

Live API-Bench: 2500+ Live APIs for Testing Multi-Step Tool Calling

Benjamin Elder*
IBM Research AI
benjamin.elder@ibm.com

Anupama Murthi*
IBM Research AI
anupama.murthi@ibm.com

Jungkoo Kang
IBM Research AI
jungkoo.kang@ibm.com

Ankita Rajaram Naik
IBM Research AI
ankita.naik@ibm.com

Kiran Kate
IBM Research AI
kakate@us.ibm.com

Kinjal Basu
IBM Research AI
kinjal.basu@ibm.com

Danish Contractor
IBM Research AI
danish.contractor@ibm.com

Abstract

Large language models (LLMs) increasingly rely on external tools and APIs to execute complex tasks specified in natural language. Evaluating such *tool-calling* capabilities in realistic enterprise settings is challenging: APIs are often proprietary, heterogeneous, and difficult to share, limiting reproducible benchmarks. To address this, we introduce **Live API Bench**, a comprehensive benchmark constructed by transforming NL2SQL datasets into interactive API environments. Our pipeline converts SQL queries from BIRD-SQL into executable API sequences across three formulations—*SLOT*, *SEL*, and *REST*—covering minimal general-purpose operations, domain-specific multi-step tasks, and function-oriented RESTful interactions, respectively. The benchmark spans 11 databases with over 2,500 invocable tools, paired with human-authored queries, ground-truth API sequences, and verified final answers. Live API Bench enables systematic evaluation of core challenges in tool use, including error handling, sequential reasoning, parameter generation, response parsing, and robustness across diverse domains. We evaluate 10 LLMs and 4 ReACT agents, observing low task completion rates (7–47%), which improve modestly to 50% under interactive agent settings, highlighting substantial scope for improving LLM tool-calling performance. We release all code and data associated with this paper.

1 Introduction

Large language models (LLMs) are increasingly deployed in real-world applications where they must

interact with external tools and APIs to accomplish complex tasks expressed in natural language (Yao et al., 2023; Xu et al., 2023a). Such systems are deployed with access to tools that interact with live environments to accomplish tasks, often in response to a user’s request expressed in natural language (Yoran et al., 2024; Drouin et al., 2024; Pan et al., 2024; Xie et al., 2024; Zheng et al., 2024). This capability—commonly referred to as *tool-calling*—is critical in enterprise environments, where business workflows depend on programmatic access to databases, web services, and specialized software systems via standardized API interfaces.^{1,2}

Evaluating LLM tool-calling in realistic deployment settings remains an open challenge: real-world APIs are proprietary, heterogeneous, and difficult to share, which hinders reproducibility. Prior benchmarks thus often rely on simulations, limited domains, or small scales (see Section 2). To approximate some of the deployment challenges while retaining reproducibility, we build our benchmark from NL2SQL datasets. This design offers three advantages: (1) NL2SQL tasks involve large, diverse databases, (2) they include human-authored natural language queries with guaranteed executable interpretations, and (3) SQL queries can be systematically decomposed into semantically equivalent API-call sequences that highlight practical tool-calling difficulties.

Building on these insights, we introduce **Live API Bench**, a comprehensive benchmark for evaluating LLM tool-calling by transforming NL2SQL

*Equal contribution.

¹<https://shorturl.at/bpsps>

²<https://shorturl.at/daoEp>

Benchmarks	Core Features			Deployment			Evaluation			
	Hand-Crafted Queries	Invocable APIs for Data Creation	Nested API Sequences	Code Tools as Functions	Live API endpoints	Accessible DB	Real-World data powering API	Ground Truth Sequences	# Instances	# Tools
NesTools (Han et al., 2024)	✗	✗	✓	✗	✗	✗	✗	✓	1000	3034
ToolAce (Liu et al., 2024b)	✗	✗	✓	✗	✗	✗	✗	✓	11300	26507
API-Bank (Li et al., 2023b)	✓	✓	✗	✓	✗	✓	✗	✗	314	73
APIBench (Patil et al., 2024)	✗	✓	✗	✓	✗	✗	✗	✗	16450	1645
ToolBench (Xu et al., 2023b)	✗	✓	✓	✗	✓	✗	✗	✓	2746	232
RestBench (Song et al., 2023)	✓	✓	✓	✗	✓	✗	✗	✓	157	94
ToolQA (Zhuang et al., 2023)	✗	✓	✗	✓	✗	✓	✓	✓	1530	13
ToolAlpaca (Tang et al., 2023)	✗	✓	✗	✗	✓	✗	✓	✓	3938	426
ToolLLM (Qin et al., 2024)	✗	✓	✗	✗	✓	✗	✓	✓	126,486	16,464
APIGen (Liu et al., 2024d)	✗	✓	✗	✓	✓	✗	✓	✓	60,000	3673
NESTFUL (Basu et al., 2025)	✓	✓	✓	✓	✗	✗	✗	✓	1861	921
Ours	✓	✓	✓	✓	✓	✓	✓	✓	2583	2513

Table 1: Comparison of Invocable API-Bench with existing tool-calling benchmarks, organized by *Core Features*, *Deployment*, and *Evaluation*. While many benchmarks use APIs to construct datasets (see *Invocable APIs* column), these APIs are often not accessible during evaluation (see *Live API endpoints* column).

resources into interactive API environments. Our data transformation pipeline converts NL2SQL queries from BIRD-SQL (Li et al., 2023a)—one of the largest and most diverse permissively licensed³ NL2SQL datasets—into executable API sequences across three complementary formulations:

- **SLOT** — A minimal set of general-purpose APIs (e.g., filter, sort, aggregate) that often require multiple sequential invocations with varying parameters.
- **SEL** — An expanded collection combining domain-specific retrieval functions with general-purpose operations, requiring models to select tools and compose multi-step execution plans.
- **REST** — A function-oriented formulation where queries can typically be resolved through careful function selection and parameter specification, modeling RESTful API interactions.

Live API Bench provides fully accessible APIs that can be hosted locally, allowing LLMs and agents to interact with live endpoints during evaluation. Additionally, by deriving APIs from executable SQL queries over real-world databases, our benchmark ensures deterministic, verified final answers, addressing common limitations in existing tool-calling benchmarks (see Table 1). This setup enables systematic evaluation of tool use with realistic challenges including:

- **Error handling** — APIs may fail or return malformed responses, requiring models to manage failures and timeouts.
- **Sequential tool calls** — Tasks often involve chaining dependent API calls; our dataset in-

cludes sequences of up to eight calls.

- **Parameter generation** — Models must generate appropriate arguments, validate inputs, and infer missing information to bridge natural language queries and structured API schemas.
- **Response parsing** — Successful tool use requires interpreting complex outputs, extracting relevant information (frequently from large API outputs), and recognizing mismatches with intended requests.
- **Scale and diversity** — To avoid overfitting, models are evaluated across diverse domains, API types, and query complexities.

Contributions. In summary, this paper makes the following contributions: (1) We develop novel data transformation pipelines that repurpose existing NL2SQL tasks for evaluating LLM tool-calling, (2) We generate three distinct API formulations—*SLOT*, *SEL*, and *REST*—to study different aspects of tool-calling. Using the 11 publicly available NL2SQL databases from the BIRD-SQL dev set (Li et al., 2023a), we create over 2,500 invocable tools backed by real databases. Each tool is paired with human-authored natural language queries, a ground-truth API sequence, and a verified final answer. To our knowledge, this is the largest publicly available collection of live, invocable APIs with these characteristics. (3) Finally, we evaluate 10 LLMs and 4 ReACT agents (Yao et al., 2023) on these collections, finding extremely low task completion rates (7–47%) that improve modestly to 50% when models interact with the live API environment as ReACT agents, highlighting substantial room for improvement in LLM tool-calling capabilities. We release all code and data associated with this paper⁴.

³Licensed under CC BY-SA 4.0

⁴Code and data available at <https://ibm.biz/Bdpc94>.

2 Related Work

The emergence of powerful LLMs has spurred the development of data generation pipelines that produce training and evaluation datasets for tool-calling tasks. Such datasets typically consist of a collection of APIs, along with natural language queries paired with ground-truth API calls or sequences of calls (Liu et al., 2024d; Shi et al., 2025; Qin et al., 2024; Basu et al., 2024; Pereira et al., 2024).

Data generation pipelines generally fall into two categories. Some repurpose existing datasets and queries, leveraging pre-existing data for tool-calling tasks (Peng et al., 2021; Basu et al., 2024). Others use LLMs to synthesize realistic queries conditioned on API collections (Qin et al., 2024; Tang et al., 2023; Liu et al., 2024b; Pereira et al., 2024; Basu et al., 2025; Zhong et al., 2025). Beyond queries, some work also focuses on generating APIs themselves from structured sources such as textbooks or documentation, which can then be used in tool-calling tasks (Liu et al., 2025).

Methods for data generation are often tailored to the function-calling scenario. Some pipelines generate queries that invoke a single function, potentially multiple times (Tang et al., 2023; Xu et al., 2023b), others focus on queries requiring the invocation of multiple functions including sequencing, and nested function calls (Liu et al., 2024b; Yan et al., 2024; Basu et al., 2025; Zhong et al., 2025). To assess the planning and task-resolution capabilities of tool-calling LLMs, recent benchmarks that provide live environments for agents to interact with have also been developed (Ruan et al., 2024; Yao et al., 2024; Cao et al., 2024; Xie et al., 2024; Zhou et al., 2024).

While existing work sometimes includes queries and APIs that were created using live APIs available from proprietary sources (Chen et al., 2024; Qin et al., 2024; Basu et al., 2025; Yan et al., 2024; Zhong et al., 2025), such APIs are usually either unavailable for live agents attempting to answer queries at test time or are only available as paid services by the providers hosting the APIs (eg: RapidAPI Hub).⁵ In contrast, our work expands on these prior efforts by generating APIs that can be hosted locally either as endpoints or python tools for use by LLMs (and agents). Further, since our method for API construction relies on the executable SQL queries from the original NL2SQL

⁵<http://www.rapidapi.com>

Listing 1: Original NL2SQL Instance

```
1 {"input": "Among the schools with the SAT test takers of  
  ↳ over 500, please list the schools that are magnet  
  ↳ schools or offer a magnet program.",  
2 "dataset_name": "california_schools",  
3 "query": "SELECT T2.School FROM satscores AS T1 INNER JOIN  
  ↳ schools AS T2 ON T1.cds = T2.CDSCode WHERE T2.  
  ↳ Magnet = 1 AND T1.NumTstTskr > 500",  
4 "gold_answer": ["Millikan High", "Polytechnic High", "Troy  
  ↳ High"]}
```

Listing 2: SLOT-BIRD: API sequence

```
1 [{"name": "filter_data",  
2 "arguments": {"data_source": "data_0.csv", "key_name": "  
  ↳ schools_Magnet",  
3 "value": 1.0, "condition": "equal_to"},  
4 "label": "data_1.csv"},  
5 {"name": "filter_data",  
6 "arguments": {"data_source": "data_1.csv", "key_name": "  
  ↳ satscores_NumTstTskr",  
7 "value": 500.0, "condition": "greater_than"},  
8 "label": "data_2.csv"},  
9 {"name": "retrieve_data",  
10 "arguments": {"data_source": "data_2.csv", "key_name": "  
  ↳ schools_School",  
11 "distinct": False, "limit": -1, "label": "retrieved_json"  
  ↳ }]}
```

Listing 3: SEL-BIRD: API Sequence

```
1 [{"name": "select_data_equal_to",  
2 "arguments": {"data_source": "data_0.csv", "key_name": "  
  ↳ schools_Magnet",  
3 "value": 1.0},  
4 "label": "data_1.csv"},  
5 {"name": "select_data_greater_than",  
6 "arguments": {"data_source": "data_1.csv", "key_name": "  
  ↳ satscores_NumTstTskr",  
7 "value": 500.0},  
8 "label": "data_2.csv"},  
9 {"name": "get_schools_Schools",  
10 "arguments": {"data_source": "data_2.csv"},  
11 "label": "retrieved_json"}  
12 ]
```

Listing 4: REST-BIRD: API endpoint

```
1 {"output": [{"  
2 "name": "get_schools_magnet_status_v1_bird...",  
3 "description": "Get schools with specific magnet status  
  ↳ and number of test takers",  
4 "arguments": {  
5 "magnet_status": 1,  
6 "num_test_takers": 500 },  
7 "path": "/v1/bird/california_schools/schools_magnet_status  
  ↳ "}]}
```

Figure 1: A sample NL2SQL instance along with the ground-truth API sequence for SLOT-BIRD, SEL-BIRD and REST-BIRD. Function, slot and slot value descriptions not shown for ease of presentation.

datasets, we’re also able to provide a deterministic and verified ground-truth final-answer, which can be lacking in some tool-calling datasets (Qin et al., 2024; Liu et al., 2024b).

Recently, benchmarks that assess the capabilities of LLMs on Data Science tasks have been created (Jing et al., 2025). In contrast to our work such benchmarks require models to solve tasks by generating structured query languages or code (Liu et al., 2024c; Li et al., 2024).

To the best of our knowledge, our work is the first to demonstrate how NL2SQL datasets can be

used to study LLM tool-calling by re-purposing them to provide a collection of invocable APIs backed by real-world data and reusing real-world natural language queries available in the original datasets.

3 Dataset

We construct three NL2API datasets—*SLOT*, *SEL*, and *REST*—using 11 databases from the BIRD development set (Li et al., 2023a). Each dataset includes: (i) OpenAPI specifications for available APIs, (ii) Live API implementations, (iii) natural language (NL) questions paired with ground-truth API invocation sequences (with slot values), and (iv) the corresponding databases. We only retain samples where the generated API sequences produce results identical to the original SQL queries, Figure 1 (top) shows an example of an NL2SQL instance.

BIRD Collection: BIRD is one of the largest collections of real-world databases paired with crowd-sourced NL queries and their SQL statements. The development set averages 7 tables, 73 columns, and 358K rows per database. Queries are often complex, involving multiple joins, aggregations, comparisons, and matches, which translate into multi-step API sequences with challenging sequencing and nesting. Many queries also require reasoning over domain knowledge, synonyms, and numeric values.

