香港中文大學
The Chinese University of Hong Kong

Final Year Project

# Large Language Models for Software Systems: Directions, Opportunities, and Challenges


*Author:*

Chang CHEN

Houtian ZHU

*Supervisor:*

Prof. Michael R. LYU


LYU2407


Department of Computer Science and Engineering

Faculty of Engineering


December 17, 2024

# Acknowledgement

The work described in this report was done under the supervision of Prof. Michael Lyu and Shuqing Li.

# Abstract

In the field of modern software engineering, many tasks now involve inputs that go beyond plain code text, incorporating multiple modalities such as images, audio, and video. This shift introduces significant challenges in handling these diverse inputs. The emergence of large multimodal models (LMMs) offers a promising solution to this issue. However, as an emerging technology, systematic research on multimodal large models within the software engineering domain remains scarce. There is still a lack of clarity regarding the specific tasks LMMs can accomplish and their performance across these tasks.

This report conducts a comprehensive and systematic survey, categorizing and summarizing all multimodal-related problems in software engineering over the past five years, and finally constructs a complete task tree. Subsequently, we develop a modular testing framework capable of automatically measuring LMM performance based on configuration files. Within the scope of input modalities currently supported by LMMs, we select several representative tasks and evaluate their capabilities.

Our findings reveal that LMMs demonstrate surprisingly strong performance in the field of software engineering. In certain tasks, they are capable of achieving results comparable to specialized models fine-tuned for specific tasks, even without any additional fine-tuning. This highlights their significant potential for development and application in the software engineering domain.

# Contents

# 1 Introduction

## 1.1 Introduction

The integration of large multimodal models (LMMs) into software systems research represents a novel frontier in artificial intelligence, blending the linguistic proficiency of large language models (LLMs) with sophisticated vision models to process and generate multimodal content. This synthesis allows LMMs to handle diverse input modalities—such as text, images, and potentially audio and video—and to produce outputs that bridge these modalities. Recent advancements in LMMs have demonstrated their capacity to perform complex reasoning tasks, often achieving strong results even in 1-shot or 0-shot scenarios. Despite these promising developments, their application within the realm of software systems remains limited, with current uses largely focused on areas like text enhancement, artwork creation, and basic summarization.

The software systems domain, encompassing software engineering, systems security, human-computer interaction (HCI), artificial intelligence, and computer graphics, involves intricate tasks that often require multimodal understanding and processing. From code analysis and software testing to user experience evaluation and cybersecurity threat detection, many of these tasks could potentially benefit from the advanced capabilities of LMMs. However, the specific challenges and opportunities presented by LMMs in these contexts have yet to be fully explored. To address this gap, a comprehensive study is necessary to understand how LMMs can enhance tasks within software systems, and to identify the architectural hurdles that must be overcome for their effective deployment.

In this vision paper, we explore the potential of LMMs to transform software systems research. We conduct a thorough review of literature spanning the past decade across related fields, constructing a task taxonomy that categorizes tasks likely to benefit from LMM capabilities. From this taxonomy, we identify representative tasks and evaluate their feasibility using a range of LMMs, such as GPT-4 Vision and Gemini Vision. Through systematic experimentation with prompt engineering, we assess the performance of these models and investigate the underlying challenges that limit their efficacy. Our work contributes to the field in the following ways:

- Task Taxonomy for LMMs in Software Systems: We develop a comprehensive taxonomy that identifies and categorizes tasks across software engineering, system security, HCI, and related fields that stand to benefit from LMM integration. This taxonomy offers a roadmap for researchers and practitioners seeking to leverage LMMs in their respective domains.
- Evaluation Framework for LMMs on Software Systems Tasks: We propose a set of evaluation criteria and experimental methods tailored to assess LMM performance on software systems tasks. By selecting representative tasks, we provide a framework for systematically test-

ing the capabilities of LMMs in real-world scenarios, including code analysis, software testing, and user experience assessment.

- Cross-Model Performance Analysis with Prompt Engineering: Using a range of LMMs, we perform a comparative analysis to understand how different models tackle similar tasks and the effectiveness of prompt engineering in enhancing their performance. This analysis sheds light on the strengths and limitations of current models, providing insights into how prompt engineering can be optimized for diverse tasks.

- Opportunities and Challenges of LMMs for Software Systems: Based on our empirical findings, we discuss the unique opportunities LMMs offer for advancing software systems research, particularly in multimodal environments such as Extended Reality (XR). We also identify the architectural challenges that hinder LMM performance, including issues related to multimodal data fusion, interpretability, and resource constraints, and propose directions for future research.

Our findings highlight the transformative potential of LMMs in software systems, paving the way for innovative applications and inspiring further exploration into this emerging field. By outlining both the current capabilities and limitations of LMMs, we aim to provide a foundation for future work that will drive the development of more intelligent, adaptive, and multimodal software solutions.

## 1.2 Background

Large Multimodal Models (LMMs) represent an evolution beyond traditional text-based Large Language Models (LLMs). In addition to supporting text input and output, LMMs can process inputs from multiple modalities such as images, audio, and video, generating corresponding multimodal outputs. Leading models, like GPT-4o[55], already support inputs from audio, images, and multiple image sources, while models such as Gemini[70] even handle video inputs. By extending the functionality of large models to cover multiple modalities, the scope and variety of tasks they can perform have significantly increased, including video summarization, image comprehension, and speech recognition.

The software engineering field encompasses a wide range of tasks, with the primary objective of ensuring high-quality software development and stable operation. Over the past few decades, software has evolved beyond simple command-line interfaces to incorporate graphical user interfaces (GUIs), animations, and voiceovers, which have become standard features. Consequently, relying solely on text-based code inputs has become increasingly insufficient for addressing the diverse needs of modern software systems, prompting the rise of multimodal inputs. For example, screenshots can be used to detect GUI issues[43], and video data can be analyzed to extract user gestures[7].

However, the integration of multimodal inputs into software engineering has been relatively slow. The primary challenge lies in the complexity and diversity of multimodal data, which has made it difficult for researchers to develop a unified and generalizable approach. Recent advances in machine learning have spurred efforts to combine machine learning models with multimodal input processing, such as using computer vision for object recognition in images. A key limitation of earlier approaches is that models were often task-specific, limiting their reusability across different contexts. The advent of multimodal large models offers a potential breakthrough. Numerous studiesneed citation have demonstrated the strong generalization capabilities of large models, showing that they can maintain high accuracy even with previously unseen tasks. As such, integrating LMMs into software engineering tasks is a logical next step. However, given that LMMs are still an emerging technology, there have been few attempts to explore their application in this domain. The goal of this paper is to address this gap and evaluate the performance of LMMs in software engineering tasks.

# 2 Related Work

In this section, we will provide an overview of how previous works utilize multimodal capabilities for problem-solving in software engineering. Then, we will discuss how LLMs can help address challenges in the software engineering domain. Finally, we will review the existing test benchmarks and evaluation criteria for assessing LMMs.

## 2.1 Ultizing Multimodal Ability in Software Engineering

Integrating multimodal capabilities, such as voice, gesture, and sentiment analysis, has emerged as a promising approach to enhancing software development processes and user experiences. Guglielmi et al. conducted automated tests on virtual personal assistants that use voice for interaction [25] . Qi et al. summarized recent research on gesture recognition through sensors and the analysis of image information [57] . Gandhi et al. investigated previous work on sentiment analysis, a domain encompassing the three modalities of text, vision, and audio working together to produce an effect [23] . However, these studies rely on specific mini-models or other traditional data analysis methods, which only perform relatively well on particular tasks or datasets. Those specialized models may lose good performance after migrating to other datasets or task settings of the same type [73] [58] . Our work reduces the expense of training different models for a specific problem by introducing LMMs with good generalization capabilities to handle different issues simultaneously.

## 2.2 LMM for Software

A branch of previous work has demonstrated that integrating LLM into the production and research of soft engineering has been a scorching trend [26] , from generating [45] [69] and pre-processing [81] [80] experimental data to using LLM as an agent for automated testing [68] [38] , all of which show that LLM has a solid potential to enhance existing soft engineering processes. Moreover, Jin et al. also illustrate the contribution that LLM can make in software design, testing, and maintenance [29]. In contrast to these studies, which only focus on specific tasks in specific domains of soft engineering and lack knowledge of what valuable tasks exist now, our work presents a systematic framework that defines what tasks are available to help optimize efficiency using LLM or LMM.

## 2.3 LMM Benchmark & Evaluation

LMMs combine information from different modalities, including text, vision, audio, and tactile, and analyze them to solve more complex real-world problems [78] [6] [79] . As a result, testing and evaluating LMMs' performance from different perspectives become a recent research interest. For instance, Wu et al. used the visual comprehension and language processing capabilities of GPT4v to test whether today's LMMs can support practical medical applications [76] . Cao et al. constructed Spider2-V, a test benchmark for LMM's ability to automate professional data science engineering workflows [9] . Cai et al. tested and improved the problem of the robustness of LMM's output when facing different styles of pictures [8] . These benchmarks and assessments have all achieved good performance in a single domain and can point out the shortcomings of LMM in the corresponding domain. Our work can complement the testing domains, bridging the gap of needing help harmonizing testing across domains and conducting migration tests.

# 3 Methodology

## 3.1 Task Taxonomy Construction

**Building the prototype of taxonomy** To build the prototype of the taxonomies, we have conducted systematic research on multimodality-related papers from four conferences (ICSE, FSE, ASE, and ISSTA) and two journals (TSE and TOSEM) in soft engineering over the last seven years [1]. We build a multimodality-related keyword list to screen the papers from these sources and manually collect 135 papers. To describe what types of tasks these papers in soft engineering are focusing on, we analyzed the papers according to the open coding procedures [19] used for qualitative data analysis. Specifically, we conducted a 5-round iterative manual analysis session involving three analysts with at least several years of development experience in the soft engineering field. In each iteration, every analyst separately summarizes what technical aspect of the paper belonged to the design, development, testing, maintenance, and repair [2] process of software from the Software Waterfall Model [53]. Each analyst analyzed two-thirds of the whole paper to ensure that each paper had been seen by at least two different analysts for cross-validation. Consequently, in the final iteration, we merge the research topics extracted in the previous rounds to form a prototype task tree resulting from our taxonomy. Finally, We used the results from 95 papers to build our taxonomy. At the top of our task tree are the

[1] from 2018.01.01 to 2024.05.15

[2] we add the ”repair” process to extend our taxonomy, which initially did not exist in the Waterfall Model

five software-building processes, followed by whether they are functional or non-functional [3]. The third level of categorization is based on modal information, such as “Vision” and “Vision with Audio.” At the bottom are up to four layers of progressively more detailed descriptions of specific technical aspects.

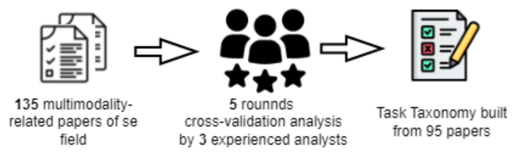


(a) Stage 1: Building task taxonomy protype

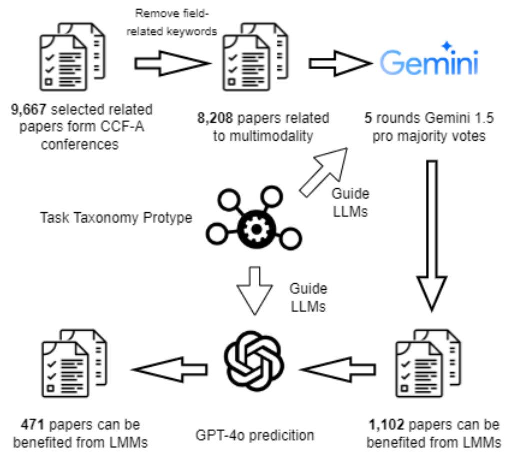


(b) Stage 2: Guiding LLM to analyze large-scale papers and extend our task taxonomy

Figure 1: Two stages workflow of building our task taxonomy

**Extending the list of papers** We expand the scope of our study to encompass all 37 A-level conferences and journals as classified by the China Computer Federation[4], with the same period considered. This inclusion covers five key domains: Computer Networks, Computer Graphics and Multimedia, Artificial Intelligence, Human-Computer Interaction, and Cross-cutting/Integrated/Emerging. Subsequently, we add software-engineering-related keywords to the search keyword list to cover a broader range of papers. We also remove the field-related keywords from the list for some specific domains. For example, we remove the keyword ”visual” from the

[3] standard is followed by ISO/IEC 25002:2024
[4] https://www.ccf.org.cn/Academic_Evaluation/By_category/

list of vision-related conferences, forming 8,208 pieces of paper. To reduce the number of papers and get a more concrete result, we involve the Gemini-1.5 as a judge to perform a 5-round majority vote, where we send the paper's title and guide it to predict whether this paper may focus on multimodal tasks using following prompt().

prompt for construct task tree

System : You are a helpful assistant designed to output JSON. You will be given a task tree generated from papers and a paper with its title and abstract. You are designed to answer the question:

What kind of task does the research task in the paper benefit from multi-modal AI to help process the target software/applications.

