

Interpreting User Requests in the Context of Natural Language Standing Instructions

Nikita Moghe^{γ*} Patrick Xia^Φ Jacob Andreas^Φ
Jason Eisner^Φ Benjamin Van Durme^Φ Harsh Jhamtani^Φ

^γSchool of Informatics, University of Edinburgh

^ΦMicrosoft Semantic Machines

nikita.moghe@ed.ac.uk hjhamtani@microsoft.com

Abstract

Users of natural language interfaces, generally powered by Large Language Models (LLMs), often must repeat their preferences each time they make a similar request. To alleviate this, we propose including some of a user’s preferences and instructions in natural language – collectively termed *standing instructions* – as additional context for such interfaces. For example, when a user states *I’m hungry*, their previously expressed preference for Persian food will be automatically added to the LLM prompt, so as to influence the search for relevant restaurants. We develop NLSI, a language-to-program dataset consisting of over 2.4K dialogues spanning 17 domains, where each dialogue is paired with a user profile (a set of user-specific standing instructions) and corresponding structured representations (API calls). A key challenge in NLSI is to identify which subset of the standing instructions is applicable to a given dialogue. NLSI contains diverse phenomena, from simple preferences to interdependent instructions such as triggering a hotel search whenever the user is booking tickets to an event. We conduct experiments on NLSI using prompting with large language models and various retrieval approaches, achieving a maximum of 44.7% exact match on API prediction. Our results demonstrate the challenges in identifying the relevant standing instructions and their interpretation into API calls.

1 Introduction

Large Language models (LLMs) such as GPT-3 (Brown et al., 2020), GPT-4 (OpenAI, 2023), and Llama-2 (Touvron et al., 2023) are increasingly being used with tools and APIs (Schick et al., 2023; Qin et al., 2023) to unlock additional functionalities for users. For example, ChatGPT allows several external plugins such as OpenTable for searching and reserving restaurants, booking

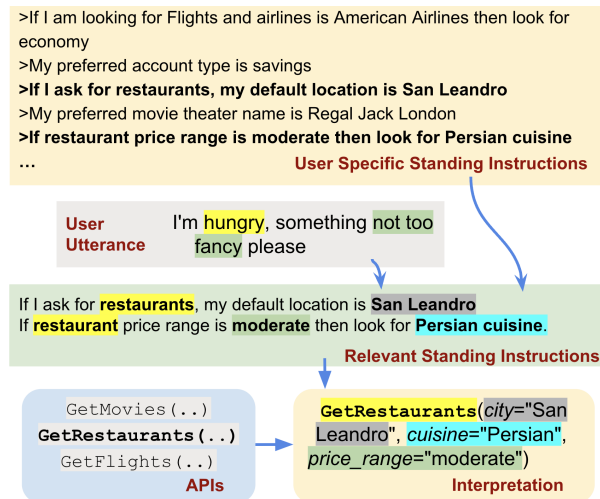


Figure 1: Parsing an utterance into a structured output, in the presence of a *user-specific* set of *standing instructions*. A model for the task needs to identify (explicitly or implicitly) the subset of instructions applicable to the utterance and interpret the utterance into API calls.

travel through Expedia or solving math problems with Wolfram.¹

As the same interface provides multiple services, these applications must learn to identify which service the user is seeking while maintaining preferences across diverse domains that are unique to each user. Understanding such preferences can aid in personalising the user experience by providing tailored responses, increased accuracy in recommendations and saving user time.

However, in most cases, users have to verbalise their preferences in detail during the interaction, including for repeated requests. For example, a user trying to find a restaurant might have to interact for multiple turns with an LLM-powered dialogue system to arrive at their preferred restaurant cuisine and location. Such interactions can become tedious, leading to a poor user experience.

Past work has explored learning preferences

*Work done while interning at Microsoft

¹<https://openai.com/blog/chatgpt-plugins>

from user-system interactions over time (Salemi et al., 2023; Micarelli et al., 2007). However, preferences may be hard to learn, or a system might need a lot of data to learn from. Moreover, these learnt preferences are *implicit* and usually cannot be interpreted or edited by the user.

We propose incorporating personalised *standing instructions* as additional context while interpreting a user’s request. Standing instructions are user-provided² natural language statements to change or prescribe system behaviour under certain circumstances. For example, in Fig. 1, the user wants to look for some restaurants around them. In the absence of standing instructions, the system will likely proceed by asking where the user is located followed by their favourite cuisine. By looking up the relevant standing instructions for restaurants, the system can directly search for *Persian restaurants* in *San Leandro*, saving the user’s time as well as providing customised/localised recommendations. Using explicit NL instructions provides both control and interpretability. A user can inspect and edit their standing instructions, especially for preferences that change over time. Further, the generated outputs can be directly linked to the relevant standing instructions, improving the user’s trust in the system (Liu et al., 2023).

Our work is related to Gupta et al. (2022), which conditions a dialogue model’s response on a set of developer guidelines. Their work focuses on controlling response generation in open-domain dialogue systems with a focus on reducing toxicity and enhancing safety. More recently, Open AI released “Custom Instructions”,³ which lets users set preferences for all their future conversations. However, not much is known about how it operates, and no reported evaluations of its usage have been documented.

This paper makes the following contributions:

1. We systematically study the incorporation of standing instructions in a task-oriented setup. We develop and introduce **NLSI** (Natural Language Standing Instructions), where every example consists of a conversation between the user and a dialogue agent and a collection of standing instructions (*user profile*) that are associated with API calls whose arguments are

inferred from the context of the conversation and the relevant standing instructions.