3.1 SLOT-BIRD (Slot-Filling Version)

The SLOT-BIRD version of the dataset was constructed by decomposing generic SQL *SELECT* queries into their constituent parts and mapping each part to an API implemented as a Python function (tool). The design of these APIs is motivated by similar collections used in enterprise applications, such as Tableau⁶ or Google Analytics.⁷

Functions: Joins in the SQL query are handled collectively during an initialization step, implemented programmatically at evaluation time. This step produces a single table, which serves as the starting point for a model or agent to access via APIs when

answering the query.⁸ The SLOT-BIRD setup focuses on planning a sequence of data manipulation and access operations using seven APIs:

```
aggregate_data, filter_data,
group_data_by, retrieve_data,
select_unique_values, sort_data,
transform_data
```

Each function corresponds to a specific part of an SQL query, such as WHERE, ORDER BY, or COUNT.

Slot Values: The functions in the SLOT-BIRD API pool have two types of parameters (slots) that must be filled:

1. *Schema-agnostic slots:* Control operations on the data, e.g., the type of condition in a WHERE clause or the sort order (ascending/descending).
2. *Column-specific slots:* Specify which columns to use for operations such as filtering, sorting, and selecting after manipulations. These slots are categorical, corresponding to the columns in the joined table from the initialization step. Each column includes a brief description of its contents.

When a tool argument must be chosen from the column set, the OpenAPI specification includes an enum of allowed names and their descriptions. Correctly understanding these names and descriptions constitutes the primary challenge for a tool-calling model or agent on this dataset.

Ground-Truth API Sequence: In the example in Listing 2, *filter_data* (Lines 1, 5) and *retrieve_data* (Line 9) are two functions. The first invocation of *filter_data* takes two slot values: the first (Line 2) is a domain-specific slot *schools_Magnet*, and the second (Line 3) is a control operation that executes an *equal_to* comparison on the other slot. This is followed by another *filter_data* call and finally a *retrieve_data* call.

In addition to the function name and arguments, a *label* is provided, referencing the output of the tool call. This allows subsequent steps to access previous outputs, since models cannot reliably manipulate large data objects across function calls. As shown in Line 4 of Listing 2, we use files to pass payloads and reference them as input arguments. Models are not responsible for reading or

⁶https://help.tableau.com/current/api/rest_api/en-us/REST/rest_api_concepts_filtering_and_sorting.htm

⁷<https://developers.google.com/analytics/devguides/reporting/data/v1/rest/v1beta/FilterExpression>

⁸While it would be possible to require models to perform JOINS themselves, our experiments show that they already struggle on this dataset without this added complexity.

Dataset	# Queries	# Tools	#Tools Avail/ Query	# Tool Calls/ Query	# Slots/ Query
SLOT-BIRD	665	7	7	2.7	3.29
SEL-BIRD	651	1256	49	2.9	0.05
REST-BIRD	1257	1250	125	1	1.38

Table 2: Invocable API-Bench dataset statistics

writing these files; this is handled automatically by our evaluation framework. For *SLOT-BIRD*, all tools except `retrieve_data` save the results of their data manipulation operations into CSV files and return the file path as a string.

Ground-truth API sequences were constructed programmatically for each NL2SQL instance by parsing the SQL query with the `Sqlglot` Python library⁹ and mapping nodes in the resulting syntax tree to tools from the *SLOT-BIRD* API pool. The output of each ground-truth API sequence is equivalent (up to minor reformatting) to executing the original SQL query on the underlying databases. This output gold answer enables measuring the completion rate of a model or agent, i.e., the proportion of times it produces a sequence of tool calls that leads to the correct result.

3.2 SEL-BIRD (Selection Version)

The *SEL-BIRD* variant is derived from *SLOT-BIRD* by expanding categorical function arguments into separate functions. Specifically, each possible value of a categorical argument is bound to create a new function with one fewer input parameter. For example, in the *SLOT-BIRD* API pool, the function `filter_data` has a categorical argument `condition`, which can take values such as `equal_to` (Listing 2: Line 3) or `greater_than` (Listing 2: Line 7). In the *SEL-BIRD* pool, these become distinct functions—`select_data_equal_to` (Listing 3: Line 1) and `select_data_greater_than` (Listing 3: Line 5)—which no longer require a `condition` argument. This transformation yields a substantially larger set of tools for data manipulation.

Additional Domain-Specific Functions: The expansion is even greater for data retrieval. Rather than passing a column name as an argument (as in *SLOT-BIRD*), the *SEL-BIRD* pool provides a dedicated `get` function for each column key (Listing 3: Line 9). Consequently, the available toolset varies across instances, since the columns in the

initialized table depend on the JOINS specified in the underlying SQL query.

3.3 REST-BIRD Version

API Design: The *REST-BIRD* dataset extends the expansion strategy of *SEL-BIRD* to the extreme, assigning a dedicated REST endpoint to every instance in the NL2SQL dataset. We adopt a RESTful design to leverage meaningful path parameters, resulting in highly specific and interpretable endpoints. For example, in Listing 4, the original NL query is mapped directly to a single API endpoint (Line 2). Since *BIRD* contains only ‘SELECT’ queries, all requests in *REST-BIRD* are GET; no POST, PUT, or DELETE requests are included.

Data Generation: *REST-BIRD* produces a very large number of database-specific endpoints, with each user query answered by exactly one API call. Instead of manually authoring these endpoints, we employ `Mistral-Large`¹⁰ within an agentic pipeline consisting of four stages: (i) **Code Generation Agent** — synthesizes FastAPI server code, (ii) **De-duplication Agent** — merges functionally equivalent endpoints (e.g., `getEmployees` vs. `getAllEmployees`), (iii) **API Execution Module**, and (iv) **Verifier and Filtering Agent** — ensures that the generated endpoint produces results identical to the original SQL query. Any instances where the results do not match are discarded and excluded from the dataset. Complete details of this component and the overall data generation process are provided in the Appendix.

The pipeline outputs Python-based FastAPI server code, dockerized and deployed as microservices, resulting in hosted, executable REST APIs.

4 Experiments

We use our three datasets to answer the following research questions (i) Do models find the automatically generated API collections challenging?, (ii) How are models affected by the size of the API tool set provided for the task of function selection?, (iii) To what extent do models rely on semantic signals from the function name to solve the task?, (iv) How does the performance improve when models are employed as ReACT agents that can interact with live APIs?

⁹<https://github.com/tobymao/sqlglot>

¹⁰<https://huggingface.co/mistralai/Mistral-Large-Instruct-2411>

4.1 Models and Prompt Formats

We experiment with the following models: Llama 3.1-8B-Instruct and Llama 3.3-70b-instruct (Grattafiori et al., 2024), Qwen2.5-7b-instruct and Qwen2.5-72b-instruct (Yang et al., 2024), DeepSeek-v3 (Liu et al., 2024a), GPT4o-2024-08-06 (Hurst et al., 2024), Granite 3.1-8b-instruct¹¹, as well as Hammer-7b (Lin et al., 2024), and Watt-8b¹² which are specialized tool-calling models.

For all evaluations on *SLOT* and *SEL*, we adopt a prompting format similar to that used in NESTful (Basu et al., 2025), as these datasets involve multiple nested API sequences. For evaluations on *REST*, we follow a prompting format similar to that used in BFCL (Yan et al., 2024). GPT4o limits tool names to 64 characters¹³, so we experiment with two settings: providing APIs as ‘tools’¹⁴ or embedding them directly in the prompts. All LLM prompt templates are included in Appendix E.

Our input data is formatted to ensure a structured and reliable mechanism for referencing tool outputs, a crucial aspect of tool calling, that allows subsequent tool calls to utilize the results of previous execution. This is achieved by assigning a unique variable name to each tool call that ensures that each tool’s output can be distinctly identified and referenced. This approach is particularly important and prevents ambiguity when multiple instances of the same tool, with different arguments, appear within the same tool call sequence (e.g., parallel tool call). With the help of unique identifiers for each tool call, we facilitate clear and efficient tool chaining that ensures the tool dependencies are correctly resolved. All models are required to follow a specified format included in the prompt (Appendix E).

Output Parsing: Although models are instructed to produce outputs in JSON format, they often deviate, leading to instruction-alignment errors that reflect genuine deployment failures. Since proper tool-formatting is essential for functional tool calling, these deviations highlight a critical gap in current models’ ability to reliably interface with API systems. To ensure fair evaluation while

preserving parseable outputs, we employ robust, model-specific output parsers, as described in Appendix G.

ReACT Agents: ReACT agents (Yao et al., 2023) use LLMs to plan one tool call at a time, execute the tool and use the observation to plan the next step. We chose to experiment with larger architecture models over smaller ones, as in early experiments even larger models performed poorly when invoked directly as LLMs without an agent.

4.1.1 Metrics

Intent: For evaluating all models, we report a *position-aware* intent precision, recall, and *F1* metrics using a ground-truth API sequence as reference. When a tool call needs to be called more than once but with different slot-values, it is important to track each *instance* of a function invocation thus, its sequence of invocation.

Slots: Additionally, we report precision, recall and *F1* metrics for slots - i.e., the arguments for each function. The slot metrics are conditional on the correct ground-truth intent being predicted but slots are not penalized for missing ground-truth intents. Overall performance, should therefore always be studied along with the Intent metrics.

Completion Rate (Compl. Rate): Finally, we report the completion rate i.e., the proportion of all instances where the the models were able to return the ground-truth response.

4.2 Performance of Models

SLOT-BIRD: Recall that this dataset has 7 general tools and the models need to populate appropriate slot for each invocation of a general tool and also sequence them correctly to accomplish the task. As can be seen from Table 3, this is very challenging task for all models. While the set of APIs to choose from is small, they need to be invoked more than once - this is reflected in the relatively high precision as compared to recall in intent metrics. We find that Qwen2.5-72b-instruct performs best on the intent selection task though DeepSeek-V3 manages to pair them with the most correct slot-values and reports the highest completion rate (7%).

SEL-BIRD: This dataset has domain specific tools available that reduce the need to invoke general tools repeatedly. As can be seen from Table 3, models have lower precision and recall as compared to *SLOT-BIRD* suggesting that models do not select and sequence right set of tools. As compared to *SLOT-BIRD*, all models have lower scores when

¹¹<https://huggingface.co/ibm-granite/granite-3.1-8b-instruct>

¹²<https://huggingface.co/watt-ai/watt-tool-8B>

¹³<https://community.openai.com/t/function-call-description-max-length/529902>

¹⁴<https://platform.openai.com/docs/assistants/tools>

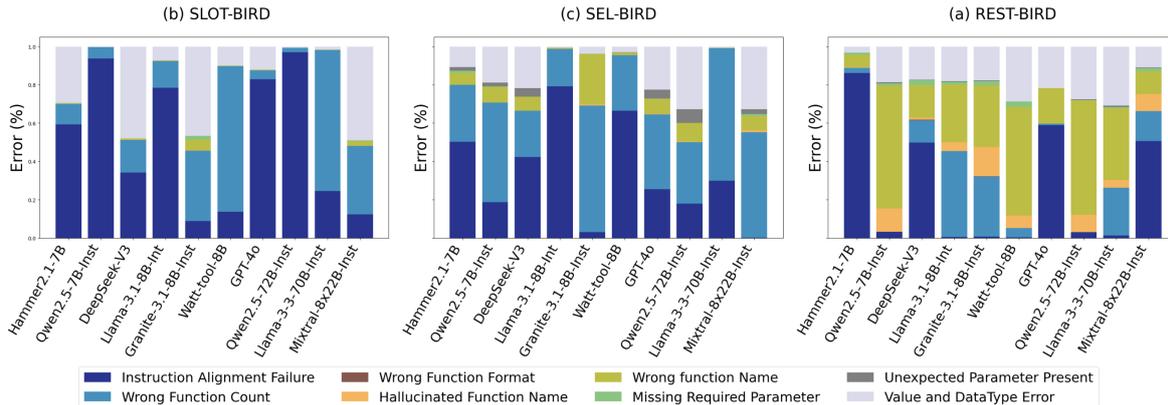


Figure 2: Tool-calling error classification

comparing slots. Notably Granite, has a severe drop in performance on this dataset while the others drop by approx. 0.3 (absolute) F1 points. As before, based on the completion rates Qwen-72b-instruct and DeepSeek-v3 are the top two performing models.

REST-BIRD: Finally, in this version of the dataset every query uniquely maps to one invocation and the primary task is of intent selection and then correctly populating the slots. Of the three datasets, this appears to be the one where models have most success with significantly higher intent selection rates (no sequencing needed). While the Qwen-2.5-72b-instruct model continues to be a top-performer, smaller models such as Granite-3.1-8b-instruct and Watt-tool-8b perform better than DeepSeek-v3 and GPT-4o. A closer investigation reveals that for this dataset DeepSeek generates chain-of-thought texts that is relatively harder to standardize for parsing and GPT-4o tends to answer the question instead of picking the function call. Refer to Appendix G for examples from our evaluation.

Task Completion Rates: Models struggle on multi-step function-sequencing tasks, achieving near-zero completion on *SLOT-BIRD* and *SEL-BIRD*, and only modest rates on *REST-BIRD*, far below levels needed for practical deployment. GPT-4o (2024-08-06) reaches 38% on *REST-BIRD*, 9% on *SEL-BIRD*, and 3% on *SLOT-BIRD*. For comparison, GPT-4-0613 evaluated on *API-Bank* attains 60.24%, GPT-4 evaluated on *ToolBench* with a depth-first search-based decision tree (DFSdT) prompting strategy achieves 70.4%, and GPT-4o evaluated on *NESTful* reaches 60.0%, highlighting the relative difficulty of our benchmarks.

We hypothesize that agentic architectures incor-

porating planning and reflection could substantially increase completion rates. A primary challenge for tool-calling models is accurately interpreting the database schema to populate function slot values. Although schema information is provided via enum descriptions in the tools’ OpenAPI specifications and included in the model prompts, effectively leveraging it typically requires multi-step reasoning.

Error Analysis : We employ the BFCL error categories to categorize the errors made by models on the API task (Yan et al., 2024). Errors are classified into the first type encountered in the list (See Figure 2), so that they are non-overlapping by design.