Please ensure the following rules while answering this question:

1. You have two kinds of action choices: output Matched if there is a node on the task tree matched the new task described in the paper. Otherwise output Add and the new task name in 1-5 levels to add a new node to the current task tree.
2. The first level of the task tree should use a combination (using with to connect) of terms from the modalities Vision, Text, Audio, and Tactile to describe the target modality the paper focuses on.
3. The second level of the task tree should be a broader technical concept term within its modality, avoiding the use of any specific software terms like AR, VR, or any specific software platform names (e.g., Android, Web, iOS).
4. The 'Function' in the tree describing the task should address either functional aspects, such as improvement, or non-functional aspects, such as accessibility.

**prompt for construct task tree**

System: Only output in 'Action': (Matched / Add), 'Function':(Functional / Non-functional) ,'1st':,'2nd':,'3rd':,'4th':,'5th:'(NA if not suitable) format.

User: The Task Tree:{SubTaskTree}. Paper: {TITLE} Abstract: {ABSTRACT}

**consolidating the taxonomy** This process helps to reduce the potential paper number to 1,102. Then, we perform another single-round GPT-4o prediction, where we prompt the LLM with the remaining paper's title and abstract to let the model know more about the details of the paper and make a more concrete prediction. Finally, we formulate an additional target multimodal related paper list with a size of 471, and the total paper list's size is 564.

Given that LLMs can identify latent patterns, [71] [52] we automate the expansion of our taxonomy by leveraging GPT-4o to learn these patterns from its prototypes. This process involves two distinct prompts. The first prompt instructs the LLM to analyze which software processes related to the technical aspect addressed in the paper may be relevant. In the subsequent step, we provide the subtask tree of the identified process, enabling the LLM to determine whether the aspect aligns with an existing child node or if a new child node is needed to describe it adequately. In each stage, the paper's title and abstract serve as user inputs, while the specific guidance for each part serves as system prompts. Additionally, we conduct manual checks to prune and merge misclassified results, ultimately consolidating the multimodal task taxonomy. Part of the task tree is shown in Figure. 2.

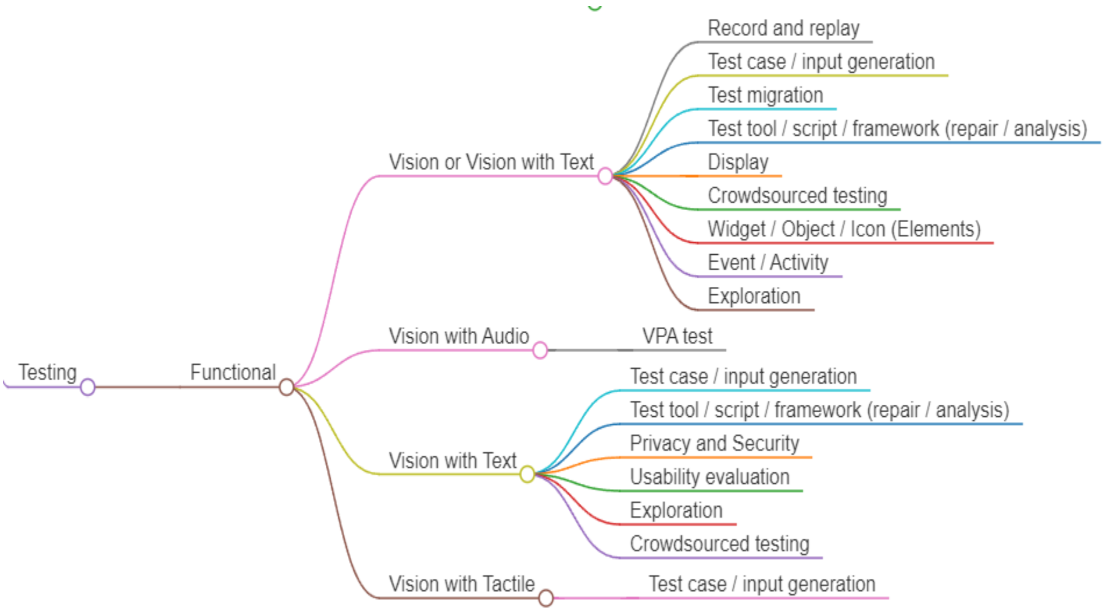


(a) Functional Testing Task Tree

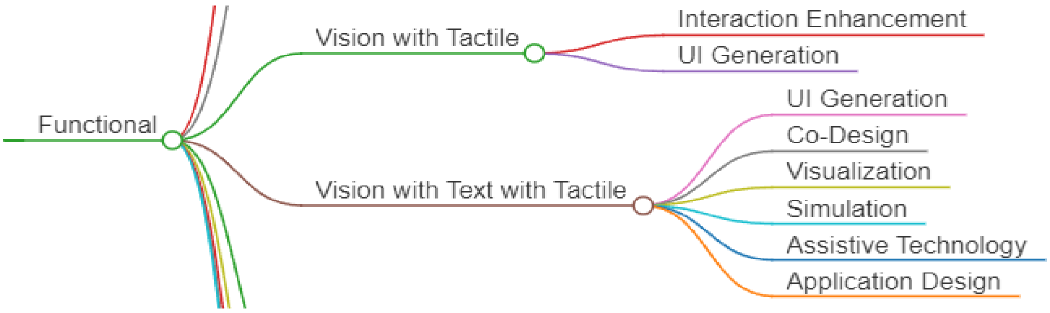


(b) Functional Design Task Tree

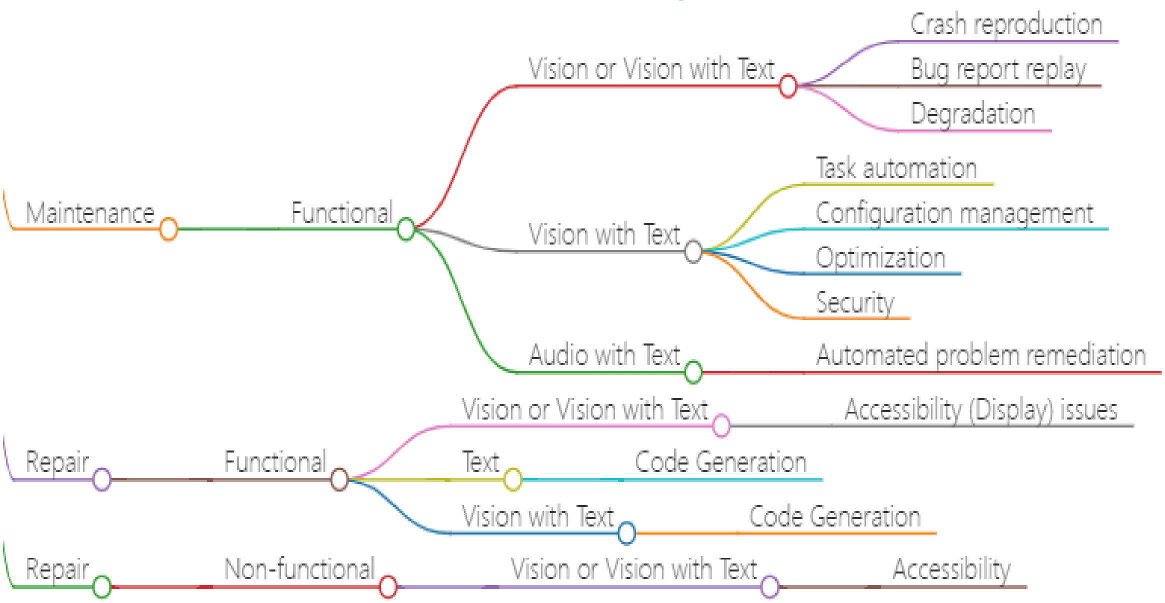


(c) Functional Maintenance, Functional and Non-Functional Repair Task Tree

Figure 2: Overview of our task tree up to 3rd level

## 3.2 Testing Framework

**Building the framework** The entire framework is built based on Python. To ensure the framework's high scalability—i.e., to ensure that our framework remains applicable as tasks in the multimodal field evolve—we have separated all task-related code, making the entire framework highly modular. This way, when new tasks need to be added or existing ones need to be modified in the future, only the corresponding task code requires adjustment. This significantly reduces the coupling between different code components, facilitating future modifications.

We have structured the workflow of the entire framework into three main components: data loading, model loading, and result evaluation. In the first component, data loading, we have developed specific loading functions tailored to different types of databases. Since various datasets may contain diverse data types, such as text, images, videos, or audio, we have implemented appropriate data processing in the Python scripts to ensure seamless integration with the models under test. For the model loading component, we have designed functions for both model initialization and request-response handling. These functions enable the model to select the appropriate data processing method based on the input data type and configuration file. For example, according to our standards, method 1 corresponds to pure image input combined with a system prompt, and more standard can be found in our released source code. Finally, in the result evaluation component, we have developed task-specific evaluation functions that efficiently and accurately assess the model's output, ensuring that the evaluation process aligns with the requirements of each task. Figure 3 is a flowchart about this framework.

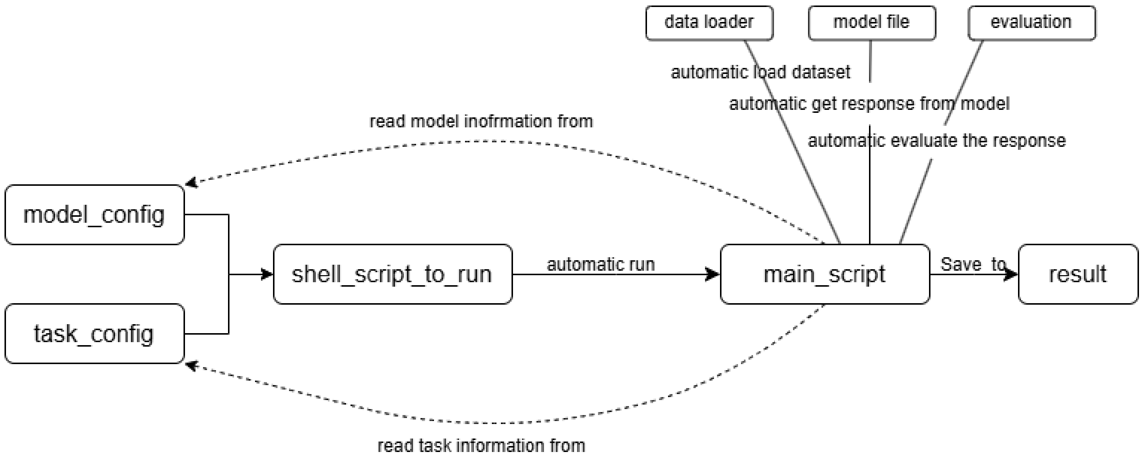


Figure 3: Framework's workflow

**Use the framework** We tried to minimize the complexity of our framework to ensure ease of use of it. Specifically, to initiate the framework, user only needs to fill in the corresponding configuration in the task config file. For example, user need to specify the task name, dataset list, model list and some other parameters in the task config file. Our framework will automatically read the corresponding parameters and perform the evaluation based on the specified task. If the user wants to add his own model for evaluation, all he has to do is to write the corresponding python file for the model to implement the relevant functions, then add the basic model information in the model config file(Algorithm 1).Similarly, if a user wants to use a new evaluation method or dataset, then only the corresponding documentation needs to be written. We have shown the sample code that needs to be implemented in each folder.

task config

task_name=YourTaskName

call_method=TheModality

system_prompt=”...”

max_token_length=MaxToken

dataset_name=’[”listOfDataLoaderPythonFile”]’

dataset_class=’[”className ofDataloader”]’

dataset_path=(”../path/to/dataset”)

batch_size=EvaluationBatchSize

eval_method=’[”evaluationPythonFileList”]’

eval_class=(”EvaluationClassName”)

middleDoc=true/false

middle_extension=txt/html....

model_list=(”listOfUsedModel”)

device=cuda

output_dir=/path/to/output

**Algorithm 1** Configuring `config.ini`

```
procedure Configure
  [SectionName] ← "name_that_will_be_used_in_the_task_config"
  call_type ← "api" or "local"
  if call_type is "api" then
    api_key ← "your_api_key_here"
    base_url ← "your_base_url_here"
    model_name ← "your_model_name_here"
  end if
  if call_type is "local" then
    conda_env_name ← "your_conda_environment_name"
    pretrained_path ← "path_to_pretrained_model"
  end if
  model_file_name ← "your_python_file_name_to_run_model_here"
  model_class ← "your_model_class_name_in_the_python_file_here"
end procedure=0
```

# 4 Experiments

In this section, we will illustrate our experiment settings. In section 4.1, we will show our target models, benchmarks, tasks, and evaluation metrics in this experiment. In section 4.2, we will demonstrate the prompts we used to guide LMMs in each task.

## 4.1 Setup

**Models** We selected 12 different LMMs as our experimental subjects. Each model can accept specific non-textual modalities as inputs and quiz the corresponding modal tasks. The information on all models is presented in Table 1. Due to time and financial constraints, in this report, we only selected 5 LMMs for our experiments, including four open-source models: LLaVA-NeXT-7B, Llama3.2 11B, 90B, InternVL-8B, and a closed-source LLM: GPT4o.