2. We investigate six reasoning types for using standing instructions that range from the inclusion of simple attributes to more complex situations like the user making new choices over existing preferences, the user proposing multiple preferences, *etc.* These reasoning types introduce challenges pertaining to
 - (a) identifying which subset of standing instructions is relevant for the conversational context. Whether an instruction is relevant or not is a function of user utterance and past turns, and requires multi-hop and cross-domain reasoning.
 - (b) effectively incorporating standing instructions while inferring the structured API calls and their attributes. This may require joint reasoning over the dialogue and the relevant standing instructions, as well as dealing with any conflicts between user utterance and instructions.
3. We benchmark the dataset on a combination of methods involving the selection and interpretation of standing instructions. We observe that our LLM-based methods are far from perfect, raising new questions in retrieval, reasoning, and semantic parsing.

2 Task Overview

We are interested in interpreting a user utterance into API calls in light of user-specific standing instructions (Figure 1). Consider a conversational context x , which consists of dialogue history between the user and the agent (if any) and the user’s current utterance. We assume a user profile u consisting of a sequence of natural language instructions u_1, u_2, \dots, u_M . We consider a *selection* task to obtain a set of relevant (to the given context x) standing instructions z from the user profile u ($z \subseteq u$). The *interpretation* task then is to predict API calls y based on the conversational context and the relevant subset of standing instructions z . We assume a schema s that lists the valid API method names and their keyword arguments (slots). The tasks can be modeled as follows:

$$z \sim p(\cdot \mid x, u)$$

$$y \sim p(\cdot \mid x, z, s)$$

²Such standing instructions could also be inferred from user-agent interactions. However, models for inferring instructions are out of the scope of this paper.

³<https://openai.com/blog/custom-instructions-for-chatgpt>

	PLAIN	MULTIHOP	MULTIPREFERENCE
Relevant Standing Instructions (z)	<p>>I always go to Santa Rosa if I'm looking for Movies.</p> <p>>I like fantasy movies the best.</p>	<p>>If I'm looking for a flight, American Airlines is my go-to.</p> <p>>If I'm flying American Airlines, check for Economy seating class.</p>	<p>>If I ask for Events, my preferred event type is Music.</p> <p>>When the event type is Music, search for Blues as the category.</p> <p>>Search for the event name Greensky Bluegrass if the category is Blues.</p> <p>>If I ask for Events, my preferred event type is Sports.</p>
Conversation (x)	<p>User: I want to go out to watch a movie, please help me find a good one.</p>	<p>User: Can you get on and get me a round trip ticket?</p> <p>Agent: Where will you go? Where are you coming from?</p> <p>User: I'm going to SFO from New York City.</p>	<p>User: My schedule is free today and I plan to go to an event in Seattle, WA. I want to look for events in that area.</p>
API calls (y)	<pre>GetMovies(genre="fantasy", location="Santa Rosa")</pre>	<pre>GetFlights(destination="SFO", origin="New York", airlines="American Airlines", seating_class="Economy")</pre>	<pre>GetEvents(city="Seattle, WA", event_type="Music", category="Blues", event_name="Greensky Bluegrass") GetEvents(city="Seattle, WA", event_type="Sports")</pre>
	CONFLICT	NONEAPPLICABLE	MULTIDOMAIN
User Profile (u)	<p>>When I request Restaurants, I prefer Italian cuisine.</p> <p>>If I'm looking for a doctor, I'd rather have a General Practitioner.</p> <p>>If I'm opening a bank account, I want it to be a savings account.</p> <p>>I'd like to get a Doctor in San Rafael if I can.</p> <p>...</p>	<p>>Request Restaurants with Filipino cuisine as my preference.</p> <p>>Request Music by Iggy Azalea as my preferred artist.</p> <p>>If I'm looking to go to the movies, my go-to theatre is Airport Stadium Cinemas.</p> <p>>If I'm looking for a flight, my go-to airline is Alaska Airlines.</p> <p>>Request Events, specifically Sports events.</p>	<p>>When I request Movies, I typically enjoy ones that are comedic.</p> <p>>My first choice when requesting Travel is Vegas</p> <p>>When it comes to Hotels, I prefer ones that are rated 1-star.</p> <p>>My go-to theater for Movies is AMC Bay Street.</p> <p>>If I'm looking into Travel, I should also check out Hotels</p> <p>>I'd like my travel to be kid-friendly.</p> <p>...</p> <p>>My first choice when requesting Travel is Vegas</p> <p>>If I'm looking into Travel, I should also check out Hotels.</p> <p>>When it comes to Hotels, I prefer ones that are rated 1-star.</p> <p>I'd like my travel to be kid-friendly.</p>
Relevant Standing Instructions (z)	<p>>I'd like to get a Doctor in San Rafael if I can.</p>	None	
Conversation (x)	<p>User: I need to find a Gynecologist</p>	<p>User: Can you help me find some attractions to see?</p> <p>Agent: Where should I look?</p> <p>User: How about in KL?</p>	<p>User: User: Any good tourist traps out there?</p>
API calls (y)	<pre>GetDoctors(type="Gynecologist", location="San Rafael")</pre>	<pre>GetTravel(location="KL")</pre>	<pre>GetTravel(good_for_kids="True" location="Vegas") GetHotels(average_rating="1", location="Vegas")</pre>

Table 1: Some examples from NLSI. User profile is not shown in all examples for brevity. (1) In PLAIN, the instructions usually represent a domain matching problem. (2) In MULTIHOP, note that the seating class attribute *Economy* is dependent on choosing the instruction with *American Airlines*. (3) For the example for MULTIPREFERENCE, as there are two preferences for the same attribute *event_type*, there are two separate API calls. Further, the API call with *event_type Music* has additional attributes. (4) In CONFLICT, user requests for an attribute that is against the standing instructions (“Gynecologist” v/s “General Practitioner”). (5) In NONEAPPLICABLE, the user makes a request which is not affected by the standing instructions. (6) In MULTIDOMAIN, the examples contain an instruction which requires invoking a hotel search for the same location when user requests for places to visit.