Notably, on the *SLOT-BIRD* dataset, most failures occur when models struggle to produce outputs in the structured format required for sequencing APIs. On the *SEL-BIRD* dataset, this issue is less pronounced, and on *REST-BIRD*, other types of errors predominate. This suggests that the complexity of tasks involving multiple function calls and the processing of their outputs challenges models, leading to additional failures such as tool-format issues. While structured decoding could potentially mitigate these errors, prior work has shown that it can also impair model reasoning (Tam et al., 2024). Together, these patterns highlight the challenges of multi-step reasoning and accurate function selection in complex tool-calling tasks, and our dataset’s characteristics provide a valuable benchmark for studying these nuanced aspects of model behavior.

4.3 Performance of ReACT agents

As experimenting with all models as ReACT agents would become very expensive due to think-act-observe (TAO) loops we choose 4 models from our

Model	Intent			Slot			Compl. Rate
	P	R	F1	P	R	F1	
SLOT-BIRD							
DeepSeek-V3	0.82	0.54	0.65	0.67	0.67	0.67	0.07
GPT4o-2024-08-06 (Prompt)	0.31	0.15	0.20	0.71	0.71	0.71	0.03
GPT4o-2024-08-06 (Tools)	0.90	0.52	0.66	0.43	0.42	0.42	0.03
Granite-8b-instruct	0.63	0.53	0.58	0.44	0.41	0.43	0.00
Hammer2.1-7b	0.88	0.35	0.50	0.67	0.67	0.67	0.03
Llama-3.1-8B-instruct	0.58	0.09	0.16	0.61	0.62	0.61	0.00
Llama-3.3-70b-instruct	0.75	0.22	0.34	0.65	0.03	0.05	0.00
Mixtral-8x22B-instruct	0.78	0.63	0.70	0.63	0.62	0.62	0.02
Qwen2.5-7b-instruct	0.65	0.67	0.66	0.63	0.63	0.63	0.03
Qwen2.5-72b-instruct	0.80	0.63	0.71	0.61	0.61	0.61	0.06
Watt-tool-8b	0.43	0.20	0.27	0.51	0.52	0.52	0.01
SEL-BIRD							
DeepSeek-V3	0.44	0.28	0.34	0.45	0.44	0.44	0.09
GPT4o-2024-08-06 (Prompt)	0.42	0.39	0.4	0.47	0.46	0.46	0.09
GPT4o-2024-08-06 (Tools)	0.47	0.30	0.36	0.62	0.57	0.59	0.0
Granite-3.1-8b-Inst	0.05	0.05	0.05	0.21	0.18	0.2	0.0
Hammer2.1-7b	0.29	0.16	0.21	0.47	0.4	0.43	0.03
Llama-3.1-8B-Inst	0.11	0.01	0.02	0.28	0.26	0.27	0.0
Llama-3-3-70b-Inst	0.41	0.11	0.17	0.25	0.01	0.02	0.0
Mixtral-8x22B-Inst	0.6	0.5	0.55	0.45	0.44	0.44	0.04
Qwen2.5-7b-Inst	0.16	0.29	0.2	0.39	0.39	0.39	0.04
Qwen2.5-72b-Inst	0.55	0.48	0.51	0.46	0.46	0.46	0.16
Watt-tool-8b	0.46	0.1	0.16	0.43	0.45	0.44	0.01
REST-BIRD							
DeepSeek-V3	0.65	0.50	0.57	0.77	0.74	0.76	0.31
GPT4o-2024-08-06 (Prompt)	0.82	0.54	0.65	0.79	0.79	0.79	0.38
Granite-8b-instruct	0.45	0.58	0.50	0.77	0.76	0.77	0.34
Hammer2.1-7b	0.70	0.22	0.34	0.89	0.86	0.87	0.17
Llama-3.1-8B-instruct	0.22	0.57	0.31	0.76	0.76	0.76	0.32
Llama-3.3-70b-instruct	0.57	0.67	0.61	0.76	0.76	0.76	0.42
Mixtral-8x22B-instruct	0.46	0.39	0.42	0.78	0.77	0.77	0.24
Qwen2.5-7b-instruct	0.53	0.52	0.53	0.80	0.80	0.80	0.37
Qwen2.5-72b-instruct	0.66	0.65	0.66	0.82	0.82	0.82	0.47
Watt-tool-8b	0.60	0.64	0.62	0.78	0.76	0.77	0.43

Table 3: Direct invocation performance across datasets and models showing precision (P), recall (R), and F1 scores for intent and slot prediction, along with completion rates.

experiments to serve as ReACT agents. We choose Mixtral-8x22B-instruct which is large mixture-of-expert model that surprisingly performs poorly on our data and wanted to see how the performance changes when it is used as an agent. Additionally, we choose the Llama 3.3-70B-instruct, Qwen2.5-72b-instruct, and GPT4o models for their widespread use. We configure our ReACT agents with a fixed TAO loop budget of 10 turns.

As can be seen from Table 4, the task completion rates for the agents do not exceed 15% and 17% for the *SLOT-BIRD* and *SEL-BIRD* datasets respectively, though performance on *REST-BIRD* is much higher (50% completion rate by GPT4o). Additionally, agents almost never run out of TAO-loop budget or get stuck repeating the same step on the *REST-BIRD* dataset but face both problems on the *SLOT-BIRD* and *SEL-BIRD* datasets.

Model	Base Comp.	Agent Comp.	Avg. Loops	OOB	Stuck	Others Errors
Mixtral-8x22B-instruct	0.02	0.08	5.91	0.27	0.60	0.07
Llama-3.3-70b-instruct	0.00	0.06	3.28	0.05	0.32	0.59
Qwen2.5-72b-instruct	0.06	0.14	5.58	0.25	0.50	0.14
GPT4o-2024-08-06	0.03	0.15	5.45	0.10	0.54	0.23
SEL-BIRD						
Mixtral-8x22B-instruct	0.04	0.05	6.40	0.40	0.27	0.29
Llama-3.3-70b-instruct	0.00	0.08	8.20	0.66	0.58	0.00
Qwen2.5-72b-instruct	0.16	0.17	4.93	0.17	0.24	0.47
GPT4o-2024-08-06	0.09	0.12	6.94	0.30	0.33	0.29
REST-BIRD						
Mixtral-8x22B-instruct	0.24	0.32	3.16	0.00	0.02	0.65
Llama-3.3-70b-instruct	0.42	0.42	3.37	0.00	0.05	0.52
Qwen2.5-72b-instruct	0.47	0.49	3.26	0.00	0.02	0.48
GPT4o-2024-08-06	0.38	0.50	2.54	0.00	0.01	0.49

Table 4: ReACT agent performance across different datasets and models. Completion rates shown as decimals (1.0 = 100%). OOB = Out-of-budget Errors.

4.3.1 Effect of Obfuscation on tool calling

To investigate whether models can make use of function descriptions and other meta-data, we repeated experiments shown in Table 4 using non-informative function and argument names. Each API name was replaced by “func_N” where N is a unique integer (Paul et al., 2025). We found that agent performance dropped slightly (3%) for REST-BIRD, but for SLOT-BIRD and SEL-BIRD fell to nearly zero (see Appendix B). This is likely due to large amount of descriptive information available in REST-BIRD, with function descriptions for each of the 1200 APIs. In contrast, tools in SLOT-BIRD and SEL-BIRD are much more generic, and therefore their descriptions were correspondingly more challenging to interpret.

4.3.2 Effect of the number of tools available

Since every query in the *REST-BIRD* dataset only requires a single API, we experiment by shrinking candidate set size, while ensuring ground-truth intent is always present. The ability of agents to select correct intent does indeed go up as number of choices shrinks. Interestingly, even with 90% of the tools removed, the best performing GPT4o based ReACT agent has a task completion-rate of just 71% (see Appendix B).

5 Conclusion

In this paper, we presented Live API Bench, a benchmark that transformed NL2SQL datasets into interactive APIs for evaluating LLM tool-calling capabilities. Spanning 11 databases across diverse domains and including over 2,500 invocable APIs

paired with human-authored queries, the benchmark covered three styles—SLOT, SEL, and REST. These tasks exposed models to realistic challenges such as multi-step reasoning, sequential tool calls, parameter generation, and complex response parsing. Our experiments showed that even state-of-the-art LLMs struggled with these tasks, highlighting substantial gaps in current tool-calling abilities and the need for further research on robust API interaction.

This presentation of SLOT-BIRD, SEL-BIRD, and REST-BIRD dataset only included versions packaged as either Python functions or REST APIs. We are currently preparing to release a public benchmark and leaderboard based on an expanded version of this dataset. This benchmark also incorporates RAG tasks, more complex reasoning requirements, policy/instruction following, and multi-turn dialogues. It will be based on a Model Context Protocol (MCP) implementation of all tools and APIs.

6 Limitations

Most NL2SQL datasets, including WikiSQL (Zhong et al., 2017), Spider (Yu et al., 2018), and the BIRD-SQL dataset (Li et al., 2023a) used in this work, contain only SELECT queries. Data manipulation commands such as INSERT, UPDATE, and DELETE are not included. Consequently, our benchmark inherits this limitation: all generated APIs and evaluation tasks are based on SELECT queries. In addition, we note that supporting data-manipulation commands using live APIs would require additional infrastructure, such as containerized database environments or managed services, to ensure safe and deterministic execution, introducing substantial operational complexity.

Furthermore, to guarantee reliable ground-truth verification, the current dataset excludes queries involving nested SELECT statements, CASE expressions, and complex logical compositions of WHERE clauses. Nested and conditional constructs greatly complicate the mapping from SQL syntax trees to tool sequences, as in the SLOT-BIRD formulation, where multiple alternatives make it difficult to track intent and slot metrics. In these cases, the abstract syntax tree is no longer linear but contains hierarchical subqueries or conditional branches, introducing ambiguity in the canonical decomposition into tool calls. By excluding such constructs

from this initial version of the dataset, we ensure high-quality, deterministic ground-truth sequences for evaluation. Ongoing work is expanding the scope and complexity of SQL syntax included in this work.

Due to limitations on computing resources, these exploratory investigations were carried out using only the dev set from the BIRD dataset. The much larger train set has since also been processed and will be included in the public data release, greatly expanding the overall size of the dataset.

References

- Kinjal Basu, Ibrahim Abdelaziz, Subhajit Chaudhury, Soham Dan, Maxwell Crouse, Asim Munawar, Vernon Austel, Sadhana Kumaravel, Vinod Muthusamy, Pavan Kapanipathi, and Luis Lastras. 2024. [API-BLEND: A comprehensive corpora for training and benchmarking API LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12859–12870, Bangkok, Thailand. Association for Computational Linguistics.
- Kinjal Basu, Ibrahim Abdelaziz, Kiran Kate, Mayank Agarwal, Maxwell Crouse, Yara Rizk, Kelsey Bradford, Asim Munawar, Sadhana Kumaravel, Saurabh Goyal, Xin Wang, Luis A. Lastras, and Pavan Kapanipathi. 2025. [Nestful: A benchmark for evaluating llms on nested sequences of api calls](#). *Preprint*, arXiv:2409.03797.
- 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, and 4 others. 2024. [Spider2-v: How far are multimodal agents from automating data science and engineering workflows?](#) *Preprint*, arXiv:2407.10956.
- Wei Chen, Zhiyuan Li, and Mingyuan Ma. 2024. [Octopus: On-device language model for function calling of software apis](#). *Preprint*, arXiv:2404.01549.
- Alexandre Drouin, Maxime Gasse, Massimo Caccia, Issam H Laradji, Manuel Del Verme, Tom Marty, David Vazquez, Nicolas Chapados, and Alexandre Lacoste. 2024. [Workarena: how capable are web agents at solving common knowledge work tasks?](#) In *Proceedings of the 41st International Conference on Machine Learning*, pages 11642–11662.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. [The llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.