Table 1: An overview of our target model list

| Models | Parameters | Open Source? | Support Modalities |
|---|---|---|---|
| GPT-4-turbo-2024-04-09 | Not published | No | Text, Vision(image), Vision(video) |
| GPT-4o-2024-05-13 | Not published | No | Text, Vision(image),Vision(Video) |
| GPT-4o-audio-preview | Not published | No | Text, Audio |
| Claude-3-5-sonnet-20240620 | Not published | No | Text, Vision(image),Vision(video) |
| Gemini-1.5-pro | Not published | No | Text, Vision(image), Vision(video), Audio |
| InternVL2-pro | Not published | No | Text, Vision(image), Vision(video) |
| Qwen-turbo-2024-09-19 | Not published | No | Text, Vision(image), Vision(video), Audio |
| Llama-3.2-90B | 90B | Yes | Text, Vision(image) |
| Llama-3.2-11B | 11B | Yes | Text, Vision(image) |
| InternVL2-8B | 8B | Yes | Text, Vision(image), Vision(video) |
| LLaVA-NeXT-7B | 7B | Yes | Text, Vision(image) |
| Qwen VL | 8B | Yes | Text, Vision(image), Vision(video) |

**Datasets** To test the LMM more comprehensively, we extracted 53 usable datasets from our collection of 471 papers as our benchmarks. For each dataset, we summarize the modality involved, which stage of the Water-Fall model is of concern, and what type of software is targeted. Detailed information for each dataset is available in Table 2 and Table 3.

**Tasks List** In order to validate the capability of LMM on our test benchmark, we summarized 11 tasks based on previous work, each involving multimodal inputs. The details of the tasks are presented in Table 4. In this report, we selected five sub-tasks from the total task list to present our findings, as shown in Table 6. These five tasks cover four input modalities: text, single image, multiple images (video), and audio. To realize the tasks, we picked a subset of 5 datasets from our test benchmarks to experiment with, each subset containing about 100 inputs. The information on the subsets can be found in Table 5.

Table 2: An overview of our benchmark, TOOL indicates a tool that can generate dataset

| Dataset Source (Used in which paper) | Dataset Link | Component |
|---|---|---|
| Uibert [3] | Here | UI Image |
| Wukong-reader [4] | Here | Document Image |
| Defects4J [10] | Here | Spectrums |
| Evosuite [10] | Here | Spectrums |
| Rico [16] [20] | Here | UI Image |
| Gestonhmd [15] | Here | Gesture Movement Description |
| PLUR [18] | Here | Graph |
| Design What You Desire [17] | Here | Icon Image |
| DI-drive [22] | Here | images + RL |
| CMU Panoptic Studio [28] | Here | 3D skeletons, Sequences |
| Prose-benchmarks [30] | Here | Text, Table |
| Silentspeller [31] | Here | GT2k(HTK) style HMM |
| Marvis [33] | Here | Text |
| Nbsearch [37] | Here | Jupyter notebook |
| Sysevr [39] | Here | SeVCs |
| Sheetcopilot [34] | Here | Excel |
| Multiviz [40] | Here | Image |
| Poseexaminer [42] | Here | Image, JSON |
| StyleGAN [54] | Here | Image |
| ImageNet [56] | Here | Image |
| SparkBraille [60] | Here | braille charts |
| DroidBench [65] | Here | TOOL |
| Head Gestures Dystonia [66] | Here | Text |
| Screen2Words [72] | Here | Text, UI actions |
| Vetter [77] | Here | TOOL |
| Seenomaly [82] | Here | UI GIF |
| DroidGem [49] | Here | TOOL |
| FraudDroid [21] | Here | UI state transition graphs (UTG) |
| GUIGAN [83] | Here | UI image |
| Combodroid [74] | Here | TOOL |
| Themis Benchmarks [63] | Here | TOOL |
| Deep Q-network Testing [32] | Here | TOOL |
| Ape [24] | Here | TOOL |
| ωDroid [27] | Here | WebView-induced bugs |
| Video2Scenario [7] | Here | Image |
| ROUTE [41] | Here | TOOL |
| DatAndroid [2] | Here | Image, xml |
| Semantic Matching [51] | Here | GUI Image, event record |
| PSC2CODE dataset [5] | Here | Text, Video |

Table 3: An overview of our benchmark, TOOL indicates a tool that can generate dataset Cont.

| Dataset Source (Used in which paper) | Dataset Link | Component |
|---|---|---|
| Annotated MYST dataset [44] | Here | Text |
| EGFE [14] | Here | UI Image, Text Label |
| VITAS [36] | Here | Text |
| Asgaardlab [47] | Here | Image, canvas json file |
| AidUI [50] | Here | UI Image, DP label |
| Webevo [61] | Here | TOOL |
| Canvas Issues [48] | Here | URL, issue class |
| Vid2Xml [1] | Here | Video, xml |
| dVermin [64] | Here | UI Image |
| IconSeer [35] | Here | Icon Image |
| GLIB [13] | Here | Game UI Image |
| Owleye [43] | Here | UI Image |
| LabelDroid [11] | Here | Web Component |
| Design2Code [62] | Here | UI Image, Text |
| Glitchbench [67] | Here | Image |

Table 4: An overview of our target tasks list

| Task Name | Input Modalities | Output Modalities |
|---|---|---|
| UI to Code | Text, Visioin | Text |
| Display Bug/Glitch Detection | Text, Visioin | Text |
| Interactable UI Element Detection | Text, Visioin | Text |
| UI to Code Optimization | Text, Visioin | Text |
| UI Code transfer | Text, Visioin | Text |
| Image Based Agent / Interaction | Text, Visioin | Text |
| Cross-application interaction | Text, Visioin | Text |
| Voice Based Agent / Interaction | Text, Audio | Text |
| Completeness Exploration | Text, Visioin | Text |
| Event Detection | Text, Visioin | Text |
| Video Display Detection | Text, Video | Text |

Table 5: An overview of our sub-dataset list

| Dataset Name | Size | Component |
|---|---|---|
| Design2Code dataset [62] | 100 | Image, HTML |
| OwlEye dataset [43] | 102 | Image |
| Annotated RICO dataset [12] | 100 | Image, Text |
| PSC2CODE dataset [5] | 74 | Text,Video |
| VITAS dataset [36] | 100 | Text |

Table 6: An overview of our current tasks list

| Task Name | Input Modalities | Output Modalities |
|---|---|---|
| UI to Code | Text, Visioin | Text |
| Display Bug/Glitch Detection | Text, Visioin | Text |
| Interactable UI Element Detection | Text, Visioin | Text |
| Voice Based Agent / Interaction | Text, Audio | Text |
| Video Display Detection | Text, Video | Text |

**Evaluation Metrics** We followed the evaluation metrics set in the original paper to evaluate our experimental results. The details of the evaluation criteria can be found in Table 7.

Table 7: An overview of evaluation metrics in our experiment

| Task Name | Eval Metics |
|---|---|
| UI to Code | Design2Code Metric [62] |
| Display Bug/Glitch Detection | OwlEye Metric [43] |
| Interactable UI Element Detection | IoU (threshold 0.6) [12] |
| Voice Based Agent / Interaction | SeMaScore [59] |
| Video Display Detection | video display detect Metric [5] |

## 4.2 Prompt Engineering

As one of the most direct and critical factors influencing model performance, prompts need to be meticulously refined to ensure the model delivers its best performance on a given task. However, a significant challenge we currently face is the lack of a suitable and direct metric for quantifying the quality of a prompt. We can only make rough assessments based on the model's responses. Therefore, despite our best efforts in prompt engineering, there remains the possibility of better prompts existing than the ones we have crafted.

Nevertheless, for evaluation purposes, as long as we apply the same prompt across all models, fairness is maintained, and the data we obtain can still be considered meaningful and reliable. During our prompt engineering process, we made several interesting observations:

1. Including the word `REMEMBER` in the prompt helps the model better

adhere to our instructions, particularly when we expect the model to output in a specific format.

2. For more complex tasks, utilizing a *chain of thought*[75] approach improves the quality of responses.
3. Providing a detailed task description and considering all potential outputs, along with explicitly stating whether they are acceptable, leads to better performance.

Below are the prompts for all our tasks.

**Prompt used for conducting UI2Code**

System:You are an expert web developer who specializes in HTML and CSS.

A user will provide you with screenshot of a webpage.

You need to return a single html file that uses HTML and CSS to reproduce the given website.

Include all CSS code in the HTML file itself.

If it involves any images, use \rick.jpg as the placeholder.

Some images on the webpage are replaced with a blue rectangle as the placeholder, use \rick.jpg for those as well.

Do not hallucinate any dependencies to external files. You do not need to include JavaScript scripts for dynamic interactions.

Pay attention to things like size, text, position, and color of all the elements, as well as the overall layout.

Respond with the content of the HTML+CSS file:

**Prompt used for conducting Display Bug Detection**

System:You are an expert UI developer. A user will provide you with screenshot of a GUI.

You only need to return a result of 0 or 1

If the screenshot shows GUI display issues, you need to response 1, otherwise 0.

You do not need to include any other answer or explanation.

Pay attention to things like text overlap, blurred screen, missing image always occur during GUI rendering on different devices due to the software or hardware compatibility. It is the things negatively influence the app usability, resulting in poor user experience. Also, you also need to distinguish between normal GUI effects such as shadows and animations and GUI effects that are not expected to appear such as strange text and incorrect overlays.REMEMBER, you should never output words other than 0 or 1, or the program will collapse!!

Respond with the 0 or 1:

**Prompt used for conducting Interactable UI Element Detection**

System:You are an expert Android developer who specializes in UI design and will answer question in JSON format.

A user will provide you with screenshot of an application.

You need to return an object detections result that including all the bouding boxes of UI elements.

You can safely ignore those bounding boxes with too small region, i.e. region < 100

REMEMBER, respond in JSON format: [{‘id’:(the index of UI element you have detected), ‘bbox’:(the bounding boxes you have found. You should output the bounding boxes location in pixels level digitals. REMEMBER In format:[x_start, y_start, X_length, y_length])}...], and DO NOT output any comment other than json code.

**Prompt used for conducting Automatic Speech Recognization**

System:You are a helpful assistant that can understand audio recordings and preform automatic speech recognition and output the recognition result in text format.

User: What is in this recording? Only output the text you heard in the recording.

**Prompt used for conducting Video display issue detection**

System:You are an expert programmer. A user wants to get the code in a video, but some parts of this code are noisy (e.g. masking, blurring). So you need to identify the video frame by frame, marking the noisy frames as 0 and the clean frames as 1.

Your filtered video content will allow the user to extract all the code content in subsequent steps with the help of a simple screen recognition program.

You only need to return a result of 0 or 1,split with space. Remember, you should always responds with 0 or 1 or the program will crash!!!The total number of 1 and 0 should be ex actly same as the number of frames the user provide.

For example: 1 1 1 1 0 1 1 0 0... Respond with the 0 or 1:

# 5 Evaluation

In this section, we will detail our evaluation process and empirically explore the following two main research questions (RQs).

- **RQ1:** Where can software system development process and research benefit from large multimodal models?
- **RQ2:** To what extent do the LMMs have sufficient capabilities to help the multimodal software system development process and research?
  - **RQ2-1:** At **Text, Image** level, do the LMMs have sufficient capabilities to help the multimodal software system development process and research?
  - **RQ2-2:** At **Text, Video** level, do the LMMs have sufficient capabilities to help the multimodal software system development process and research?
  - **RQ2-3:** At **Text, Audio** level, do the LMMs have sufficient capabilities to help the multimodal software system development process and research?

## 5.1 RQ1: Where can software system development process and research benefit from large multimodal models?

Software system processes and research often involve analyzing multimodal information, and LMM, which combines the text comprehension capability of LLM with the ability to analyze multimodal information, is undoubtedly quite capable of optimizing this process. Therefore, in this section, we examine what research directions and processes might benefit from utilizing the capabilities of LMM.

As described in Section 3.1, we predicted whether the studies in the corresponding paper could benefit from the LMM's capabilities by guiding the LLM with a prototype of our taxonomy. After obtaining the predictions from LLM, we manually merged with three experienced evaluators to remove some unsuitable classifications. Consequently, we received a task tree [5] (demonstrated through markmap [6]) covering 176 secondary classifications. Our task tree covers four modalities (text, visual, audio, tactile) and five software processes (Design, Develop, Test, Maintain, and Repair). Researchers can easily find potential, unattended problems from the AI community.

**Answer to RQ1:** Our task tree demonstrates the software system development processes and research that can benefit from LMMs.

[5] https://storage.googleapis.com/testvideocuhk/demo/markmap.html
[6] https://markmap.js.org/repl

## 5.2 RQ2: To what extent do the LMMs have sufficient capabilities to help the multimodal software system development process and research?

In Section 5.1, we verified that LMMs can help many aspects of software system development and research. Therefore, it is necessary to measure whether today's LMMs can understand the corresponding modalities and to be able to accomplish the corresponding domain tasks. In this section, we evaluate the LMM in three different modality combinations: the primary text modality plus a specific modality: single image, multiple images (video), and audio. We design at least one task for each modality combination as a measure. In Section 5.2.1, we selected three tasks to assess the LMM's comprehension of text combined with a single image. In Section 5.2.2, we developed one task to evaluate the LMM's understanding of text alongside multiple images. In Section 5.2.3, we designed one task to measure the LMM's comprehension of text combined with audio.