3 Dataset: NLSI

Existing related datasets have focused on generating safer responses in open-domain dialogue via natural language guidelines (Gupta et al., 2022), personalized text generation by conditioning on a set of past user-written documents like emails or reviews (Salemi et al., 2023), or conditioning on past user feedback for tasks such as ethical reasoning and word scrambling (Madaan et al., 2022). Due to the lack of datasets that study the use of natural language standing instructions in a language-to-program setup, we created NLSI.

3.1 Reasoning Types

In the context of standing instructions, various types of reasoning could be needed to predict the API calls. Following a single standing instruction may be easier than composing and reasoning over several instructions. Furthermore, reasoning across several instructions in the same domain, like booking hotels, may be easier than across domains. Thus, to enable comparisons at different difficulties, we designated six reasoning types for NLSI. While these are not exhaustive, they allow us to systematically study a range of situations ranging from simple domain matching to more complex reasoning (examples in Table 1):

NONEAPPLICABLE For these examples, no standing instructions from the user profile are required for interpreting the user’s utterance ($z = \emptyset$).

PLAIN These examples use the standing instructions directly: each argument can be predicted from a single standing instruction. All the relevant standing instructions, z , belong to the same domain.

MULTIHOP These examples contain at least one standing instruction in z that is deemed relevant to the dialogue x by virtue of presence of another standing instruction. These are of the form “if A then B” and “if B then C”, where A, B, and C are slot names from the same domain. These examples test multi-hop reasoning abilities of the model.

MULTIDOMAIN The examples are similar to **MULTIHOP** except that there is at least one relevant instruction in z that links two domains. These example types typically involve triggering of APIs from an additional domain while being consistent on any shared arguments such as location. For example, invoking a hotel search when user requests

for places to visit (Table 1). These examples challenge multi-domain understanding in addition to multi-hop reasoning.

MULTIPREFERENCE These examples contain standing instructions catering towards multiple preferences for the same attribute. The interpretation task for such examples requires placing multiple API calls respecting the different constraints (*Music* or *Sports* in the example in Table 1). Note that the preferences may interplay differently with other instructions, like in our example, *Blues* and *Greensky Bluegrass* are only required in the call about *Music*.

CONFLICT These examples include instructions in the profile u that conflict with the last user utterance in the dialogue x . The model should gracefully handle such situations and give preference to the user’s request.

Examples can contain standing instructions demonstrating multiple reasoning types. In our work, we associate each example with a single type as based on the above ordering (a type occurring later in the above ordering gets precedence). We provide an example of each type in Table 1.

3.2 Dataset Creation

We constructed NLSI in a semi-automatic fashion by extending Schema Guided Dataset (SGD) (Rastogi et al., 2020). SGD consists of 16K multi-turn conversations across 20 domains like airlines or restaurants. We chose SGD because the dialogues in that dataset include natural and rich conversations with diverse reasoning types, and the accompanying annotations make it possible to construct the ground truth API labels. We note that the process outlined intends to repurpose an existing dataset for studying the selection and interpretation tasks. In a real-world setting, a user might provide explicit preferences through another interface, or else such preferences would be inferred from the user’s continuous interaction with the system.

Extracting standing instructions: We first identified which slots within the SGD schema can be translated into standing instructions based on the slot descriptions provided in the original dataset. For example, while booking movie tickets, *theatre_name* is inclined to be a persistent user preference, hence it can be part of a standing instruction. In contrast, *movie_title* or *date* of booking the movie ticket should not be converted to standing in-

structions, as these are likely to change every time the user interacts with the system.

Each conversation in SGD originates from a series of rules consisting of actions that a user or agent should take. For example, `Greet()` \rightarrow `Inform(location)` \rightarrow `Request(cuisine)` \rightarrow `Inform(date)` \rightarrow `Offer(restaurant_name)`. These actions were then paraphrased into dialogues which constitute the examples in SGD. We reverse-engineer the original SGD creation process to construct the standing instructions.

To convert an SGD dialogue to an NLSI dialogue with standing instructions, we retained the first 1 or 3 turns as the conversational context x , and converted the remaining turns into the relevant standing instructions z . (We ignored any turns that could not be converted into instructions.) For example, the second NLSI example in Table 1 was derived from an SGD dialogue that had originally continued with natural language turns that specified `airlines`="American Airlines", `seating_class`="Economy". Those remaining turns are not needed to predict y from x provided that the standing instructions z can be detected as relevant to that prediction.

Revisiting the example SGD series of rules mentioned earlier, the dialogue corresponding to `Greet()` and `Inform(location)` becomes part of the conversation context x and `Request(cuisine)` and `Offer(restaurant_name)` become the standing instruction. The instruction is templated as "If I ask about *Restaurants*, my preferred *cuisine* is *Italian*". As *date* is a non-instructional slot, we exclude it. Additional details on how examples for various reasoning types are constructed are discussed in Appendix A.

Forming user profiles: The above process provides us with the *relevant* standing instructions z for the given example from SGD, but these are only part of the full user profile u . A user will have additional preferences that are not relevant to the given example. To emulate this, for the given example, we create u by augmenting z with M randomly sampled instructions from other examples.⁴ These

⁴We drew M uniformly from the range $[3, 12]$. Note that owing to this augmentation procedure, the examples in our synthetic dataset are not IID; they may be artificially similar to one another (in addition to being artificial in other ways). In particular, we drew the distractor instructions before splitting the dataset into train/dev/test, so training examples were constructed with some information from the test set. Given this dataset, however, our experiments followed the usual protocol of holding out the test set while constructing our systems.

"distractor" instructions are sampled from domains unrelated to the current domain(s).

Post-processing of instructions: We also included several rounds of pre-checks and post-processing on the dataset to remove undesirable or unrealistic situations that arise either through the noise in the base dataset or our extraction process like domain mismatch ("Play music" followed by "Book me a bus ticket"). To make the standing instructions more natural and diverse, we paraphrased the templates using LLMs (GPT-3.5). Please also see Appendix A for more details.

