonly used subset of SWE-bench, has soared from 3% to 43% (Jimenez et al., 2024b).

However, with respect to the broader landscape of software engineering, SWE-bench reflects only a fraction of real-world applications. Its 17 repositories are predominantly written in Python. Task instance codebases tend to be structured similarly because every repository is a PyPI package. Though SWE-bench features several commonly used backend and data science libraries, many other use cases are not represented. In addition, many domains of software development rely on visual assets, such as user interface design, gaming, virtual reality, and data visualizations (Foley & Van Dam, 1982; Shneiderman, 1998), but only 5.6% of SWE-bench tasks contain an image.

In light of these limitations, this paper poses the research question: “*Do autonomous programming systems generalize to new software domains?*”

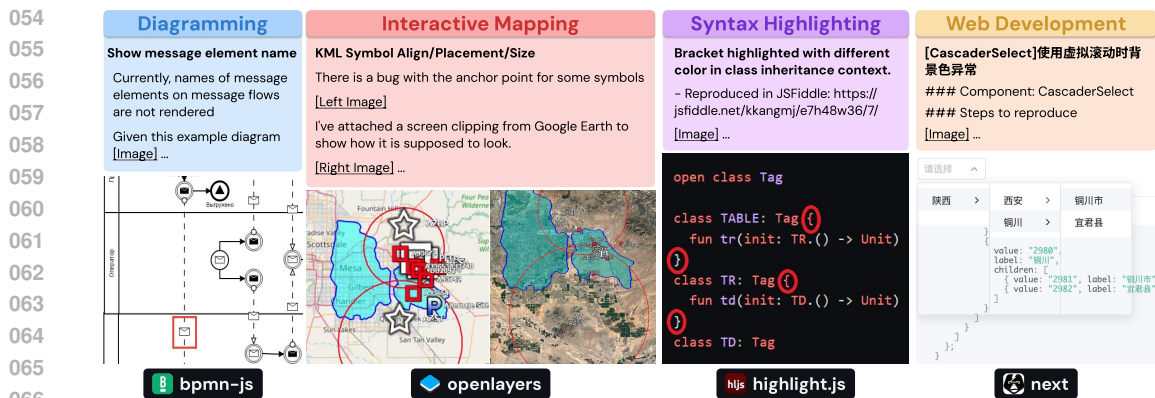


Figure 1: **Four task instances from SWE-bench Multimodal.** JavaScript repositories featured in SWE-bench M introduce new categories of software development challenges, such as the use cases shown above, along with performance profiling, creative coding, and more. In addition to multimodal comprehension, SWE-bench M contains more multi-language issues for both programming (JavaScript, TypeScript, HTML, CSS), and natural (English, Chinese) languages.

A popular sub-field capturing several of these unexplored dimensions is *front-end web development*, which focuses on building user-facing applications and websites with visual and interactive elements Si et al. (2024). Here, coding problems frequently contain both textual and visual components (Cherubini et al., 2007; Moody, 2009). As of 2023, JavaScript, the primary vehicle for front-end development, was the most popular programming language for the last decade (Daigle & GitHub, 2023). The JavaScript ecosystem encompasses a wide range of frameworks, such as Node.js, React, and Angular, each embodying unique architectural principles (Wittern et al., 2016).

In this work, we introduce **SWE-bench Multimodal** (SWE-bench M), a dataset of 619 task instances focused on visual JavaScript problems. Like SWE-bench, we derive instances from real-world issues on GitHub. Unlike SWE-bench, our collection targets JavaScript-based, user-facing applications, such as UI design systems, web app development, interactive mapping, and syntax highlighters. Furthermore, we filter exclusively for task instances that contain images or videos in their problem descriptions or testing scenarios. Finally, through verification by human experts, we find that for 83.5% of task instances, images are necessary for solving the corresponding issue.

We discover that **existing systems perform significantly worse on SWE-bench M than they do on SWE-bench**, due in large part to the challenges of visual problem solving and JavaScript’s diverse development practices. The performance results on SWE-bench M appear to vary depending on the diverse types of images and challenges presented. Different visual elements, such as code snippets, website screenshots, and diagrams, seem to require distinct comprehension abilities. Furthermore, JavaScript’s support for object oriented, functional, and procedural programming introduces substantial variance in how codebase and file structures, which standardized solutions struggle with.

Our efforts to adapt existing baselines highlight *generalizability* as a desirable but overlooked consideration for building LM systems. Current systems evaluated on SWE-bench rely heavily on Python-only parsers to perform fault localization steps independent of an LM (Örwall, 2024; Xia et al., 2024; Zhang et al., 2024b). These approaches either force system builders to re-engineer tools for other programming languages or fail entirely when similar tools are not available. Among our takeaways, we offer suggestions for building agent systems that operate efficiently in numerous programming languages and repositories with visual content.

2 SWE-BENCH MULTIMODAL

We first review SWE-bench’s task formulation and limitations (Section 2.1) and describe how SWE-bench M builds upon this work. We then discuss SWE-bench M’s data collection heuristics (Section 2.2). Finally, we thoroughly characterize SWE-bench M (Section 2.3), highlighting the novel challenges it poses to LM agent-based software development.

Table 1: **Comparison of repository-level coding benchmarks.** We characterize benchmarks by their number of repositories (**# Repos**), programming languages (**Lang.**), whether they employ execution-based evaluation (**Exec.**), and include tasks with image content (**Images**). Programming languages for front-end apps and visual assets are novel attributes introduced by SWE-bench M.

Dataset	# Repos	Lang.	Exec.	Images
RepoEval (Zhang et al., 2023b)	14	Python	✗	✗
RepoBench (Liu et al., 2023b)	10 K+	Python, Java	✗	✗
SWE-bench (Jimenez et al., 2024a)	17	Python	✓	✗
SWE-bench Multimodal	17	JavaScript	✓	✓

2.1 PRELIMINARIES

Formulations. SWE-bench (Jimenez et al., 2024a) has emerged as a popular LM agent benchmark. By drawing from real GitHub workflows, task instances reflect diverse, practical software challenges and require systems to ingest meaningful, long-form inputs. Repository-level coding also involves meticulous refactoring of interwoven modules, each with dependency trees spanning multiple source files. SWE-bench’s collection strategy yields human-written unit tests that ensure robust evaluation.

SWE-bench consists of 2,294 task instances derived from pull requests (PRs) collected across 12 open source Python repositories. Each task instance corresponds to a PR and one or more resolved *issues*; issue(s) describe a bug or feature request, and the PR contains the corresponding solution code along with unit tests validating its correctness. The unit tests fail before the solution code is applied to the codebase but pass afterwards, also referred to as *fail-to-pass* (F2P) tests. Additional auxiliary *pass-to-pass* (P2P) tests verify that the codebase’s existing functionality is maintained.

A task worker is shown the codebase and the issue, also called the *problem statement*. The worker must then modify the codebase to solve the problem. The proposed modification is run against both F2P and P2P tests to check if (1) the issue is fixed and (2) prior working behavior is not broken. If all unit tests pass, the task instance is considered resolved.

Limitations. Within the SWE-bench task formulation, several facets of software development remain unexplored. This work primarily investigates two such facets. First, SWE-bench task instances are predominantly *text only*, and there are no discussions of the implications of the interplay between images and videos with software development. For the 5.6% of SWE-bench task instances with an image, it is unclear what these images portray and whether they are necessary to solving the task. To fill this gap, we focus on task instances with visual elements and demonstrate the significance of multi-modal reasoning in software engineering evaluations. Second, SWE-bench is *comprised exclusively of Python repositories*. We show that a benchmark containing SWE-bench-style task instances from an alternative programming language, i.e., JavaScript, highlights previously unconsidered complexities (e.g., web development, asynchronous programming, DOM/state manipulation).

2.2 COLLECTION

JavaScript repositories have a high concentration of visual assets due to its popularity for full stack web development and browser manipulation. We summarize our modifications to SWE-bench’s task collection pipeline to identify task instances with visual components. See Appendix B for details.

1. Find user-facing JavaScript repositories. Using GitHub’s search feature, we look for JavaScript repositories with 5,000 or more stars and 500 or more pull requests. We then manually pick 17 repositories, filtering for user-facing libraries that have a visual aspect, such as mapping, plotting, or syntax highlighting. We scrape 135k PRs from these selected repositories. Every repository’s source code is at least 70% JavaScript or TypeScript code, with the remainder usually being HTML or Cascading Style Sheets (CSS). None of the selected repositories have Python code.

2. Filter for issues that have visuals in their description or unit tests. We first look for PRs with one or more issues. We then filter these [issue(s), PR] pairs down to ones with visual assets in either the issue or testing code. Specifically, we inspect the issue text and test patch for working hyperlinks

162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215

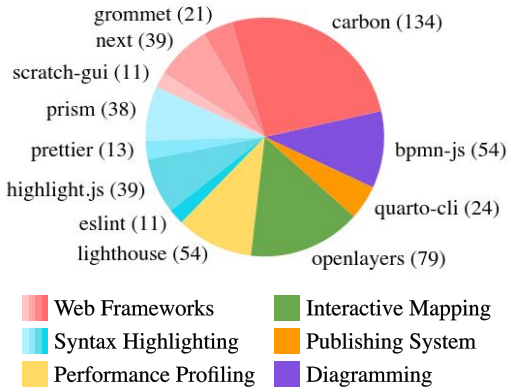


Figure 2: SWE-bench M Test Split Distribution of 517 task instances from 12 open source GitHub repositories written mainly in JavaScript.

Table 2: Median values of different attributes of a SWE-bench M task instance.

		Median
Issue Text	Length (Words)	105
Codebase	# Lines (non-test)	549K
	# Files (non-test)	1799
Gold Patch	# Lines edited	27
	# Files edited	2
	# Funcs edited	3
Tests	# Fail to Pass	1
	# Pass to Pass	5
Images	Aspect Ratio	5:3
	File Size (KB)	42.95
	Resolution (Pixels)	262K

that point to an image (e.g., `jpg`, `png`) or video (e.g., `gif`, `mov`). From 135k PRs, this new filtering criteria yields 1,478 candidate task instances.

3. Set up the environment for each repository. SWE-bench’s Docker containers do not support JavaScript out of the box. Therefore, we first add foundational infrastructure, such as Node.js and Chrome, to support JavaScript execution, visual testing, and webpage rendering in-browser. Next, for each repository, we read its public contribution guidelines and then write tailored installation and testing scripts. This process requires several rounds of manual trial and error not only to validate the setup procedure, but also to overcome undocumented obstacles. Designing environments suited to run the majority of task instances took an average of ten hours of manual labor per repository. We are able to install the repository and execute tests for 679 of the 1478 task instances.

4. Remove inconsistent tests. During manual inspection, we found that a small percentage of tests exhibited inconsistency, i.e., given the same patch, a test’s output (*pass* or *fail*) was inconsistent across multiple evaluation runs. This phenomenon was also reported in SWE-bench (Brown et al., 2024). To eliminate such cases, we run validation per task instance 10 times and remove tests with inconsistent results, reducing 679 candidates to 643 viable instances.

5. Human validation. To remove impossible task instances and perform dataset characterization that is difficult to conduct automatically, the authors manually inspect each task instance. We identify the kinds of image(s) provided and assess how necessary image(s) are to solving a task. We include the annotation procedure and queries in Appendix D. Annotators removed 24 impossible tasks for a final total of 619 task instances. We discuss takeaways from annotations in Section 2.3.

2.3 FEATURES

SWE-bench M is a dataset of 619 multimodal, JavaScript task instances collected from 17 open-source GitHub repositories. It contains a diverse set of libraries, including ones for data visualization, building diagrams, website UI components, displaying maps, and syntax highlighting. As shown in Figure 2, the test split consists of 517 task instances from 12 repositories. The development split contains 102 task instances from 5 repositories.

Diversity of images. Across all SWE-bench M task instances, there are 862 “problem statement” images, meaning the image hyperlink is in the issue text. These images feature a multitude of visual processing challenges. Our annotation procedure grouped these images into seven distinct categories. Interpreting UI issues in *website screenshots* (401 images), such as malformed front-end layouts or accessibility issues, involves spotting design flaws and assessing spatial relationships between elements. Grounding visual cues from *code screenshots* (194) and *error messages* (54) to specific codebase entities can provide hints for bug localization. Images of *diagrams* (107), *art* (38), *maps* (35), and *data visualizations* (28) depict repository-specific challenges that address properly generating complex statistical, geographic, and creative content.

Videos, actual/expected reference images. SWE-bench M task instances often have multiple pictures (221 instances) or videos (70) to succinctly communicate helpful context and reproduction steps. For instance, `gifs` illustrate issues for 43 instances from `bpmn-js`, a diagramming toolkit. Unexpected behaviors are easily depicted by recording diagram interactions, drag-and-drop actions, and zooming in/out. 67 instances from `carbon`, a UI design system, contain “actual” and “expected” image pairs that highlight visual discrepancies and incorrect implementations.

Visual testing. A subset of 69 SWE-bench M task instances test submissions’ functional correctness with pixel-level visual testing. This approach renders web pages and then compares screenshots pixel-by-pixel, identifying visual differences that functional unit tests cannot detect.

Necessity of images. To gauge the importance of problem statements’ images, we pose two questions to human annotators. First, *does an image convey more information than the text in the image?* For instance, a syntax-highlighting bug screenshot would show several lines of code, but the text coloring is an important visual element. Annotators answer “Yes” to this question for 80% of problem statement images. Second, *is it feasible to solve a task instance without its associated image(s)?* Of 557 instances with image(s), 83.5% were considered crucial for resolving the corresponding task. When provided, images are essential for diagnosing software bugs, and the information they convey is necessarily visual. See Appendices D.2 and D.3 for details.

Difficulty curve. Human annotators estimated the time it would take for a developer to solve the task instances in the benchmark. This analysis shows that SWE-bench M features a range of difficulty levels, with tasks classified as follows: 13% take less than 15 minutes, 43% take 15 minutes to an hour, 38% take 1-4 hours, and 6% take more than 4 hours. Comparing these ratios to a similar study performed on SWE-bench (Chowdhury et al., 2024) shows that SWE-bench M tasks are longer on average overall. Reference solutions in SWE-bench M also edit more files, functions, and lines than in SWE-bench. See Appendices A.3 and D.4 for details.

3 EVALUATING ON SWE-BENCH M

We now discuss the challenge of generalization for existing open-source solutions for SWE-bench and describe how we have adapted these methods for evaluation on SWE-bench M (Section 3.1). Additionally, we provide details about our the setup of our experiments (Section 3.2).

3.1 DO EXISTING SYSTEMS GENERALIZE?

We attempted to run existing open-source systems that perform well on SWE-bench (Jimenez et al., 2024b) on SWE-bench M. However, when adapting these systems to SWE-bench M, we found that several solutions were so heavily tailored to Python and SWE-bench to the point that they were unusable for JavaScript repositories and SWE-bench M evaluation.

Our observations from exploring a number of existing solutions led us to identify *generalizability* as a desirable but overlooked property of automated software engineering systems. Specifically, do existing agent systems perform well on bugs that are not similar to the ones found in SWE-bench (i.e., non-Python or issues that include images)?