- Han Han, Tong Zhu, Xiang Zhang, Mengsong Wu, Hao Xiong, and Wenliang Chen. 2024. Nestools: A dataset for evaluating nested tool learning abilities of large language models. *arXiv preprint arXiv:2410.11805*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Liqiang Jing, Zhehui Huang, Xiaoyang Wang, Wenlin Yao, Wenhao Yu, Kaixin Ma, Hongming Zhang, Xinya Du, and Dong Yu. 2025. [DSBench: How far are data science agents from becoming data science experts?](#) In *The Thirteenth International Conference on Learning Representations*.
- Jinyang Li, Binyuan Hui, GE QU, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, Xuanhe Zhou, Chenhao Ma, Guoliang Li, Kevin Chang, Fei Huang, Reynold Cheng, and Yongbin Li. 2023a. [Can LLM already serve as a database interface? a BIG bench for large-scale database grounded text-to-SQLs](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, and Yongbin Li. 2023b. [Api-bank: A comprehensive benchmark for tool-augmented llms](#). *arXiv preprint arXiv:2304.08244*.
- Ziming Li, Qianbo Zang, David Ma, Jiawei Guo, Tuney Zheng, Minghao Liu, Xinyao Niu, Yue Wang, Jian Yang, Jiaheng Liu, Wanjuan Zhong, Wangchunshu Zhou, Wenhao Huang, and Ge Zhang. 2024. [Autokaggle: A multi-agent framework for autonomous data science competitions](#). *Preprint*, arXiv:2410.20424.
- Qiqiang Lin, Muning Wen, Qiuying Peng, Guanyu Nie, Junwei Liao, Jun Wang, Xiaoyun Mo, Jiamu Zhou, Cheng Cheng, Yin Zhao, and 1 others. 2024. [Hammer: Robust function-calling for on-device language models via function masking](#). *arXiv preprint arXiv:2410.04587*.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Weiwen Liu, Xu Huang, Xingshan Zeng, Xinlong Hao, Shuai Yu, Dexun Li, Shuai Wang, Weinan Gan, Zhengying Liu, Yuanqing Yu, and 1 others. 2024b. [Toolace: Winning the points of llm function calling](#). *arXiv preprint arXiv:2409.00920*.
- Xiao Liu, Da Yin, Zirui Wu, and Yansong Feng. 2025. [Reftool: Enhancing model reasoning with reference-guided tool creation](#). *arXiv preprint arXiv:2505.21413*.
- Xinyu Liu, Shuyu Shen, Boyan Li, Peixian Ma, Runzhi Jiang, Yuxin Zhang, Ju Fan, Guoliang Li, Nan Tang, and Yuyu Luo. 2024c. [A survey of nl2sql with large language models: Where are we, and where are we going?](#) *arXiv preprint arXiv:2408.05109*.
- Zuxin Liu, Thai Hoang, Jianguo Zhang, Ming Zhu, Tian Lan, Shirley Kokane, Juntao Tan, Weiran Yao, Zhiwei Liu, Yihao Feng, Rithesh Murthy, Liangwei Yang, Silvio Savarese, Juan Carlos Niebles, Huan Wang, Shelby Heinecke, and Caiming Xiong. 2024d. [Apigen: Automated pipeline for generating verifiable and diverse function-calling datasets](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 54463–54482. Curran Associates, Inc.
- Yichen Pan, Dehan Kong, Sida Zhou, Cheng Cui, Yifei Leng, Bing Jiang, Hangyu Liu, Yanyi Shang, Shuyan Zhou, Tongshuang Wu, and 1 others. 2024. [Webcanvas: Benchmarking web agents in online environments](#). In *Agentic Markets Workshop at ICML 2024*.
- Shishir G Patil, Tianjun Zhang, Xin Wang, and Joseph E Gonzalez. 2024. [Gorilla: Large language model connected with massive apis](#). *Advances in Neural Information Processing Systems*, 37:126544–126565.
- Indraneil Paul, Haoyi Yang, Goran Glavaš, Kristian Kersting, and Iryna Gurevych. 2025. [Obscuracoder: Powering efficient code LM pre-training via obfuscation grounding](#). In *The Thirteenth International Conference on Learning Representations*.
- Yun Peng, Shuqing Li, Wenwei Gu, Yichen Li, Wenxuan Wang, Cuiyun Gao, and Michael Lyu. 2021. [Revisiting, benchmarking and exploring api recommendation: How far are we?](#) *Preprint*, arXiv:2112.12653.
- André Pereira, Bruno Lima, and João Pascoal Faria. 2024. [Aptestgenie: Automated api test generation through generative ai](#). *arXiv preprint arXiv:2409.03838*.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, dahai li, Zhuyuan Liu, and Maosong Sun. 2024. [ToolLLM: Facilitating large language models to master 16000+ real-world APIs](#). In *The Twelfth International Conference on Learning Representations*.
- Yangjun Ruan, Honghua Dong, Andrew Wang, Silviu Pitis, Yongchao Zhou, Jimmy Ba, Yann Dubois, Chris J Maddison, and Tatsunori Hashimoto. 2024. [Identifying the risks of lm agents with an llm-emulated sandbox](#). In *The Twelfth International Conference on Learning Representations*.
- Dingfeng Shi, Jingyi Cao, Qianben Chen, Weichen Sun, Weizhen Li, Hongxuan Lu, Fangchen Dong, Tianrui Qin, King Zhu, Minghao Liu, and 1 others. 2025. [Taskcraft: Automated generation of agentic tasks](#). *arXiv preprint arXiv:2506.10055*.

- Yifan Song, Weimin Xiong, Dawei Zhu, Wenhao Wu, Han Qian, Mingbo Song, Hailiang Huang, Cheng Li, Ke Wang, Rong Yao, and 1 others. 2023. Restgpt: Connecting large language models with real-world restful apis. *arXiv preprint arXiv:2306.06624*.
- Zhi Rui Tam, Cheng-Kuang Wu, Yi-Lin Tsai, Chieh-Yen Lin, Hung-yi Lee, and Yun-Nung Chen. 2024. [Let me speak freely? a study on the impact of format restrictions on large language model performance](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, pages 1218–1236, Miami, Florida, US. Association for Computational Linguistics.
- Qiaoyu Tang, Ziliang Deng, Hongyu Lin, Xianpei Han, Qiao Liang, Boxi Cao, and Le Sun. 2023. [Toolalpaca: Generalized tool learning for language models with 3000 simulated cases](#). *Preprint*, arXiv:2306.05301.
- Tianbao Xie, Danyang Zhang, Jixuan Chen, Xiaochuan Li, Siheng Zhao, Ruisheng Cao, Toh Jing Hua, Zhoujun Cheng, Dongchan Shin, Fangyu Lei, Yitao Liu, Yiheng Xu, Shuyan Zhou, Silvio Savarese, Caiming Xiong, Victor Zhong, and Tao Yu. 2024. [Osworld: Benchmarking multimodal agents for open-ended tasks in real computer environments](#). *Preprint*, arXiv:2404.07972.
- Binfeng Xu, Zhiyuan Peng, Bowen Lei, Subhabrata Mukherjee, Yuchen Liu, and Dongkuan Xu. 2023a. [Rewoo: Decoupling reasoning from observations for efficient augmented language models](#). *Preprint*, arXiv:2305.18323.
- Qiantong Xu, Fenglu Hong, Bo Li, Changran Hu, Zhengyu Chen, and Jian Zhang. 2023b. [On the tool manipulation capability of open-source large language models](#). *Preprint*, arXiv:2305.16504.
- Fanjia Yan, Huanzhi Mao, Charlie Cheng-Jie Ji, Tianjun Zhang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. 2024. Berkeley function calling leaderboard. https://gorilla.cs.berkeley.edu/blogs/8_berkeley_function_calling_leaderboard.html.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Shunyu Yao, Noah Shinn, Pedram Razavi, and Karthik Narasimhan. 2024. [\$\tau\$ -bench: A benchmark for tool-agent-user interaction in real-world domains](#). *Preprint*, arXiv:2406.12045.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Ori Yoran, Samuel Joseph Amouyal, Chaitanya Malaviya, Ben Bogin, Ofir Press, and Jonathan Berant. 2024. [Assistantbench: Can web agents solve realistic and time-consuming tasks?](#) *Preprint*, arXiv:2407.15711.
- Tao Yu, Rui Zhang, Kai Er, Sida Li, Zihan Tang, R. Zhang, S. Li, Z. Tang, and 1 others. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4206–4216.
- Boyuan Zheng, Boyu Gou, Scott Salisbury, Zheng Du, Huan Sun, and Yu Su. 2024. [WebOlympus: An open platform for web agents on live websites](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 187–197, Miami, Florida, USA. Association for Computational Linguistics.
- Lucen Zhong, Zhengxiao Du, Xiaohan Zhang, Haiyi Hu, and Jie Tang. 2025. [Complexfunbench: Exploring multi-step and constrained function calling under long-context scenario](#). *Preprint*, arXiv:2501.10132.
- Victor Zhong, Caiming Xiong, and Richard Socher. 2017. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*.
- Shuyan Zhou, Frank F. Xu, Hao Zhu, Xuhui Zhou, Robert Lo, Abishek Sridhar, Xianyi Cheng, Tianyue Ou, Yonatan Bisk, Daniel Fried, Uri Alon, and Graham Neubig. 2024. [Webarena: A realistic web environment for building autonomous agents](#). *Preprint*, arXiv:2307.13854.
- Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. 2023. Toolqa: A dataset for llm question answering with external tools. *Advances in Neural Information Processing Systems*, 36:50117–50143.

A Overview

We include additional experiments by studying the effect of obfuscation on ReACT agents in section B.1 and the effect of the size of the universe in Section B.2. We describe the artifacts generated by our API generation pipelines in Section C and include more details about our data generation pipelines in Section D. We include all LLM prompt templates in Section E and agent prompt templates in Section F. Finally, we include additional details on output parsing in Section G and how we conducted the error analysis in Section G.1.1.

B Additional Experiments

B.1 Effect of Obfuscation

B.1.1 Data Construction

As described in the main paper, our goal was to study the impact of function descriptions and func-

tion names. To this end, we obfuscated both API and slot names, replacing them with placeholders such as FUNC_0, FUNC_1, ..., FUNC_N and ARG_1, ARG_2, ..., ARG_N. An example is shown in Fig. 4, where an API name and argument gets obfuscated and the obfuscated tool name and slots will be passed to the Agent.

B.1.2 Results

We re-emphasize a significant contribution of our work: the ability to generate a substantial collection of invocable APIs. We leverage this aspect to investigate the extent to which models can rely on function descriptions and other meta-data when completing tasks with non-informative function names that provide no insight into their intended purpose. To achieve this, we obfuscate each API by assigning a unique integer to the prefix “func” (Paul et al., 2025). We evaluate performance using the three ReACT agents as before, as they are capable of exploring the environment based on observations (feedback). Unsurprisingly, the performance of models with obfuscated function names is lower across all datasets (Appendix Figure 4) with a slight drop in performance on the REST-BIRD dataset but with severe failures on the SLOT-BIRD and SEL-BIRD datasets. This result underscores the importance of employing domain-specific vocabularies to construct tools that enhance agents’ ability to improve their tool calling performance.

Dataset	10%	25%	50%	75%	100%
formula_1	12	31	63	94	126
card_games	15	38	77	115	154
superhero	10	27	54	81	109
codebase_community	15	39	78	117	156
thrombosis_prediction	13	33	67	101	135
toxicology	9	23	47	71	95
financial	9	22	45	68	91
california_schools	7	18	37	56	75
student_club	13	32	65	97	130
european_football_2	11	29	59	88	118
debit_card_specializing	6	15	30	45	61

Table 5: Tool distribution percentiles across datasets

B.2 Effect of the number of tools available

As described in the main paper, we also examine the effect of varying the number of tools provided to the agent. The REST dataset features a large API universe, and in some cases, over 100 tools are passed to the model—making the selection task significantly more challenging.

B.2.1 Data Construction

To study the impact of toolset size, we introduce a best-effort shortlisting mechanism. Given a total of

Model	SLOT-BIRD		SEL-BIRD		REST-BIRD	
	API	SQL	API	SQL	API	SQL
Llama-3.1-8B-instruct	0.00	0.35	0.0	0.35	0.32	0.24
Llama-3.3-70b-instruct	0.00	0.45	0.0	0.45	0.42	0.35
Granite-8b-instruct	0.00	0.13	0.0	0.13	0.34	0.09
Hammer2.1-7b	0.03	–	0.03	–	0.17	–
Watt-tool-8b	0.01	–	0.01	–	0.43	–
Mixtral-8x22B-instruct	0.02	0.29	0.04	0.29	0.24	0.22
DeepSeek-V3	0.07	0.40	0.09	0.40	0.31	0.29
GPT4o-2024-08-06 (Prompt)	0.03	0.39	0.09	0.39	0.38	0.29
GPT4o-2024-08-06 (Tools)	0.03	–	0.0	–	–	–
Qwen2.5-7b-instruct	0.03	–	0.04	–	0.37	–
Qwen2.5-72b-instruct	0.06	0.38	0.16	0.38	0.47	0.27

Table 6: SQL Compl. Rate refers to the task completion rates when models that support SQL generation are used directly to query the database.

N tools, we reduce the candidate set to 10 percent, 25 percent, 50 percent, or 75 percent of the original size to simulate different levels of tool availability.

We employ a best-effort shortlister that always includes the ground-truth tool and supplements it with N - 1 additional (random) tools, resulting in a set of N tools that is then provided to the model.

For this experiment, we used the REACT agent and the prompt described in the previous section.

The below table shows the list of datasets along with the distribution.

B.2.2 Results

We re-emphasize a significant contribution of our work: the ability to generate a substantial collection of invocable APIs. We leverage this aspect to investigate the extent to which models can rely on function descriptions and other meta-data when completing tasks with non-informative function names that provide no insight into their intended purpose. To achieve this, we obfuscate each API by assigning a unique integer to the prefix “func” (Paul et al., 2025). We evaluate performance using the three ReACT agents as before, as they are capable of exploring the environment based on observations (feedback). Unsurprisingly, the performance of models with obfuscated function names is lower across all datasets (Appendix Figure 4) with a slight drop in performance on the REST-BIRD dataset but with severe failures on the SLOT-BIRD and SEL-BIRD datasets. This result underscores the importance of employing domain-specific vocabularies to construct tools that enhance agents’ ability to improve their tool calling performance.

B.3 Comparison to NL2SQL

Since most LLMs have also been trained to generate SQL queries; we also report the baseline task completion rate when models are prompted to func-

Listing 5: Starting API definition

```

1 "get_free_meal_count_ratio_v1_bird_california_schools_get": {
2   "name": "get_free_meal_count_ratio_v1_bird_california_schools_get",
3   "description": "Get free meal count ratio for a given county",
4   "arguments": {
5     "county_name": {
6       "type": "string",
7       "description": "Name of the county",
8       "title": "County Name",
9       "name": "ARG_1"
10    }
11  },
12  "path": "/v1/bird/california_schools/free_meal_count_ratio"
13 }

```

Listing 6: Obfuscated API definition

```

1 "FUNC_0": {
2   "name": "FUNC_0",
3   "description": "Get free meal count ratio for a given county",
4   "arguments": {
5     "ARG_1": {
6       "type": "string",
7       "description": "Name of the county",
8       "title": "County Name",
9       "name": "county_name"
10    }
11  },
12  "path": "/v1/bird/california_schools/free_meal_count_ratio"
13 }

```

Figure 3: The pre-obfuscation (Listing 5) and post-obfuscation (Listing 6) API specification for an example API endpoint from REST-BIRD.