API calls: The outputs of the interpretation task are API calls y , in line with the recent works of integrating LLMs with tools and plugins (Schick et al., 2023; Qin et al., 2023). The API calls are of the format `GetDomain(slot_1=value_1, slot_2=value_2)`. The argument names and values are derived from dialogue states in the SGD examples, which are either explicitly mentioned in the user’s utterance or provided in the standing instructions.

3.3 Dataset Statistics

We obtained 2441 examples through the above process spanning 17 domains. We construct a balanced test set based on the different reasoning types – 340 per reasoning type type, leading to a total of 2040 examples in test. The train set contains at most 10 examples per domain with a minimum of five examples per reasoning type type, for a total of 150 examples. The remaining examples are a part of the development set (251). There are 10.4 ± 3.0 instructions on average in a user profile (min: 3, max: 22) and there are 2.1 ± 1.7 relevant standing instructions per example in the dataset (min: 0, max: 10). As there are 17 domains, there are 17 function calls corresponding to it.

4 Methods

Given the recent success of using LLMs to generate outputs in structured prediction tasks (Roy et al., 2023; Schick et al., 2023; Heck et al., 2023), we resort to using an LLM-based method to interpret a user utterance into a structured API call. We use in-context learning (Dong et al., 2023) by providing K demonstration examples, where K is tuned on the dev set. These demonstration examples are obtained by retrieving examples from the training set that are most similar to the current

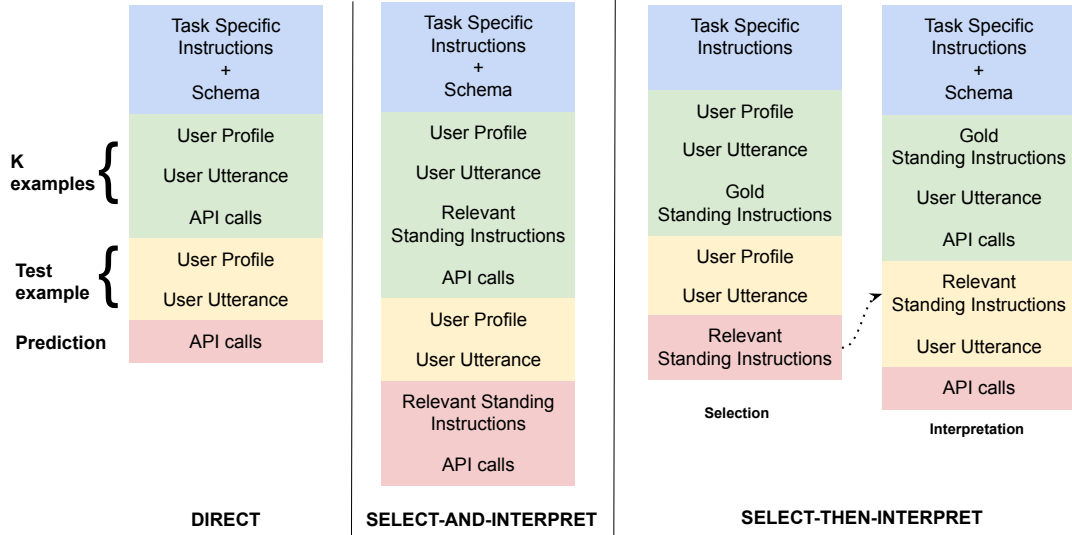


Figure 2: Illustration of different prompting methods. The blocks in red are the expected output generation and every other block is part of the input. The green bits are repeated K times, providing K demonstrations for in-context learning. **DIRECT** Interpretation conditions the generation of API calls on the user profile and user utterance. **SELECT-AND-INTERPRET** requires the generation of the appropriate standing instructions based on user profile and user utterance followed by API generation. **SELECT-THEN-INTERPRET** receives the predicted standing instructions from a separate Selection Model (see left) in addition to the user utterance and then generates the API calls. The selection step only generates the standing instructions based on the user profile and the user utterance.

dialogue of the test example utterance using the BM25 similarity measure (Robertson et al., 1994) as in Rubin et al. (2022); Roy et al. (2023). The examples are arranged in a best-first order. We describe the different paradigms (Fig. 2) used for the interpretation task by selecting the instructions implicitly (**DIRECT** Interpretation), jointly (**SELECT-AND-INTERPRET**) or explicitly (**SELECT-THEN-INTERPRET**).

4.1 Direct Interpretation

In **DIRECT** method, we do not have any explicit selection of standing instructions from the user profile, and directly interpret the dialogue context into API calls. The input to the LLM (Fig. 2) consists of (i) instructions about the interpretation task including the information about using standing instructions, (ii) the schema of the dataset (list of functions and arguments that can be used when generating API calls) s , (iii) user profile u , (iv) user’s dialogue x , and (v) API calls y . Of these, (iii)-(v) are repeated for every demonstration example and the test example only consists of the user profile and the dialogue. We also include the list of categorical slots and their categories as well as a list of boolean slots while describing the schema. This setup allows us to evaluate the ability of implicit selection of the relevant standing instructions for

the interpretation task.

4.2 Joint Selection and Interpretation

Inspired by the effectiveness of techniques like *Chain-of-Thought* prompting (Wei et al., 2022) across several tasks (Chu et al., 2023), we also treat the direct interpretation task with a two-step approach: generate the relevant standing instructions $z \subseteq u$ and then generate the corresponding API calls y .

In addition to potentially improving the accuracy, such explicit selection can enhance the transparency of the system by exposing the relevant subset of instructions to the user. To implement the method, the input prompt to the LLM is modified such that the demonstration examples include the set of all standing instructions u , the relevant standing instructions z , and then the API calls y (Fig. 2). We refer to this method as **SELECT-AND-INTERPRET**.