Though current LMs perform well in multiple programming languages (Cassano et al., 2022), a rigidly defined system can force LMs to follow a particular problem solving pipeline that restricts its capabilities. Although such workflow-oriented systems may work well for a specific type of repository or even programming language, they can fail when confronted with minor distribution shifts outside their original design parameters. Simply put, the shift from Python to JavaScript or the addition of the image modality exceeds the abilities of many existing approaches.

We attempt to adapt the top four open-source systems on the SWE-bench leaderboard (Jimenez et al., 2024b) to work on SWE-bench M tasks. For each system, we discuss below whether adaptations were feasible based on the preceding questions and the changes we made, if applicable.

SWE-agent (Yang et al., 2024) is a lightweight framework connecting an LM to an operating system and shell process. It integrates a text-based agent-computer interface (ACI) for LMs to edit files in addition to the ability to execute shell commands. The primary modifications necessary for SWE-bench was adapting the message processing and history functionality to work with images

for multimodal LMs. To explore the effects of further adaptation, we evaluate three ACI configurations for SWE-agent: (1) SWE-agent Base, the original SWE-agent ACI, initially designed for evaluation on Python and SWE-bench tasks; (2) SWE-agent JS, which converts SWE-agent Base’s editing interface to detect JavaScript edit errors, mirroring the linting feature used by SWE-agent for Python files; and (3) SWE-agent M, which extends the SWE-agent JS ACI to provide a simple web browser, screenshot, and image viewing capabilities, allowing models to visually reproduce image-based issues and verify changes or generated images.

Agentless (Xia et al., 2024) proposes a two stage localize-then-repair pipeline. A repository is first pre-processed with the Python `ast` module into a simplified overview of the repository’s structure. Based on this syntax tree, the LM can then view files’ content and select which files to fix. In the repair stage, based on the identified files, the LM then generates multiple candidate patches. The most commonly repeated solution among the candidates is its final submission.

The main adaptation needed is swapping out the Python `ast` module with a custom JavaScript parser written from scratch. Based on feedback from one of the authors, we also update Python-centric prompts with more JavaScript-oriented instructions. Without these changes, Agentless gets 0% resolved on SWE-bench M’s development split. We denote our adapted version Agentless JS.

AutoCodeRover (Zhang et al., 2024b) employs a two-phase approach similar to Agentless. The first phase involves a customized retrieval process where an LM performs multiple iterative searches. This is followed by a patch generation and testing phase, where an LM, provided with the search history and results, generates and refines a code patch using results from executing tests.

We found that most tools AutoCodeRover provides to the agent rely heavily on Python-specific program analysis features and even prior knowledge of the particular repositories included in SWE-bench. We ultimately chose not to benchmark AutoCodeRover since its specialized tools would require extensive redesign for SWE-bench M, likely resulting in a fundamentally different system.

Moatless (Örwall, 2024) subscribes to the localize-then-repair workflow, but does not use an LM for localization. Instead, code files are first converted in abstract syntax trees (AST). These ASTs are then aggregated into a single, searchable code graph represented as a Faiss index (Johnson et al., 2019). The LM then queries this index to generate a repair.

We worked off the author’s in-progress implementation for generating ASTs from JavaScript and TypeScript files¹. Although the code graph representation is language agnostic, the input AST representation needed to generate the Faiss index is heavily based on Python’s object oriented design. As reflected by the magnitude of the pull request’s changes, writing an equivalent parser for JavaScript that (1) reflects the code and (2) subscribes to the index’s input format is non-trivial due to usage of functional and declarative programs being more commonplace. We do not benchmark Moatless.

RAG. Jimenez et al. (2024a) proposed a retrieval augmented baseline using BM25 (Robertson & Zaragoza, 2009) from which we inherit the document formatting, retrieval method, and prompt structure. We adapt the prompt template used in Jimenez et al. (2024a) to use a JavaScript example for patch generation instead of Python. We also include a new section in the prompt for inserting reproduction code collected from links in the problem statement, as described in B.3.

3.2 EXPERIMENT SETUP

Models. Resolving SWE-bench M issues with existing systems requires handling very long contexts, processing images and text simultaneously, and producing sophisticated structured outputs. Given these constraints, we focus all of our evaluations on GPT-4o (`gpt-4o-2024-08-06`) (OpenAI, 2024) and Claude 3.5 Sonnet (`claude-3-5-sonnet-20240620`) (Anthropic, 2024), the two most well-supported mulitmodal LMs for long-context RAG and agent systems. Though alternative models continue to improve, they currently lack the combination of long-context handling, mulitmodal processing, and structured prediction abilities required to work with existing autonomous software engineering systems.

Baselines. Based on our Section 3.1 findings, we benchmark the performance of five baselines: Retrieval Augmented Generation (RAG), SWE-agent (Base, JS, M), and Agentless JS. Each of our baseline systems were adapted and developed using the development split of SWE-bench M. We

¹<https://github.com/aorwall/moatless-tools/pull/34>

Table 3: **Performance comparison of various baselines on SWE-bench M.** The table shows results for different software development agent frameworks, including SWE-agent (with multimodal and JavaScript-specific variations) and a retrieval augmented generation (RAG) approach. Each system’s success rate (% Resolved) and average cost (\$ Avg. Cost) per task are reported.

System	Model	% Resolved	\$ Avg. Cost
SWE-agent M	GPT-4o	12.2	2.94
	Claude 3.5 Sonnet	11.4	3.11
SWE-agent JS	GPT-4o	9.2	0.99
	Claude 3.5 Sonnet	12.0	3.11
SWE-agent Base	GPT-4o	12.0	2.07
	Claude 3.5 Sonnet	12.2	1.52
Agentless JS	GPT-4o	3.1	0.38
	Claude 3.5 Sonnet	6.2	0.42
RAG	GPT-4o	6.0	0.17
	Claude 3.5 Sonnet	5.0	0.15

describe more about how the final test configurations, including hyperparameters, were selected in Appendices C.1 and C.2.

Metrics. We report two evaluation metrics: (1) **% Resolved**, the main performance metric, which represents the proportion of successfully resolved task instances and (2) **Avg. \$ Cost**, the mean per-instance inference cost incurred from running the baseline, as used in Yang et al. (2024).

4 RESULTS

We compare the performance of each baseline system in Table 3. While overall performance on SWE-bench M is relatively low, we observe a substantial performance gap between the interactive SWE-agent systems (11.5 % resolved on average) and the Agentless and RAG baselines, which achieve 3.9 and 5.5 % resolved on average respectively.

Across all SWE-agent configurations, we observe similar absolute performance, suggesting that the JavaScript-specific customizations in SWE-agent JS and SWE-agent M had minimal impact on performance. Additionally, we also observe minimal performance differences when swapping the underlying LM (GPT-4o or Claude 3.5 Sonnet) within a system. However, a study of SWE-agent ablations on the development set reveals that the added multimodal tooling can improve agent performance in some cases, though the overall picture remains ambiguous.

An analysis of solution rates segmented by the date each task instance was originally solved, is performed in Appendix C.3. It finds no evidence for a test set advantage from solutions leaking into the LM training sets: we show for example, that SWE-agent M using GPT-4o performs better on task instances based on issues that were resolved *after* its knowledge cutoff date.

4.1 ANALYSIS

We present insights into how images influence system performance, evaluate the impact of multimodal actions in SWE-agent M, and explore the generalization of different approaches to automated software engineering. All analyses are conducted on the development set of SWE-bench M.

Resolving SWE-bench M issues requires improved visual understanding. Table 4 compares the performance of the RAG and SWE-agent JS systems, with and without images. The results reveal a significant performance decrease when models lack access to visual information. Table 5 analyzes performance based on the two key questions from our human validation: “Do the images contain more information than just text?” and “Are the images necessary to solve this task?” (For detail, see Appendices D.2 and D.3, respectively.) As expected, system performance generally improves when images are provided with the problem statement. In addition, for instances where annotators deemed images necessary to solve the task, performance plummets without visual input. Performance changes vary based on image content. When images contain primarily text, SWE-

agent JS maintains 23.1% performance with or without images. Larger performance drops occur when images contain non-textual elements. For non-textual visuals, performance drops from 13.0% to 8.7% without images for SWE-agent JS. Visual elements, especially those that are non-textual, appear to provide additional context that substantially improves system performance.

Table 4: **Performance of RAG and SWE-agent JS with (W/) and without (W/o) Images.** The consistent improvement across both models and systems emphasizes the critical role of multimodal inputs, particularly for complex tasks requiring visual reasoning.

System	Model	W/ Images	W/o Images
SWE-agent JS	GPT-4o	11.0	8.0
	Claude Sonnet 3.5	16.0	13.0
RAG	GPT-4o	10.0	8.0
	Claude Sonnet 3.5	14.1	11.2

Table 5: **System performance segmented by annotator response.** We show performance for SWE-agent JS and RAG using Claude 3.5 Sonnet broken down by response to human validation questions detailed in Appendix D.2 and D.3. Each performance value corresponds to the performance of the configuration (e.g. SWE-agent JS with images) on the subset of instances corresponding to the annotation class for each question. The questions in the tables have been simplified from the original for readability.

Does the image contain more information than just text?				Are the images necessary to solve this task?			
System	With Images?	Yes	No	System	With Images?	Yes	No
SWE-agent JS	✓	13.0	23.1	SWE-agent JS	✓	17.6	11.1
	✗	8.7	23.1		✗	8.8	22.2
RAG	✓	11.6	13.8	RAG	✓	12.4	11.9
	✗	6.1	16.2		✗	9.1	8.1

Localization modules are overly engineered for Python. As discussed in Section 3.1, except for SWE-agent, the systems that we study (Agentless, Moatless, and AutoCodeRover) impose fixed, procedural problem-solving workflows. Every system starts with a bug localization step that relies on abstract syntax tree (AST) parsing libraries to identify programmatic symbols.

These Python-specific modules do not work for other languages. Overcoming this design is contingent upon either finding or creating a similarly reliable tool. For Moatless and AutoCodeRover, no such replacement tool exists; reimplementation was prohibitively laborious and would have resulted in a system with little resemblance to the original. For Agentless, replacing invocations of the Python `ast` package with the JavaScript `tree-sitter` library still yielded a 0% resolve rate on the development split; this is due to Python-centric assumptions the localization module makes about the AST representations. Therefore, we wrote a custom JavaScript parser from scratch compatible with the Agentless logic, which took 15 hours of labor.

However, because the original design of Agentless is influenced by Python, fundamental differences between Python and JavaScript’s design patterns (e.g., multiple programming styles, object prototypes, arrow functions) lead to many failure modes during localization, resulting in Agentless JS’s low performance (4.6% in Table 3). For instance, for `grommet_grommet-6749`, the localization module, designed based on Python’s *imperative* and object-oriented style, does not recognize `Tab`, the buggy component, because it was defined *declaratively* as an arrow function. In the subsequent repair phase, Agentless JS writes a completely new `Tab` imperatively, while the reference solution makes a minor change to its existing declarative implementation. We also found that given the raw code, GPT-4o correctly identified `Tab` as the buggy component.

We calculate the F1 Score for file localization on the development split with Claude Sonnet 3.5 as the base LM. Agentless JS achieves a score of 0.142, while SWE-agent gets 0.367. Flaws with

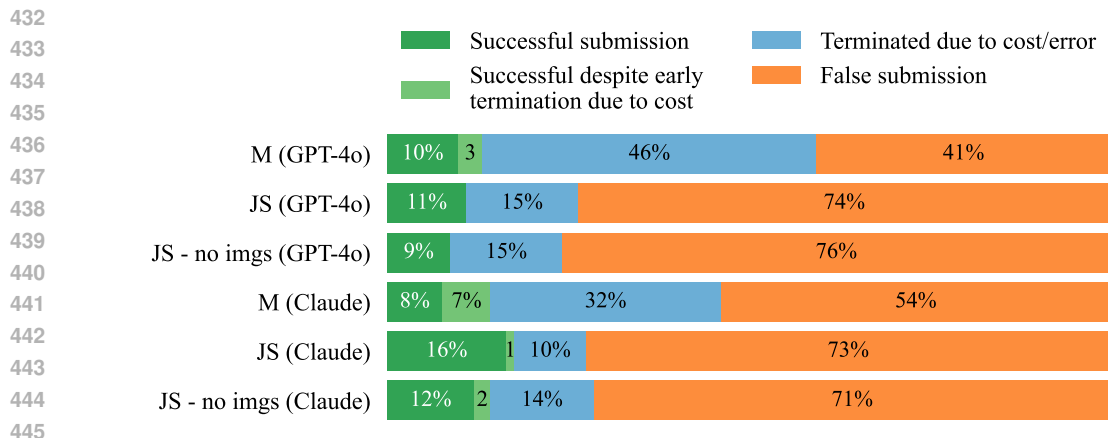


Figure 3: **SWE-agent performance across configurations.** The frequency of outcomes by success and exit status of SWE-agent under different configurations on the development set. Configurations include multimodal (M), JavaScript-specific (JS), and no images, for both GPT-4o and Claude 3.5 Sonnet models. A similar analysis by repository is shown in Figure 8.

language-specific design are only amplified in SWE-bench M. Unlike SWE-bench solutions, which edit only Python files, 28% of reference solutions in SWE-bench M edit multiple files types (e.g., TypeScript, HTML, CSS, Lua). Adapting current solutions requires writing unique AST generator and navigator functions to pre-process code files for many languages.

From these observations, *we hypothesize that generalizable LM-based systems for software engineering will emphasize interaction, not problem solving.* Instead of inserting LMs at specific points of a workflow, our results show the superiority of an LM-first approach that focuses on building tools that enhance an LM’s ability to navigate and manipulate an environment. SWE-agent’s superior performance and Agentless’s localization struggles suggest that *the burden of problem solving should remain on the LM instead of being offloaded to manually engineered pipelines.*

Multimodal tools increase task complexity but can improve agent performance. To evaluate the impact of the presence of images on agent behavior, we conduct an analysis of the SWE-agent trajectories generated while resolving the development set. We run SWE-agent in three different configurations: SWE-agent M, SWE-agent JS, and SWE-agent JS without image inputs. The web-specific tools introduced in SWE-agent M lead to approximately 20% of actions being dedicated to building and taking screenshots of websites, though this proportion varies significantly across repositories (see Figures 6 and 7). SWE-agent M with GPT-4o builds websites and takes screenshots for 38.3% of the instances. For these instances, an average of 7.5 screenshots are taken, showing that the multimodal tooling is used as part of an iterative problem-solving process for some instances.

However, reproducing issues and verifying solutions by building and taking screenshots of websites makes the task for SWE-agent M more complex than directly performing and submitting edits as in SWE-agent JS. This is reflected in an almost threefold increase of attempts that are terminated because they exceed their cost limit (see Figure 3). The web-specific tools lead to degraded performance with Claude 3.5 Sonnet, while GPT-4o improves, even though some of its successes are from runs that are terminated prematurely due to costs.

Another key figure for agent performance is the *fraction of correct submissions*, as incorrect submissions increase manual reviewing workload. To increase this fraction, submissions from SWE-agent runs that were terminated prematurely due to cost can be excluded, as they have a lower probability to be correct. In this setup, *adding multimodal tools almost doubles the fraction of correct submissions with GPT-4o to 19.6% compared to 10.4% without image inputs.* This result is not observed in the Claude 3.5 Sonnet-based SWE-agent.