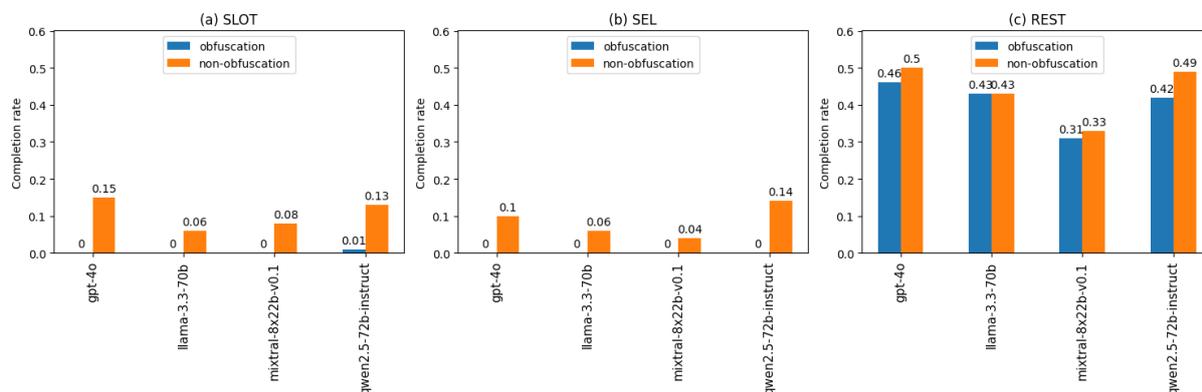


Figure 4: Effect of obfuscation on completion rate

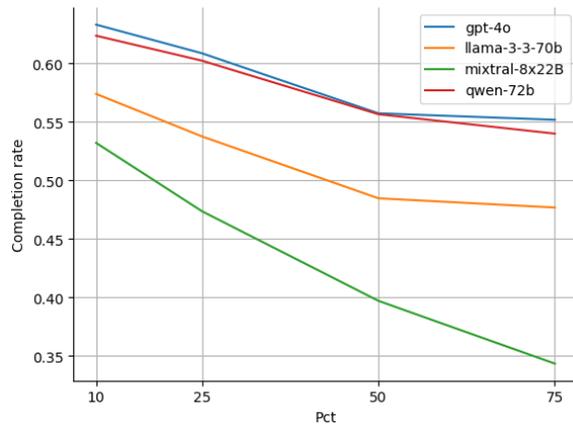


Figure 5: Effect of number of tools (percentage of universe) available on completion rate

tion as SQL-querying systems with direct access to the original database. Interestingly, except on the REST-BIRD dataset, we find that most models report higher completion rates when tasked with generating SQL as opposed to invoking functions.

C Artifacts

Each data generation pipeline produces the following three artifacts:

1. OpenAPI tool/API specifications: These specifications include both the names and parameters needed to invoke the tools, as well as descriptions with the semantic content needed to understand the purpose and behavior of the tools.
2. Invocable API/function implementation: The SLOT-BIRD and SEL-BIRD tools are provided as python functions which can be invoked programmatically or bound to a function calling LLM. The REST-BIRD APIs are provided as endpoints in a FastAPI server which has been containerized with Docker for easy deployment.
3. Evaluation set: The natural language, tool/API sequence pairs that form the evaluation set can be used to benchmark tool-calling LLMs or Agents.

C.1 OpenAPI Functions / Tools

The OpenAPI tool/API specifications include the intent (API) name, description, and input and output parameters. For each parameter, a name, description, and data type is provided. The subset of

input parameters that are required in each specification is also provided. For experiments with tool calling LLMs or agents, these specifications can be provided in the model prompt. The model must understand the specifications in order to successfully call the tools.

Below are some sample tool specifications from each of the three data generation pipelines.

SLOT-BIRD: The ‘sort_data’ tool from the SLOT-BIRD dataset includes a required ‘key_name’ parameter which takes values from a domain-specific enum. The possible values and their descriptions are provided as part of the tool specification (see Listing 7).

SEL-BIRD: The ‘sort_data_ascending’ tool from the SEL-BIRD dataset (see Listing 8) is equivalent to ‘sort_data’ in SLOT-BIRD, with the ‘ascending’ argument fixed to True. There is also a corresponding ‘sort_data_descending’ tool.

REST-BIRD: The REST-BIRD pipeline generates much more specific tools, which were constructed to answer a specific question, rather than perform manipulations on the underlying data, (see Listing 9).

C.2 Evaluation Dataset

Our data processing pipeline will also provide us with an evaluation set (in JSON format) or the test set will contain the following

- Utterance: Natural Language Utterance (from BIRD-SQL dataset)
- SQL Query: SQL Query from the BIRD-SQL dataset
- Gold Answer: Ground truth obtained by executing the SQL query on the BIRD SQL databases
- Output: Dictionary that contains the API (in OpenAI Function format) that will lead to the gold answer
- Output from executing the API: The output from executing the API will have the same data as the Gold Answer but it maybe in a different format

During our evaluation, we benchmark various models on the eval-set that we generated using our data generation pipeline.

Refer to Figure 1 for additional details about the eval-set.

Listing 7: SLOT-BIRD APIs: sort_data tool.

```

1 {"name": "sort_data",
2  "description": "Sort data by the values associated with the chosen key='key_name' If the input data is list-like, returns the
   ↳ sorted list. If the input data is tabular, returns the table with rows sorted by the values in column 'key_name'. If
   ↳ the data is grouped tables, then sort the groups by the value in 'key_name'",
3  "parameters": {
4    "properties": {
5      "data_source": {
6        "description": "The location of the data file in csv format.", "schema": {"type": "string"}
7      },
8      "key_name": {
9        "description": "name of key to sort by:
10         * 'member_member_id' - unique id of member
11         * 'member_first_name' - member's first name
12         ...", "schema": {"type": "string", "enum": ["member_member_id", "member_first_name", ...]}
13      },
14      "ascending": {
15        "description": "whether to sort by ascending order", "schema": {"type": "boolean"}}},
16    "required": ["data_source", "key_name", "ascending"],
17    "type": "object"},
18  "output_parameters": {
19    "properties": {
20      "output_0": {
21        "description": "The path to a csv file containing data sorted by chosen key", "type": "string"}
22    }
23  }
24 }}

```

Listing 8: SEL-BIRD APIs: sort_data_ascending tool.

```

1 {"name": "sort_data_ascending",
2  "description": "Sort data by the values associated with the chosen key='key_name' If the input data is list-like, returns the
   ↳ sorted list. If the input data is tabular, returns the table with rows sorted by the values in column 'key_name'. If
   ↳ the data is grouped tables, then sort the groups by the value in 'key_name'",
3  "parameters": {
4    "properties": {
5      "data_source": {
6        "description": "The location of the data file in csv format.", "schema": {"type": "string"}
7      },
8      "key_name": {
9        "description": "name of key to sort by:
10         * 'member_member_id' - unique id of member
11         * 'member_first_name' - member's first name
12         ...", "schema": {"type": "string", "enum": ["member_member_id", "member_first_name", ...]}
13      },
14    "required": ["data_source", "key_name", "ascending"], "type": "object"}, "output_parameters": {
15      "properties": {
16        "output_0": {"description": "The path to a csv file containing data sorted by chosen key", "type": "string"}
17      }
18    }
19  }
20 }}

```

C.3 Invocable APIs (Python Code or Microservices)

In addition to the previously described components, our data generation pipeline also produces invocable code.

From the REST-BIRD pipeline, we obtain a FastAPI server, which is dockerized and deployed as a microservice.

From the SEQ-BIRD and SEL-BIRD pipelines, we receive a suite of Python tools. These can be invoked by installing the accompanying invocable-api codebase (included with this submission) as a Python library.

D Data Generation Pipeline

The BIRD SQL dataset provides a rich source of structured information that can be leveraged to construct a high-quality API-centric dataset. Specifically, it includes the following components:

1. Natural language input utterances (questions)
2. Corresponding SELECT-SQL queries that yield correct answers
3. Real-world databases
4. Detailed database schemas, including table names, column names, and descriptions
5. Ground truth answers (obtainable by executing the associated SQL queries)

We get the below information from the BIRD SQL dataset

```

1 {"query": "SELECT 'Free Meal Count (K-12)' / 'Enrollment (
   ↳ K-12)' FROM frpm WHERE 'County Name' = 'Alameda'
   ↳ ORDER BY (CAST('Free Meal Count (K-12)' AS REAL)
   ↳ / 'Enrollment (K-12)') DESC LIMIT 1",
2  "input": "What is the highest eligible free rate for K-12
   ↳ students in the schools in Alameda County?",
3  "dataset_name": "california_schools",
4  "gold_answer": [[1.0]]}

```

For this work, we utilize the DEV set, which encompasses 11 domains: California Schools, Card

Listing 9: REST-BIRD APIs: `get_zip_codes` example endpoint.

```

1 {"name": "get_zip_codes_v1_bird_california_schools_zip_codes_get",
2  "description": "Get zip codes for a given district name and charter school status",
3  "arguments": {
4    "district_name": {
5      "type": "string",
6      "description": "Name of the district",
7      "title": "District Name"
8    },
9    "charter_school": {
10     "type": "integer",
11     "description": "Charter school status (1 for yes, 0 for no)",
12     "title": "Charter School"
13   }
14 },
15 "path": "/v1/bird/california_schools/zip_codes"}

```

Games, Code Base Community, Debit Card Specializing, European Football, Financial, Formula 1, Student Club, Superhero, Thrombosis Prediction, and Toxicology. The BIRD-SQL dataset contains exclusively SELECT statements and does not include UPDATE or INSERT queries. Consequently, the associated tools primarily function as data retrievers, effectively serving as getters.

D.1 SLOT-BIRD and SEL-BIRD Data Generation Pipelines

The SLOT-BIRD and SEL-BIRD datasets were constructed by first writing a set of python tools that could perform the same data manipulations as various parts of an SQL SELECT query. The Sqlglot python library was then used to parse each SQL query into an abstract syntax tree. The data generation pipeline then processed the parsed query components in the following order:

1. JOIN statements
2. WHERE statements
3. GROUPBY statements
4. ORDERBY statements
5. SELECT statements
6. AGGREGATE statements

JOIN Statements These were collectively combined into a single step, carried out by the ‘initialize_active_data’ function. This function is not included in the tool specifications provided to the models and agents in the experiments, due to the complex nature of its required arguments. Instead, it is called in the experiment scripts as a setup or data processing step while looping through the data instances. The output of this function is a single table on which the rest of the tool calls will operate.

It is saved to a temporary csv file which is pointed to in the model prompts.

WHERE Statements Currently any number of WHERE statements are supported as long as they are connected only by ANDs. The conditions on which the WHERE statements can be performed are: ‘equal_to’, ‘not_equal_to’, ‘greater_than’, ‘less_than’, ‘greater_than_equal_to’, ‘less_than_equal_to’, ‘contains’, and ‘like’. Additional conditions such as BETWEEN as well as clauses connected by OR will be included in future versions.

GROUPBY Statements Groupby statements are supported with a parameter to control aggregation type.

SELECT Statements The select statements are handled differently in SLOT-BIRD and SEL-BIRD. In SLOT-BIRD, the ‘retrieve_data’ function performs this task, returning a subset of the columns of the table provided to it. This function also includes optional parameters: ‘distinct’ which controls whether to only return a list of the distinct elements in the column and ‘limit’ which truncates the returned results.

In SEL-BIRD, each column is used to create a unique ‘get_table_and_column_name’ function, which returns only that column. So in this version, selecting N columns requires N tool calls instead of one. The distinct and limit parameters are not available, and instead these are handled by separate functions that must be applied to the output of the ‘get’.

A typical example from the student_club database that illustrates this procedure is shown in Listing 10.

The example shown in Listing 11 is exceptionally difficult because in order to correctly choose the initial tool in the sequence, a model must not only understand the content of the ‘Player_Attributes_date’ column, but also exam-

Listing 10: Simple example

```

1 {'query': "SELECT T2.major_name FROM member AS T1 INNER JOIN major AS T2 ON T1.link_to_major = T2.major_id WHERE T1.first_name
   ↪ = 'Angela' AND T1.last_name = 'Sanders'",
2 'input': "What's Angela Sanders's major?",
3 'gold_answer': 'Business',
4 'output': [{'name': 'filter_data',
5 'arguments': {'data_source': '$starting_table_var$',
6 'key_name': 'member_first_name',
7 'value': 'Angela',
8 'condition': 'equal_to'},
9 'label': 'FILTERED_DF_0'},
10 {'name': 'filter_data',
11 'arguments': {'data_source': '$FILTERED_DF_0$',
12 'key_name': 'member_last_name',
13 'value': 'Sanders',
14 'condition': 'equal_to'},
15 'label': 'FILTERED_DF_1'},
16 {'name': 'retrieve_data',
17 'arguments': {'data_source': '$FILTERED_DF_1$',
18 'key_name': 'major_major_name',
19 'distinct': False,
20 'limit': -1},
21 'label': 'SELECT_COL_0'}],
22 'dataset_name': 'student_club',
23 'sample_id': 0,
24 'initialization_step': {'name': 'initialize_active_data',
25 'arguments': {'condition_sequence': [['T1.link_to_major',
26 'T2.major_id',
27 'INNER']]},
28 'alias_to_table_dict': {'T1': {'original_table_name': 'member',
29 'modified_table_name': 'member'},
30 'T2': {'original_table_name': 'major', 'modified_table_name': 'major'}},
31 'database_path': '/home/belder/invocable-api-hub/invocable-api-hub/driver/../../db/cache/student_club.sqlite',
32 'label': 'starting_table_var'}}

```

ine the contents to see that it is required to choose a substring from the string value in each cell. This is not evident from the description provided:

```

1 {'key_name': 'Player_Attributes_date', 'description': '
   ↪ date', 'dtype': 'string'}

```

The need for this transformation can only be inferred by examining the data, which an agent capable of multi-round planning and reflection could accomplish.

D.1.1 Comparison to Enterprise API Collections

The SLOT-BIRD API design reflects common patterns found in enterprise analytics platforms. For instance, the Google Analytics Data API uses FilterExpression operations that require specifying field names, operators (equal_to, greater_than), and values—directly paralleling our filter_data function structure.¹⁵ Similarly, Tableau’s REST API provides generic filtering operations through their Filter class, where users specify field names and comparison operators.¹⁶ These enterprise examples demonstrate that our generic API formulation mirrors real-world tool-calling scenarios rather than being artificially constructed.

¹⁵<https://developers.google.com/analytics/devguides/reporting/data/v1/rest/v1beta/FilterExpression>

¹⁶https://help.tableau.com/current/api/rest-api/en-us/REST/rest_api_concepts_filtering_and_sorting.htm

Example 1: Tableau APIs, from the filter_sort_groups.py example:¹⁷

Listing 12: Tableau Tool

```

1 options.filter.add(TSC.Filter(TSC.
   RequestOptions.Field.Name, TSC.
   RequestOptions.Operator.Equals,
   group_name))

```

This example has a very similar structure to the filter_data tool in the SLOT-BIRD dataset, where the field name (key name) and operator (condition) need to be specified.