4.3 Selection Then Interpretation

In this method, we treat selection and interpretation with two separate models (see Section 2). The selection model is not limited to an LLM-based approach. The interpretation model is similar to the one described for **DIRECT**, except that instead of user profile, the relevant standing instructions

are used directly. By decoupling the selection task from the interpretation task, we can explore popular methods of information retrieval for selection. Additionally, as the user profile size increases, and the instructions no longer fit into the prompt, a separate selection step can be convenient. We now describe various approaches for the selection step.

ORACLE: In this setup, the selection step simply returns the true z . This setup measures the standalone performance of the interpretation task when given the correct standing instructions.

BM25: The selection step sets z to the N instructions from the user profile u that are most similar to the dialogue x , where N is tuned on the dev set. As the similarity measure, we use BM25 (Robertson et al., 1994). To compute the corpus statistics used by BM25 to define similarity, each instruction in u is considered a document, and so too is each standing instruction from the train examples.

CONTRIEVER: As above, but replace BM25 with cosine similarity. The dialogue x and each standing instruction in u is embedded into \mathbb{R}^{768} with a pretrained sentence encoder, CONTRIEVER (Izacard et al., 2022). Both BM25 and CONTRIEVER have been used as baselines in similar past work (Gupta et al., 2022; Salemi et al., 2023).

ICL: We also experiment with using LLMs for the selection task. The fixed input prompt to the LLM consists of instructions for the selection task, followed by exactly six demonstration examples, each consisting of a dialogue x , user profile u , and relevant standing instructions z and then the test example (see Fig. 2, Selection). We randomly sampled the six demonstration examples from the training set, one per reasoning type, and used the same demonstration examples for all the test examples.

ICL-DYNAMIC: Similar to ICL, except that now K demonstration examples are dynamically retrieved from the train split by using the ones that are similar to the dialogue in the current example through BM25.

MULTI-PASS: In our preliminary experiments with the previous LLM-based selection methods, we observed that the LLMs consistently missed a subset of relevant instructions in the MULTIHOP and MULTIDOMAIN reasoning types. We use the standing instructions selected in the first pass of the selection process from ICL as part of the prompt to

perform a new selection step. We instruct the model to find the standing instructions that are missing from the current selection set. Though the process can be iterated across multiple steps, in the preliminary experiments we found the best results with only one additional round of selection.

5 Experiments

We benchmark the dataset on the methods discussed above to explain the various challenges on the benchmark. We used GPT-3.5 (text-davinci-003) and GPT-4 as the base LLMs.

5.1 Evaluation

For the interpretation task, we report exact match and slot-F1 of the predicted API call. For the selection task, we report the exact match and sentence-F1. For the interpretation task, the exact match requires getting every function call and its arguments equal to the ground truth. The slot-F1 is F1 score per example and then averaged over the test set. The order in which the API calls are generated is not important. We post-processed the outputs to make punctuation consistent and lowercase both the ground truth and prediction API strings.

For the selection task, exact match and sentence-F1 operate similarly to the above definition except that the triples are now individual standing instructions. Similar to the API strings, all instructions are converted to lowercase. Additionally, we excluded any generated or selected instructions that were not present in the user profile.

We provide some additional details about evaluation in Appendix B.

5.2 Results

We report the results for the different methods in Table 2. Overall, across all the methods, using GPT-4 as the LLM outperforms using GPT-3.5. Within the different ways of incorporating the selection task with the interpretation task, we find that DIRECT interpretation gives the best result, closely followed by the SELECT-AND-INTERPRET and then the ICL results when using GPT-3.5 as base LLM. This trend shifts for GPT-4 where ICL has the best result followed by DIRECT and then SELECT-AND-INTERPRET. Despite the success of chain-of-thought methods in tasks like mathematical reasoning (Wei et al., 2022) and multi-hop question answering (Yoran et al.,

Instruction Subset Selection	GPT-3.5				GPT-4			
	Selection		Interpretation		Selection		Interpretation	
	EM↑	Instruction F1↑	EM↑	Slot F1↑	EM↑	Instruction F1↑	EM↑	Slot F1↑
DIRECT	N/A	N/A	32.0	66.4	N/A	N/A	42.0	67.9
SELECT-AND-INTERPRET	25.9	50.3	28.0	65.9	46.5	67.6	40.2	73.2
SELECT-THEN-INTERPRET								
BM25	17.3	3.0	11.2	39.7	17.3	3.0	11.8	40.8
CONTRIEVER	14.6	51.5	17.2	57.5	14.6	51.5	25.4	62.7
ICL	33.5	48.1	24.7	61.6	50.7	67.7	44.7	75.5
ICL-DYNAMIC	29.0	32.2	19.5	54.9	44.7	61.3	40.7	73.4
MULTI-PASS	24.3	52.1	20.6	57.2	67.3	70.0	*	*
ORACLE	N/A	N/A	55.9	82.8	N/A	N/A	58.5	84.1

Table 2: Results of the different methods on the NLSI dataset for the interpretation task and selection task evaluated on (Instruction/Slot) F1 and Exact Match (EM). DIRECT has the highest score on exact match followed by SELECT-AND-INTERPRET for GPT-3.5 while ICL is best followed by DIRECT for GPT-4. ICL is the best method within SELECT-THEN-INTERPRET and LLM based selection methods are better. * denotes that experiment was couldn’t complete by the time of the submission.

2023), we find that generating selection step and generating API call within the same prompt may not be suitable for incorporating standing instructions.

Models struggle to effectively incorporate standing instructions The best-performing configuration across all the methods only has an exact match of 44.7%. Considering the ORACLE method has an exact match of 58.5%, there is a considerable gap in performance. Incorporating standing instructions to interpret the user’s context is not a trivial problem and would require approaches beyond contemporary prompting methods. Further, even with the gold standing instructions in ORACLE, the models fail to reach maximum exact match for interpretation, suggesting the difficulty of the interpretation task as well. We attribute this to the examples in our dataset that require understanding from different contexts - standing instructions, list of valid APIs, and the current dialogue. Further, the standing instructions are also dependent on each other, hence, models need to reason over the order in which instructions need to be processed to generate the correct API call (see an example of MULTIHOP).