Together, these findings show that LM agents can use web-based multimodal tools to reproduce issues and verify solutions, in some cases increasing success rates or solution correctness. However,

486 using these tools to their full potential requires complex workflows that are still challenging for
487 agents and result in an ambiguous record overall.
488

489 5 RELATED WORK 490

491 **Multimodal Code Benchmarks.** Code generation, the task of synthesizing code from natural lan-
492 guage, has served as a long-standing measure of LM performance (Austin et al., 2021; Chen et al.,
493 2021; Hendrycks et al., 2021). As performance on these benchmarks has plateaued—Claude Sonnet
494 3.5 achieves 92% on HumanEval—subsequent works have extended this task along different di-
495 mensions to increase difficulty, such as more robust evaluation (Liu et al., 2023a; Jain et al., 2024),
496 multilinguality (Cassano et al., 2022; Wang et al., 2023a; Zheng et al., 2024), open-domain gener-
497 ation (Wang et al., 2023b; Zhuo et al., 2024), data science (Lai et al., 2022; Yin et al., 2022; Cao
498 et al., 2024), understanding execution traces (Gu et al., 2024; Muennighoff et al., 2024), repository
499 comprehension (Liu et al., 2023b; Zhang et al., 2023b), cybersecurity (Yang et al., 2023b; Zhang
500 et al., 2024a; Shao et al., 2024; Abramovich et al., 2024), and efficiency (Huang et al., 2024; Liu
501 et al., 2024; Waghjale et al., 2024). SWE-bench M represents a merger of two promising directions:
502 software engineering (Jimenez et al., 2024a) and multimodal code generation (Li et al., 2024; Si
503 et al., 2024; Wu et al., 2024; Nishina & Matsui, 2024). By drawing on real-world problems from
504 GitHub, SWE-bench M overcomes the synthetic and short-form limitations inherent in code gener-
505 ation task formulation. In addition, SWE-bench M represents a significant multimodal upgrade to
506 SWE-bench, which has already become a popular benchmark for LM evaluation.

507 **LM Agents for Web and Code.** Prior to the proliferation of LMs, earlier works explored translating
508 between user interfaces and front-end code (HTML/CSS/JavaScript) (Beltramelli, 2018; Robinson,
509 2019). More recently, LM agents (Yao et al., 2023b; Sumers et al., 2024) have been increasingly
510 employed and evaluated on various tasks (Wang et al., 2024a; Xie et al., 2024). Two popular use
511 cases, web or application navigation (Hong et al., 2023; Soselia et al., 2023; Yan et al., 2023; Yao
512 et al., 2023a; Zhang et al., 2023a; Koh et al., 2024; Putta et al., 2024; Rawles et al., 2024; Press
513 et al., 2024; Yoran et al., 2024) and software engineering (Yang et al., 2023a; 2024; Zhang et al.,
514 2024b; Xia et al., 2024), have typically been studied separately. To the best of our knowledge, SWE-
515 bench M is the first benchmark that meaningfully couples these two tasks. Though prior approaches
516 have provided web browsing tools to programming agents (Wang et al., 2024b), such tools usually
517 represent webpages in text form and were not designed with a clear downstream objective. SWE-
518 bench M demonstrates how agents can iterate meaningfully between code updates and their effects
519 rendered as images or in a browser.

520 6 CONCLUSION 521

522 This work introduces SWE-bench Multimodal (SWE-bench M), the first benchmark to evaluate
523 coding agents on real-world software engineering tasks involving visual elements. SWE-bench M
524 contains 619 task instances from 17 user-facing JavaScript repositories, including ones for web
525 user interface design, data visualization, art and mapping. Our analysis reveals that SWE-bench
526 M contains diverse visual challenges and increases task complexity compared to SWE-bench. Fur-
527 thermore, existing systems perform poorly on it, with the top resolve rate reaching only 12.2%;
528 SWE-bench M presents repository-level programming tasks that are challenging even for state-of-
529 the-art systems built on top of the strongest LMs. Incorporating multimodality in SWE-bench M
530 not only expands the coverage of exciting, practical challenges in software engineering, but it also
531 encourages practitioners to develop more general-purpose, language-agnostic solutions that do not
532 overfit to SWE-bench or Python repositories.
533
534
535
536
537
538
539

REFERENCES

- 540
541
542 Talor Abramovich, Meet Udeshi, Minghao Shao, Kilian Lieret, Haoran Xi, Kimberly Milner, Sofija
543 Jancheska, John Yang, Carlos E. Jimenez, Farshad Khorrani, Prashanth Krishnamurthy, Brendan
544 Dolan-Gavitt, Muhammad Shafique, Karthik Narasimhan, Ramesh Karri, and Ofir Press. Enigma:
545 Enhanced interactive generative model agent for ctf challenges, 2024. URL <https://arxiv.org/abs/2409.16165>.
546
- 547 Anthropic. Introducing claude 3.5 sonnet, 2024. URL <https://www.anthropic.com/news/claude-3-5-sonnet>.
548
- 549 Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan,
550 Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and Charles Sutton. Program synthesis with large
551 language models, 2021. URL <https://arxiv.org/abs/2108.07732>.
552
- 553 Kaustubh Bagalkote. Amazon ceo andy jassy says company’s ai assistant has saved \$260m and 4.5k
554 developer-years of work, 2024. URL <https://finance.yahoo.com/news/amazon-ceo-andy-jassy-says-213018283.html>.
555
- 556 Tony Beltramelli. pix2code: Generating code from a graphical user interface screenshot. In *Pro-*
557 *ceedings of the ACM SIGCHI symposium on engineering interactive computing systems*, pp. 1–6,
558 2018.
559
- 560 Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V. Le, Christopher Ré, and
561 Azalia Mirhoseini. Large language monkeys: Scaling inference compute with repeated sampling,
562 2024. URL <https://arxiv.org/abs/2407.21787>.
- 563 Ruisheng Cao, Fangyu Lei, Haoyuan Wu, Jixuan Chen, Yeqiao Fu, Hongcheng Gao, Xinzhuang
564 Xiong, Hanchong Zhang, Yuchen Mao, Wenjing Hu, Tianbao Xie, Hongshen Xu, Danyang
565 Zhang, Sida Wang, Ruoxi Sun, Pengcheng Yin, Caiming Xiong, Ansong Ni, Qian Liu, Victor
566 Zhong, Lu Chen, Kai Yu, and Tao Yu. Spider2-v: How far are multimodal agents from
567 automating data science and engineering workflows? *CoRR*, abs/2407.10956, 2024. URL
568 <https://arxiv.org/abs/2407.10956>.
- 569 Federico Cassano, John Gouwar, Daniel Nguyen, Sydney Nguyen, Luna Phipps-Costin, Donald
570 Pinckney, Ming-Ho Yee, Yangtian Zi, Carolyn Jane Anderson, Molly Q Feldman, Arjun Guha,
571 Michael Greenberg, and Abhinav Jangda. Multipl-e: A scalable and extensible approach to bench-
572 marking neural code generation, 2022. URL <https://arxiv.org/abs/2208.08227>.
573
- 574 Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared
575 Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri,
576 Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan,
577 Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian,
578 Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fo-
579 tios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex
580 Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders,
581 Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec
582 Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob Mc-
583 Grew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large
584 language models trained on code, 2021. URL <https://arxiv.org/abs/2107.03374>.
585
- 586 Mauro Cherubini, Gina Venolia, Rob DeLine, and Amy J Ko. Let’s go to the whiteboard: how
587 and why software developers use drawings. In *Proceedings of the SIGCHI conference on Human*
588 *factors in computing systems*, pp. 557–566, 2007.
589
- 590 Neil Chowdhury, James Aung, Chan Jun Shern, Oliver Jaffe, Dane Sherburn, Giulio Starace, Evan
591 Mays, Rachel Dias, Marwan Aljubei, Mia Glaese, Carlos E. Jimenez, John Yang, Kevin Liu, and
592 Aleksander Madry. Introducing swe-bench verified, 2024. URL [https://openai.com/i](https://openai.com/index/introducing-swe-bench-verified/)
593 [ndex/introducing-swe-bench-verified/](https://openai.com/index/introducing-swe-bench-verified/).
- 594 Kyle Daigle and GitHub. Octoverse: The state of open source and rise of ai in 2023, 2023. URL
595 [https://github.blog/news-insights/research/the-state-of-open-sou](https://github.blog/news-insights/research/the-state-of-open-source-and-ai/#the-most-popular-programming-languages)
596 [rce-and-ai/#the-most-popular-programming-languages](https://github.blog/news-insights/research/the-state-of-open-source-and-ai/#the-most-popular-programming-languages).

- 594 James D Foley and Andries Van Dam. *Fundamentals of interactive computer graphics*. Addison-
595 Wesley Longman Publishing Co., Inc., 1982.
- 596
- 597 Paul Gauthier. Aider, 2024. URL <https://github.com/paul-gauthier/aider>.
- 598
- 599 Alex Gu, Baptiste Rozière, Hugh Leather, Armando Solar-Lezama, Gabriel Synnaeve, and Sida I.
600 Wang. Cruxeval: A benchmark for code reasoning, understanding and execution, 2024. URL
601 <https://arxiv.org/abs/2401.03065>.
- 602
- 603 Dan Hendrycks, Steven Basart, Saurav Kadavath, Mantas Mazeika, Akul Arora, Ethan Guo, Collin
604 Burns, Samir Puranik, Horace He, Dawn Song, and Jacob Steinhardt. Measuring coding challenge
605 competence with apps, 2021. URL <https://arxiv.org/abs/2105.09938>.
- 606
- 607 Wenyi Hong, Weihang Wang, Qingsong Lv, Jiazhen Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan
608 Wang, Yuxuan Zhang, Juanzi Li, Bin Xu, Yuxiao Dong, Ming Ding, and Jie Tang. Cogagent: A
609 visual language model for gui agents, 2023. URL <https://arxiv.org/abs/2312.08914>.
- 610
- 611 Dong Huang, Yuhao Qing, Weiyi Shang, Heming Cui, and Jie M. Zhang. Effibench: Benchmarking
612 the efficiency of automatically generated code, 2024. URL <https://arxiv.org/abs/2402.02037>.
- 613
- 614 Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando
615 Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free
616 evaluation of large language models for code, 2024. URL <https://arxiv.org/abs/2403.07974>.
- 617
- 618 Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik
619 Narasimhan. Swe-bench: Can language models resolve real-world github issues?, 2024a. URL
620 <https://arxiv.org/abs/2310.06770>.
- 621
- 622 Carlos E. Jimenez, John Yang, Alexander Wettig, Shunyu Yao, Kexin Pei, Ofir Press, and Karthik R.
623 Narasimhan. Swe-bench leaderboard, 2024b. URL <https://www.swebench.com/>.
- 624
- 625 Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with GPUs. *IEEE
626 Transactions on Big Data*, 7(3):535–547, 2019.
- 627
- 628 Jing Yu Koh, Robert Lo, Lawrence Jang, Vikram Duvvur, Ming Chong Lim, Po-Yu Huang, Graham
629 Neubig, Shuyan Zhou, Ruslan Salakhutdinov, and Daniel Fried. Visualwebarena: Evaluating
630 multimodal agents on realistic visual web tasks. *ACL*, 2024.
- 631
- 632 Yuriy Kotsarenko. Measuring perceived color difference using yiq ntsc transmission color space in
633 mobile applications. *Programacion Matematica y Software*, 2018.
- 634
- 635 Yuhang Lai, Chengxi Li, Yiming Wang, Tianyi Zhang, Ruiqi Zhong, Luke Zettlemoyer, Scott Wen
636 tau Yih, Daniel Fried, Sida Wang, and Tao Yu. Ds-1000: A natural and reliable benchmark for
637 data science code generation, 2022. URL <https://arxiv.org/abs/2211.11501>.
- 638
- 639 Kaixin Li, Yuchen Tian, Qisheng Hu, Ziyang Luo, and Jing Ma. Mmcode: Evaluating multi-modal
640 code large language models with visually rich programming problems, 2024. URL <https://arxiv.org/abs/2404.09486>.
- 641
- 642 Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by
643 chatgpt really correct? rigorous evaluation of large language models for code generation, 2023a.
644 URL <https://arxiv.org/abs/2305.01210>.
- 645
- 646 Jiawei Liu, Songrun Xie, Junhao Wang, Yuxiang Wei, Yifeng Ding, and Lingming Zhang. Evaluat-
647 ing language models for efficient code generation, 2024. URL <https://arxiv.org/abs/2408.06450>.
- 648
- 649 Tianyang Liu, Canwen Xu, and Julian McAuley. Repobench: Benchmarking repository-level code
650 auto-completion systems, 2023b. URL <https://arxiv.org/abs/2306.03091>.