### 5.2.1 RQ2-1: At Text, Image level, do the LMMs have sufficient capabilities to help the multimodal software system development process and research?

We conducted experiments on five LMMs that accept text and image input: UI2Code, Display Bug/Glitch Detection, and Interactable UI Element Detection.

**UI2Code** UI2Code requires the conversion of a given UI image into working HTML code. Following Si et al.'s setup, we instruct the LMM to read the UI image and generate the HTML code through a system prompt [62]. For the evaluation of the results, we followed the configuration in the paper and evaluthis five scoreveness of the generated code in five different dimensions, where the final score is the average of these five score:

- Block-Match: computing the total sizes of all matched blocks divided by the total sizes of all blocks.
- Text: computing character-level Sørensen-Dice similarity and averaging across all matched pairs.
- Position: computing IoU between matched pairs
- Color: computing following CIEDE2000 color difference formula [46]
- CLIP: high-level visual similarity through CLIP library [7]

Table 8: Experiment Result of UI2Code [62]

| Models | Final Score | Block-Match | Text | Position | Color | CLIP |
|---|---|---|---|---|---|---|
| GPT-4o-2024-05-13 | **0.887** | **0.907** | 0.972 | **0.855** | **0.822** | **0.879** |
| Llama3.2-11b | ≈0 | ≈0 | ≈0 | ≈0 | ≈0 | ≈0 |
| Llava-Next-7b | 0.735 | 0.665 | 0.846 | 0.69 | 0.641 | 0.834 |
| InternVL-8b | 0.149 | 0.0 | 0.0 | 0.0 | 0.0 | 0.746 |
| Llama3.2-90b | 0.54 | 0.357 | 0.61 | 0.486 | 0.437 | 0.812 |
| Baseline | 0.848 | 0.858 | **0.974** | 0.805 | 0.733 | 0.869 |

The experiment result of this sub-task is shown in Table 8. Of all five metrics, gpt4o did well inside 4, and the remaining Text metrics were close

[7]https://pypi.org/project/open-clip-torch/

to the baseline. Undoubtedly, in a classic soft-engineering domain task like UI2code, the state-of-the-art LMM GPT4o demonstrates a strong capability far beyond the baseline model, proving that the LMM is capable enough to make soft-engineering-related development as well as research gainful. Among the remaining LMMs, surprisingly, Llava-NeXT, with the smallest parameter size, is the best performer, even outperforming llama3.2 90B, which has more than 10 times the parameter size, demonstrating that small and even smaller mobile models have the potential to be utilized in a wide range of applications with low hardware requirements.

**Display Bug/Glitch Detection** UIDIsplay Issue Detection focuses on detecting potential display issues in given UI screenshots, such as texture loading failures, text rendering errors, or overlapping elements. In this task, the image recognition ability of large multimodal models (LMMs) becomes critical. For evaluation, we randomly sampled 100 images from the dataset constructed by Liu et al.[43], with an equal distribution of labels: 50% representing problematic UIs and 50% representing normal UIs. In this dataset, a label of 1 (true label) indicates that the presented UI screenshot contains display issues, while a label of 0 (false label) indicates that the UI is functioning normally.

As a baseline, we adopted the best-performing evaluation results reported by Liu et al.[43], which utilized a deep learning-based model. This serves as a benchmark to assess the accuracy of large multimodal models on this task. The experiment result of this sub-task is shown in Table 9 .

Table 9: Experiment Result of Display Bug/Glitch Detection [43]

| Models | Percision | Recall | F1 | TP | FP | FN |
|---|---|---|---|---|---|---|
| GPT-4o-2024-05-13 | **0.92** | 0.597 | 0.724 | 46 | 4 | 31 |
| Llama3.2-11b | ≈0 | ≈0 | ≈0 | ≈0 | ≈0 | ≈0 |
| Llava-Next-7b | 0.02 | 1 | 0.039 | 1 | 49 | 0 |
| InternVL-8b | 0 | 0 | 0 | 0 | 50 | 0 |
| Llama3.2-90b | 0.18 | 0.45 | 0.257 | 9 | 41 | 11 |
| Baseline | 0.850 | **0.848** | **0.849** | - | - | - |

**Interactable UI Element Detection** Interactable UI Element Detection aims to detect small elements inside a UI image and generate several bounding boxes to indicate them. We use system prompts to guide LMM in finding the suitable interactable UI element and generating bounding boxes. We followed Chen et al. to verify the result and compute the IoU between the truth and the predicted bounding boxes [12]. We adjusted the threshold in this experiment from 0.9 to 0.6 to allow for more mistakes LMM made. The experiment result of this sub-task is shown in Table 10.

Table 10: Experiment Result of Interactable UI Element Detection [12]

| Models | Percision | Recall | F1 | TP | FP | FN |
|---|---|---|---|---|---|---|
| GPT-4o-2024-05-13 | 0.014 | 0.017 | 0.016 | 13 | 918 | 730 |
| Llama3.2-11b | ≈0 | ≈0 | ≈0 | ≈0 | ≈0 | ≈0 |
| Llava-Next-7b | 0.0009 | 0.004 | 0.001 | 3 | 3411 | 740 |
| InternVL-8b | 0.002 | 0.009 | 0.003 | 7 | 3288 | 736 |
| Llama3.2-90b | 0 | 0 | 0 | 0 | 2373 | 743 |
| Baseline | **0.490** | **0.557** | **0.524** | - | - | - |

Compared to the baseline model [12], all LMMs performed poorly, with only GPT4o, as well as Llava-NeXT, reaching single-digit FPs, and none of the other models failing to achieve a single correct prediction, even after lowering the standard threshold for IoU in this experiment. This result suggests that LMM is much less effective at tasks requiring fine-level analysis of image data than those requiring only high-level image analysis and significantly less effective than small models with specialized specializations.

To improve the performance of LMMs, we envisioned two possible solutions: to provide more detailed prompt guidelines or to perform more detailed preprocessing of the image to reduce the pressure on the LMM to analyze the whole image. Another approach is to let the LMM play the role of an assistant to work with a specific model, where the LMM only performs the high-level task of determining whether a specific UI element exists and then calls a specific mini-model to generate accurate results. Both of these approaches can be used to improve the performance of specific aspects of the LMM in the future.

**Conclusion** On many previously unseen tasks, LMMs have already surpassed models specifically trained for those tasks, highlighting the feasibility and potential of applying multimodal large models in this domain. However, it is worth noting that in certain specialized tasks, such as small object detection, the performance of these large models is significantly suboptimal. This reveals a highly valuable direction for future research in improving their capabilities in such scenarios.

Another interesting finding is that of the five models tested in this section, all of the llama3.2 series models demonstrated in addition to poor comprehension, as evidenced by the inability to analyze the input instructions.

One of them, the 90B version, could understand the input instructions and output them in the required format due to its larger parameter size, but the results showed a complete lack of understanding of the task requirements. For example, in the UI element detection task, the output of llama 3.2-90B repeats the coordinates of the bounding boxes of the whole picture size, which ultimately fails to understand the detection of a specific element as required by the task's prompt, while the understanding of llama 3.2-11B is even worse, as all the outputs repeat the same paragraph. All the output is a repetition of a meaningless response, showing no understanding of the task requirements. In contrast, models such as Llava-NeXT have fewer parameters but show an understanding of the task setup and give a response. This finding warrants subsequent exploration of the token level of generation.

> **Answer to RQ2-1:** At the **Text and Image** level, LMMs can be experts on some specialized pre-trained tasks but are inferior to baseline methods for other tasks.

#### 5.2.2 RQ2-2: At Text, Video level, do the LMMs have sufficient capabilities to help the multimodal software system development process and research?

Video input is a unique modality that differs significantly from simply using multiple images as input. In videos, there is strong correlation and continuity between frames, requiring contextual understanding to interpret the content. For tasks involving this modality, we selected the task of detecting whether video frames are valid, as presented by Bao et al[5]. This task is a sub-task of a larger problem—extracting code from videos. Specifically, this task involves analyzing each frame of a video to determine whether it contains useful code content that needs to be extracted. If a frame contains such content, it is labeled as valid; otherwise, it is labeled as invalid.

It is important to note that this differs from analyzing a single image to determine its validity, as the validity of a video frame often depends on its context within the sequence, rather than solely on the content of the frame itself. As such, this task is well-suited for evaluating a model's capability in video understanding.

We used the video dataset provided by Bao et al. to conduct this evaluation[5]. Again, we use the data from the original paper as the baseline. Since current large models typically process video understanding by extracting frames and treating them as a set of images, we adopted the same frame extraction approach as Bao et al. to ensure experimental rigor. The evaluation metrics also follow the methodology presented in their paper. In this task, we use the label 1 to indicate that a frame is valid and the label 0 to indicate that a frame is invalid. The experiment result of this sub-task is shown in Table 11.

Table 11: Experiment Result of Video display detect [5]

| Models | Percision | Recall | F1 | TP | FP | FN |
|---|---|---|---|---|---|---|
| GPT-4o-2024-05-13 | 0.891 | **0.891** | 0.891 | 57 | 7 | 7 |
| InternVL-8b | **0.938** | 0.857 | **0.895** | 60 | 4 | 10 |
| Baseline | 0.91 | 0.85 | 0.88 | 2459 | 256 | 445 |

It is important to note that, although InternVL-8B appears to perform exceptionally well at first glance— even surpassing GPT4o in video understanding—this is actually due to an imbalance in the dataset, where the number of invalid frames is insufficient, resulting in too few negative samples. In practice, during real testing, InternVL often merely predicts most frames as valid and, in some cases, fails to generate predictions corresponding to the number of input frames. This behavior leads to seemingly impressive performance metrics for the model.

However, it is undeniable that the multimodal large model still performs very well on this task, requiring no additional training at all while maintaining a very high accuracy.

**Answer to RQ2-2:** At the text and video level, the LMMs have very strong potential to assist in this area, even achieving performance comparable to the baseline.

#### 5.2.3 RQ2-3: At Text, Audio level, do the LMMs have sufficient capabilities to help the multimodal software system development process and research?

In order to serve the user like a Virtual personal assistant (VPA), an audio-capable LMM should be able to recognize the same meaning in different speech inputs, e.g., ”What’s up today?” and ”Tell me the news of the day” should trigger the same news-playing state of a VPA. Based on the work of Guglielmi et al., we designed a series of tests to check whether the LMM is good at detecting the corresponding trigger state in the input text [25]. However, since the evaluation criterion used in the paper is to obtain the truth label through a textual conversation with the AWS skill VPA [8] in the simulator, in this experiment, we only tested the speech recognition accuracy of the LMM and then multiplied it by the original result that the LMM can test the accuracy and comprehensiveness of the VPA. For the sake of rigor, we only report the results of the Automatic Speech Recognition (ASR) part of the experiment in this report Table 12. We used SemaScore [59] as the evaluation metric of ASR accuracy, a criterion for determining the accuracy of ASR work from the language model token level.

Table 12: Experiment Result of Automatic Speech Recognition (ASR) [59]

| **Models** | **SemaScore** |
|---|---|
| GPT-4o-audio-preview | **0.9583** |

The experiment result indicates that LMM, like GPT-4o, has good speech recognition capabilities, and we believe that future LMMs can be used as VPAs to provide a broader range of services. However, the current support for speech input and the ability to analyze the non-textual information of speech are still lacking, and only some of the fine-tuned mini-models [85] [84] have good performance in this area. There is still a certain distance from the performance of solving text-level problems like LLM.

**Answer to RQ2-3:** LMM can understand text information inside audio, so LMM has sufficient capabilities to help the multimodal software system development process and research.

In summary, large multimodal models have demonstrated exceptional capabilities in integrating with the field of software engineering across cur-

[8] https://explore.skillbuilder.aws/learn

rently supported input modalities, including images, videos, and audio. For previously unseen tasks, these models can provide accurate and rapid responses based solely on prompts, often without requiring additional fine-tuning. Large multimodal models hold tremendous potential in the software engineering domain, offering researchers an incredibly practical tool for handling multimodal inputs without the need for extensive effort and time spent on fine-tuning or debugging.

**Answer to RQ2:** Large multimodal models exhibit exceptional performance within currently supported modalities, highlighting significant potential for integration into the field of software engineering.

# 6 Conclusion & Future Work

## 6.1 Conclusion

This report proposes a task taxonomy and a corresponding task tree of classification results, bridging the gap of the lack of an explicit specification for the application of LMMs in the field of software engineering. We construct a testing framework for evaluating LMMs that allows developers to test LMMs more flexibly by combining different datasets and evaluation criteria. We perform systematic qualitative testing of existing LMMs by constructing a cross-modal test benchmark. Our experimental results show that existing LMMs perform well on certain specific, extensively researched tasks but perform much worse than the baseline on underappreciated tasks. This finding emphasizes the importance of a more comprehensive and nuanced assessment of LMMs. In addition, the LMM's ability to understand and execute multiple multimodal tasks is encouraging, and we look forward to expanding the LMM's capabilities to more complex environments, such as analyzing simultaneous inputs from more than two modalities in XR software, with the collaborative efforts of the entire AI community.