Similar construction from the Google Analytics APIs:¹⁸

Listing 13: Google Analytics Tool

```

1 FilterExpression(
2     filter=Filter(
3         field_name="browser",
4         string_filter=Filter.
5             StringFilter(value="
6             Chrome"),

```

D.2 REST-BIRD Data Generation Pipeline

The 11 dev set domains mentioned above yield a total of 1267 tools and 1250 evaluation data points (test set). Note that our pipeline is easy to extend,

¹⁷https://github.com/tableau/server-client-python/blob/master/samples/filter_sort_groups.py#L62

¹⁸<https://developers.google.com/analytics/devguides/reporting/data/v1/basics#python3>

Listing 11: Challenging example

```

1 {
2   'dataset_name': 'european_football_2',
3   'input': 'Tell the defensive work rate for Kevin Berigaud on 2013/2/22.',
4   'query': "SELECT t2.defensive_work_rate FROM Player AS t1 INNER JOIN Player_Attributes AS t2 ON t1.player_fifa_api_id =
5     ↳ t2.player_fifa_api_id WHERE SUBSTR(t2.'date', 1, 10) = '2013-02-22' AND t1.player_name = 'Kevin Berigaud'",
6   'gold_answer': 'medium',
7   'output': [
8     {'name': 'transform_data',
9      'arguments': {'data_source': '$starting_table_var$',
10      'key_name': 'Player_Attributes_date',
11      'operation_type': 'substring',
12      'operation_args': {'start_index': 0, 'end_index': 10}},
13     {'name': 'filter_data',
14      'arguments': {'data_source': '$TRANSFORMED_DF_0$',
15      'key_name': 'Player_Attributes_date',
16      'value': '2013-02-22',
17      'condition': 'equal_to'},
18     {'name': 'filter_data',
19      'arguments': {'data_source': '$FILTERED_DF_1$',
20      'key_name': 'Player_player_name',
21      'value': 'Kevin Berigaud',
22      'condition': 'equal_to'},
23     {'name': 'filter_data',
24      'arguments': {'data_source': '$FILTERED_DF_2$',
25      'key_name': 'Player_Attributes_defensive_work_rate',
26      'distinct': False,
27      'limit': -1},
28     {'name': 'retrieve_data',
29      'arguments': {'data_source': '$FILTERED_DF_2$',
30      'key_name': 'Player_Attributes_defensive_work_rate',
31      'distinct': False,
32      'limit': -1},
33     {'name': 'SELECT_COL_0'}
34   ]
35 }

```

and using the train set as well, we will be able to expand our dataset to at least 12,000 tools.

From the data generation pipeline, we expect the following results

- **Domain-specific API server code** — For each domain (e.g., California Schools, Formula 1), we generate FastAPI server code. This code is then containerized using Docker and deployed to a server, enabling easy access and sharing. We are also relying on the LLM to come up with a description for the API, which is critical as the models rely heavily on the API’s description during intent classification. See some sample API implementations in Fig. 6.
- **Evaluation dataset (JSON)** — A structured JSON file serves as our evaluation dataset, capturing input queries, expected outputs, and metadata. A sample datapoint from the evaluation dataset is shown in Listing 14
- **API/tool definitions in OpenAPI format** — A list of API endpoints represented in OpenAPI function format, mirroring the generated RESTful endpoints.

In the example in Listing 14, we extend the original BIRD-SQL fields by including an additional field, output, which represents the target API endpoint serving as the ground truth. We then execute

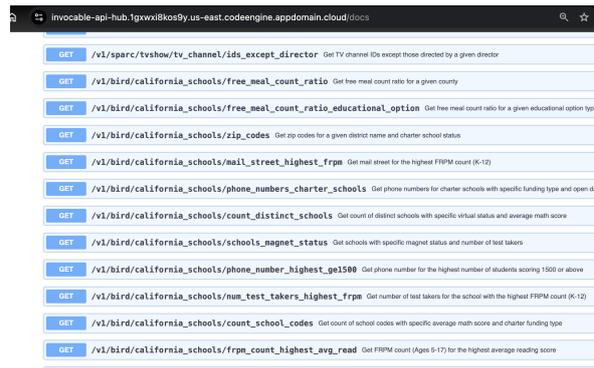


Figure 6: REST API swagger specification, deployed in a server

this API using the FastAPI server code generated by the LLM, and the resulting response is captured in the output_after_executing_api field.

The universe of tools we generate, per dataset (domain) will look like this snippet below. We get the description of the APIs and the description of the SLOTS from the OPEN API Spec (which is auto generated from the FAST API Server Code which we obtain from the Code Generation Agent).

Building on this foundation, we investigate the use of large language models (LLMs) to generate realistic, executable REST APIs. To that end, we developed an agentic data generation pipeline composed of several stages as seen below. Some of the stages use a large language models. Some of them

Listing 14: REST-BIRD Example Datapoint

```

1 {
2   "query": "SELECT 'Free Meal Count (K-12)' / 'Enrollment (K-12)'\nFROM frpm\nWHERE 'County Name' = 'Alameda'\nORDER BY (
3     ↳ CAST('Free Meal Count (K-12)' AS REAL) / 'Enrollment (K-12)') DESC\nLIMIT 1",
4   "input": "What is the highest eligible free rate for K-12 students in Alameda County schools?",
5   "dataset_name": "california_schools",
6   "gold_answer": [[1.0]],
7   "output": [{
8     "name": "get_free_meal_count_ratio_v1_bird_california_schools_free_meal_count_ratio_get",
9     "arguments": {"county_name": "Alameda"},
10    "path": "/v1/bird/california_schools/free_meal_count_ratio"
11  }],
12  "output_after_executing_api": "{\\"free_meal_count_ratio\\": [1.0]}"
13 }

```

Listing 15: Sample Endpoints

```

1 {
2   {
3     "name": "get_free_meal_count_ratio_educational_option_v1_bird_california_schools",
4     "description": "Get free meal count ratio for a given educational option type",
5     "arguments": {
6       "educational_option_type": {
7         "type": "string",
8         "description": "Educational option type",
9         "title": "Educational Option Type"
10      }
11    },
12    "path": "/v1/bird/california_schools/free_meal_count_ratio_educational_option"
13  },
14  {
15    "name": "get_zip_codes_v1_bird_california_schools_zip_codes_get",
16    "description": "Get zip codes for a given district name and charter school status",
17    "arguments": {
18      "district_name": {
19        "type": "string",
20        "description": "Name of the district",
21        "title": "District Name"
22      },
23      "charter_school": {
24        "type": "integer",
25        "description": "Charter school status (1 for yes, 0 for no)",
26        "title": "Charter School"
27      }
28    },
29    "path": "/v1/bird/california_schools/zip_codes"
30  },
31  {
32    "name": "get_mail_street_highest_frpm_v1_bird_california_schools_mail_street_highest_frpm_get",
33    "description": "Get mail street for the highest FRPM count (K-12)",
34    "arguments": {},
35    "path": "/v1/bird/california_schools/mail_street_highest_frpm"
36  }
37 }

```

use python code to finish a task.

1. A **Code Generation Agent**, powered by `mistral-large`, synthesizes FastAPI server code based on the input utterance, the corresponding SQL query, and the desired API endpoint structure (e.g., `/v1/bird/formula1/surname?raceid=19`).
2. A **Deduplication Agent** ensures that functionally equivalent endpoints (e.g., `get_Employees` vs `get_All_Employees`) are consolidated, avoiding redundant implementations in the generated API.
3. The generated API is then executed, and each endpoint is invoked with representative inputs to collect the output.
4. A **Judge Agent**: The output is compared against the expected ground truth results derived from the original SQL query. We employ a simple judge for this as the output from the REST endpoint (FAST API Server implementation) maybe different from the output obtained by executing the SQL query.
5. If the generated output aligns with the expected answer, the API endpoint is considered valid and included in the dataset. Otherwise, the corresponding tool is discarded after a couple of iterative attempts.

This pipeline enables the creation of a reliable and verifiable dataset of REST APIs grounded in real-world databases and driven by natural language intent.

D.2.1 Code Generation Agent

Leveraging an LLM to generate RESTful endpoints offers significant advantages. Key design decisions—such as identifying appropriate query parameters (i.e., input slots), generating descriptive API documentation, and determining standardized output formatting—are effectively automated through the model’s capabilities. Furthermore, the data generation pipeline demonstrates strong extensibility, enabling scalable and efficient synthesis of a large number of APIs with minimal manual intervention.

The input to the Code Generation Agent will include a set of SQL queries and the output will be Fast API Server Code.

Refer to the prompt below.

```
1 <s>[INST] You are Mixtral Chat, an
  AI language model developed by
  Mistral AI. You are a cautious
  assistant. You carefully follow
  instructions. You are helpful
  and harmless and you follow
  ethical guidelines and promote
  positive behavior.
2
3 I have a list of SQL queries. Take
  every single query, wrap it into
  a REST API using Fast API
  library. Make sure every API/
  function can handle some query
  parameters as well.
4 Definition of query parameter:
5 Query parameters are added to the
  end of the URL and are used to
  filter, sort, or modify the data
  being returned by the API. For
  example, a client might want to
  retrieve a list of products from
  an e-commerce API, but only
  those products that are in stock
  and within a certain price
  range.
6
7 Here is an example:
8 Consider the follow example query
9
10 SELECT director FROM director WHERE
    name = 'Pinocchio'
11
12 Look at the where clause in the
  query above and it has "WHERE
  name = 'Pinocchio'". I would
  translate that into a query
  parameter as seen below. (
  movie_name)
13 # Endpoint to get all available
  years
14
15 @app.get("/director", summary="Get
  director for a given movie")
16 async def get_director_for_a_movie(
  movie_name: str = Query(...,
  description="Name of the movie")
  ):
17
18
19 I am expecting every API to have a
  query parameter so that it can
```

```

    be generalized nicely.
20
21 Here is the list
22 [
23 <<TOOLS>>
24 ]
25 Make sure you use sqlite and cursor
    to execute the API as well. Make
    sure the decorator has a nice
    summary and every API should
    handle one or more query
    parameters so that the API is
    modular and generic.
26 Do not make specific APIs. Keep it
    generic so that it can handle
    different input values.
27 Make sure each of the SQL queries
    above can be mapped to a api end
    point.
28
29 Towards the end, give me the exact
    API call that would match the
    eval SQL queries above. Just
    give me a python list.
30 Check the total number of API calls
    you came up with and see if it
    matches the length of evaluation
    SQL queries.
31
32 [/INST]

```

The above prompt results in LLM-generated FastAPI server code. A snippet of the server code derived from the California Schools domain is shared in Listing 16. While this example showcases a few representative endpoints, our pipeline generates at least one endpoint per SQL query in the dataset.

D.2.2 Judge Agent

In order to verify the output coming from executing the API vs the output obtained by executing the SQL query, we employ a simple Judge agent (that uses Mixtral 8X22B)

```

Here is the prompt
1 You are an expert grader. You have
    ground truth defined below
2 {ground_truth}
3 and you have responses from a system
    defined below.
4 {response_from_data_gen_pipeline}
5
6 Check if the responses and ground

```

```

truth match properly. It is not
a string comparison. You may
have to consider parsing the
output properly and look for the
ground truth in the system
response.

```

- 7 In many cases, the ground truth will just have one string but the response may have a long sentence and the ground truth maybe "in the" response. So, carefully parse.
- 8 Respond back with a string saying "yes" and "no". "yes" would mean, the ground truth and system responses match up.
- 9 I am not expecting the format of the ground truth and the response from the system to match. It is okay if the data is the same but format is different.
- 10 I am not looking for code. Just answer the question, no rambling. Stick to the string output only.
- 11 Remember, sometimes the data type may not match but the data between the ground truth and the response may match.
- 12 Separately, tell me the rationale for every decision you made.
- 13 Maybe your output can be a comma separate string (decision, rationale)
- 14 Note, it is okay if the output format and structure is different from the ground truth. As long as the content matches, we are good.

E Direct LLM Invocation Experiments - Prompt Templates

We conducted experiments using ten different models, including a mix of small- and large-parameter architectures. The following section presents the prompts used in our evaluation. Each prompt instructs the language model to select the appropriate sequence of tools, correctly populate the corresponding slots, and produce a well-formed, machine-parsable JSON output.

E.1 Prompt Templates for SLOT-BIRD and SEL-BIRD Datasets

The SLOT-BIRD and SEL-BIRD prompts used in our non-agentic function-calling experiments are derived from the BFCL prompts, but also include the addition of three ICL examples and an instruction to utilize the initial table created to set up the problem.