- 648 Daniel Moody. The “physics” of notations: toward a scientific basis for constructing visual notations
649 in software engineering. *IEEE Transactions on software engineering*, 35(6):756–779, 2009.
650
- 651 Niklas Muennighoff, Qian Liu, Armel Zebaze, Qinkai Zheng, Binyuan Hui, Terry Yue Zhuo,
652 Swayam Singh, Xiangru Tang, Leandro von Werra, and Shayne Longpre. Octopack: Instruction
653 tuning code large language models, 2024. URL <https://arxiv.org/abs/2308.07124>.
- 654 Kunato Nishina and Yusuke Matsui. Svgeditbench: A benchmark dataset for quantitative assessment
655 of llm’s svg editing capabilities, 2024. URL <https://arxiv.org/abs/2404.13710>.
- 656
- 657 OpenAI. Hello gpt-4o, 2024. URL <https://openai.com/index/hello-gpt-4o/>.
- 658 Albert Örwall. Moatless tools, 2024. URL [https://github.com/aorwall/moatless-t
659 ools](https://github.com/aorwall/moatless-tools).
- 660
- 661 Ori Press, Andreas Hochlehnert, Ameya Prabhu, Vishaal Udandaraao, Ofir Press, and Matthias
662 Bethge. Citeme: Can language models accurately cite scientific claims?, 2024. URL <https://arxiv.org/abs/2407.12861>.
- 663
- 664 Pranav Putta, Edmund Mills, Naman Garg, Sumeet Motwani, Chelsea Finn, Divyansh Garg, and
665 Rafael Rafailov. Agent q: Advanced reasoning and learning for autonomous ai agents, 2024.
666 URL <https://arxiv.org/abs/2408.07199>.
- 667
- 668 Christopher Rawles, Sarah Clinckemaiillie, Yifan Chang, Jonathan Waltz, Gabrielle Lau, Marybeth
669 Fair, Alice Li, William Bishop, Wei Li, Folawiyi Campbell-Ajala, Daniel Toyama, Robert Berry,
670 Divya Tyamagundlu, Timothy Lillicrap, and Oriana Riva. Androidworld: A dynamic benchmark-
671 ing environment for autonomous agents, 2024.
- 672
- 673 Stephen E. Robertson and Hugo Zaragoza. The probabilistic relevance framework: Bm25 and be-
674 yond. *Found. Trends Inf. Retr.*, 3:333–389, 2009. URL [https://api.semanticschola
675 r.org/CorpusID:207178704](https://api.semanticscholar.org/CorpusID:207178704).
- 676
- 677 Alex Robinson. Sketch2code: Generating a website from a paper mockup, 2019. URL <https://arxiv.org/abs/1905.13750>.
- 678
- 679 Minghao Shao, Sofija Jancheska, Meet Udeshi, Brendan Dolan-Gavitt, Haoran Xi, Kimberly Milner,
680 Boyuan Chen, Max Yin, Siddharth Garg, Prashanth Krishnamurthy, Farshad Khorrami, Ramesh
681 Karri, and Muhammad Shafique. Nyu ctf dataset: A scalable open-source benchmark dataset for
682 evaluating llms in offensive security, 2024. URL <https://arxiv.org/abs/2406.05590>.
- 683
- 684 Ben Shneiderman. *Designing the User Interface: Strategies for Effective Human-Computer Inter-*
685 *action*. Pearson, 1998. URL [https://api.semanticscholar.org/CorpusID:
54156244](https://api.semanticscholar.org/CorpusID:54156244).
- 686
- 687 Chenglei Si, Yanzhe Zhang, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. Design2code: How far
688 are we from automating front-end engineering?, 2024. URL [https://arxiv.org/abs/24
03.03163](https://arxiv.org/abs/2403.03163).
- 689
- 690 Davit Soselia, Khalid Saifullah, and Tianyi Zhou. Learning ui-to-code reverse generator using visual
691 critic without rendering, 2023. URL <https://arxiv.org/abs/2305.14637>.
- 692
- 693 Theodore R. Sumers, Shunyu Yao, Karthik Narasimhan, and Thomas L. Griffiths. Cognitive archi-
694 tectures for language agents, 2024. URL <https://arxiv.org/abs/2309.02427>.
- 695
- 696 Siddhant Waghjale, Vishruth Veerendranath, Zora Zhiruo Wang, and Daniel Fried. Ecco: Can we
697 improve model-generated code efficiency without sacrificing functional correctness?, 2024. URL
<https://arxiv.org/abs/2407.14044>.
- 698
- 699 Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai
700 Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Jirong Wen. A survey on large
701 language model based autonomous agents. *Frontiers of Computer Science*, 18(6), March 2024a.
ISSN 2095-2236. doi: 10.1007/s11704-024-40231-1. URL [http://dx.doi.org/10.10
07/s11704-024-40231-1](http://dx.doi.org/10.1007/s11704-024-40231-1).

- 702 Xingyao Wang, Boxuan Li, Yufan Song, Frank F. Xu, Xiangru Tang, Mingchen Zhuge, Jiayi Pan,
703 Yueqi Song, Bowen Li, Jaskirat Singh, Hoang H. Tran, Fuqiang Li, Ren Ma, Mingzhang Zheng,
704 Bill Qian, Yanjun Shao, Niklas Muennighoff, Yizhe Zhang, Binyuan Hui, Junyang Lin, Robert
705 Brennan, Hao Peng, Heng Ji, and Graham Neubig. Opendevin: An open platform for ai software
706 developers as generalist agents, 2024b. URL <https://arxiv.org/abs/2407.16741>.
- 707 Zhiruo Wang, Grace Cuenca, Shuyan Zhou, Frank F. Xu, and Graham Neubig. Mconala: A bench-
708 mark for code generation from multiple natural languages, 2023a. URL <https://arxiv.org/abs/2203.08388>.
- 709
710 Zhiruo Wang, Shuyan Zhou, Daniel Fried, and Graham Neubig. Execution-based evaluation for
711 open-domain code generation, 2023b. URL <https://arxiv.org/abs/2212.10481>.
- 712
713 Erik Wittern, Philippe Suter, and Shriram Rajagopalan. A look at the dynamics of the javascript
714 package ecosystem. In *Proceedings of the 13th International Conference on Mining Software*
715 *Repositories*, MSR '16, pp. 351–361, New York, NY, USA, 2016. Association for Computing
716 Machinery. ISBN 9781450341868. doi: 10.1145/2901739.2901743. URL <https://doi.org/10.1145/2901739.2901743>.
- 717
718 Chengyue Wu, Yixiao Ge, Qiushan Guo, Jiahao Wang, Zhixuan Liang, Zeyu Lu, Ying Shan, and
719 Ping Luo. Plot2code: A comprehensive benchmark for evaluating multi-modal large language
720 models in code generation from scientific plots, 2024. URL <https://arxiv.org/abs/2405.07990>.
- 721
722 Chunqiu Steven Xia, Yinlin Deng, Soren Dunn, and Lingming Zhang. Agentless: Demystifying llm-
723 based software engineering agents, 2024. URL <https://arxiv.org/abs/2407.01489>.
- 724
725 Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing
726 Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio
727 Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. Osworld: Benchmarking multimodal agents
728 for open-ended tasks in real computer environments, 2024. URL <https://arxiv.org/abs/2404.07972>.
- 729
730 An Yan, Zhengyuan Yang, Wanrong Zhu, Kevin Lin, Linjie Li, Jianfeng Wang, Jianwei Yang, Yiwu
731 Zhong, Julian McAuley, Jianfeng Gao, Zicheng Liu, and Lijuan Wang. Gpt-4v in wonderland:
732 Large multimodal models for zero-shot smartphone gui navigation, 2023. URL <https://arxiv.org/abs/2311.07562>.
- 733
734 John Yang, Akshara Prabhakar, Karthik Narasimhan, and Shunyu Yao. Intercode: Standardizing
735 and benchmarking interactive coding with execution feedback, 2023a. URL <https://arxiv.org/abs/2306.14898>.
- 736
737 John Yang, Akshara Prabhakar, Shunyu Yao, Kexin Pei, and Karthik R Narasimhan. Language
738 agents as hackers: Evaluating cybersecurity skills with capture the flag. In *Multi-Agent Security*
739 *Workshop@ NeurIPS'23*, 2023b.
- 740
741 John Yang, Carlos E. Jimenez, Alexander Wettig, Kilian Lieret, Shunyu Yao, Karthik Narasimhan,
742 and Ofir Press. Swe-agent: Agent-computer interfaces enable automated software engineering,
743 2024. URL <https://arxiv.org/abs/2405.15793>.
- 744
745 Shunyu Yao, Howard Chen, John Yang, and Karthik Narasimhan. Webshop: Towards scalable real-
746 world web interaction with grounded language agents, 2023a. URL <https://arxiv.org/abs/2207.01206>.
- 747
748 Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao.
749 React: Synergizing reasoning and acting in language models, 2023b. URL <https://arxiv.org/abs/2210.03629>.
- 750
751 Erin Yepis and StackOverflow. Developers get by with a little help from ai: Stack overflow knows
752 code assistant pulse survey results, 2024. URL <https://stackoverflow.blog/2024/05/29/developers-get-by-with-a-little-help-from-ai-stack-overflow-knows-code-assistant-pulse-survey-results/>.
- 753
754
755

- 756 Pengcheng Yin, Wen-Ding Li, Kefan Xiao, Abhishek Rao, Yeming Wen, Kensen Shi, Joshua
757 Howland, Paige Bailey, Michele Catasta, Henryk Michalewski, Alex Polozov, and Charles Sut-
758 ton. Natural language to code generation in interactive data science notebooks, 2022. URL
759 <https://arxiv.org/abs/2212.09248>.
- 760 Ori Yoran, Samuel Joseph Amouyal, Chaitanya Malaviya, Ben Bogin, Ofir Press, and Jonathan
761 Berant. Assistantbench: Can web agents solve realistic and time-consuming tasks?, 2024. URL
762 <https://arxiv.org/abs/2407.15711>.
- 763
764 Andy K. Zhang, Neil Perry, Riya Dulepet, Eliot Jones, Justin W. Lin, Joey Ji, Celeste Menders,
765 Gashon Hussein, Samantha Liu, Donovan Jasper, Pura Peetathawatchai, Ari Glenn, Vikram
766 Sivashankar, Daniel Zamoshchin, Leo Glikbarg, Derek Askaryar, Mike Yang, Teddy Zhang, Rishi
767 Alluri, Nathan Tran, Rinnara Sangpisit, Polycarpos Yiorkadjis, Kenny Osele, Gautham Raghupathi,
768 Dan Boneh, Daniel E. Ho, and Percy Liang. Cybench: A framework for evaluating cyber-
769 security capabilities and risk of language models, 2024a. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2408.08926)
770 [2408.08926](https://arxiv.org/abs/2408.08926).
- 771 Chi Zhang, Zhao Yang, Jiakuan Liu, Yucheng Han, Xin Chen, Zebiao Huang, Bin Fu, and Gang Yu.
772 Appagent: Multimodal agents as smartphone users, 2023a. URL [https://arxiv.org/abs/](https://arxiv.org/abs/2312.13771)
773 [2312.13771](https://arxiv.org/abs/2312.13771).
- 774
775 Fengji Zhang, Bei Chen, Yue Zhang, Jacky Keung, Jin Liu, Daoguang Zan, Yi Mao, Jian-Guang
776 Lou, and Weizhu Chen. Repocoder: Repository-level code completion through iterative retrieval
777 and generation, 2023b. URL <https://arxiv.org/abs/2303.12570>.
- 778 Yuntong Zhang, Haifeng Ruan, Zhiyu Fan, and Abhik Roychoudhury. Autocoderover: Autonomous
779 program improvement, 2024b. URL <https://arxiv.org/abs/2404.05427>.
- 780
781 Qinkai Zheng, Xiao Xia, Xu Zou, Yuxiao Dong, Shan Wang, Yufei Xue, Zihan Wang, Lei Shen,
782 Andi Wang, Yang Li, Teng Su, Zhilin Yang, and Jie Tang. Codegeex: A pre-trained model for
783 code generation with multilingual benchmarking on humaneval-x, 2024. URL [https://arxi](https://arxiv.org/abs/2303.17568)
784 [v.org/abs/2303.17568](https://arxiv.org/abs/2303.17568).
- 785 Terry Yue Zhuo, Minh Chien Vu, Jenny Chim, Han Hu, Wenhao Yu, Ratnadira Widyasari, Imam
786 Nur Bani Yusuf, Haolan Zhan, Junda He, Indraneil Paul, Simon Brunner, Chen Gong, Thong
787 Hoang, Armel Randy Zebaze, Xiaoheng Hong, Wen-Ding Li, Jean Kaddour, Ming Xu, Zhihan
788 Zhang, Prateek Yadav, Naman Jain, Alex Gu, Zhoujun Cheng, Jiawei Liu, Qian Liu, Zijian Wang,
789 David Lo, Binyuan Hui, Niklas Muennighoff, Daniel Fried, Xiaoning Du, Harm de Vries, and
790 Leandro Von Werra. Bigcodebench: Benchmarking code generation with diverse function calls
791 and complex instructions, 2024. URL <https://arxiv.org/abs/2406.15877>.
- 792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809

APPENDIX

In the appendix, we provide additional details about the collection process for curating SWE-bench M, more characterizations of the SWE-bench M dataset, supplementary experiments and ablations of LM agent performance on SWE-bench M, limitations and more.

A DATASET

A.1 DEVELOPMENT SPLIT

As shown in Figure 4, the SWE-bench M development split consists of 100 tasks instances from 5 open source repositories largely written in the JavaScript family of programming languages (e.g. `js`, `jsx`, `ts`, `tsx`). Repositories were chosen for the development split based on two criteria. First, the number of task instances is about one-fifth the size of the test split. Second, each repository has an “equivalent” test repository, meaning for each repository in the development set, there is a repository with a similar purpose in the test set. This parallelism is meant to make solutions crafted on the development set more transferable to the test set.

We have observed that, in practice, practitioners developing solutions to the SWE-bench task have often iterated on task instances from the *test* split of the dataset. For those interested in working on SWE-bench M, we encourage practitioners to follow the preferable practice of iterating on the *development* split, and then only evaluate on the test split once the approach is finalized.




Figure 4: Distribution of SWE-bench M development set tasks (in parenthesis) across 5 open source GitHub repositories.

Table 6: Median values of different attributes of a SWE-bench M task instance.

		Median
Issue Text	Length (Words)	145
Codebase	# Lines (non-test)	228 K
	# Files (non-test)	1,324
Gold Patch	# Lines edited	35
	# Files edited	2
	# Funcs edited	3
Tests	# Fail to Pass	2
	# Pass to Pass	7

As reflected in Table 2 compared to Table 6, a task instance from the development split has a slightly longer issue text, but also a smaller codebase in terms of number of lines and files. The median number of files (2) and functions (3) edited is identical. Development split task instances change slightly more lines, and have generally have somewhat more Fail-to-Pass and Pass-to-Pass tests.

A.2 ADDITIONAL CHARACTERIZATIONS

We provide further characterizations about the SWE-bench M dataset. In this section, we include both extended discussions of filtering criteria and features discussed in Sections 2.2 and 2.3 as well as additional, noteworthy details not mentioned in the main paper.

Visual Testing. A total of 69 task instances from `Chart.js`, a data visualization framework, and `openlayers`, an interactive mapping library, verify functional correctness with pixel-level visual testing. Across this subset of instances, there are 273 reference screenshots for visual testing.

Within the JavaScript development ecosystem, visual testing like this is commonplace, with `Puppeteer`² and `Pixelmatch`³ being two popular libraries. Such libraries are necessary because they provide several sensible advantages over a simple binary comparison of two images to check if

²<https://github.com/puppeteer/puppeteer>

³<https://github.com/mapbox/pixelmatch>

they're identical. Generally, the overarching reason is that a direct `diff` is too strict and doesn't account for the discrepancies of how different browser render user interfaces and webpages.

There are several additional benefits. First, such libraries provide the ability to specify tolerances for minor pixel level differences that might arise from anti-aliasing or sub-pixel rendering (Kotsarenko, 2018). Second, visual testing libraries allow for targeting certain components of a webpage. Specifically, Puppeteer has tools for taking screenshots of specified sub-parts of a page, and Pixelmatch allows for region-specific comparisons of images. This use case is particularly important for selectively ignoring dynamic website content, such as timestamps or ads, that is irrelevant to the code.

Task instances by year. In the following Table 7, we show the number of task instances for each repository across different years. Task instances are distributed fairly uniformly, with the majority of tasks being from 2019 to 2022. Some repositories have a higher distribution of older task instances, such as `lighthouse` and `wp-calypso`.

Table 7: Number of task instances per year for each repository represented in SWE-bench M.

Repository	≤ 2018	2019	2020	2021	2022	2023	2024
<code>bpmn-js</code>	0	22	6	4	18	4	0
<code>carbon</code>	0	36	30	30	24	12	2
<code>eslint</code>	4	3	0	3	0	1	0
<code>grommet</code>	3	1	0	1	11	5	0
<code>highlight.js</code>	0	0	14	21	4	0	0
<code>lighthouse</code>	25	12	7	3	3	4	0
<code>next</code>	3	9	11	10	4	0	2
<code>openlayers</code>	3	4	12	16	22	14	8
<code>prettier</code>	3	1	3	1	1	3	1
<code>prism</code>	5	6	8	14	5	0	0
<code>quarto-cli</code>	0	0	0	0	8	16	0
<code>scratch-gui</code>	7	2	0	0	1	1	0
<code>Chart.js</code>	0	0	2	17	3	2	0
<code>marked</code>	2	2	6	1	2	1	0
<code>p5.js</code>	1	4	1	1	6	3	0
<code>react-pdf</code>	1	1	0	8	0	1	0
<code>wp-calypso</code>	26	11	0	0	0	0	0
Total	83	114	100	130	112	67	13

Asset collection. For each task instance, we download all image, video, and reproduction code assets referenced by hyperlinks in the issue. This step, not done in SWE-bench, ensures that problem statements are reproducible in the future even if these hyperlinks expire.