Through our work, we aim to inspire more software developers to explore using large multimodal models (LMMs) to accelerate project development and research. We believe that integrating LMMs into software engineering will unlock new possibilities, setting the stage for groundbreaking developments in research and application.

## 6.2 Future Work

In the future, we will further refine our taxonomy to check for instances of LLM misclassification due to random numbers, such as when an error excludes a potential research direction. We will also generalize and propose more tasks from the existing datasets and task trees and tasks that cover more modalities. In addition, we will further expand the number of test models and benchmark size to conduct a complete set of experiments covering the entire dataset.

## References

[1] Mohammad D Alahmadi. Vid2xml: Automatic extraction of a complete xml data from mobile programming screencasts. *IEEE Transactions on Software Engineering*, 49(4):1726–1740, 2022. 27

[2] Luca Ardito, Andrea Bottino, Riccardo Coppola, Fabrizio Lamberti, Francesco Manigrasso, Lia Morra, and Marco Torchiano. Feature matching-based approaches to improve the robustness of android visual gui testing. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 31(2):1–32, 2021. 26

[3] Chongyang Bai, Xiaoxue Zang, Ying Xu, Srinivas Sunkara, Abhinav Rastogi, Jindong Chen, et al. Uibert: Learning generic multimodal representations for ui understanding. *arXiv preprint arXiv:2107.13731*, 2021. 26

[4] Haoli Bai, Zhiguang Liu, Xiaojun Meng, Wentao Li, Shuang Liu, Nian Xie, Rongfu Zheng, Liangwei Wang, Lu Hou, Jiansheng Wei, et al. Wukong-reader: Multi-modal pre-training for fine-grained visual document understanding. *arXiv preprint arXiv:2212.09621*, 2022. 26

[5] Lingfeng Bao, Zhenchang Xing, Xin Xia, David Lo, Minghui Wu, and Xiaohu Yang. psc2code: Denoising code extraction from programming screencasts. *ACM Transactions on Software Engineering and Methodology*, 29(3):1–38, June 2020. 26, 27, 28, 41, 42

[6] Arnab Barua, Mobyen Uddin Ahmed, and Shahina Begum. A systematic literature review on multimodal machine learning: Applications, challenges, gaps and future directions. *IEEE Access*, 11:14804–14831, 2023. 12

[7] Carlos Bernal-Cárdenas, Nathan Cooper, Madeleine Havranek, Kevin Moran, Oscar Chaparro, Denys Poshyvanyk, and Andrian Marcus. Translating video recordings of complex mobile app ui gestures into replayable scenarios. *IEEE Transactions on Software Engineering*, 49(4):1782–1803, 2022. 9, 26

[8] Rizhao Cai, Zirui Song, Dayan Guan, Zhenhao Chen, Yaohang Li, Xing Luo, Chenyu Yi, and Alex Kot. Benchlmm: Benchmarking cross-style visual capability of large multimodal models. In *European Conference on Computer Vision*, pages 340–358. Springer, 2025. 13

[9] Ruisheng Cao, Fangyu Lei, Haoyuan Wu, Jixuan Chen, Yeqiao Fu, Hongcheng Gao, Xinzhuang Xiong, Hanchong Zhang, Yuchen Mao, Wenjing Hu, Tianbao Xie, Hongshen Xu, Danyang Zhang, Sida Wang, Ruoxi Sun, Pengcheng Yin, Caiming Xiong, Ansong Ni, Qian Liu, Victor Zhong, Lu Chen, Kai Yu, and Tao Yu. Spider2-v: How far are multimodal agents from automating data science and engineering workflows?, 2024. 13

[10] Prantik Chatterjee, José Campos, Rui Abreu, and Subhajit Roy. Augmenting automated spectrum based fault localization for multiple faults. In *IJCAI*, pages 3140–3148, 2023. 26

[11] Jieshan Chen, Chunyang Chen, Zhenchang Xing, Xiwei Xu, Liming Zhu, Guoqiang Li, and Jinshui Wang. Unblind your apps: Predicting natural-language labels for mobile gui components by deep learning. In *Proceedings of the ACM/IEEE 42nd international conference on software engineering*, pages 322–334, 2020. 27

[12] Jieshan Chen, Mulong Xie, Zhenchang Xing, Chunyang Chen, Zhu Liming Xu, Xiwei, and Guoqiang Li. Object detection for graphical user interface: Old fashioned or deep learning or a combination? In *Proceedings of the 2020 28th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, New York, NY, 2020. ACM. 27, 28, 39

[13] Ke Chen, Yufei Li, Yingfeng Chen, Changjie Fan, Zhipeng Hu, and Wei Yang. Glib: towards automated test oracle for graphically-rich applications. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 1093–1104, 2021. 27

[14] Liuqing Chen, Yunnong Chen, Shuhong Xiao, Yaxuan Song, Lingyun Sun, Yankun Zhen, Tingting Zhou, and Yanfang Chang. Egfe: End-to-end grouping of fragmented elements in ui designs with multimodal learning. In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*, pages 1–12, 2024. 27

[15] Taizhou Chen, Lantian Xu, Xianshan Xu, and Kening Zhu. Gestonhmd: Enabling gesture-based interaction on low-cost vr head-mounted display. *IEEE Transactions on Visualization and Computer Graphics*, 27(5):2597–2607, 2021. 26

[16] Yan Chen, Maulishree Pandey, Jean Y Song, Walter S Lasecki, and Steve Oney. Improving crowd-supported gui testing with structural guidance. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2020. 26

[17] Yinpeng Chen, Zhiyu Pan, Min Shi, Hao Lu, Zhiguo Cao, and Weicai Zhong. Design what you desire: Icon generation from orthogonal application and theme labels. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 2536–2546, 2022. 26

[18] Zimin Chen, Vincent J Hellendoorn, Pascal Lamblin, Petros Maniatis, Pierre-Antoine Manzagol, Daniel Tarlow, and Subhodeep Moitra. Plur: A unifying, graph-based view of program learning, understanding, and repair. *Advances in Neural Information Processing Systems*, 34:23089–23101, 2021. 26

[19] John W Creswell and Cheryl N Poth. *Qualitative inquiry and research design: Choosing among five approaches*. Sage publications, 2016. 14

[20] Biplab Deka, Zifeng Huang, Chad Franzen, Joshua Hibschman, Daniel Afergan, Yang Li, Jeffrey Nichols, and Ranjitha Kumar. Rico: A mobile app dataset for building data-driven design applications. In

*Proceedings of the 30th annual ACM symposium on user interface software and technology*, pages 845–854, 2017. 26

[21] Feng Dong, Haoyu Wang, Li Li, Yao Guo, Tegawendé F Bissyandé, Tianming Liu, Guoai Xu, and Jacques Klein. Frauddroid: Automated ad fraud detection for android apps. In *Proceedings of the 2018 26th ACM joint meeting on European software engineering conference and symposium on the foundations of software engineering*, pages 257–268, 2018. 26

[22] Sevinc Eroglu, Arthur Voigt, Benjamin Weyers, and Torsten W Kuhlen. Vrscenariobuilder: Free-hand immersive authoring tool for scenario-based testing of automated vehicles. In *2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*, pages 196–202. IEEE, 2024. 26

[23] Ankita Gandhi, Kinjal Adhvaryu, Soujanya Poria, Erik Cambria, and Amir Hussain. Multimodal sentiment analysis: A systematic review of history, datasets, multimodal fusion methods, applications, challenges and future directions. *Information Fusion*, 91:424–444, 2023. 11

[24] Tianxiao Gu, Chengnian Sun, Xiaoxing Ma, Chun Cao, Chang Xu, Yuan Yao, Qirun Zhang, Jian Lu, and Zhendong Su. Practical gui testing of android applications via model abstraction and refinement. In *2019 IEEE/ACM 41st International Conference on Software Engineering (ICSE)*, pages 269–280. IEEE, 2019. 26

[25] Emanuela Guglielmi, Giovanni Rosa, Simone Scalabrino, Gabriele Bavota, and Rocco Oliveto. Help them understand: Testing and improving voice user interfaces. *ACM Transactions on Software Engineering and Methodology*, 33(6):1–33, 2024. 11, 43

[26] Xinyi Hou, Yanjie Zhao, Yue Liu, Zhou Yang, Kailong Wang, Li Li, Xiapu Luo, David Lo, John Grundy, and Haoyu Wang. Large language models for software engineering: A systematic literature review. *ACM Trans. Softw. Eng. Methodol.*, September 2024. Just Accepted. 12

[27] Jiajun Hu, Lili Wei, Yepang Liu, Shing-Chi Cheung, and Huaxun Huang. A tale of two cities: How webview induces bugs to android applications. In *Proceedings of the 33rd ACM/IEEE International Conference on Automated Software Engineering*, pages 702–713, 2018. 26

[28] Pascal Jansen, Julian Britten, Alexander Häusele, Thilo Segschneider, Mark Colley, and Enrico Rukzio. Autovis: Enabling mixed-immersive analysis of automotive user interface interaction studies. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–23, 2023. 26

[29] Haolin Jin, Linghan Huang, Haipeng Cai, Jun Yan, Bo Li, and Huaming Chen. From llms to llm-based agents for software engineering: A survey of current, challenges and future. *arXiv preprint arXiv:2408.02479*, 2024. 12

[30] Harshit Joshi, José Cambronero Sanchez, Sumit Gulwani, Vu Le, Gust Verbruggen, and Ivan Radiček. Repair is nearly generation: Multilingual program repair with llms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 5131–5140, 2023. 26

[31] Naoki Kimura, Tan Gemicioglu, Jonathan Womack, Richard Li, Yuhui Zhao, Abdelkareem Bedri, Zixiong Su, Alex Olwal, Jun Rekimoto, and Thad Starner. Silentspeller: Towards mobile, hands-free, silent speech text entry using electropalatography. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–19, 2022. 26

[32] Yuanhong Lan, Yifei Lu, Zhong Li, Minxue Pan, Wenhua Yang, Tian Zhang, and Xuandong Li. Deeply reinforcing android gui testing with deep reinforcement learning. In *Proceedings of the 46th IEEE/ACM International Conference on Software Engineering*, pages 1–13, 2024. 26

[33] Ricardo Langner, Marc Satkowski, Wolfgang Büschel, and Raimund Dachselt. Marvis: Combining mobile devices and augmented reality for visual data analysis. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–17, 2021. 26

[34] Hongxin Li, Jingran Su, Yuntao Chen, Qing Li, and ZHAO-XIANG ZHANG. Sheetcopilot: Bringing software productivity to the next

level through large language models. *Advances in Neural Information Processing Systems*, 36, 2024. 26

[35] Linlin Li, Ruifeng Wang, Xian Zhan, Ying Wang, Cuiyun Gao, Sinan Wang, and Yepang Liu. What you see is what you get? it is not the case! detecting misleading icons for mobile applications. In *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 538–550, 2023. 27

[36] Suwan Li, Lei Bu, Guangdong Bai, Zhixiu Guo, Kai Chen, and Hanlin Wei. Vitas: Guided model-based vui testing of vpa apps. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, ASE '22, New York, NY, USA, 2023. Association for Computing Machinery. 27

[37] Xingjun Li, Yuanxin Wang, Hong Wang, Yang Wang, and Jian Zhao. Nbsearch: Semantic search and visual exploration of computational notebooks. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2021. 26

[38] Yuanchun Li, Hao Wen, Weijun Wang, Xiangyu Li, Yizhen Yuan, Guohong Liu, Jiacheng Liu, Wenxing Xu, Xiang Wang, Yi Sun, et al. Personal llm agents: Insights and survey about the capability, efficiency and security. *arXiv preprint arXiv:2401.05459*, 2024. 12

[39] Zhen Li, Deqing Zou, Shouhuai Xu, Hai Jin, Yawei Zhu, and Zhaoxuan Chen. Sysevr: A framework for using deep learning to detect

software vulnerabilities. *IEEE Transactions on Dependable and Secure Computing*, 19(4):2244–2258, 2021. 26

[40] Paul Pu Liang, Yiwei Lyu, Gunjan Chhablani, Nihal Jain, Zihao Deng, Xingbo Wang, Louis-Philippe Morency, and Ruslan Salakhutdinov. Multiviz: Towards user-centric visualizations and interpretations of multimodal models. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–21, 2023. 26

[41] Jun-Wei Lin, Navid Salehnamadi, and Sam Malek. Route: Roads not taken in ui testing. *ACM Transactions on Software Engineering and Methodology*, 32(3):1–25, 2023. 26

[42] Qihao Liu, Adam Kortylewski, and Alan L Yuille. Poseexaminer: Automated testing of out-of-distribution robustness in human pose and shape estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 672–681, 2023. 26

[43] Zhe Liu, Chunyang Chen, Junjie Wang, Yuekai Huang, Jun Hu, and Qing Wang. Owl eyes: spotting ui display issues via visual understanding. In *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering*, ASE '20, page 398–409. ACM, December 2020. 9, 27, 28, 38, 39

[44] Zhijie Liu, Liang Feng Zhang, and Yutian Tang. Enhancing malware detection for android apps: Detecting fine-granularity malicious com-

ponents. In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 1212–1224. IEEE, 2023. 27