Comparison across selection methods: We find that LLM-based selection methods surpass traditional methods based on lexical statistics and embedding similarity as also seen in (Sun et al., 2023). However, the gap between the ORACLE setting in the selection module and the best-performing configuration is also huge (on

both exact match and F1) suggesting that selecting the relevant standing instructions explicitly from the user profile in the context of the conversation is challenging. We find that ICL has a significant improvement over ICL-DYNAMIC selection method including the per reasoning type distribution in Table 3. This indicates that using simple lexical overlap is not sufficient for demonstration example selection when performing in-context learning.

5.3 Results by reasoning type

We now break down the examples by reasoning type in Table 3, and investigate the accuracy of different methods, using GPT-3.5 as the base LLM. We include the ORACLE for understanding the gaps in these results.

- We observe that different methods display varying trends across different reasoning types and there is no one consistent *winner* among these methods.
- We find that PLAIN is the easiest reasoning type for all the methods, suggesting that LLMs do have the capacity to follow simple standing instructions. Methods perform worse on more complex MULTIDOMAIN examples (<5%). These examples require sharing arguments across multiple domains, following individual standing instructions under respective domains, and reasoning across different standing instructions.
- From Table 2, we note that the MULTI-PASS setup has an overall exact match lower than

Type	ORACLE	DIRECT	JOINT	ICL-D	ICL	MULTI-P
NONEAPPLICABLE	65.3	45.9	37.9	54.4	58.5	29.4
PLAIN	80.3	56.2	56.5	41.8	28.5	36.5
MULTIHOP	65.3	41.8	34.1	27.6	19.1	34.1
MULTIPREFERENCE	40.0	11.5	11.5	8.8	4.1	9.7
MULTIDOMAIN	23.2	3.5	3.2	0.6	0.3	1.2
CONFLICT	70.3	34.1	26.2	17.1	6.8	14.7

Table 3: Per reasoning type exact match on the interpretation task (GPT-3.5). ICL-D is ICL-DYNAMIC and MULTI-P is MULTI-PASS. All the methods find PLAIN easiest and struggle on MULTIDOMAIN. Different methods show different trends without a consistent winner.

ICL. However, the improvement in MULTIHOP, MULTIPREFERENCE, and MULTIDOMAIN reasoning type types over the ICL setup suggests that another round of standing instruction selection can benefit the reasoning types where some complex reasoning over the instructions is required.

These trends also hold for GPT-4 and we report this result in Table 5 in Appendix C.

5.4 Qualitative Analysis

We analyse 100 erroneous examples each from the DIRECT, the best-performing model and ICL, the best model under SELECT-THEN-INTERPRET setup (excluding ORACLE). We identify the most prominent error in an example and discuss trends of errors across these three experiments (for GPT-3.5). We list these error examples in Table 4. The errors that are common to both these methods include

- The hallucination of slot names and slot values while generating the API calls (Example 1) as well as missing some arguments (Example 3).
- Within the MULTIPREFERENCE reasoning type, the models tend to exclude the second API call. Further, if one of the repeating argument/slot has a standing instruction dependent on its value, the model does not include this conditional dependence when generating the API call (Example 2).
- Within the MULTIDOMAIN reasoning type, the models tend to exclude the API call for one of the two domains (Example 3).

In addition to the above errors, DIRECT INTERPRETATION also suffers from over-generation of API calls (Example 1). This is partly because the model may confuse demonstration examples

from PLAIN or CONFLICT with MULTIDOMAIN or MULTIPREFERENCE. Another possible reason is that the model incorrectly considers many irrelevant instructions in the profile while generating the API calls.

5.5 OpenAI’s Custom Instructions

OpenAI also recently reported the introduction of custom instructions⁵ that allow the users to add requirements or preferences that ChatGPT should consider when generating the responses. This is similar to our notion of standing instructions. To test the effectiveness of this feature (free version), we use the instructions from the user profile as “custom instructions”. We pose the API generation task as a standalone task and hope for the model to directly incorporate the standing instructions from the custom instructions. We also use the ICL setup to provide examples about the task as discussed in Section 4.3. As this effort requires manual copy-pasting of examples, we randomly selected and evaluated 17 examples per type, amounting to 102 test examples. While not directly comparable with Table 2, the exact match for the interpretation task on this subset is 15.6 and the slot F1 score is 45.5. Thus, the model does not necessarily incorporate the correct custom instructions every time. It is prone to copying arguments from the demonstration example as well as hallucinating the arguments and their values. For some examples, the model is prone to over-generation of API calls and other unrelated text. We remark that due to the opacity of the “custom instructions” UI, we do not know the prompt or exact LLM API call and we’ve observed nondeterminism in the outputs.