Multilingual problem statements. 55 task instances have languages other than English in the problem statement text or images, with Mandarin Chinese being the second most frequently occurring language (38 task instances). The `next` repository, a design component library based on React, contributes the most of these instances, with 26 task instances. This is in large part due to the fact that the repository's owner is the Alibaba Design team. Although not explicitly emphasized in SWE-bench M's collection process, we find that collecting from repositories owned by maintainers whose primary language is not English is a fruitful source of multilingual SWE-bench style task instances.

References solutions edit multiple file types. In SWE-bench, all reference solutions edit exclusively Python files. In SWE-bench M, 174 task instances have references solutions that modify two or more file types. We did not consider non-code files such as text, Markdown, or images files when calculating the counts. We list the frequencies for the different sets of file types edited by reference solutions in Table 8.

A.3 ADDITIONAL ANALYSES

We provide further details and breakdowns of the SWE-bench M dataset, such as repository specific statistics and additional information from scraping task instance information on GitHub.

Table 8: Counts for the number of task instances with reference solutions that edit a specific set of file types. 28% of SWE-bench M task instances modify two or more file types.

File Types Edited	Count	File Types Edited	Count
js	400	js, scss, jsx	8
js, scss	52	lua	8
jsx	19	js, hbs, scss	6
ts, js	18	html, js, css	4
js, jsx	16	scss, jsx	4
html, js	14	js, lock	3
ts	10	js, css	3

Repository Statistics. Following Tables 2 and 6, we show the same median statistics for each repository in Table 9. The discrepancies in the issue text length, codebase size, and gold patch changes can be directly attributed to the codebases themselves - relative to the original online sources for the code and issue text, the SWE-bench M collection process does not affect these values at all. On the other hand, the differences in the number of tests per repository, particularly Pass-to-Pass tests, is an outcome of how the testing specifications and test log parsing are manually specified. Two factors lead to these differences. First, for some repositories, the entire testing suite is run (e.g., `highlight.js`, `p5.js`, `quarto-cli`), while for the majority of repositories, only specific test files, usually derived from file paths modified by the test patch, are run (e.g., `carbon`, `openlayers`, `wp-calypso`). Consequently, for the latter set of repositories, the number of Pass-to-Pass tests is usually few to none. Second, the granularity of test logging is also different across codebases. The primary distinction is that while some test logs show Pass/Fail on a per-file basis, other test logs will show Pass/Fail on a per-test case basis where one files has multiple test cases.

Table 9: **Median values for metrics about each task instance**, mirroring the statistics shown in Tables 2 and 6. Repositories from the test set are listed above the separator, while development set repositories are listed below. “Fail to Pass” is abbreviated as “F2P.” “Pass to Pass” is “P2P.”

Repository	Issue Text	Codebase		Gold Patch			Tests	
	Length	Lines	Files	Lines	Files	Funcs	F2P	P2P
GoogleChrome/lighthouse	94	1076K	834	29	2	3	1	1
PrismJS/prism	79	184K	2,721	12	2	2	1	9
alibaba-fusion/next	14	200K	1,731	57	5	5	3	25
bpmn-io/bpmn-js	92	99K	593	15	1	2	1	0
carbon-design-system/carbon	116	822K	4,250	45	3	4	1	33
eslint/eslint	324	398K	1,567	23	1	2	2	0
grommet/grommet	105	612K	1,186	8	1	1	1	77
highlightjs/highlight.js	139	105K	1,476	16	2	2	2	1,641
openlayers/openlayers	128	544K	1,761	32	2	3	1	0
prettier/prettier	105	348K	6,855	48	2	5	1	0
quarto-dev/quarto-cli	202	1399K	3,639	12	2	2	1	434
scratchfoundation/scratch-gui	148	100K	514	562	12	9	1	0
Automatic/wp-calypso	121	562K	7,990	76	3	5	2	11
chartjs/Chart.js	148	154K	1,376	18	1	2	1	0
diegomura/react-pdf	88	188K	806	39	2	2	1	210
markedjs/marked	89	38K	283	11	1	1	1	0
processing/p5.js	254	436K	1,009	30	2	4	12	2,373

As shown in Figure 5, we include plots of the cumulative distribution functions for the statistics presented in Table 9 across all task instances from both the development and test splits. We also do the same calculations for the SWE-bench dataset (2519 task instances total) and overlay the CDFs for each statistic. The frequencies are normalized to a range of zero to one.

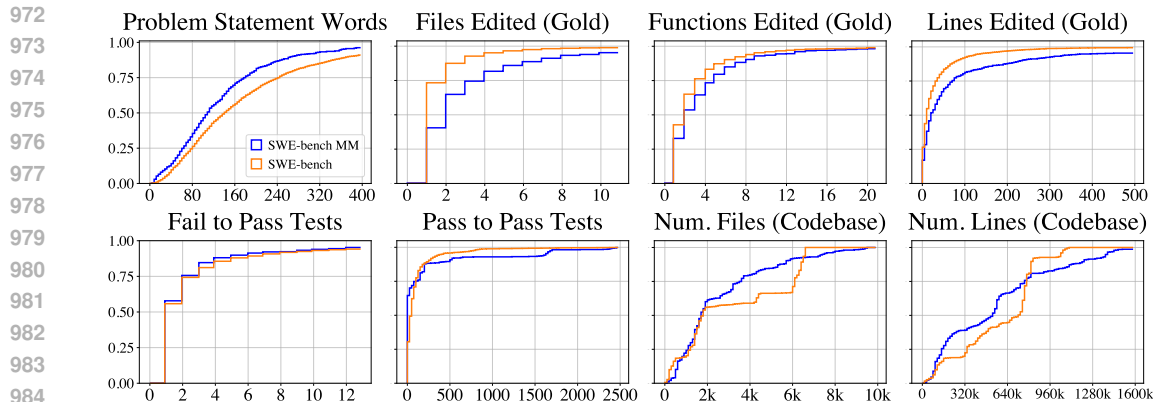


Figure 5: Normalized cumulative distribution functions for different task instance statistics. We compare these statistics between SWE-bench M (blue) and SWE-bench (orange).

While the median number of files and functions is low, with a majority of changes being local and limited to a few lines, there is a long tail of tasks that edit multiple files and functions. The distribution of edits is unlike SWE-bench’s. In SWE-bench, 83% of task instances edit one file, and 65% edit one function. In SWE-bench M, just 40% of task instances change one file, and 32.5% change one function. On average, the changes by reference solutions in SWE-bench M are larger than those in SWE-bench, with multi-file edits being more commonplace.

The average SWE-bench M codebase is somewhat smaller than a SWE-bench repository. The median SWE-bench M codebase has 200k fewer lines of code than its SWE-bench counterpart (535k compared to 734k). However, in the long tail, SWE-bench M has codebases that are larger in terms of both number of files and lines.

Testing compared to SWE-bench is quite similar. For SWE-bench, 80% of task instances have two or fewer fail to pass tests compared to 75% for SWE-bench M. Finally, another noticeable trend is that SWE-bench M issues tend to be shorter than SWE-bench problem statements. This trend can be attributed to the inclusion of images as a complementary form of communication.

GitHub Issue Tags. For every issue collected in the SWE-bench M dataset, we check what labels, if any, were assigned to each issue. These human-annotated labels provide further insight into the kinds of fixes and contributions that a SWE-bench M task instance makes to the codebase. Table 10 summarizes the categories of tags along with the counts of individual tags.

Table 10: **Categories of 1018 tags associated with 535 issues associated with one of the 517 task instances** from the test split of SWE-bench M (499 task instances have 1 issue, and 18 have 2).

Category	Count	Examples
Bug	267	“bug” (207); “type: bug” (56)
Feature	79	“type: enhancement” (44); “enhancement” (19); “feature request” (10); “feature” (6);
Other	672	“a11y” (40); “language” (37); “help welcome” (35); “language-definitions” (32); “role: dev” (32); “good first issue” (28); “package: react” (21); “package: @carbon/react” (14); modeling (13);

While bug fixes are the most popular type of tag in terms of raw frequency, they only constitute a quarter of all tags. In addition to features, there are a variety of more repository specific bug fixes that highlight the diversity of the types of fixes corresponding to SWE-bench M task instances. For example, the “a11y” tag, shorthand for “accessibility”, is used in both the `carbon` and `bpmn-js` repositories. Issues with this tag usually describe how the accessibility features of a UI component from the JavaScript library is either missing or defective. The “language” and “language-definitions” tags are used by syntax highlighting libraries including `highlight.js` and `prism`. Issues with this tag refer to bugs where language-specific symbols are not being highlighted correctly in a code

editor. The “modeling” tag, used exclusively in the `bpmn-js` repository, refers to unexpected behavior the user encounters when manipulating `bpmn-js` diagrams. Every SWE-bench M task instance associated with an issue that has this tag contains either a video or annotated picture that captures the phenomenon via screen recording and/or describes what the expected action should be. Such problem statement require not only visual reasoning acumen, but also the ability to ground accompanying issue description text in the visual elements.

Table 11: **Categories of 186 tags associated with 109 issues associated with one of the 100 task instances** from the development split of SWE-bench M (91 task instances have 1 issue, 9 have 2).

Category	Count	Examples
Bug	52	“type: bug” (21); “[Type] Bug” (17); “bug” (14);
Feature	45	“[Type] Enhancement” (9); “[Feature] Signup & Account Creation” (6); “[Feature] Plans & Upgrades” (6);
Other	89	“Area:WebGL” (13); “L1 - broken” (7); “released” (6); “L2 - annoying” (5); “Store” (5); “FixTheFlows” (5);

Mirroring Table 10, we redo the analysis of the distribution of tags for issues associated with task instances from the development set in Table 11. The findings for the development set are similar to the conclusions drawn from the test set. Bug fixes constitute a quarter of all found tags. The majority of tags, categorized under “Other”, are specifically to subsets of 1+ repositories and reflect more fine-grained, interesting purposes. For example, the “Area:WebGL” tag, which is specific to `p5.js`, refers to problems that arise when using `WebGL`, a popular API for rendering 2/3-D graphics, to render `p5.js` components. Such problems require visual reasoning to fully comprehend; usually, the problem is related to how the component is created properly when using the default `P2D` renderer, but not when `WebGL` is used instead.

A.4 LICENSING

We fully list the license associated with each repository included in the SWE-bench M dataset in Table 12. These licenses all allow for non-commercial use of a repository’s code and content, which the collection and evaluation processes of SWE-bench M respects.

Table 12: **Summary and licenses for all GitHub repositories represented in SWE-bench M.** Test split repositories are listed above the delimiter, while development split repositories are listed below.

Repository	Summary	License
carbon-design-system/carbon	A design system built by IBM	Apache-2.0
GoogleChrome/lighthouse	Automated auditing, perf. metrics for web	Apache-2.0
grommet/grommet	React-based framework for web app dev.	Apache-2.0
openlayers/openlayers	High-performance, feature-packed library for all your mapping needs	BSD-2-Clause
highlightjs/highlight.js	JS syntax highlighter with lang. auto-detection	BSD-3-Clause
scratchfoundation/scratch-gui	GUI to create + run Scratch 3.0 projects	BSD-3-Clause
bpmn-io/bpmn-js	BPMN 2.0 rendering toolkit + web modeler	Custom
quarto-dev/quarto-cli	Open-source scientific and technical publishing system built on Pandoc.	Custom
alibaba-fusion/next	Configurable component library for web	MIT
eslint/eslint	Find and fix problems in JS code.	MIT
PrismJS/prism	Lightweight, robust syntax highlighter	MIT
prettier/prettier	Prettier is an opinionated code formatter.	MIT
markedjs/marked	Markdown parser and compiler.	Custom
Automattic/wp-calypso	JavaScript and API powered WordPress.com	GPL-2.0
processing/p5.js	JS library for learning to code, create art	LGPL-2.1
chartjs/Chart.js	Simple HTML5 Charts w/ <code><canvas></code> tag	MIT
diplomura/react-pdf	Create PDF files using React	MIT

B COLLECTION

Additional information about new dataset curation processes introduced for SWE-bench M are included here, such as the reasons for and number of task instances removed due to inconsistency testing, challenges of setting up automatic evaluations for JavaScript repositories, and the procedure with which digital assets were automatically downloaded.

B.1 STATISTICS

In Table 13, we provide a summary of the number of pull requests that were eventually successfully converted into viable SWE-bench M task instances along with how many pull requests were removed per step of the filtering process.

Table 13: **Per repository, the number of pull requests eventually converted into viable SWE-bench M task instances and task instances kept per stage of the filtering process.** The red subscript denotes the number of task instances that were filtered out at this stage with respect to the total from the previous stage.

	PRs Crawled	Conversion	Validation	Inconsistent Test	Manual Filter
GoogleChrome/lighthouse	6,022	116 _{↓5906}	55 _{↓61}	55 _{↓0}	54 _{↓1}
PrismJS/prism	1,942	52 _{↓1890}	42 _{↓10}	38 _{↓4}	38 _{↓0}
alibaba-fusion/next	2,133	85 _{↓2048}	50 _{↓35}	50 _{↓0}	39 _{↓11}
bpmn-io/bpmn-js	820	86 _{↓734}	54 _{↓32}	54 _{↓0}	54 _{↓0}
carbon-design-system/carbon	8,096	196 _{↓7900}	143 _{↓53}	134 _{↓9}	134 _{↓0}
eslint/eslint	7,454	42 _{↓7412}	11 _{↓31}	11 _{↓0}	11 _{↓0}
grommet/grommet	3,848	68 _{↓3780}	22 _{↓46}	22 _{↓0}	21 _{↓1}
highlightjs/highlight.js	1,955	60 _{↓1895}	39 _{↓21}	39 _{↓0}	39 _{↓0}
openlayers/openlayers	9,466	143 _{↓9323}	88 _{↓55}	80 _{↓8}	79 _{↓1}
prettier/prettier	9,424	34 _{↓9390}	13 _{↓21}	13 _{↓0}	13 _{↓0}
quarto-dev/quarto-cli	1,620	60 _{↓1560}	25 _{↓35}	25 _{↓0}	24 _{↓1}
scratchfoundation/scratch-gui	6,347	38 _{↓6309}	11 _{↓27}	11 _{↓0}	11 _{↓0}
Automatic/wp-calypso	61,165	214 _{↓60951}	37 _{↓177}	37 _{↓0}	37 _{↓0}
chartjs/Chart.js	3,792	200 _{↓3592}	33 _{↓167}	24 _{↓9}	24 _{↓0}
diegomura/react-pdf	784	22 _{↓762}	11 _{↓11}	11 _{↓0}	11 _{↓0}
markedjs/marked	1,816	16 _{↓1800}	14 _{↓2}	14 _{↓0}	14 _{↓0}
processing/p5.js	3,186	42 _{↓3144}	27 _{↓15}	16 _{↓11}	16 _{↓0}
Total	134,866	1,478 _{↓133388}	679 _{↓799}	643 _{↓36}	619 _{↓24}

B.2 INCONSISTENCY TESTING

All task instances in SWE-bench M should yield consistent test case results, so we use a simple scheme to filter out candidate task instances that are inconsistent. For each task instance, we run the test cases five consecutive times and check whether 1) all runs contain the same test cases, and 2) all runs contain the same passing test cases and failing test cases. If either condition is false, we filter out the task instance.