[45] Lin Long, Rui Wang, Ruixuan Xiao, Junbo Zhao, Xiao Ding, Gang Chen, and Haobo Wang. On llms-driven synthetic data generation, curation, and evaluation: A survey, 2024. 12

[46] M Ronnier Luo, Guihua Cui, and Bryan Rigg. The development of the cie 2000 colour-difference formula: Ciede2000. *Color Research & Application: Endorsed by Inter-Society Color Council, The Colour Group (Great Britain), Canadian Society for Color, Color Science Association of Japan, Dutch Society for the Study of Color, The Swedish Colour Centre Foundation, Colour Society of Australia, Centre Français de la Couleur*, 26(5):340–350, 2001. 37

[47] Finlay Macklon, Mohammad Reza Taesiri, Markos Viggiato, Stefan Antoszko, Natalia Romanova, Dale Paas, and Cor-Paul Bezemer. Automatically detecting visual bugs in html5 canvas games. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, pages 1–11, 2022. 27

[48] Finlay Macklon, Markos Viggiato, Natalia Romanova, Chris Buzon, Dale Paas, and Cor-Paul Bezemer. A taxonomy of testable html5 canvas issues. *IEEE Transactions on Software Engineering*, 49(6):3647–3659, 2023. 27

[49] Vikas K Malviya, Yan Naing Tun, Chee Wei Leow, Ailys Tee Xynyn, Lwin Khin Shar, and Lingxiao Jiang. Fine-grained in-context permission classification for android apps using control-flow graph embedding. In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, pages 1225–1237. IEEE, 2023. 26

[50] SM Hasan Mansur, Sabiha Salma, Damilola Awofisayo, and Kevin Moran. Aidui: Toward automated recognition of dark patterns in user interfaces. In *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*, pages 1958–1970. IEEE, 2023. 27

[51] Leonardo Mariani, Ali Mohebbi, Mauro Pezzè, and Valerio Terragni. Semantic matching of gui events for test reuse: are we there yet? In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 177–190, 2021. 26

[52] Suvir Mirchandani, Fei Xia, Pete Florence, Brian Ichter, Danny Driess, Montserrat Gonzalez Arenas, Kanishka Rao, Dorsa Sadigh, and Andy Zeng. Large language models as general pattern machines. *arXiv preprint arXiv:2307.04721*, 2023. 18

[53] Waterfall Model. Waterfall model. *Luettavissa: http://www. waterfall-model. com/. Luettu*, 3, 2015. 14

[54] Mohammad Amin Mozaffari, Xinyuan Zhang, Jinghui Cheng, and Jin LC Guo. Ganspiration: Balancing targeted and serendipitous inspiration in user interface design with style-based generative adversar-

ial network. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2022. 26

[55] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele,

Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O'Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass,

Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. 9

[56] Sara Pieri, Jose Restom, Samuel Horvath, and Hisham Cholakkal. Handling data heterogeneity via architectural design for federated visual recognition. *Advances in Neural Information Processing Sys-*

*tems*, 36:4115–4136, 2023. 26

[57] Jing Qi, Li Ma, Zhenchao Cui, and Yushu Yu. Computer vision-based hand gesture recognition for human-robot interaction: a review. *Complex & Intelligent Systems*, 10(1):1581–1606, 2024. 11

[58] Patrik Reizinger, Szilvia Ujváry, Anna Mészáros, Anna Kerekes, Wieland Brendel, and Ferenc Huszár. Position: Understanding llms requires more than statistical generalization, 2024. 11

[59] Zitha Sasindran, Harsha Yelchuri, and T. V. Prabhakar. Semascore: A new evaluation metric for automatic speech recognition tasks. In *Interspeech 2024*, interspeech$_2024, page 4558–4562. ISCA, September$ 2024. 28, 44

[60] JooYoung Seo, Yilin Xia, Bongshin Lee, Sean McCurry, and Yu Jun Yam. Maidr: Making statistical visualizations accessible with multimodal data representation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–22, 2024. 26

[61] Fei Shao, Rui Xu, Wasif Haque, Jingwei Xu, Ying Zhang, Wei Yang, Yanfang Ye, and Xusheng Xiao. Webevo: taming web application evolution via detecting semantic structure changes. In *Proceedings of the 30th ACM SIGSOFT International Symposium on Software Testing and Analysis*, pages 16–28, 2021. 27

[62] Chenglei Si, Yanzhe Zhang, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. Design2code: How far are we from automating front-end engineering?,

2024. 27, 28, 37

[63] Ting Su, Jue Wang, and Zhendong Su. Benchmarking automated gui testing for android against real-world bugs. In *Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering*, pages 119–130, 2021. 26

[64] Yuhui Su, Chunyang Chen, Junjie Wang, Zhe Liu, Dandan Wang, Shoubin Li, and Qing Wang. The metamorphosis: Automatic detection of scaling issues for mobile apps. In *Proceedings of the 37th IEEE/ACM International Conference on Automated Software Engineering*, pages 1–12, 2022. 27

[65] Cong Sun, Yuwan Ma, Dongrui Zeng, Gang Tan, Siqi Ma, and Yafei Wu. $\mu$dep: Mutation-based dependency generation for precise taint analysis on android native code. *IEEE Transactions on Dependable and Secure Computing*, 20(2):1461–1475, 2022. 26

[66] Qin Sun, Yunqi Hu, Mingming Fan, Jingting Li, and Su-Jing Wang. "can it be customized according to my motor abilities?": Toward designing user-defined head gestures for people with dystonia. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pages 1–11, 2024. 26

[67] Mohammad Reza Taesiri, Tianjun Feng, Cor-Paul Bezemer, and Anh Nguyen. Glitchbench: Can large multimodal models detect video game glitches? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22444–22455, 2024. 27

[68] Yashar Talebirad and Amirhossein Nadiri. Multi-agent collaboration: Harnessing the power of intelligent llm agents. *arXiv preprint arXiv:2306.03314*, 2023. 12

[69] Ruixiang Tang, Xiaotian Han, Xiaoqian Jiang, and Xia Hu. Does synthetic data generation of llms help clinical text mining?, 2023. 12

[70] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, Orhan Firat, James Molloy, Michael Isard, Paul R. Barham, Tom Hennigan, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, Ryan Doherty, Eli Collins, Clemens Meyer, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Jack Krawczyk, Cosmo Du, Ed Chi, Heng-Tze Cheng, Eric Ni, Purvi Shah, Patrick Kane, Betty Chan, Manaal Faruqui, Aliaksei Severyn, Hanzhao Lin, YaGuang Li, Yong Cheng, Abe Ittycheriah, Mahdis Mahdieh, Mia Chen, Pei Sun, Dustin Tran, Sumit Bagri, Balaji Lakshminarayanan, Jeremiah Liu, Andras Orban, Fabian Güra, Hao Zhou, Xinying Song, Aurelien Boffy, Harish Ganapathy, Steven Zheng, HyunJeong Choe, Ágoston Weisz, Tao Zhu, Yifeng Lu, Siddharth Gopal, Jarrod Kahn, Maciej Kula, Jeff Pitman, Rushin Shah, Emanuel Taropa, Majd Al Merey, Martin Baeuml, Zhifeng Chen, Laurent El Shafey, Yujing Zhang, Olcan Sercinoglu, George Tucker, Enrique Piqueras, Maxim Krikun, Iain Barr, Nikolay Savinov, Ivo Danihelka,

Becca Roelofs, Anaïs White, Anders Andreassen, Tamara von Glehn, Lakshman Yagati, Mehran Kazemi, Lucas Gonzalez, Misha Khalman, Jakub Sygnowski, Alexandre Frechette, Charlotte Smith, Laura Culp, Lev Proleev, Yi Luan, Xi Chen, James Lottes, Nathan Schucher, Federico Lebron, Alban Rrustemi, Natalie Clay, Phil Crone, Tomas Kocisky, Jeffrey Zhao, Bartek Perz, Dian Yu, Heidi Howard, Adam Bloniarz, Jack W. Rae, Han Lu, Laurent Sifre, Marcello Maggioni, Fred Alcober, Dan Garrette, Megan Barnes, Shantanu Thakoor, Jacob Austin, Gabriel Barth-Maron, William Wong, Rishabh Joshi, Rahma Chaabouni, Deeni Fatiha, Arun Ahuja, Gaurav Singh Tomar, Evan Senter, Martin Chadwick, Ilya Kornakov, Nithya Attaluri, Iñaki Iturrate, Ruibo Liu, Yunxuan Li, Sarah Cogan, Jeremy Chen, Chao Jia, Chenjie Gu, Qiao Zhang, Jordan Grimstad, Ale Jakse Hartman, Xavier Garcia, Thanumalayan Sankaranarayana Pillai, Jacob Devlin, Michael Laskin, Diego de Las Casas, Dasha Valter, Connie Tao, Lorenzo Blanco, Adrià Puigdomènech Badia, David Reitter, Mianna Chen, Jenny Brennan, Clara Rivera, Sergey Brin, Shariq Iqbal, Gabriela Surita, Jane Labanowski, Abhi Rao, Stephanie Winkler, Emilio Parisotto, Yiming Gu, Kate Olszewska, Ravi Addanki, Antoine Miech, Annie Louis, Denis Teplyashin, Geoff Brown, Elliot Catt, Jan Balaguer, Jackie Xiang, Pidong Wang, Zoe Ashwood, Anton Briukhov, Albert Webson, Sanjay Ganapathy, Smit Sanghavi, Ajay Kannan, Ming-Wei Chang, Axel Stjerngren, Josip Djolonga, Yuting Sun, Ankur Bapna, Matthew Aitchison, Pedram Pejman, Henryk Michalewski, Tianhe Yu, Cindy Wang, Juliette Love, Junwhan Ahn, Dawn Bloxwich, Kehang Han, Peter Humphreys, Thibault Sel-

lam, James Bradbury, Varun Godbole, Sina Samangooei, Bogdan Damoc, Alex Kaskasoli, Sébastien M. R. Arnold, Vijay Vasudevan, Shubham Agrawal, Jason Riesa, Dmitry Lepikhin, Richard Tanburn, Srivatsan Srinivasan, Hyeontaek Lim, Sarah Hodkinson, Pranav Shyam, Johan Ferret, Steven Hand, Ankush Garg, Tom Le Paine, Jian Li, Yujia Li, Minh Giang, Alexander Neitz, Zaheer Abbas, Sarah York, Machel Reid, Elizabeth Cole, Aakanksha Chowdhery, Dipanjan Das, Dominika Rogozińska, Vitaliy Nikolaev, Pablo Sprechmann, Zachary Nado, Lukas Zilka, Flavien Prost, Luheng He, Marianne Monteiro, Gaurav Mishra, Chris Welty, Josh Newlan, Dawei Jia, Miltiadis Allamanis, Clara Huiyi Hu, Raoul de Liedekerke, Justin Gilmer, Carl Saroufim, Shruti Rijhwani, Shaobo Hou, Disha Shrivastava, Anirudh Baddepudi, Alex Goldin, Adnan Ozturel, Albin Cassirer, Yunhan Xu, Daniel Sohn, Devendra Sachan, Reinald Kim Amplayo, Craig Swanson, Dessie Petrova, Shashi Narayan, Arthur Guez, Siddhartha Brahma, Jessica Landon, Miteyan Patel, Ruizhe Zhao, Kevin Villela, Luyu Wang, Wenhao Jia, Matthew Rahtz, Mai Giménez, Legg Yeung, James Keeling, Petko Georgiev, Diana Mincu, Boxi Wu, Salem Haykal, Rachel Saputro, Kiran Vodrahalli, James Qin, Zeynep Cankara, Abhanshu Sharma, Nick Fernando, Will Hawkins, Behnam Neyshabur, Solomon Kim, Adrian Hutter, Priyanka Agrawal, Alex Castro-Ros, George van den Driessche, Tao Wang, Fan Yang, Shuo yiin Chang, Paul Komarek, Ross McIlroy, Mario Lučić, Guodong Zhang, Wael Farhan, Michael Sharman, Paul Natsev, Paul Michel, Yamini Bansal, Siyuan Qiao, Kris Cao, Siamak Shakeri, Christina Butterfield, Justin Chung, Paul Kishan Rubenstein,