6 Discussion

Identifying the subset of relevant instructions is challenging: We compared and contrasted several baseline methods on the NLSI dataset. One peculiar observation is that all methods need to perform better on the selection task for our dataset especially the ones only making a single pass of selection. The generation of the correct API call requires understanding of the user’s context, the schema, the set of relevant standing instructions, as well as the dependence between standing instructions. Further, the order in which these instructions are retrieved can be important, which means the

⁵<https://openai.com/blog/custom-instructions-for-chatgpt>

Reasoning Type	Conversation	Gold SI	Gold API	Predicted API	Error
PLAIN	User: I am looking for a place of interest to visit.	>If I'm looking to go somewhere, my top choice is New York City.	GetTravel(category="Tourist Attraction", location="New York City")	GetTravel(category="Place of Interest", location="New York City") GetMusic(genre="Country") GetMedia(genre="Drama", directed_by="Qui Sheng") ...	Hallucination
MULTI PREFERENCE	User: I fancy watching a movie at home this evening	>When I request Media, my favorite type is adventure. >If you're looking for an adventure movie, search for one directed by Joel Coen. >When I request Media, my favorite type is biographical.	GetMedia(directed_by="Joel Coen", genre="adventure") GetMedia(genre="biographical")	GetMedia(genre="adventure") GetMedia(genre="biographical")	Skipped Instruction
MULTI DOMAIN	User: I'm looking for some interesting attractions to visit, specifically a Museum. Can you help me find any?	>I prefer London as my destination when I travel. >If I'm looking into Travel, I should also check out Hotels. >I prefer to stay in hotels that have a two-star rating when I am making reservations.	GetTravel(category="Museum", location="London") GetHotels(average_rating="2", location="London")	GetTravel(category="Museum")	Missing API Prediction

Table 4: Examples of prominent errors across the DIRECT and ICL methods (with GPT-3). The incorrectly predicted or missing arguments and function calls are marked in red. The DIRECT models tend to produce several unrelated API calls (first example). Both DIRECT and ICL have a tendency to miss an argument that is only dependent on one of the attributes in MULTIPREFERENCE, in this case missing the director *Joel Coen*. Majority of predictions in MULTIDOMAIN fail at generating the API calls for the second domain.

instruction selection task is no longer a retrieval of independent facts but a reasoning over attributes from the user’s context. This may explain why we found that standard retrieval approaches fail at this task. Our findings align with the observations made in other tasks that find the retrieval of some form of context from a separate memory to be challenging (Weir et al., 2023; Majumder et al., 2023).

Impact of size of user profile: The maximum size of the user profile in our dataset is 22 standing instructions. Our current methods include all the standing instructions explicitly as part of the user profile in the respective prompts. Over time, we envision the capability to add new standing instructions to user profiles, which might exceed the prompt’s capacity. We anticipate that our benchmark can be useful for evaluating interesting questions in LLMs augmented with external memory (Lewis et al., 2020). Further, decoupling the selection step would provide more transparency, as it would allow users to see their individual standing instructions that influenced the generated output (Liu et al., 2023)

Interface to incorporate standing instructions: Our current dataset assumes that the instructions are already provided and the user has consented to the use of the same. The CONFLICT reasoning type also assumes that the user’s request is preferred over standing instructions. In the future, standing instructions can be extracted from user’s interac-

tions with the system. As standing instructions become a component of a larger interface, UX design must include the user’s consent to include or update such existing and inferred standing instructions. Our dataset only provides a starting point on how standing instructions can be considered by LLM-based systems.

7 Related Work

NL guidelines: Gupta et al. (2022) collected and released a dataset of NL guidelines that govern the safe response generation in dialogue systems. Compared to theirs, we showcase a more challenging retrieval setup: we have more applicable instructions on average, with rich phenomena such as MULTIHOP or MULTIPREFERENCE (which is also highlighted in their Limitations). Moreover, we are concerned with generating structured representations as a more complex final task. Irfan et al. (2021) consider a variant of standing instructions in a barista setting where the instruction consists of the favourite drink and snack of the corresponding user. Our work offers more diverse scenarios and domains. We also explore the complexity of selecting relevant standing instructions.

The use of declarative NL specifications has been explored in past work. For example, Ye et al. (2023) use an LLM to generate a declarative task specification, coupled with an off-the-shelf automated theorem prover to derive the final answer.

Weir et al. (2023) discuss methods to generate user-NPC dialogues based on game quest specifications. Constitutional AI (Bai et al., 2022) identifies whether some model response violates a given rule, and then revises the response accordingly.

Closely related to the use of standing instructions is also learning from feedback (Labutov et al., 2018; Tandon et al., 2022; Madaan et al., 2022), where the goal is to maintain a memory of user-provided feedback and use it to augment the knowledge used by question-answering models at test time. Analogously, standing instructions can also be seen as a form of memory.

Personalisation: Personalisation in dialogue has been extensively studied (Li et al. (2016); Zhang et al. (2018); Majumder et al. (2020); *inter-alia*) where the personality traits are provided through NL statements. Closer to our work, Joshi et al. (2017) provide a user profile consisting of age, gender, and favourite food item structured as a dictionary. However, all these works focus on providing a persona to the bot to generate more engaging responses rather than assisting the users in completing their request.

In a broader sense, learning from preferences has been fundamental to improving user experience. These include personalised review generation (Li et al., 2020), personalised search results through collaborative filtering (Micarelli et al., 2007) or leveraging a profile of user interests (Speretta and Gauch, 2005). More recently, Salemi et al. (2023) explored personalised text generation with LLMs on tasks such as article generation given past articles authored by the user. Our work provides incorporation of preferences explicitly through standing instructions allowing better understanding of a generated result.

8 Conclusion

We proposed the use of standing instructions - a set of natural language statements that contain the user’s preferences to enhance the interpretation of the user’s requests. To facilitate this, we created NLSI, a dataset based on the SGD dataset. We explored the interpretation task of generating API calls which are conditioned on the user’s current interaction with the system and selecting the relevant standing instructions from a list of pre-defined preferences. We experimented with several methods for the selection and interpretation tasks. Our results show that while LLMs are somewhat ca-

pable of incorporating standing instructions as an additional context, their usage of standing instructions is far from perfect. The models struggle at selecting the instructions in the user profile that were relevant for the given dialogue, which in turn affects the interpretation task. Moreover, as reasoning types become more intricate and involve complex reasoning or interactions among the respective standing instructions, the interpretation of these instructions becomes increasingly challenging for the methods. This calls for the development of new approaches in incorporating standing instructions, reasoning-based retrieval, and memory-augmented representations.