B.3 RESOURCE COLLECTION

Besides text, issue descriptions can contain

1. Images, often in the form of web browser screenshots detailing bugs or specifications of new features. For issues describing unexpected user interface behaviors that involve complex mouse actions (dragging, etc.) or a longer sequence of reproduction steps, animated images (GIFs) are often used.

- 1134 2. Links to online integrated development environments (IDEs) for testing and showcasing
 1135 HTML, CSS, and JavaScript code snippets. Around 17% of the instances contain at least
 1136 one such link, and 15% contain more than one.

1137
 1138 While images can be directly downloaded, copying the resources from the online IDEs is more
 1139 involved. We handle four different such IDEs, *codesandbox.io* (63), *jsfiddle.net* (25), *codepen.io*
 1140 (21), *stackblitz.com* (5), and *editor.p5js.org* where numbers indicate the number of occurrences in
 1141 the final dataset.

1142 The IDEs typically consist of three panes for the code in the three respective languages and an
 1143 additional preview pane. Furthermore, menus provide options to link external resources, such as
 1144 JavaScript and CSS libraries, from content delivery networks. Depending on the occurrence of
 1145 the IDE, we have written extraction scripts (downloading and parsing source code or using web
 1146 browser automation tools) or manually downloaded content to obtain all resources. Most snippets
 1147 are a combination of HTML, CSS, and JavaScript, while others take the form of React apps. For
 1148 the former category, the online IDEs apply a variety of heuristics to interpret the HTML (linking
 1149 external resources, adding missing tags, etc.), so we similarly apply a series of postprocessing steps
 1150 to ensure that the websites can be immediately served without additional steps.

1151 C EXPERIMENTS

1152
 1153 In this section, we include additional details concerning the experiments and ablations we ran on
 1154 the SWE-bench M dataset. We provide information such as how the optimal context window for
 1155 retrieval augmented generation was selected and the refactoring efforts required for making software
 1156 development agent scaffolds compatible with SWE-bench M task instances. We also show more
 1157 detailed analyses of experimental results, such as solve rates by repository and year.

1158 C.1 BASELINE ADAPTATIONS

1159
 1160 Here we provide further technical details on how we adapted each of the software development
 1161 scaffolds for evaluation on SWE-bench MM.

1162
 1163 **Retrieval Augmented Generation.** We largely inherit the retrieval pipeline, prompts, and format-
 1164 ting for RAG from those used in Jimenez et al. (2024a). This includes building a standard BM25
 1165 index using *pysnerini*. Adaptations include the inclusion of “reproduction code” where present ref-
 1166 erenced in links as described in Appendix B.3 and adapting the patch generation example in the
 1167 prompt from a Python example to an equivalent JavaScript example.

1168
 1169 Lastly, we perform some minor reformatting of issue text. We convert all images to Markdown-style
 1170 links with the following format:

```
1171 ![alt text - or image url if not provided](image url)
```

1172
 1173 **Agentless** (Xia et al., 2024) is split into two main phases, localization and repair, the former of
 1174 which is heavily dependent on Python-specific constructs. The original localization phase relies on
 1175 the Python *ast* module to collect the line numbers of all function and class declarations. Unfor-
 1176 tunately, *ast* is specific to the Python abstract syntax grammar, and therefore cannot be applied
 1177 to other languages such as Typescript and Javascript. We also found that the *Esprima* library for
 1178 parsing Javascript code had issues with Typescript code, so we used the *Tree-sitter* parsing library,
 1179 which supports a wide variety of languages. Due to differences in function variants in Python vs.
 1180 Javascript (e.g. arrow functions in Javascript), we also added extra parsing steps in the localization
 1181 step. Furthermore, Agentless uses several in-context Python code snippets for both the localization
 1182 and repair phases – we replace all prompts with Javascript equivalents. Finally, when we provide
 1183 the repository structure as input to the model, we replace their tab-indented nested repository struc-
 1184 ture with full file paths – we found this to be extremely important in ensuring the models do not
 1185 hallucinate false file paths.

1186 The original Agentless work generates a list of candidate solutions and manually checks whether
 1187 they 1) pass syntax checks and 2) pass all *pass-to-pass* P2P tests. We do not include these filters
 for our implementation of Agentless because 1) we found syntax parsing to be inconsistent between

1188 different versions of Javascript, leading to many valid candidate solutions being filtered unless a
1189 particular parser was used per instance and 2) we were interested in a fully automated framework.
1190 Finally, the original Agentless implementation did not implement support for Claude models, in-
1191 cluding how to handle extraneous outputs when querying for structured outputs. We added extra
1192 functionality to support Claude outputs for our experiments while remaining faithful to the original
1193 implementation.

1194 **Aider** (Gauthier, 2024) similarly parses an abstract syntax tree for Python code to help the language
1195 model localize code with a limited context window – however, they use Tree-sitter instead of the
1196 `ast` module, enabling their framework to generalize easier to other languages. Other than adding
1197 the Javascript / Typescript versions of Tree-sitter, we do not modify Aider. Unlike Agentless, which
1198 simultaneously generates multiple candidates and does not require a syntax-checking step, Aider
1199 uses syntax-checking and linting in an iterative fashion to produce plausible generations. We found
1200 that due to syntax differences in Javascript versions, the parser would often invalidate correct or
1201 syntactically valid solutions generated by the language model, leading us not to use the baseline in
1202 its current form.

1203 **AutoCodeRover** (Zhang et al., 2024b) also uses the `ast` module to generate an abstract syntax
1204 tree (AST) for Python code in its *context-retrieval phase*. Furthermore, it defines a set of APIs that
1205 the agent can use to traverse the AST. In our attempted implementation, we replace the `ast` module
1206 with the Tree-sitter library and some extra custom functions, and re-define the APIs to interface with
1207 the Javascript AST while retaining the same stratified search functionality as the original method.
1208 We found the traversal APIs to be too specific to the `ast` module, as the structure of the Tree-sitter
1209 Javascript AST was entirely different, leading to faulty localization.

1210 To re-purpose AutoCodeRover for SWE-Bench M, a faithful implementation would require entirely
1211 swapping out the localization / code-searching APIs.

1212 **Moatless** (Örwall, 2024) employs a multi stage pipeline for solving task instances. Similar to Agent-
1213 less, Moatless also generates an AST of the desired language, and converts this AST into a custom
1214 code graph object that they use for their search and localization APIs. Because this code graph object
1215 is, in theory, language agnostic, it enables flexibility in applying Moatless to different programming
1216 languages. However, while the authors wrote code for converting a Python-generated AST to their
1217 code graph, we noticed that performing an analogous conversion for Javascript/Typescript does not
1218 capture all of the relevant design patterns (e.g. arrow functions). Thus, we chose not to use Moatless
1219 as a baseline.

1220 **SWE-agent Base.** SWE-agent involves relatively few Python specific tools and components. The
1221 most notable exception to this is that for tasks in the original SWE-bench dataset, it sets up the OS
1222 environment with the necessary dependencies and installs the packages locally to provide the agent
1223 with a local testbed for experimenting with changes. To enable similar behavior in Javascript, we
1224 also install package dependencies and build the project locally when applicable. We also globally
1225 link the package name to the local repository when the instance contains an npm package.

1226 New for SWE-bench M, we supply agents with “reproduction code” from links available in the
1227 problem statement. As described in Appendix B.3, these links and assets can contain useful code for
1228 reproducing issues locally. For each link containing reproduction code, we create a directory in the
1229 repository directory where the agent starts, and paste the code into the corresponding files (usually
1230 ‘index.html’, ‘script.js’, and ‘style.css’). Finally, we initialize those repositories as an npm package
1231 and link the package from main repository directory so the agent may use the package as installed
1232 locally.

1233 **SWE-agent JS.** Extending SWE-agent Base, SWE-agent JS adapts the base edit command to work
1234 more effectively with JavaScript and related programming languages. In particular, the original
1235 SWE-agent edit command includes a component that performs a linting operation after every edit
1236 an agent applies to a Python file. This linting operation checks whether the resulting file would
1237 be syntactically correct and can detect certain other mistakes such as a reference to an undefined
1238 variable. During editing, SWE-agent will “roll-back” an edit if it is found to be invalid according to
1239 the linting operation. It then reports this error to the agent and requests that the agent to correct and
1240 try its intended edit again.

1241

Table 14: Specialized tools provided to the agent for web interaction and image handling. Required arguments are enclosed in `<>` and optional arguments are in `[]`.

Command	Documentation
open_webpage <code><website_dir></code>	Opens the webpage in a new display window using Google Chrome. <code>website_dir</code> is the directory where an <code>index.html</code> file exists and will be served.
restart_webpage <code><website_dir></code>	Restarts the webpage with the given directory. <code>website_dir</code> is the directory where an <code>index.html</code> file exists and will be served.
close_webpage	Closes the open webpage.
screenshot [<code><num_images></code>] [<code><interval></code>]	Takes screenshots of the display. <code>num_images</code> (optional, default 1) is the number of screenshots to take. <code>interval</code> (optional, default 1.0) is the interval between screenshots in seconds.
open_image <code><image_path></code> [<code><image_path></code> ...]	Opens the image(s) in base64 format to be fed to the LM. Multiple image paths can be provided.

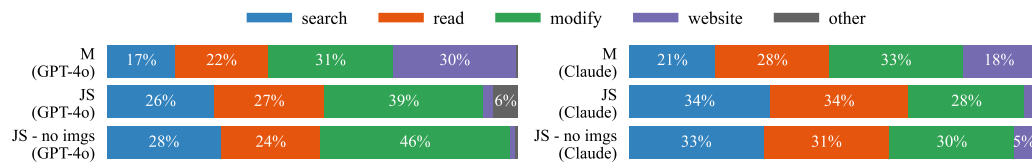


Figure 6: Action space of SWE-agent with different configurations for Claude 3.5 Sonnet and GPT-4o using the three configurations mentioned in Section 4 on the development set. For a version of this figure stratified by repositories, see Figure 7.

To adapt this for SWE-bench M, we use ESLint with the local repositories configuration files, to primarily detect fundamental syntax errors. While typically ESLint is capable of enforcing a wide variety of rules, we found that relying too heavily on these rules from the projects’ configuration substantially degraded performance.

SWE-agent M. Further extending SWE-agent JS, SWE-agent M introduces a variety of new commands and features of the ACI, which we show in Table 14. These new commands require some environmental support for simulating displays. We simulate a display for the X window system using Xvfb, and take screenshots with xwd.

C.2 BASELINE CONFIGURATIONS

SWE-agent. Due to the expense and flexibility in developing new SWE-agent interfaces, we developed new features and for SWE-agent JS and SWE-agent M primarily iterating while validating performance on a small subset of the development set. By default, for long trajectories, SWE-agent “collapses” past environment observations by replacing the content of old observations with a single line: “n lines omitted”. SWE-agent will usually show the model the last five observations, and collapse all observations prior to that except the initial problem statement and demonstration. For the final configuration evaluated for each system, we perform a very small grid search over two different options for the number of past observations to show in $\{5, 9\}$. We perform this search on a set of 50 instances from the development set and show the results of this search in Table 15.

Retrieval Augmented Generation. For RAG systems, the amount of context we provide, either in terms of the number of documents or simply the absolute length of the context retrieved, is an important hyperparameter that may affect a model’s performance. As in Jimenez et al. (2024a), we determine the final RAG system to evaluate by performing a grid search over three possible context lengths in $\{32K, 64K, 100K\}$ and the inclusion or not of images as input with the problem

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

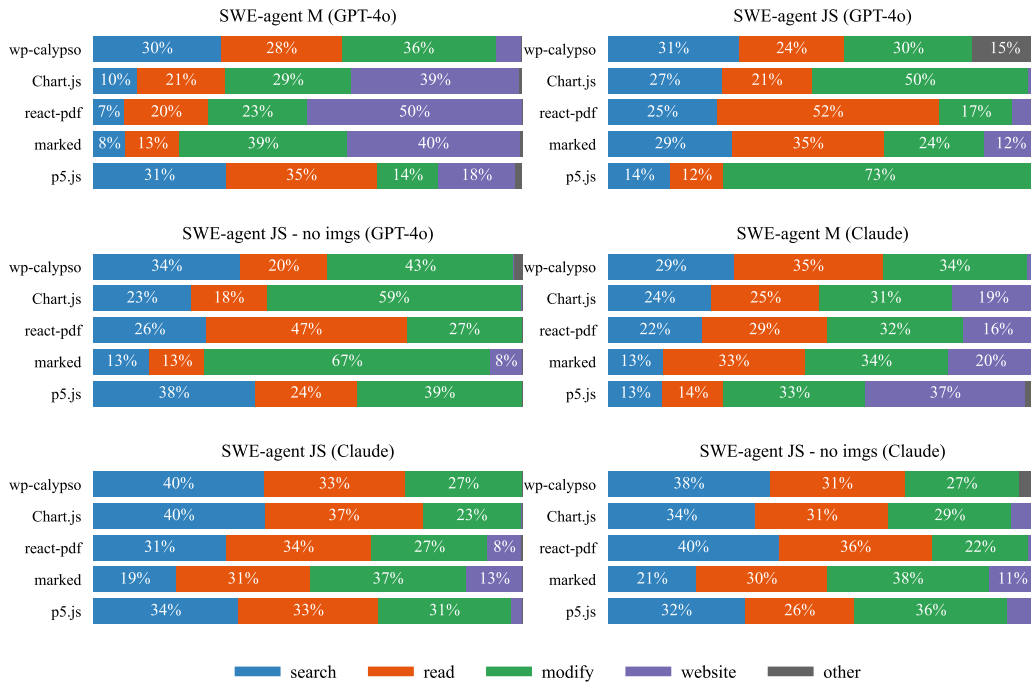


Figure 7: Action space of SWE-agent with different configurations for Claude 3.5 Sonnet and GPT-4o using the three configurations mentioned in Section 4 on the development set broken down by repository.