Shivani Agrawal, Arthur Mensch, Kedar Soparkar, Karel Lenc, Timothy Chung, Aedan Pope, Loren Maggiore, Jackie Kay, Priya Jhakra, Shibo Wang, Joshua Maynez, Mary Phuong, Taylor Tobin, Andrea Tacchetti, Maja Trebacz, Kevin Robinson, Yash Katariya, Sebastian Riedel, Paige Bailey, Kefan Xiao, Nimesh Ghelani, Lora Aroyo, Ambrose Slone, Neil Houlsby, Xuehan Xiong, Zhen Yang, Elena Gribovskaya, Jonas Adler, Mateo Wirth, Lisa Lee, Music Li, Thais Kagohara, Jay Pavagadhi, Sophie Bridgers, Anna Bortsova, Sanjay Ghemawat, Zafarali Ahmed, Tianqi Liu, Richard Powell, Vijay Bolina, Mariko Iinuma, Polina Zablotskaia, James Besley, Da-Woon Chung, Timothy Dozat, Ramona Comanescu, Xiance Si, Jeremy Greer, Guolong Su, Martin Polacek, Raphaël Lopez Kaufman, Simon Tokumine, Hexiang Hu, Elena Buchatskaya, Yingjie Miao, Mohamed Elhawaty, Aditya Siddhant, Nenad Tomasev, Jinwei Xing, Christina Greer, Helen Miller, Shereen Ashraf, Aurko Roy, Zizhao Zhang, Ada Ma, Angelos Filos, Milos Besta, Rory Blevins, Ted Klimenko, Chih-Kuan Yeh, Soravit Changpinyo, Jiaqi Mu, Oscar Chang, Mantas Pajarskas, Carrie Muir, Vered Cohen, Charline Le Lan, Krishna Haridasan, Amit Marathe, Steven Hansen, Sholto Douglas, Rajkumar Samuel, Mingqiu Wang, Sophia Austin, Chang Lan, Jiepu Jiang, Justin Chiu, Jaime Alonso Lorenzo, Lars Lowe Sjösund, Sébastien Cevey, Zach Gleicher, Thi Avrahami, Anudhyan Boral, Hansa Srinivasan, Vittorio Selo, Rhys May, Konstantinos Aisopos, Léonard Hussenot, Livio Baldini Soares, Kate Baumli, Michael B. Chang, Adrià Recasens, Ben Caine, Alexander Pritzel, Filip Pavetic, Fabio Pardo, Anita Gergely, Justin Frye, Vinay Ramasesh, Dan

Horgan, Kartikeya Badola, Nora Kassner, Subhrajit Roy, Ethan Dyer, Víctor Campos Campos, Alex Tomala, Yunhao Tang, Dalia El Badawy, Elspeth White, Basil Mustafa, Oran Lang, Abhishek Jindal, Sharad Vikram, Zhitao Gong, Sergi Caelles, Ross Hemsley, Gregory Thornton, Fangxiaoyu Feng, Wojciech Stokowiec, Ce Zheng, Phoebe Thacker, Çağlar Ünlü, Zhishuai Zhang, Mohammad Saleh, James Svensson, Max Bileschi, Piyush Patil, Ankesh Anand, Roman Ring, Katerina Tsihlas, Arpi Vezer, Marco Selvi, Toby Shevlane, Mikel Rodriguez, Tom Kwiatkowski, Samira Daruki, Keran Rong, Allan Dafoe, Nicholas FitzGerald, Keren Gu-Lemberg, Mina Khan, Lisa Anne Hendricks, Marie Pellat, Vladimir Feinberg, James Cobon-Kerr, Tara Sainath, Maribeth Rauh, Sayed Hadi Hashemi, Richard Ives, Yana Hasson, Eric Noland, Yuan Cao, Nathan Byrd, Le Hou, Qingze Wang, Thibault Sottiaux, Michela Paganini, Jean-Baptiste Lespiau, Alexandre Moufarek, Samer Hassan, Kaushik Shivakumar, Joost van Amersfoort, Amol Mandhane, Pratik Joshi, Anirudh Goyal, Matthew Tung, Andrew Brock, Hannah Sheahan, Vedant Misra, Cheng Li, Nemanja Rakićević, Mostafa Dehghani, Fangyu Liu, Sid Mittal, Junhyuk Oh, Seb Noury, Eren Sezener, Fantine Huot, Matthew Lamm, Nicola De Cao, Charlie Chen, Sidharth Mudgal, Romina Stella, Kevin Brooks, Gautam Vasudevan, Chenxi Liu, Mainak Chain, Nivedita Melinkeri, Aaron Cohen, Venus Wang, Kristie Seymore, Sergey Zubkov, Rahul Goel, Summer Yue, Sai Krishnakumaran, Brian Albert, Nate Hurley, Motoki Sano, Anhad Mohananey, Jonah Joughin, Egor Filonov, Tomasz Kępa, Yomna Eldawy, Jiawern Lim, Rahul Rishi, Shirin Badiezadegan, Taylor Bos, Jerry

Chang, Sanil Jain, Sri Gayatri Sundara Padmanabhan, Subha Puttagunta, Kalpesh Krishna, Leslie Baker, Norbert Kalb, Vamsi Bedapudi, Adam Kurzrok, Shuntong Lei, Anthony Yu, Oren Litvin, Xiang Zhou, Zhichun Wu, Sam Sobell, Andrea Siciliano, Alan Papir, Robby Neale, Jonas Bragagnolo, Tej Toor, Tina Chen, Valentin Anklin, Feiran Wang, Richie Feng, Milad Gholami, Kevin Ling, Lijuan Liu, Jules Walter, Hamid Moghaddam, Arun Kishore, Jakub Adamek, Tyler Mercado, Jonathan Mallinson, Siddhinita Wandekar, Stephen Cagle, Eran Ofek, Guillermo Garrido, Clemens Lombriser, Maksim Mukha, Botu Sun, Hafeezul Rahman Mohammad, Josip Matak, Yadi Qian, Vikas Peswani, Pawel Janus, Quan Yuan, Leif Schelin, Oana David, Ankur Garg, Yifan He, Oleksii Duzhyi, Anton Älgmyr, Timothée Lottaz, Qi Li, Vikas Yadav, Luyao Xu, Alex Chinien, Rakesh Shivanna, Aleksandr Chuklin, Josie Li, Carrie Spadine, Travis Wolfe, Kareem Mohamed, Subhabrata Das, Zihang Dai, Kyle He, Daniel von Dincklage, Shyam Upadhyay, Akanksha Maurya, Luyan Chi, Sebastian Krause, Khalid Salama, Pam G Rabinovitch, Pavan Kumar Reddy M, Aarush Selvan, Mikhail Dektiarev, Golnaz Ghiasi, Erdem Guven, Himanshu Gupta, Boyi Liu, Deepak Sharma, Idan Heimlich Shtacher, Shachi Paul, Oscar Akerlund, François-Xavier Aubet, Terry Huang, Chen Zhu, Eric Zhu, Elico Teixeira, Matthew Fritze, Francesco Bertolini, Liana-Eleonora Marinescu, Martin Bölle, Dominik Paulus, Khyatti Gupta, Tejasi Latkar, Max Chang, Jason Sanders, Roopa Wilson, Xuewei Wu, Yi-Xuan Tan, Lam Nguyen Thiet, Tulsee Doshi, Sid Lall, Swaroop Mishra, Wanming Chen, Thang Luong, Seth Benjamin, Jasmine Lee, Ewa Andrejczuk, Dominik Rabiej, Vipul Ranjan, Krzysztof Styrc, Pengcheng Yin, Jon Simon, Malcolm Rose Harriott, Mudit Bansal, Alexei Robsky, Geoff Bacon, David Greene, Daniil Mirylenka, Chen Zhou, Obaid Sarvana, Abhimanyu Goyal, Samuel Andermatt, Patrick Siegler, Ben Horn, Assaf Israel, Francesco Pongetti, Chih-Wei "Louis" Chen, Marco Selvatici, Pedro Silva, Kathie Wang, Jackson Tolins, Kelvin Guu, Roey Yogev, Xiaochen Cai, Alessandro Agostini, Maulik Shah, Hung Nguyen, Noah Ó Donnaile, Sébastien Pereira, Linda Friso, Adam Stambler, Adam Kurzrok, Chenkai Kuang, Yan Romanikhin, Mark Geller, ZJ Yan, Kane Jang, Cheng-Chun Lee, Wojciech Fica, Eric Malmi, Qijun Tan, Dan Banica, Daniel Balle, Ryan Pham, Yanping Huang, Diana Avram, Hongzhi Shi, Jasjot Singh, Chris Hidey, Niharika Ahuja, Pranab Saxena, Dan Dooley, Srividya Pranavi Potharaju, Eileen O'Neill, Anand Gokulchandran, Ryan Foley, Kai Zhao, Mike Dusenberry, Yuan Liu, Pulkit Mehta, Ragha Kotikalapudi, Chalence Safranek-Shrader, Andrew Goodman, Joshua Kessinger, Eran Globen, Prateek Kolhar, Chris Gorgolewski, Ali Ibrahim, Yang Song, Ali Eichenbaum, Thomas Brovelli, Sahitya Potluri, Preethi Lahoti, Cip Baetu, Ali Ghorbani, Charles Chen, Andy Crawford, Shalini Pal, Mukund Sridhar, Petru Gurita, Asier Mujika, Igor Petrovski, Pierre-Louis Cedoz, Chenmei Li, Shiyuan Chen, Niccolò Dal Santo, Siddharth Goyal, Jitesh Punjabi, Karthik Kappaganthu, Chester Kwak, Pallavi LV, Sarmishta Velury, Himadri Choudhury, Jamie Hall, Premal Shah, Ricardo Figueira, Matt Thomas, Minjie Lu, Ting Zhou, Chintu Kumar, Thomas Jurdi, Sharat Chikkerur, Yenai Ma, Adams Yu, Soo Kwak, Victor Ähdel, Sujeevan Rajayogam, Travis Choma, Fei Liu, Aditya Barua, Colin Ji, Ji Ho Park, Vincent Hellendoorn, Alex Bailey, Taylan Bilal, Huanjie Zhou, Mehrdad Khatir, Charles Sutton, Wojciech Rzadkowski, Fiona Macintosh, Konstantin Shagin, Paul Medina, Chen Liang, Jinjing Zhou, Pararth Shah, Yingying Bi, Attila Dankovics, Shipra Banga, Sabine Lehmann, Marissa Bredesen, Zifan Lin, John Eric Hoffmann, Jonathan Lai, Raynald Chung, Kai Yang, Nihal Balani, Arthur Bražinskas, Andrei Sozanschi, Matthew Hayes, Héctor Fernández Alcalde, Peter Makarov, Will Chen, Antonio Stella, Liselotte Snijders, Michael Mandl, Ante Kärrman, Paweł Nowak, Xinyi Wu, Alex Dyck, Krishnan Vaidyanathan, Raghavender R, Jessica Mallet, Mitch Rudominer, Eric Johnston, Sushil Mittal, Akhil Udathu, Janara Christensen, Vishal Verma, Zach Irving, Andreas Santucci, Gamaleldin Elsayed, Elnaz Davoodi, Marin Georgiev, Ian Tenney, Nan Hua, Geoffrey Cideron, Edouard Leurent, Mahmoud Alnahlawi, Ionut Georgescu, Nan Wei, Ivy Zheng, Dylan Scandinaro, Heinrich Jiang, Jasper Snoek, Mukund Sundararajan, Xuezhi Wang, Zack Ontiveros, Itay Karo, Jeremy Cole, Vinu Rajashekhar, Lara Tumeh, Eyal Ben-David, Rishub Jain, Jonathan Uesato, Romina Datta, Oskar Bunyan, Shimu Wu, John Zhang, Piotr Stanczyk, Ye Zhang, David Steiner, Subhajit Naskar, Michael Azzam, Matthew Johnson, Adam Paszke, Chung-Cheng Chiu, Jaume Sanchez Elias, Afroz Mohiuddin, Faizan Muhammad, Jin Miao, Andrew Lee, Nino Vieillard, Jane Park, Jiageng Zhang, Jeff Stanway, Drew Garmon, Abhijit Karmarkar, Zhe Dong, Jong Lee, Aviral Kumar, Luowei Zhou, Jonathan Evens, William Isaac, Geoffrey Irving, Edward