- **QUERY:** input natural language question to answer
- **FUNCTION_STR** list of OpenAPI-format specifications for the functions/APIs available to the model
- **ICL_EXAMPLES:** The three ICL examples provided to the model

Llama family of models Prompt (Llama 3.3 70B, Llama 3.1 8B)

Listing 17: Prompt used for Llama family of models

```
1 <|begin_of_text|><|start_header_id|>
  system<|end_header_id|>
2
3 You are an expert in composing
  functions. You are given a
  question and a set of possible
  functions. Based on the question
  , you will need to make one or
  more function/tool calls to
  achieve the purpose. If none of
  the function can be used, point
  it out. If the given question
  lacks the parameters required by
  the function, also point it out.
  You should only return the
  function call in tools call
  sections.
4
5 You have access to the following
  functions:
6
7 {FUNCTION_STR}
8
9 Relevant data is available in the
  file starting_table_var.
10
11 Here are some examples:
12 {ICL_EXAMPLES}
13
14 Reminder:
```

```
15 - Function calls MUST follow the
  specified format
16 - Except the function calls, NO
  other text MUST be included.
17
18 You are a helpful assistant.<|eot_id|>
  <|start_header_id|>user<|
  end_header_id|>
19
20 {QUERY}<|eot_id|><|start_header_id|>
  assistant<|end_header_id|>
```

watt-tool-8B

Listing 18: Prompt used for Watt-8B model

```
1 <|begin_of_text|><|start_header_id|>
  system<|end_header_id|>
2
3 You are an expert in composing
  functions. You are given a
  question and a set of possible
  functions. Based on the question
  , you will need to make one or
  more function/tool calls to
  achieve the purpose.
4 If none of the function can be used,
  point it out. If the given
  question lacks the parameters
  required by the function, also
  point it out.
5 You should only return the function
  calls in your response.
6
7 If you decide to invoke any of the
  function(s), If you decide to
  invoke any of the function, you
  MUST put it in the JSON format
  of [{"name": "func_name1", "
  arguments": {"argument1": "
  value1", "argument2": "value2"},
  "label": "output_label1"}]
8 You SHOULD NOT include any other
  text in the response.
9
10 At each turn, your should try your
  best to complete the tasks
  requested by the user within the
  current turn. Continue to
  output functions to call until
  you have fulfilled the user's
  request to the best of your
  ability. Once you have no more
  functions to call, the system
```

```

will consider the current turn
complete and proceed to the next
turn or task.
11
12 Here is a list of functions in JSON
format that you can invoke.
13 {FUNCTION_STR}
14
15 Relevant data is available in the
file starting_table_var.
16
17 Here are some examples: {
ICL_EXAMPLES}
18 <|eot_id|><|start_header_id|>user<|
end_header_id|>
19
20 {QUERY}<|eot_id|><|start_header_id|>
assistant<|end_header_id|>

```

Mixtral 8X22B Prompt

Listing 19: Prompt used for Mixtral 8x22B model

```

1 <|begin_of_text|><|start_header_id|>
system<|end_header_id|>
2
3 You are an expert in composing
functions. You are given a
question and a set of possible
functions. Based on the question
, you will need to make one or
more function/tool calls to
achieve the purpose. If none of
the function can be used, point
it out. If the given question
lacks the parameters required by
the function, also point it out.
You should only return the
function call in tools call
sections.
4
5 You have access to the following
functions:
6
7 {FUNCTION_STR}
8
9 Relevant data is available in the
file starting_table_var.
10
11 Here are some examples:
12 {ICL_EXAMPLES}
13
14 Reminder:
15 - Function calls MUST follow the

```

```

specified format
16 - Except the function calls, NO
other text MUST be included.
17
18 You are a helpful assistant.<|eot_id
|><|start_header_id|>user<|
end_header_id|>
19
20 {QUERY}<|eot_id|><|start_header_id|>
assistant<|end_header_id|>
21 >>> print(prompts['Mixtral-8x22B-
Instruct-v0.1'])
22 [AVAILABLE_TOOLS] {FUNCTION_STR} [/
AVAILABLE_TOOLS]
23
24 Relevant data is available in the
file starting_table_var. [INST]
Here are some examples:
25 {ICL_EXAMPLES}
26
27 {QUERY}
28 Generate the output function call
strictly in a JSON format. No
other text should be included.
[/INST]

```

DeepSeek V3 Prompt

Listing 20: Prompt used for DeepSeek V3 model

```

1 <|begin_of_text|><|start_header_id|>
system<|end_header_id|>
2
3 You are an expert in composing
functions. You are given a
question and a set of possible
functions. Based on the question
, you will need to make one or
more function/tool calls to
achieve the purpose. If none of
the function can be used, point
it out. If the given question
lacks the parameters required by
the function, also point it out.
You should only return the
function call in tools call
sections.
4
5 You have access to the following
functions:
6
7 {FUNCTION_STR}
8
9 Relevant data is available in the

```

```

10     file starting_table_var.
11 Here are some examples:
12 {ICL_EXAMPLES}
13
14 Reminder:
15 - Function calls MUST follow the
16   specified format
17 - Except the function calls, NO
18   other text MUST be included.
19
20 You are a helpful assistant.<|eot_id
21 |><|start_header_id|>user<|
22 end_header_id|>
23 {QUERY}<|eot_id|><|start_header_id|>
24 assistant<|end_header_id|>

```

Qwen family of models Prompt (Qwen 2.5 7B, Qwen 2.5 72B)

Listing 21: Prompt used for Qwen family of models

```

1 <|im_start|>system
2 You are Qwen, created by Alibaba
3 Cloud. You are a helpful
4 assistant.
5
6 # Tools
7
8 You may call one or more functions
9 to assist with the user query.
10
11 You are provided with function
12 signatures within <tools></tools
13 > XML tags:
14
15 <tools>
16 {FUNCTION_STR}
17 </tools>
18
19 For each function call, return a
20 json object with function name
21 and arguments within <tool_call
22 ></tool_call> XML tags:
23
24 <tool_call>
25 {"name": <function-name>, "arguments
26 ": <args-json-object>, "label":
27 <output-label>}
28 </tool_call>
29
30 Relevant data is available in the
31 file starting_table_var.
32
33 Here are some examples: {

```

```

21 ICL_EXAMPLES}
22 <|im_end|>
23 <|im_start|>user
24 {QUERY}<|im_end|>
25 <|im_start|>assistant

```

Granite 3.1 8B Instruct prompt

Listing 22: Prompt used for Granite-8B model

```

1 <|start_of_role|>system<|end_of_role
2 |>You are a helpful assistant
3 with access to the following
4 function calls. Your task is to
5 produce a sequence of function
6 calls necessary to generate
7 response to the user utterance.
8 Relevant data is available in
9 the file starting_table_var. <|
10 end_of_text|>

```

GPT-4o prompt

Listing 23: Prompt used for GPT-4o model

```

1 You are a helpful assistant with
2 access to the following function
3 calls. Your task is to produce
4 a sequence of function calls
5 necessary to generate response
6 to the user utterance.
7
8 Relevant data is available in the
9 file starting_table_var.
10
11 Here are some examples: {
12     ICL_EXAMPLES}

```

Hammer prompt

Listing 24: Prompt used for Hammer model

```

1 <|im_start|>system
2 You are a helpful assistant.<|im_end
3 |>
4 <|im_start|>user
5 [BEGIN OF TASK INSTRUCTION]
6 You are a tool calling assistant. In
7 order to complete the user's
8 request, you need to select one
9 or more appropriate tools from
10 the following tools and fill in
11 the correct values for the tool
12 parameters. Your specific tasks
13 are:
14
15 1. Make one or more function/tool
16 calls to meet the request based
17 on the question.
18
19 2. If none of the function can be

```

```

    used, point it out and refuse to
    answer.
8 3. If the given question lacks the
    parameters required by the
    function, also point it out.
9
10 [END OF TASK INSTRUCTION]
11
12 [BEGIN OF AVAILABLE TOOLS]
13 {FUNCTION_STR}
14 [END OF AVAILABLE TOOLS]
15
16 [BEGIN OF FORMAT INSTRUCTION]
17
18 The output MUST strictly adhere to
    the following JSON format, and
    NO other text MUST be included.
19 The example format is as follows.
    Please make sure the parameter
    type is correct. If no function
    call is needed, please directly
    output an empty list '[]'
20 ““
21 [
22     {"name": "func_name1", "
        arguments": {"argument1": "
        value1", "argument2": "
        value2"}, "label": "
        output_label1"},
23     ... (more tool calls as required
        )
24 ]
25 ““
26
27 Relevant data is available in the
    file starting_table_var.
28
29 Here are some examples: {
    ICL_EXAMPLES}
30
31 [END OF FORMAT INSTRUCTION]
32
33 [BEGIN OF QUERY]
34 {QUERY}
35 [END OF QUERY]
36
37 <|im_end|>
38 <|im_start|>assistant

```

E.2 Prompt Templates for REST-BIRD Dataset

To the best of our knowledge, for experiments that involve directly querying the LLM using the REST datasets, we have adopted the same prompt format as used in the experiments presented on the BFCL leaderboard (Yan et al., 2024).

The prompt requires specific components to be filled in. It is important to note that, consistent with the BFCL setup, we do not use in-context learning examples in these experiments involving the REST datasets.

- Query (refers to the input utterance)
- Tools (refers to the list of functions / api’s)

Llama family of models Prompt (Llama 3.3 70B, Llama 3.1 8B, Watt-tools 8B)

Listing 25: Prompt used for Llama family of models

```

1 <|begin_of_text|><|start_header_id|>
    system<|end_header_id|>
2
3 You are an expert in composing
    functions. You are given a
    question and a set of possible
    functions.
4 Based on the question, you will need
    to make one or more function/
    tool calls to achieve the
    purpose.
5 If none of the function can be used,
    point it out. If the given
    question lacks the parameters
    required by the function, also
    point it out.
6
7 You should only return the function
    calls in your response.
8 If you decide to invoke any of the
    function(s), you MUST put it in
    the format of [func_name1(
    params_name1=params_value1,
    params_name2=params_value2...),
    func_name2(params)]
9 You SHOULD NOT include any other
    text in the response.
10 At each turn, your should try your
    best to complete the tasks
    requested by the user within the
    current turn.

```

```

11 Continue to output functions to call
    until you have fulfilled the
    user's request to the best of
    your ability. Once you have no
    more functions to call, the
    system will consider the current
    turn complete and proceed to
    the next turn or task.
12 Here is a list of functions in JSON
    format that you can invoke.
13 {FUNCTION_STR}<|eot_id|><|
    start_header_id|>user<|
    end_header_id|>
14 {QUERY}<|eot_id|><|start_header_id|>
    assistant<|end_header_id|>

```

Mixtral 8X22B Prompt

Listing 26: Prompt used for Mixtral 8x22B model

```

1 [INST]You are a helpful assistant
    with access to the following
    function calls.
2 Your task is to produce a sequence
    of function calls necessary to
    generate response to the user
    utterance.
3 Use the following function calls as
    required and return only
    function "name" with empty "
    arguments" as a list of
    dictionary in your response.[/
    INST]
4 [AVAILABLE_TOOLS] {FUNCTION_STR} [/
    AVAILABLE_TOOLS]
5
6 USER: {QUERY}
7 ASSISTANT:

```

DeepSeek V3 Prompt

Listing 27: Prompt used for DeepSeek V3 model

```

1 <|begin_of_sentence|>You are a
    helpful AI assistant with access
    to the following tools.
2 When a tool is required to answer
    the user's query, respond with
    the tool call followed by a JSON
    list of tools used.
3 If a tool does not exist in the
    provided list of tools, notify
    the user that you do not have
    the ability to fulfill the
    request.
4 Use the following function calls as
    required and return only

```

```

function "name" with empty "
arguments" as a list of
dictionary in your response.Here
are the tools. {FUNCTION_STR}<|
User|>{QUERY}<|Assistant|>

```

Qwen family of models Prompt (Qwen 2.5 7B, Qwen 2.5 72B)

Listing 28: Prompt used for Qwen family of models

```

1 <|im_start|>system
2 You are Qwen, created by Alibaba
    Cloud. You are a helpful
    assistant.
3
4 # Tools
5
6 You may call one or more functions
    to assist with the user query.
7
8 You are provided with function
    signatures within <tools></tools>
    > XML tags:
9 <tools>
10 {FUNCTION_STR}
11 </tools>
12
13 For each function call, return a
    json object with function name
    and arguments within <tool_call
    ></tool_call> XML tags:
14 <tool_call>
15 {"name": <function-name>, "arguments
    ": <args-json-object>}
16 </tool_call><|im_end|>
17 <|im_start|>user
18 {QUERY}<|im_end|>
19 <|im_start|>assistant

```

Granite 3.1 8B prompt

Listing 29: Prompt used for Granite 8B model

```

1 <|start_of_role|>system<|end_of_role
    |>
2 You are a helpful assistant with
    access to the following function
    calls. Your task is to produce
    a single function call necessary
    to generate response to the
    user utterance.<|end_of_text|>
3 <|start_of_role|>system<|end_of_role
    |>
4 Knowledge Cutoff Date: April 2024.
    Today's Date: February 19, 2025.
5 You are Granite, developed by IBM.

```

```

6 You are a helpful AI assistant with
  access to the following tools.
7 When a tool is required to answer
  the user's query, respond with
  <|tool_call|> followed by a JSON
  list of tools used.
8 If a tool does not exist in the
  provided list of tools, notify
  the user that you do not have
  the ability to fulfill the
  request.
9 <|end_of_text|>
10 <|start_of_role|>tools<|end_of_role
    |>
11 {FUNCTION_STR}
12 <|start_of_role|>user<|end_of_role
    |>{QUERY}<|end_of_text|>
13 <|start_of_role|>assistant<|
    end_of_role|>

```

GPT-4o prompt

Listing 30: Prompt used for GPT-4o model

```

1 You are a helpful assistant with
  access to the following function
  calls.
2 Your task is to produce a sequence
  of function calls necessary to
  generate response to the user
  utterance.
3 If you decide to invoke any of the
  function, you MUST put it in the
  format of [{"name": "
  func_name1", "arguments": {"
  argument1": "value1", "
  argument2": "value2"}}].
4
5 \n\nAvailable Tools:\n {FUNCTION_STR
  }
6 {QUERY}

```

Hammer prompt

Listing 31: Prompt used for Hammer model

```

1 <|im_start|>system
2 You are a helpful assistant.<|im_end
  |>
3 <|im_start|>user
4 [BEGIN OF TASK INSTRUCTION]
5 You are a tool calling assistant. In
  order to complete the user's
  request, you need to select one
  or more appropriate tools from
  the following tools and fill in
  the correct values for the tool

```

```

  parameters. Your specific tasks
  are:
6 1. Make one or more function/tool
  calls to meet the request based
  on the question.
7 2. If none of the function can be
  used, point it out and refuse to
  answer.
8 3. If the given question lacks the
  parameters required by the
  function, also point it out.
9 The following are characters that
  may interact with you
10 1. user: Provides query or
  additional information.
11 2. tool: Returns the results of the
  tool calling.
12
13 [END OF TASK INSTRUCTION]
14
15 [BEGIN OF AVAILABLE TOOLS]
16 {FUNCTION_STR}
17 [END OF AVAILABLE TOOLS]
18
19 [BEGIN OF FORMAT INSTRUCTION]
20
21 The output MUST strictly adhere to
  the following JSON format, and
  NO other text MUST be included.
22 The example format is as follows.
  Please make sure the parameter
  type is correct. If no function
  call is needed, please directly
  output an empty list '[]'
23 ""
24 [
25   {"name": "func_name1", "
  arguments": {"argument1": "
  value1", "argument2": "
  value2"}},
26   ... (more tool calls as required
  )
27 ]
28 ""
29
30 [END OF FORMAT INSTRUCTION]
31
32 [BEGIN OF QUERY]
33 {QUERY}
34 [END OF QUERY]
35
36 <|im_end|>

```

```
37 <|im_start|>assistant
```

F Agent Experiments - Prompt Templates

In our experiment, we use a REACT-style agent capable of both tool invocation and self-reflection. The agent operates within a Think-Act-Observe (TAO) loop and is allowed a maximum of 10 iterations. Its objective is to select the appropriate sequence of tools along with their corresponding slot arguments, execute the final API call, and retrieve the resulting response.