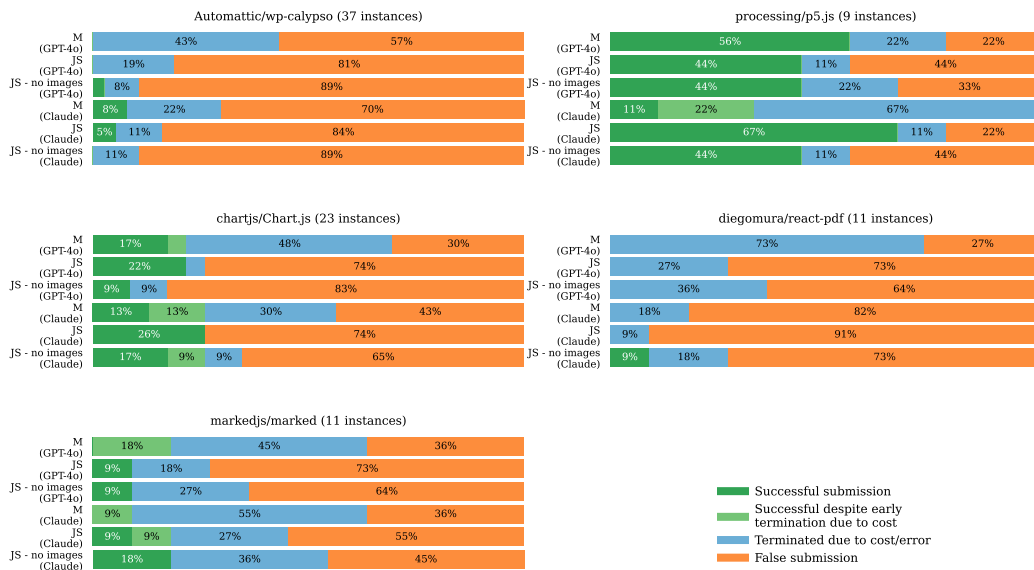


Figure 8: Frequency of outcomes by success and exit status of SWE-agent under different configurations on the development set.

Table 15: Hyperparameter search for SWE-agent system for non-collapsed history length.

Model	System	History Length	% Resolved
Claude 3.5 Sonnet	SWE-agent Base	5	18%
		9	14%
	SWE-agent JS	5	28%
		9	18%
	SWE-agent M	5	24%
		9	16%
GPT-4o	SWE-agent Base	5	18%
		9	12%
	SWE-agent JS	5	24%
		9	10%
	SWE-agent M	5	26%
		9	16%

Table 16: Hyperparameter search for RAG models with differing lengths of context and image input.

Model	Retrieval Context	With Images	% Resolved
Claude 3.5 Sonnet	32K	✗	9.4 \pm 1.1
		✓	7.8 \pm 2.6
	64K	✗	11.2 \pm 1.3
		✓	14.1 \pm 2.6
	100K	✗	7.6 \pm 2.7
		✓	13.7 \pm 1.8
GPT-4o	32K	✗	6.7 \pm 2.1
		✓	6.9 \pm 2.0
	64K	✗	6.7 \pm 1.9
		✓	7.6 \pm 2.2
	100K	✗	8.0 \pm 1.5
		✓	10.0 \pm 2.2

statement. We run each configuration 5 times on the development set and select the configuration with the best mean performance. We report the results of this grid search in Table 16, highlighting the performance of our selected configuration.

C.3 FURTHER ANALYSES

Here, we include additional analyses of the experimental results and performance by different baselines on SWE-bench M.

Resolved by Repository. We show the resolved rates per repository accomplished by different baselines in Table 17.

1404
 1405
 1406
 1407
 1408
 1409
 1410
 1411
 1412
 1413
 1414
 1415
 1416
 1417
 1418
 1419
 1420
 1421
 1422
 1423
 1424
 1425
 1426
 1427
 1428
 1429
 1430
 1431
 1432
 1433
 1434
 1435
 1436
 1437
 1438
 1439
 1440
 1441
 1442
 1443
 1444
 1445
 1446
 1447
 1448
 1449
 1450
 1451
 1452
 1453
 1454
 1455
 1456
 1457

Table 17: **% Resolved Rates by repository for different baselines.** All performance numbers listed here use GPT-4o as the base model. Performance for the test split repositories (above delimiter) are pass@1, while performance for the development split repositories (below delimiter) represents the best result achieved from a hyperparameter search.

Repository	Count	SWE-agent M	Agentless JS	RAG
GoogleChrome/lighthouse	54	5.6	7.4	1.9
PrismJS/prism	38	0.0	0.0	0.0
alibaba-fusion/next	39	0.0	2.6	0.0
bpmn-io/bpmn-js	54	27.8	11.1	13.0
carbon-design-system/carbon	134	1.5	0.7	0.0
eslint/eslint	11	0.0	0.0	0.0
grommet/grommet	21	0.0	0.0	0.0
highlightjs/highlight.js	39	2.6	2.6	0.0
openlayers/openlayers	79	51.9	3.8	30.4
prettier/prettier	13	7.7	0.0	0.0
quarto-dev/quarto-cli	24	0.0	0.0	0.0
scratchfoundation/scratch-gui	11	0.0	0.0	0.0
Automattic/wp-calypso	37	0.0	0.0	2.2
chartjs/Chart.js	24	20.8	0.0	11.7
diegomura/react-pdf	11	0.0	0.0	0.0
markedjs/marked	14	7.1	0.0	20.0
processing/p5.js	16	25.0	6.2	23.8

1458 **Temporal analysis does not reveal any indication of solution leakage.** Table 18 shows resolution
 1459 rates for different baselines by year. While agentless with GPT-4o shows somewhat increased per-
 1460 formance for older task instances, all other system shows significantly higher performance on newer
 1461 instances with a pronounced peak in 2024. However, the instances of the different repositories have
 1462 uneven temporal distributions and average resolution rates (see Tables 7 and 17), making it difficult
 1463 to interpret the yearly solution rates. For example, 8 out of the 13 instances from 2024 are from
 1464 openlayers/openlayers, which has by far the highest solution rates with SWE-agent M and
 1465 agentless across all years.

1466 To account for this, we calculate average resolution rates before and after the GPT-4o training cutoff
 1467 (October 2023). We then compare the post resolution rates to the pre-resolution rates reweighted to
 1468 the post-cutoff repository distribution⁴. The results are shown in Table 19 and show that the post-
 1469 cutoff performance of all tested systems exceeds that of the pre-cutoff even when accounting for the
 1470 shift in repository distribution.

1471 Since Claude 3.5 Sonnet’s training cutoff of April 2024 would leave too few instances, a similar
 1472 analysis is not feasible. However, SWE-agent M with Claude 3.5 Sonnet follows a very similar
 1473 trend as SWE-agent M with GPT-4o and shows an even more pronounced performance jump in
 1474 2024.

1475 Table 18: **% Resolved performance for SWE-bench M test split task instances from different**
 1476 **years.** Each row corresponds to a year, while each column is the agent’s performance for task
 1477 instances from that year.

Year	Count	SWE-agent M (Claude)	SWE-agent M (GPT-4o)	Agentless JS (GPT-4o)	RAG (GPT-4o)
2017	22	4.5	4.5	9.1	4.5
2018	31	0.0	0.0	3.2	0.0
2019	96	9.4	7.3	3.1	1.0
2020	91	8.8	4.4	5.5	2.2
2021	103	13.6	11.7	2.9	3.9
2022	101	12.9	19.8	1.0	13.9
2023	60	13.3	20.0	0.0	6.7
2024	13	46.2	53.8	7.7	46.2

1490 Table 19: **% Resolved Rates with GPT-4o before and after training cutoff.** This table shows
 1491 the performance for the SWE-bench M test split task instances before and after the GPT-4o training
 1492 cutoff date (October 2023). Pre-cutoff (reweighted) is the pre-cutoff performance reweighted to the
 1493 post-cutoff repository distribution (see text).

	SWE-agent M	Agentless JS	RAG
Pre-cutoff	11.0	3.0	5.0
Pre-cutoff (reweighted)	27.6	3.1	13.6
Post-cutoff	47.1	5.9	41.2

1510 ⁴i.e., we calculate $\sum_{\text{repo}} f_{\text{post}}^{\text{repo}} \eta_{\text{pre}}^{\text{repo}}$, where $f_{\text{post}}^{\text{repo}}$ is the fraction of instances of a specific repository among the
 1511 post-cutoff instances and $\eta_{\text{pre}}^{\text{repo}}$ is the corresponding pre-cutoff solution rate.

D HUMAN VALIDATION

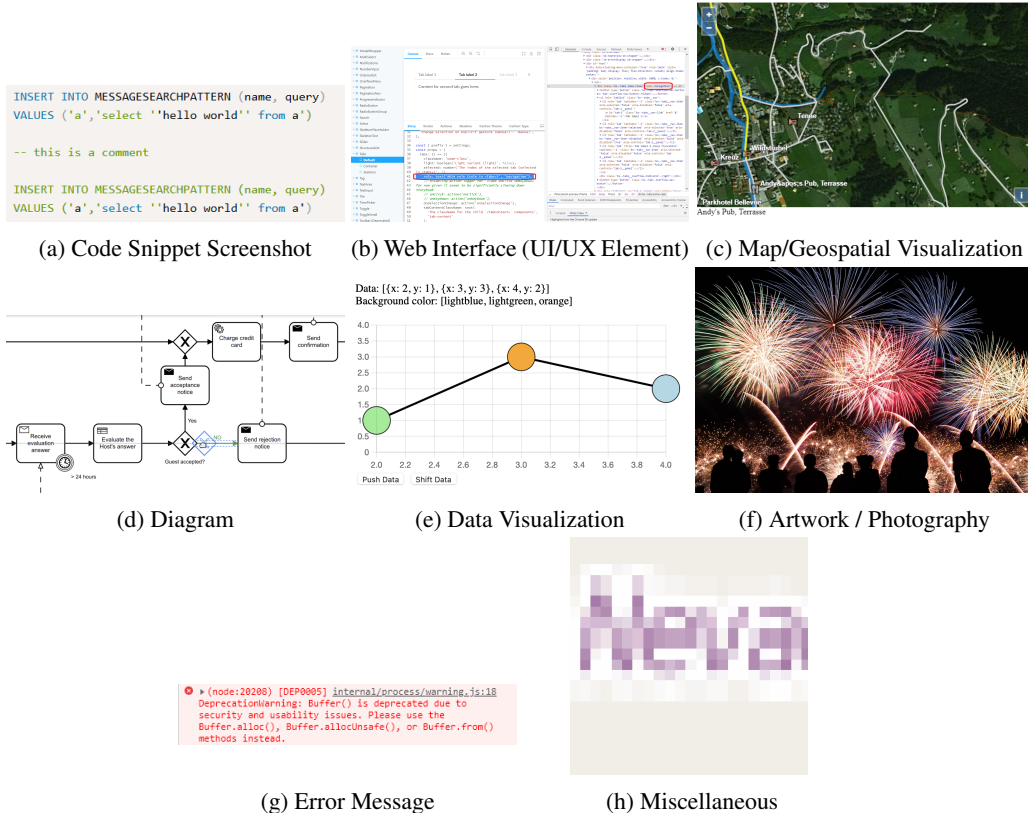
We provide full details about the human validation process for SWE-bench M task instances. We include each question with corresponding examples that were provided to the papers' authors. We also show additional analyses about the authors' answers and agreement.

D.1 IMAGE CATEGORIZATION

This question aims to categorize the types of visual content commonly associated with open-source software issues. By classifying images across a diverse range of repositories, we aim to analyze the visual ways in which developers convey information about software development. To accomplish this, we randomly sampled 50 images from the dataset and manually derived 8 categories from these samples. Human participants then manually inspected and labeled all images with one of the categories.

Prompt: Classify the image as one of these categories: (1) Code Snippet Screenshot (2) Web Interface (UI/UX Element) (3) Map/Geospatial Visualization (4) Diagram (5) Data Visualization (Plots) (6) Artwork / Photography (7) Error Message (8) Miscellaneous.

Here are examples of each category of images.



The images are from the following task instances:

1. Code: PrismJS_prism-1500
2. Webpage: carbon-design-system_carbon-6964
3. Geospatial: openlayers_openlayers-10545
4. Diagram: bpmn-io_bpmn-js-1542
5. Data Viz.: chartjs_Chart.js-9101
6. Art: quarto-dev_quarto-cli-5547
7. Error Trace: diegomura_react-pdf-1285
8. Other: openlayers_openlayers-12393

Table 20: Response to *Q1: Image Categorization* indicate the majority of images are code and website screenshots. Other categories tend to be more repository-specific.

	Code	Website	Geospatial	Diagram	Data Viz.	Art	Errors	Other	Total
bpmn-js	1	13	1	54	1	1	1	1	69
carbon	21	136	0	36	0	0	30	1	224
Chart.js	5	2	0	0	18	0	1	0	26
eslint	10	0	0	1	0	0	2	0	13
grommet	3	25	0	2	2	0	0	0	32
highlight.js	62	4	0	4	0	0	0	0	70
lighthouse	5	55	0	0	5	0	3	0	68
marked	1	22	0	0	0	2	0	0	25
next	11	36	0	0	0	0	6	0	53
openlayers	6	6	35	0	0	0	2	0	49
p5.js	5	0	0	3	0	22	2	0	32
prettier	9	3	0	0	1	3	0	0	16
prism	41	3	0	3	0	0	0	0	47
quarto-cli	7	22	0	2	1	10	0	1	43
react-pdf	6	7	0	2	1	0	4	0	20
scratch-gui	0	12	0	0	0	1	0	0	13
wp-calypso	1	55	0	0	0	0	3	0	59
Total	194	401	35	107	28	38	54	2	859

Results: We applied this labeling procedure to the 862 problem statement images across all SWE-bench M task instances. The counts for each of the categories are presented below in Table 20.

The majority of SWE-bench M images are screenshots of webpages or code. These categories dominate because they are most often used to point out problems with linting (e.g. `eslint`, `highlight.js`, `prism`, prominently feature code-related images due to their focus on code analysis and syntax highlighting) and incorrectly rendered web elements (e.g. `carbon`, `lighthouse`, `next`, `wp-calypso`, as these tools are used to address issues in web design, layout and performance). The representation for other categories correlates heavily with specific repositories. In other words, images falling under a certain category are usually distributed across at most two to three repositories. For example, diagrams tend to be found in `bpmn-js`, a tool for creating BPMN (Business Process Model and Notation) diagrams, and `carbon`, which often uses design component specifications to communicate details about web elements, such as spacing and size.

Geospatial images, including maps or location-based data, are exclusive found in `openlayers`. Data visualization, such as charts or graphs, are predominantly found in the `Chart.js` repository, which specializes in creating data-driven visual content. Additionally, creative coding repositories like `p5.js` often include images related to artistic outputs, those images typically include visual glitches or unexpected rendering behaviors. Lastly, categories such as errors are found sporadically across repositories and represent instances where tools fail to operate correctly, such as compilation failures. Based on these findings, we conclude that SWE-bench M features a wide range of images that capture various aspects of software development, with code and webpage screenshots dominating the dataset. At the same time, certain types of images are tightly coupled with the specific functionalities of individual repositories, reflecting the unique focus and purpose of each project.

D.2 IS AN IMAGE REPRESENTABLE AS TEXT?

How important is it for a SWE-bench M image to be an image? To answer this, we attempt to quantify the proportion of images that are solely text. For these images, it may not be necessary for the image’s information to be communicated visually. For instance, it’s possible that a large proportion of images are screenshots of code or error messages, some of which could be represented

perfectly as text. If this is the case, the fact that the information is presented as image becomes somewhat trivial.