Loper, Michael Fink, Isha Arkatkar, Nanxin Chen, Izhak Shafran, Ivan Petrychenko, Zhe Chen, Johnson Jia, Anselm Levskaya, Zhenkai Zhu, Peter Grabowski, Yu Mao, Alberto Magni, Kaisheng Yao, Javier Snaider, Norman Casagrande, Evan Palmer, Paul Suganthan, Alfonso Castaño, Irene Giannoumis, Wooyeol Kim, Mikołaj Rybiński, Ashwin Sreevatsa, Jennifer Prendki, David Soergel, Adrian Goedeckemeyer, Willi Gierke, Mohsen Jafari, Meenu Gaba, Jeremy Wiesner, Diana Gage Wright, Yawen Wei, Harsha Vashisht, Yana Kulizhskaya, Jay Hoover, Maigo Le, Lu Li, Chimezie Iwuanyanwu, Lu Liu, Kevin Ramirez, Andrey Khorlin, Albert Cui, Tian LIN, Marcus Wu, Ricardo Aguilar, Keith Pallo, Abhishek Chakladar, Ginger Perng, Elena Allica Abellan, Mingyang Zhang, Ishita Dasgupta, Nate Kushman, Ivo Penchev, Alena Repina, Xihui Wu, Tom van der Weide, Priya Ponnapalli, Caroline Kaplan, Jiri Simsa, Shuangfeng Li, Olivier Dousse, Fan Yang, Jeff Piper, Nathan Ie, Rama Pasumarthi, Nathan Lintz, Anitha Vijayakumar, Daniel Andor, Pedro Valenzuela, Minnie Lui, Cosmin Paduraru, Daiyi Peng, Katherine Lee, Shuyuan Zhang, Somer Greene, Duc Dung Nguyen, Paula Kurylowicz, Cassidy Hardin, Lucas Dixon, Lili Janzer, Kiam Choo, Ziqiang Feng, Biao Zhang, Achintya Singhal, Dayou Du, Dan McKinnon, Natasha Antropova, Tolga Bolukbasi, Orgad Keller, David Reid, Daniel Finchelstein, Maria Abi Raad, Remi Crocker, Peter Hawkins, Robert Dadashi, Colin Gaffney, Ken Franko, Anna Bulanova, Rémi Leblond, Shirley Chung, Harry Askham, Luis C. Cobo, Kelvin Xu, Felix Fischer, Jun Xu, Christina Sorokin, Chris Alberti, Chu-Cheng Lin, Colin Evans, Alek Dimitriev, Hannah Forbes, Dylan Banarse, Zora Tung, Mark Omernick, Colton Bishop, Rachel Sterneck, Rohan Jain, Jiawei Xia, Ehsan Amid, Francesco Piccinno, Xingyu Wang, Praseem Banzal, Daniel J. Mankowitz, Alex Polozov, Victoria Krakovna, Sasha Brown, MohammadHossein Bateni, Dennis Duan, Vlad Firoiu, Meghana Thotakuri, Tom Natan, Matthieu Geist, Ser tan Girgin, Hui Li, Jiayu Ye, Ofir Roval, Reiko Tojo, Michael Kwong, James Lee-Thorp, Christopher Yew, Danila Sinopalnikov, Sabela Ramos, John Mellor, Abhishek Sharma, Kathy Wu, David Miller, Nicolas Sonnerat, Denis Vnukov, Rory Greig, Jennifer Beattie, Emily Caveness, Libin Bai, Julian Eisenschlos, Alex Korchemniy, Tomy Tsai, Mimi Jasarevic, Weize Kong, Phuong Dao, Zeyu Zheng, Frederick Liu, Fan Yang, Rui Zhu, Tian Huey Teh, Jason Sanmiya, Evgeny Gladchenko, Nejc Trdin, Daniel Toyama, Evan Rosen, Sasan Tavakkol, Linting Xue, Chen Elkind, Oliver Woodman, John Carpenter, George Papamakarios, Rupert Kemp, Sushant Kafle, Tanya Grunina, Rishika Sinha, Alice Talbert, Diane Wu, Denese Owusu-Afriyie, Cosmo Du, Chloe Thornton, Jordi Pont-Tuset, Pradyumna Narayana, Jing Li, Saaber Fatehi, John Wieting, Omar Ajmeri, Benigno Uria, Yeongil Ko, Laura Knight, Amélie Héliou, Ning Niu, Shane Gu, Chenxi Pang, Yeqing Li, Nir Levine, Ariel Stolovich, Rebeca Santamaria-Fernandez, Sonam Goenka, Wenny Yustalim, Robin Strudel, Ali Elqursh, Charlie Deck, Hyo Lee, Zonglin Li, Kyle Levin, Raphael Hoffmann, Dan Holtmann-Rice, Olivier Bachem, Sho Arora, Christy Koh, Soheil Hassas Yeganeh, Siim Põder, Mukarram Tariq, Yanhua Sun, Lucian Ionita, Mojtaba Seyedhosseini, Pouya Tafti, Zhiyu Liu, Anmol Gulati, Jasmine Liu, Xinyu Ye, Bart

Chrzaszcz, Lily Wang, Nikhil Sethi, Tianrun Li, Ben Brown, Shreya Singh, Wei Fan, Aaron Parisi, Joe Stanton, Vinod Koverkathu, Christopher A. Choquette-Choo, Yunjie Li, TJ Lu, Abe Ittycheriah, Prakash Shroff, Mani Varadarajan, Sanaz Bahargam, Rob Willoughby, David Gaddy, Guillaume Desjardins, Marco Cornero, Brona Robenek, Bhavishya Mittal, Ben Albrecht, Ashish Shenoy, Fedor Moiseev, Henrik Jacobsson, Alireza Ghaffarkhah, Morgane Rivière, Alanna Walton, Clément Crepy, Alicia Parrish, Zongwei Zhou, Clement Farabet, Carey Radebaugh, Praveen Srinivasan, Claudia van der Salm, Andreas Fidjeland, Salvatore Scellato, Eri Latorre-Chimoto, Hanna Klimczak-Plucińska, David Bridson, Dario de Cesare, Tom Hudson, Piermaria Mendolicchio, Lexi Walker, Alex Morris, Matthew Mauger, Alexey Guseynov, Alison Reid, Seth Odoom, Lucia Loher, Victor Cotruta, Madhavi Yenugula, Dominik Grewe, Anastasia Petrushkina, Tom Duerig, Antonio Sanchez, Steve Yadlowsky, Amy Shen, Amir Globerson, Lynette Webb, Sahil Dua, Dong Li, Surya Bhupatiraju, Dan Hurt, Haroon Qureshi, Ananth Agarwal, Tomer Shani, Matan Eyal, Anuj Khare, Shreyas Rammohan Belle, Lei Wang, Chetan Tekur, Mihir Sanjay Kale, Jinliang Wei, Ruoxin Sang, Brennan Saeta, Tyler Liechty, Yi Sun, Yao Zhao, Stephan Lee, Pandu Nayak, Doug Fritz, Manish Reddy Vuyyuru, John Aslanides, Nidhi Vyas, Martin Wicke, Xiao Ma, Evgenii Eltyshev, Nina Martin, Hardie Cate, James Manyika, Keyvan Amiri, Yelin Kim, Xi Xiong, Kai Kang, Florian Luisier, Nilesh Tripuraneni, David Madras, Mandy Guo, Austin Waters, Oliver Wang, Joshua Ainslie, Jason Baldridge, Han Zhang, Garima Pruthi, Jakob Bauer, Feng

Yang, Riham Mansour, Jason Gelman, Yang Xu, George Polovets, Ji Liu, Honglong Cai, Warren Chen, XiangHai Sheng, Emily Xue, Sherjil Ozair, Christof Angermueller, Xiaowei Li, Anoop Sinha, Weiren Wang, Julia Wiesinger, Emmanouil Koukoumidis, Yuan Tian, Anand Iyer, Madhu Gurumurthy, Mark Goldenson, Parashar Shah, MK Blake, Hongkun Yu, Anthony Urbanowicz, Jennimaria Palomaki, Chrisantha Fernando, Ken Durden, Harsh Mehta, Nikola Momchev, Elahe Rahimtoroghi, Maria Georgaki, Amit Raul, Sebastian Ruder, Morgan Redshaw, Jinhyuk Lee, Denny Zhou, Komal Jalan, Dinghua Li, Blake Hechtman, Parker Schuh, Milad Nasr, Kieran Milan, Vladimir Mikulik, Juliana Franco, Tim Green, Nam Nguyen, Joe Kelley, Aroma Mahendru, Andrea Hu, Joshua Howland, Ben Vargas, Jeffrey Hui, Kshitij Bansal, Vikram Rao, Rakesh Ghiya, Emma Wang, Ke Ye, Jean Michel Sarr, Melanie Moranski Preston, Madeleine Elish, Steve Li, Aakash Kaku, Jigar Gupta, Ice Pasupat, Da-Cheng Juan, Milan Someswar, Tejvi M., Xinyun Chen, Aida Amini, Alex Fabrikant, Eric Chu, Xuanyi Dong, Amruta Muthal, Senaka Buthpitiya, Sarthak Jauhari, Nan Hua, Urvashi Khandelwal, Ayal Hitron, Jie Ren, Larissa Rinaldi, Shahar Drath, Avigail Dabush, Nan-Jiang Jiang, Harshal Godhia, Uli Sachs, Anthony Chen, Yicheng Fan, Hagai Taitelbaum, Hila Noga, Zhuyun Dai, James Wang, Chen Liang, Jenny Hamer, Chun-Sung Ferng, Chenel Elkind, Aviel Atias, Paulina Lee, Vít Listík, Mathias Carlen, Jan van de Kerkhof, Marcin Pikus, Krunoslav Zaher, Paul Müller, Sasha Zykova, Richard Stefanec, Vitaly Gatsko, Christoph Hirnschall, Ashwin Sethi, Xingyu Federico Xu, Chetan Ahuja, Beth Tsai, Anca Stefanoiu, Bo Feng, Keshav

Dhandhania, Manish Katyal, Akshay Gupta, Atharva Parulekar, Divya Pitta, Jing Zhao, Vivaan Bhatia, Yashodha Bhavnani, Omar Alhadlaq, Xiaolin Li, Peter Danenberg, Dennis Tu, Alex Pine, Vera Filippova, Abhipso Ghosh, Ben Limonchik, Bhargava Urala, Chaitanya Krishna Lanka, Derik Clive, Yi Sun, Edward Li, Hao Wu, Kevin Hongtongsak, Ianna Li, Kalind Thakkar, Kuanysh Omarov, Kushal Majmundar, Michael Alverson, Michael Kucharski, Mohak Patel, Mudit Jain, Maksim Zabelin, Paolo Pelagatti, Rohan Kohli, Saurabh Kumar, Joseph Kim, Swetha Sankar, Vineet Shah, Lakshmi Ramachandruni, Xiangkai Zeng, Ben Bariach, Laura Weidinger, Tu Vu, Alek Andreev, Antoine He, Kevin Hui, Sheleem Kashem, Amar Subramanya, Sissie Hsiao, Demis Hassabis, Koray Kavukcuoglu, Adam Sadovsky, Quoc Le, Trevor Strohman, Yonghui Wu, Slav Petrov, Jeffrey Dean, and Oriol Vinyals. Gemini: A family of highly capable multimodal models, 2024. 9

[71] Johannes Treutlein, Dami Choi, Jan Betley, Samuel Marks, Cem Anil, Roger Grosse, and Owain Evans. Connecting the dots: Llms can infer and verbalize latent structure from disparate training data. *arXiv preprint arXiv:2406.14546*, 2024. 18

[72] Bryan Wang, Gang Li, Xin Zhou, Zhourong Chen, Tovi Grossman, and Yang Li. Screen2words: Automatic mobile ui summarization with multimodal learning. In *The 34th Annual ACM Symposium on User Interface Software and Technology*, pages 498–510, 2021. 26

[73] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang

Lu, Yiqiang Chen, Wenjun Zeng, and S Yu Philip. Generalizing to unseen domains: A survey on domain generalization. *IEEE transactions on knowledge and data engineering*, 35(8):8052–8072, 2022. 11

[74] Jue Wang, Yanyan Jiang, Chang Xu, Chun Cao, Xiaoxing Ma, and Jian Lu. Combodroid: generating high-quality test inputs for android apps via use case combinations. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*, pages 469–480, 2020. 26

[75] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023. 29

[76] Chaoyi Wu, Jiayu Lei, Qiaoyu Zheng, Weike Zhao, Weixiong Lin, Xiaoman Zhang, Xiao Zhou, Ziheng Zhao, Ya Zhang, Yanfeng Wang, et al. Can gpt-4v (ision) serve medical applications? case studies on gpt-4v for multimodal medical diagnosis. *arXiv preprint arXiv:2310.09909*, 2023. 13

[77] Hao Lin Zhenhua Li Feng Qian Xianlong Wang Yunhao Liu Xinlei Yang, Wei Liu and Tianyin Xu. Visual-aware testing and debugging for web performance optimization. In *Proceedings of the International World Wide Web Conference (WWW)*. ACM, 2023. 26

[78] Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on multimodal large language models. *National Science Review*, page nwae403, 2024. 12

[79] Duzhen Zhang, Yahan Yu, Jiahua Dong, Chenxing Li, Dan Su, Chenhui Chu, and Dong Yu. Mm-llms: Recent advances in multimodal large language models, 2024. 12

[80] Haochen Zhang, Yuyang Dong, Chuan Xiao, and Masafumi Oyamada. Jellyfish: A large language model for data preprocessing. *arXiv preprint arXiv:2312.01678*, 2023. 12

[81] Haochen Zhang, Yuyang Dong, Chuan Xiao, and Masafumi Oyamada. Large language models as data preprocessors. *arXiv preprint arXiv:2308.16361*, 2023. 12

[82] Dehai Zhao, Zhenchang Xing, Chunyang Chen, Xiwei Xu, Liming Zhu, Guoqiang Li, and Jinshui Wang. Seenomaly: Vision-based linting of gui animation effects against design-don't guidelines. In *Proceedings of the ACM/IEEE 42nd international conference on software engineering*, pages 1286–1297, 2020. 26

[83] Tianming Zhao, Chunyang Chen, Yuanning Liu, and Xiaodong Zhu. Guigan: Learning to generate gui designs using generative adversarial networks. In *2021 IEEE/ACM 43rd International Conference on Software Engineering (ICSE)*, pages 748–760. IEEE, 2021. 26

[84] Yang Zhao, Zhijie Lin, Daquan Zhou, Zilong Huang, Jiashi Feng, and Bingyi Kang. Bubogpt: Enabling visual grounding in multi-modal llms, 2023. 44

[85] Zhisheng Zheng, Puyuan Peng, Ziyang Ma, Xie Chen, Eunsol Choi, and David Harwath. Bat: Learning to reason about spatial sounds with large language models. *arXiv preprint arXiv:2402.01591*, 2024. 44