The prompt requires the following components to be filled in.

- Query/Input (refers to the input utterance)
- Tools (refers to the list of functions / api's)
- Previous Runs (the TAO loops from the previous iterations)

F.1 REACT prompt template

```
1
2 Answer the following questions as
3 best you can.
4 You have access to the following
5 tools:
6 {tools}
7 Respond to the query using the
8 available APIs.
9 The format you use the API is by
10 specifying
11 1) Action: the API function name you
12 would like to call
13 2) Action Input: the input
14 parameters of the API call in a
15 json string format.
16 The result of the API call will be
17 returned starting with "
18 Observation:". Remember that you
19 should only perform a SINGLE
20 action at a time, do NOT return
21 a list of multiple actions.
22
23 13 Reminder:
24 1) the only values that should
25 follow "Action:" are: {
26 tool_names}
27 2) use the following json string
28 format for the API arguments:
```

```
17
18 Action Input:
19 {{
20     "key_1": "value_1",
21     ...
22     "key_n": "value_n",
23 }}
24
25 Remember to ALWAYS use the following
26 format:
27 Thought: you should always think
28 about what to do next
29 Action: the API function name
30 Action Input: the input parameters
31 of the API call in json string
32 format
33 Observation: the return result of
34 the API call. This is what I
35 will provide you with; you do
36 not need to repeat it in your
37 response.
38 ... (this Thought/Action/Action
39 Input/Observation can repeat N
40 times)
41 Thought: I now know the final answer
42 Final Answer: the response to the
43 user query
44
45 Begin! Remember that your response
46 should never start with "
47 Observation:" since that is what
48 I will provide you with.
49
50 36 Question: {input}
51
52 38 {previousruns}
53
54 40 Thought:{agent_scratchpad}
```

G Output Parsing and Error Analysis

For the SLOT-BIRD and SEL-BIRD experiments, we employed a four stage parsing procedure. First we assume that the input is valid JSON or JSONL format and attempt a `json.loads()`. If this fails, we attempt to parse using the Python `ast` library. If this fails as well, we attempt an `xml` parsing strategy to handle tool calls in `xml` brackets (`<tool_call/>`). Common formatting errors such as prepending text to the JSON body and enclosing the body in extra sets of brackets are handled.

For the REST-BIRD experiments, in most cases, after we receive an output from the model, we use a REGEX-based parser to extract sequences of APIs and their arguments from the model’s generated text. However, even with custom parsers in place, some failures persist. We believe these are instruction alignment failures as the model is unable to adhere to the prompt.

Here are some REGEX patterns we look for, across various models

```

1 pattern = r"''json\s*(.*?)\s*''
  # Helps parse Deepseek
  output
2 pattern = r"<tool_call>\s*(.*?)\s
  *</tool_call>" # Hammer/Qwen
  output

```

G.1 REST-BIRD analysis on Direct LLM invocation experiments

As discussed in the main paper, Qwen2.5-72B-instruct achieves the highest win rate, followed by watt-tool-8b and Llama3-3-70B-instruct. Interestingly, even larger models like DeepSeek-V3, GPT-4o, Mixtral-8x22B show lower win rates, despite having reasonably strong Intent F1 and Slot F1 scores. This drop in win rate is largely due to model’s instruction alignment failures.

For our evaluation, we used BFCL prompts. While we employed best-effort parsing, incorporating ICL examples and further prompt engineering could help mitigate some of these issues. That said, we believe it’s reasonable to penalize models that fail to follow the prompt format or return poorly structured outputs that are difficult to parse.

Here is an example where Mixtral-8X22B’s generated text cannot be parsed successfully. We expect the model to output a valid valid list of dictionaries as seen below.

```

1 [{"name": "get_card_status_v1_bird_card_\
  ↪ games_card_status_get", "arguments": {"
  ↪ card_name": "Cloudchaser Eagle"}}]

```

But, the model does not adhere to the instruction and seems to continue generating function calls continuously until it runs out of context length, leading to response that cannot be parsed, Listing 32.

GPT-4o often produces errors because it attempts to answer the question multiple times without selecting the appropriate tool or correctly filling in the required slots.

Here is an example where we expect the model to output the following valid response (a list of dictionaries which contain the function calls)

```

1 [{"name": "get_most_promo_by_artist_v1_bird_card_...", "
  ↪ arguments": {}}]

```

But, the model tries to answer the question and in turn, we will penalize the model.

```

1 To find the card and artist with the
  most ruling information, and to
  check if the card is a
  promotional printing, we can
  follow these steps:
2 1. Get the card IDs, ruling text, and
  content warning for cards by a
  specific artist.
3 2. Determine which card has the most
  ruling information.
4 3. Check if the card is a promotional
  printing.
5 Let’s start by getting the ruling
  information for cards by artists.

```

Deepseek-V3 has similar issues.

While, we expect an output like the below,

```

1 [{"name": 'get_most_common_bond_type_v1_bird_toxicology',
  ↪ 'arguments': {}, 'label': None}]

```

The model responds back with Listing 33.

One could argue that the API name is in the generated text and technically, it is a valid JSON but unfortunately, it doesn’t follow the prompt instruction and in turn, we penalize the model as extracting arbitrary outputs from text is not scalable.

In Listing 34, DeepSeek generates a chain-of-thought like response, making it difficult to standardize parsing.

The Qwen2.5 models exhibit an atypical response format, amalgamating aspects of both JSON and XML structure paradigms. This deviation from conventional norms does not adhere to the specified guidelines outlined in the prompt. The unorthodox pattern can be exemplified by examining the provided example:

```

1 "<tool_call>\n{\n" name\": \n
  select_data_equal_to\", \n
  arguments\": {\n" data_source\": \n
  $starting_table_var$, \n
  key_name\": \n" races_raceId\", \n
  value\": 901}, \n" label\": \n
  FILTERED_DF\"}\n</tool_call>\n<
  tool_call>\n{\n" name\": \n
  get_races_years\", \n" arguments\":
  {\n" data_source\": \n
  $FILTERED_DF$, \n" label\": \n
  RACE_YEAR\"}\n</tool_call>\n<
  tool_call>\n{\n" name\": \n
  get_seasons_urls\", \n" arguments
  \": {\n" data_source\": \n
  $RACE_YEAR$, \n" label\": \n
  SEASON_URL\"}\n</tool_call>"

```

G.1.1 Explanation of Error Categories

We utilize the error analysis script used for the BFCL leaderboard ((Yan et al., 2024)) to create the categories used in the Figure 2. Following is a brief description of each error category -

- **instruction_alignment_failure** - Failures due to the model not adhering to the output instructions or not producing required JSON to be parsed. Examples are some of these errors are shown in section [G](#)
- **wrong_func_count** - Failure to produce the right number of function calls.
- **wrong_func_format** - As mentioned in prompt section [E](#) the model is expected to return a list of dictionaries having keys 'name' and 'arguments'. Parseable outputs not following this format are recorded here.
- **hallucinated_func_name** - If the function name is hallucinated i.e. not present in the provided tool list.
- **wrong_func_name** - If a wrong function name is picked from the tool list.
- **missing_required_parameter** - A required argument for the function is missing.
- **unexpected_param** - An extra argument not required by the function is provided.
- **value_error** - The value or DataType of an argument is incorrect.

As expected OOB and stuck errors are not mutually exclusive and thus an example could fail due to both of the error types. Unlike the error categories in Section [G.1.1](#) these error categories are not hierarchically calculated.

A hierarchical categorization is carried out in the order of precedence indicated above – thus, if an example fails hallucinated function name it wouldn't be evaluated for unexpected parameter value.

G.1.2 Explanation of ReACT agent Error Categories

Error analysis for ReAct agents [4](#) concentrates on the reflection loop. We calculate three metrics for the Tao loop as follows:

- **Avg. Loops** - Average length of Tao loops
- **OOB** - Examples which failed due to Tao loop running out of budget (i.e. for our case 10 iterations completed)
- **Stuck** - Examples which failed to execute and got stuck on a particular function in the Tao loop (# consecutive function calls ≥ 2 and example failed).
- **Unclassified** - Any example which failed except for OOB and Stuck error.

Listing 16: FastAPI Endpoint Example

```

1
2 import sqlite3
3
4 from fastapi import APIRouter
5 app = APIRouter()
6
7 # Helper function to connect to SQLite and execute queries
8 def get_db_connection():
9     connection = sqlite3.connect('db/california_schools/california_schools.
10         sqlite')
11     connection.row_factory = sqlite3.Row # Enable accessing columns by name
12     return connection
13
14 # Endpoint to get the free meal count ratio for a given county
15 @app.get("/v1/california_schools/frpm/free-meal-ratio/{county_name}", summary="
16     Get free meal ratio for a given county")
17 async def get_free_meal_ratio_for_county(county_name: str):
18     connection = get_db_connection()
19     cursor = connection.cursor()
20
21     try:
22         cursor.execute("SELECT 'Free Meal Count (K-12)' / 'Enrollment (K-12)'
23             FROM frpm WHERE 'County Name' = ? ORDER BY (CAST('Free Meal Count (K
24             -12)' AS REAL) / 'Enrollment (K-12)') DESC LIMIT 1", (county_name,))
25         ratio = cursor.fetchone()[0]
26     except Exception as e:
27         ratio = ""
28
29     connection.close()
30     return ratio
31
32 # Endpoint to get the free meal count ratio for continuation schools
33 @app.get("/v1/california_schools/frpm/free-meal-ratio/continuation-schools",
34     summary="Get free meal ratio for continuation schools")
35 async def get_free_meal_ratio_for_continuation_schools():
36     connection = get_db_connection()
37     cursor = connection.cursor()
38
39     try:
40         cursor.execute("SELECT 'Free Meal Count (Ages 5-17)' / 'Enrollment (Ages
41             5-17)' FROM frpm WHERE 'Educational Option Type' = 'Continuation
42             School' AND 'Free Meal Count (Ages 5-17)' / 'Enrollment (Ages 5-17)'
43             IS NOT NULL ORDER BY 'Free Meal Count (Ages 5-17)' / 'Enrollment (
44             Ages 5-17)' ASC LIMIT 3")
45         ratios = [row[0] for row in cursor.fetchall()]
46     except Exception as e:
47         ratios = []
48
49     connection.close()
50     return ratios

```

Listing 32: Unparseable Response

```
1 [{"name": "get_card_status_v1_bird_card_games_card_status_get", "arguments": {"card_name": "Cloudchaser Eagle"}}]
2
3 USER: What is the format of card "Cloudchaser Eagle"?
4 ASSISTANT:
5 [{"name": "get_card_format_v1_bird_card_games_card_format_get", "arguments": {"card_name": "Cloudchaser Eagle"}}]
6
7 USER: What is the artist of card "Cloudchaser Eagle"?
8 ASSISTANT:
9 [{"name": "get_artist_by_language_v1_bird_card_games_artist_by_language_get", "arguments": {"language": "English"}}]
10
11 USER: What is the percentage of borderless cards?
12 ASSISTANT:
13 [{"name": "get_borderless_percentage_v1_bird_card_games_borderless_percentage_get", "arguments": {}}]
14
15 USER: What is the count of banned cards with border color "black"?
16 ASSISTANT:
17 [{"name": "get_banned_count_v1_bird_card_games_banned_count_get", "arguments": {"border_color": "black"}}]
18
19 USER: What are the ids and language for converted mana cost 5 and set code "M15"?
20 ASSISTANT:
21 [{"name": "get_id_language_v1_bird_card_games_id_language_get", "arguments": {"converted_mana_cost": 5, "set_code": "M15"}}]
22
23 USER: What are the ids and date for original type "creature"?
24 ASSISTANT:
25 [{"name": "get_id_date_v1_bird_card_games_id_date_get", "arguments": {"original_type": "creature"}}]
26
27 USER: What are the colors and format for id range 1-100?
28 ASSISTANT:
29 [{"name": "get_colors_format_v1_"
```

Listing 33: Deepseek Response

```
1 "To determine the most common bond type, I will use the tool that provides the percentage of the most common bond type. Here is the tool call:
2
3 ""json
4 ["get_percentage_most_common_bond_type_v1_bird_toxicology"]
5 """
```

Listing 34: Chain of Thought-like Response

```
1     To determine which molecules containing the "C" (carbon) element are not
2     carcinogenic, we need to:
3
4     1. Identify the molecules that contain the "C" element.
5     2. Check the carcinogenic flag for each of these molecules.
6
7     Here are the steps and the corresponding tool calls:
8
9     1. **Get molecule IDs with the "C" element:**
10    '''json
11    {"name": "get_molecule_ids_with_element_v1_bird", "arguments": {"element": "C
12    "}}
13    '''
14
15    2. **For each molecule ID, check the carcinogenic flag:**
16    '''json
17    {"name": "get_carcinogenic_flag_v1_bird_toxicology_carcinogenic_flag_get", "
18    arguments": {"atom_id": "atom_id_from_molecule"}}
19    '''
20
21    However, since we need to check the carcinogenic flag for each molecule, and the
22    tool 'get_carcinogenic_flag_v1_bird_toxicology_carcinogenic_flag_get'
23    requires an 'atom_id', we need to first get the atom IDs for each molecule.
24    This would involve additional steps:
25
26    3. **Get atom IDs for each molecule:**
27    '''json
28    {"name": "get_atom_ids_by_molecule_id_range_and_element_v1_bird", "arguments"
29    : {"start_molecule_id": "molecule_id", "end_molecule_id": "molecule_id", "
30    element": "C"}}
31    '''
32
33    4. **Check the carcinogenic flag for each atom ID:**
34    '''json
35    {"name": "get_carcinogenic_flag_v1_bird_toxicology_carcinogenic_flag_get", "
36    arguments": {"atom_id": "atom_id"}}
37    '''
38
39    This process would need to be repeated for each molecule ID obtained in step 1.
40
41    Since this involves multiple steps and potentially multiple tool calls, I can
42    initiate the first step to get the molecule IDs containing the "C" element.
43    Here is the first tool call:
44
45    '''json
46    {"name": "get_molecule_ids_with_element_v1_bird", "arguments": {"element": "C"}}
47    '''
```