To clarify, we are not asking whether the image can be processed with Optical Character Recognition (OCR) or verbalized, as such approaches would still likely result in a loss of information and are hard to answer without considering the problem context. For instance, for syntax highlighting libraries, although many images are code screenshots, there could be color highlighting of different characters that is crucial to understanding the problem correctly. Therefore, while traditional OCR could turn the code in the image into text, it would fail to capture additional visual details. Asking whether the same image can be verbalized is also not straightforward, as the answer could likely depend on the type of problem being solved and what information about the image is actually necessary. In a small labeling trial, we also found that answers are sensitive to the annotator’s subjectivity, reflected by low inter-annotator agreement.

Similar to *Q1: Image Categorization*, we ask a human participant to look at the image and respond with either “Yes” if the image can be represented faithfully as text, and “No” otherwise. A more detailed discussion of how to decide is included in the following prompt.

Prompt: Can you assess if this image can be faithfully represented with text? A faithful representation means that the contents of the image can be represented as text with *no loss* in information. An image fits this criteria if (1) The image only contains text (2) The text in the image is mostly or entirely monochrome, and (3) The image does not contain any meaningful visuals or patterns independent of text that could be important to know for solving a problem. Please only respond with [1/0], 1 meaning “Yes” and 0 meaning “No”. Do not include any additional justification or text.


Here are examples of images that can and cannot be conveyed effectively with OCR.




Figure 10: Examples of images that cannot (top row) and can (bottom row) be processed into text by OCR software. Typically, code and webpages are OCR-able while diagrams, plots, and art are not.

- Top Left (scratchfoundation__scratch-gui-8891): The eye icons cannot be represented as text. While there is text, it does not take up much of the image.
- Top Middle (chartjs__Chart.js-8567): The bar plot cannot be converted into text.
- Top Right (processing__p5.js-3709): This is an graphic with no text in it.
- Bottom Left (openlayers__openlayers-11649): Although a significant portion of this image is made up of text, the code syntax coloring is an aspect of the image that cannot be conveyed effectively in a text representation. In addition, the black highlighting and striped background cannot be portrayed as text.
- Bottom Middle (diegomura__react-pdf-1285): Similar to the bottom left image, this is also “No” because of the text coloring.
- Bottom Right (PrismJS__prism-2782): The contents of this image can be faithfully represented as text, as the majority of it is colored red. The main purpose of this screenshot is to communicate an error message. There are no meaningful visual elements.

Results: We again answer these questions for problem statement images in SWE-bench M. Out of 862 problem statement images, 691 (80%) were labeled as “No”, while 171 (20%) were labeled as



1685 Figure 11: Responses to whether an image can be represented with text per repository.



1697 Figure 12: Responses to whether an image can be represented with text per image category.

1698
1699
1700 “Yes”. Below, we present two figures. Figure 11 shows the labeling results per repository. Figure 12
1701 shows the labeling results with respect to image category.

1702 With the exception of `prettier`, across repositories, the majority of images were labeled as not
1703 representable with text, with several repositories having no images labeled “Yes” (e.g. `bpmn-js`,
1704 `highlight.js`, `quarto-cli`). The principal component that corresponds to whether an image
1705 can be labeled is the image category. As shown in Table 12, different image categories have varying
1706 rates of whether or not the image can be represented with text. Error messages are more often labeled
1707 “Yes” than not. A smaller proportion of web and code screenshots were labeled “Yes”. Although
1708 at first glance it may seem like this proportion should be higher, through the annotation procedure,
1709 we found out that for code, text color and highlighting are often very important signals for solving
1710 task instances, particularly linting repositories such as `prism` and `eslint`. For web screenshots,
1711 while there is usually text on a webpage, the spatial arrangement of web components is a critical
1712 aspect of understanding problems in design framework libraries, such as `carbon` and `next`. For
1713 the remaining categories, low to no images are labeled as being representable in text.

1714 Based on our findings, we conclude that information from problem statement images in SWE-bench
1715 M frequently must be represented as images. The rate at which images are not just screenshots of
1716 text varies by repository and image category. Images that are errors are more often representable as
1717 text. Other categories are less often due to either a complete absence of text or the text coloring and
1718 organization being an important part of understanding the problem being communicated.

1719
1720
1721 **D.3 IMAGE NECESSITY.**

1722 We elicit human feedback on how necessary images are to solving the task at hand. This question
1723 attempts to directly determine how the visual asset(s) provided with a problem statement are actually
1724 necessary to solving a GitHub issue when provided. We label the 557 SWE-bench M task instances
1725 that have one or more images in its problem statement.

1726 For the annotation procedure, a human participant is asked to look at each task instance’s problem
1727 statement, codebase, gold patch, and test patch, and judged whether it could be solved. They should

then look at the associated problem statement image(s) and assign a label according to the procedure described below. Task instances with no images in the problem statement are not considered for this question. If a task instance is deemed generally unsolvable, it is removed from the dataset.

Prompt: For each task instance, please determine whether its associated images are necessary to solving the problem described by this task instance. The following workflow is recommended:

1. Read the problem statement. If you'd like, check the codebase, gold patch, and test patch as needed. Make a mental note of whether the task instance can be solved.
2. Then, look at the image(s) associated with the task instance. Now, re-answer whether the task instance can be solved.
3. Provide answer according to the following logic:
 - “Yes”: If the answer changed from “No” to “Yes”, then the image is necessary.
 - “No”: If the answer remained “Yes”, the image is not necessary to solve the task.
 - Remove: If the answer before and after seeing the image was both “No”, please note this and we will remove the task instance.

Below, we provide examples of task instances where the image *is* necessary to solving the task (PrismJS_prism-3442), and where it *is not* (quarto-dev_quarto-cli-6659).

PrismJS_prism-3442. A screenshot of code is provided for this task instance. It is shown together with the corresponding problem statement in the text box below.

The main contribution of this image is the syntax highlighting. Specifically, it shows the undesired effect of Prism highlighting the quotes with the incorrect color. The double quotes around “Month”, “Days”, and “Jan” should not be highlighted gray.

Without the screenshot, it is not exactly clear how the quotes are being highlighted incorrectly. For instance, it’s unclear what color corresponds to “highlighted as punctuation”. In the absence of the image, while it is still possible that an experienced maintainer of this repository could make educated guesses about the exact nature of the bug or reproduce it exactly, the image provides concrete evidence of what the user is experiencing.

PrismJS_prism-3442: [language-markup] Quotes in HTML attribute values are highlighted

****Information:****

- Prism version: Latest (reproducible on test.html)
- Plugins: none
- Environment: Browser

Quotes are highlighted as punctuation inside HTML attributes. E.g. this:

```
“html
<google-chart data='["Month", "Days"], ["Jan", 31]}'></google-chart>
“
```

is highlighted as:

```
<google-chart data='["Month", "Days"], ["Jan", 31]}'></google-chart>
```

We might decide that this is a feature, not a bug, but figured I'd flag this in case, as it looked wrong to me.

On the other hand, the image in quarto-dev_quarto-cli-6659’s problem statement is not necessary.

The provided image is an arbitrary link that is used in the reproduction code shown in the issue description. While the image is helpful in that it is readily usable by a developer for recreating the problem, there is no additional information being conveyed by the image’s content that helps with understanding the task at hand.

quarto-dev__quarto-cli-6659: Figures from R code blocks don’t render if ‘fig-cap:
!expr ...’ evaluates to ‘character(0)’

With Quarto 1.4.330, quarto-r 1.2, and knitr 1.43 on R 4.2.1, the command ‘quarto render example.qmd –to html’ renders the following Quarto document as expected:


```
““_r
#| fig-cap: !expr caption
caption = "hello world"
knitr::include_graphics("http://arfer.net/mlp/img/rara-jiggs.png")
““
```

But if ‘caption = "hello world"’ is changed to ‘caption = character(0)’, the figure disappears from the output.

I originally hit this issue with a mistaken ‘sprintf’ call, which can produce a zero-length character vector.

Results: Out of 557 task instances, 465 were labeled as the image being *necessary* to solving the problem. The remaining 93 were labeled as “No”, meaning the image is *not necessary*. The following Figure 13 shows the labeling splits for this question by repository.

Figure 13: Responses to whether an image is necessary to solve a task instance per repository.



Annotations suggest the significant majority (83.5%) of problem statement images are necessary to solving the problem at hand. Several repositories have zero issues with unessential images.

Annotators attributed several factors that affected their decision making. First, if the task instance included reproduction code, annotators were more likely to rate the task instance as not requiring the image. For these cases, the image was redundant because running the reproduction code would produce the exact effect that the image was capturing. Second, a repository’s contributing guidelines influences the role of the image. For instance, some repositories will explicitly ask for an image or screenshot that details the error message, which is why for repositories such as bpmn-js and lighthouse, there are zero occurrences of unessential images. On the other hand, images are sometimes provided under an “Additional Information” section, in which case the image is helpful, but not crucial. These observations also hold for reproduction code. Task instances from carbon, p5.js, and next will often have links to online editors reproducing the described bug.

1836 D.4 TASK DIFFICULTY

1837
1838 Following the instructions used in Chowdhury et al. (2024), we label the difficulty of a task instance.
1839 We define difficulty as the estimated amount of time it would take for an experienced developer to
1840 accomplish this task. There are four possible labels for difficulty:

- 1841 • **<15 min** fix (e.g., a trivial change adding some assertions to a function)
- 1842
- 1843 • **15 min–1 hour** (e.g., a small change that requires a bit of thought)
- 1844
- 1845 • **1–4 hours** (e.g., substantially rewriting a function or editing multiple files)
- 1846
- 1847 • **>4 hours** (e.g., a very esoteric issue that clearly requires a substantial amount of research
to fix, changing >100 lines of code)

1848 For every estimate, we assume the developer is approaching the problem from the following context:

- 1849
- 1850 • The developer has neither worked on nor used the codebase before.
- 1851
- 1852 • The developer is familiar with JavaScript code and has experience working on JS related
frameworks (e.g., React).
- 1853
- 1854 • The developer has had a couple hours to become familiar with the codebase prior to starting
work on the task instance.
- 1855

1856 To label, each annotator looked at the full details of the task instance (problem statement, images,
1857 gold patch, codebase) and made a best-effort judgment of how intensive the fix would be. Annotators
1858 were asked to *not* consider the time spent on writing or modifying tests.

1859 Because of the potential variance in annotators’ answers to this question along with the amount of
1860 time it takes to comprehend a task instance before making a judgment, we do not label the entire
1861 SWE-bench M dataset. Instead, we take 100 task instances from all of SWE-bench M, uniformly
1862 sampling with respect to the ratios of task instances in each repository. Three annotators then label
1863 each of the tasks with one of the four labels. For each task instance, the final difficulty level is
1864 assigned to be whichever label has a majority vote. If there is no majority vote, the “middle” of the
1865 three annotations (either **15 min–1 hour** or **1–4 hours** is used as the label.

1866
1867
1868 *Results:* From this labeling procedure, 100 SWE-bench M task instances are labeled with differing
1869 levels of difficulty, as shown in Table 21. The Fleiss’ kappa score is calculated to be 0.78, reflecting
1870 a substantial amount of agreement between the three annotators who carried out labeling.

1871 Table 21: Estimated level of difficulty for different splits of SWE-bench in addition to SWE-bench
1872 M. Each number is the percentage of sampled task instances classified at that level of difficulty.

1874 Dataset	1875 # Samples	<15 min	15 min-1 hour	1-4 hours	>4 hours
1876 SWE-bench	1,699	24.5	53.3	19.4	2.8
1877 Lite	231	37.7	56.3	6.1	0.0
1878 Verified	500	38.8	52.2	8.4	0.6
1879 SWE-bench M	100	13.0	43.0	38.0	6.0

1880
1881 After the labeling procedure was completed, annotators agreed that the strongest signals for estimat-
1882 ing the level of difficulty came most frequently from three indicators.

1883 First, smaller gold patches tend to be labeled as requiring less time and visa versa. For instance,
1884 for the 13 task instances labeled to take < 15 minutes, the gold patch change sizes (calculated as
1885 Lines Added + Lines Removed) are [1, 3, 4, 1, 3, 2, 163, 5, 5, 10, 7, 6, 4]. The one aberration of ± 163
1886 lines is due to updates to an auto-generated `package-lock.json`. On the other hand, for task
1887 instances requiring > 4 hours to fix, the gold patch sizes are [350, 186, 557, 367, 2682, 937].

1888 Task instances with more descriptive problem statements tend to get rated as easier. The annota-
1889 tors’ noted that more detailed descriptions and the presence of reproduction code made them more

1890 inclined to reduce the estimated amount of time required. The images in the problem statement
1891 also had an effect on the labeling. Annotators mentioned that screenshots of web elements and
1892 source code helps developers localize errant code in a codebase faster. For instance in `carbon` and
1893 `next`, many screenshots would point out spacing or coloring issues with a specific web component.
1894 A screenshot is often informative for not only describing the problem, but also directly point out
1895 which component folder the task worker should look into.

1896 Finally, annotators noticed that sometimes, a small change can still take some time, particularly if
1897 the edit is precise and requires a good amount of context to find. In addition to the size of a gold
1898 patch edit, the quality of an edit, specifically the types of entities and symbols that were used in the
1899 fix, influences the amount of time a problem takes. For instance, a larger edit that mostly consists
1900 of JavaScript primitives is usually regarded as simpler than a shorter edit mostly made up of entities
1901 from other parts of the codebase.

1902 In conclusion, SWE-bench M presents task instances at a range of difficulties, from short quick fixes
1903 to problems that require large scale refactoring of multiple modules in a codebase. The difficulty
1904 distribution suggests that SWE-bench M can be effective for tracking improvements in model and
1905 agent capabilities.

1906

1907 E LIMITATIONS

1908

1909 **Broader scope** To comprehensively evaluate multimodal AI systems, we curate 617 task instances
1910 from 17 JavaScript libraries for SWE-Bench M. However, the scope could be broadened across three
1911 key dimensions: **(1) More programming languages:** Besides JavaScript, multimodal content also
1912 appears in issues related to code in other programming languages, such as Python, C++, or Rust.
1913 **(2) More modalities:** While images and videos are likely the most commonly used media besides
1914 text in GitHub issues, users can also upload audio files or other media and file types. **(3) More**
1915 **tasks:** Our 617 tasks across 17 libraries cover various domains including web frameworks, data
1916 visualization and syntax highlighting, however, there are many more JavaScript libraries in other
1917 domains that could be added to the benchmark. We are excited about research broadening the scope
1918 across these and other axes. However, here we focus on quality over quantity. Scope extensions
1919 likely warrant separate dedicated future work, as we found that adding task instances to SWE-bench
1920 M was a labor-intensive task.

1921

1922 **Improved models and environments** While we primarily focus on developing the benchmark
1923 and evaluating current models, we also extend SWE-Agent to handle multimodal issues. There
1924 remains substantial opportunity for future enhancements, both by advancing the underlying models
1925 (e.g., leveraging future versions of GPT and Claude) and by enriching the agent’s environment with,
1926 for example, better browsing capabilities or more tools. We look forward to seeing future work
1927 along these lines to improve performance on SWE-Bench Multimodal beyond the systems that we
1928 present.

1928

1929

1930

1931

1932

1933

1934

1935

1936

1937

1938

1939

1940

1941

1942

1943