9 Limitations

Our task setup is limited to generating API calls for the current turn. In an ideal scenario, the LLM or the service should also display the results in a user-friendly format, like natural language or Markdown, and perhaps confirm with the user before executing the call. Our dataset is not accompanied by the results from respective API calls or replies from the system due to the unavailability of results from the base dataset. The different reasoning types in our dataset are not exhaustive and future work could look into expanding them.

As our dataset is derived from an existing task-oriented dialogue dataset, it is useful for testing methods, but we caution readers that it is only a synthetic dataset. Preferences stated explicitly by a human user would likely take a wider range of natural language forms. Preferences deduced from the user’s past history might take a non-linguistic form, as in recommendation systems; they might be uncertain or soft constraints that cannot be passed directly as arguments to simple search APIs.

References

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bow-

- man, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. [Constitutional AI: Harmlessness from AI feedback](#).
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Advances in Neural Information Processing systems*, 33:1877–1901.
- Zheng Chu, Jingchang Chen, Qianglong Chen, Weijiang Yu, Tao He, Haotian Wang, Weihua Peng, Ming Liu, Bing Qin, and Ting Liu. 2023. [A survey of chain of thought reasoning: Advances, frontiers and future](#).
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. 2023. [A survey on in-context learning](#).
- Prakhar Gupta, Yang Liu, Di Jin, Behnam Hedayatnia, Spandana Gella, Sijia Liu, Patrick Lange, Julia Hirschberg, and Dilek Hakkani-Tur. 2022. [Dialguide: Aligning dialogue model behavior with developer guidelines](#). *arXiv preprint arXiv:2212.10557*.
- Michael Heck, Nurul Lubis, Benjamin Ruppik, Renato Vukovic, Shutong Feng, Christian Geischauser, Hsien-chin Lin, Carel van Niekerk, and Milica Gasic. 2023. [ChatGPT for zero-shot dialogue state tracking: A solution or an opportunity?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 936–950, Toronto, Canada. Association for Computational Linguistics.
- Bahar Irfan, Mehdi Hellou, and Tony Belpaeme. 2021. [Coffee with a hint of data: Towards using data-driven approaches in personalised long-term interactions](#). *Frontiers in Robotics and AI*, 8.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Trans. Mach. Learn. Res.*, 2022.
- Chaitanya K. Joshi, Fei Mi, and Boi Faltings. 2017. [Personalization in goal-oriented dialog](#). *NeurIPS 2017 Conversational AI Workshop*.
- Igor Labutov, Bishan Yang, and Tom Mitchell. 2018. [Learning to learn semantic parsers from natural language supervision](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1676–1690, Brussels, Belgium. Association for Computational Linguistics.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios Spithourakis, Jianfeng Gao, and Bill Dolan. 2016. [A persona-based neural conversation model](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 994–1003, Berlin, Germany. Association for Computational Linguistics.
- Junyi Li, Siqing Li, Wayne Xin Zhao, Gaole He, Zhicheng Wei, Nicholas Jing Yuan, and Ji-Rong Wen. 2020. [Knowledge-enhanced personalized review generation with capsule graph neural network](#). In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 735–744.
- Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo, Hao Cheng, Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. 2023. [Trustworthy llms: a survey and guideline for evaluating large language models’ alignment](#).
- Aman Madaan, Niket Tandon, Peter Clark, and Yiming Yang. 2022. [Memory-assisted prompt editing to improve gpt-3 after deployment](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2833–2861.
- Bodhisattwa Prasad Majumder, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Julian McAuley. 2020. [Like hiking? You probably enjoy nature: Persona-grounded dialog with commonsense expansions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9194–9206.
- Bodhisattwa Prasad Majumder, Bhavana Dalvi Mishra, Peter Jansen, Oyvind Tafjord, Niket Tandon, Li Zhang, Chris Callison-Burch, and Peter Clark. 2023. [CLIN: A continually learning language agent for rapid task adaptation and generalization](#).
- Alessandro Micarelli, Fabio Gaspiretti, Filippo Sciarone, and Susan Gauch. 2007. [Personalized search on the World Wide Web](#). In Peter Brusilovsky, Alfred Kobsa, and Wolfgang Nejdl, editors, *The Adaptive Web: Methods and Strategies of Web Personalization*, volume 4321 of *Lecture Notes in Computer Science*, pages 195–230. Springer.
- OpenAI. 2023. [GPT-4 technical report](#).
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2023. [ToolLLM: Facilitating large language models to master 16000+ real-world apis](#).

- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. [Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8689–8696.
- Stephen E. Robertson, Steve Walker, Susan Jones, Micheline Hancock-Beaulieu, and Mike Gatford. 1994. [Okapi at TREC-3](#). In *Text Retrieval Conference*.
- Subhro Roy, Sam Thomson, Tongfei Chen, Richard Shin, Adam Pauls, Jason Eisner, and Benjamin Van Durme. 2023. [BenchCLAMP: A benchmark for evaluating language models on syntactic and semantic parsing](#). In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2022. [Learning to retrieve prompts for in-context learning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2671, Seattle, United States. Association for Computational Linguistics.
- Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2023. [LaMP: When large language models meet personalization](#).
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Mirco Speretta and Susan Gauch. 2005. [Personalized search based on user search histories](#). *The 2005 IEEE/WIC/ACM International Conference on Web Intelligence (WI'05)*, pages 622–628.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Shuaiqiang Wang, Pengjie Ren, Zhumin Chen, Dawei Yin, and Zhaochun Ren. 2023. [Is ChatGPT good at search? investigating large language models as re-ranking agents](#).
- Niket Tandon, Aman Madaan, Peter Clark, and Yiming Yang. 2022. [Learning to repair: Repairing model output errors after deployment using a dynamic memory of feedback](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 339–352, Seattle, United States. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *NeurIPS*.
- Nathaniel Weir, Ryan Thomas, Randolph D’Amore, Kellie Hill, Benjamin Van Durme, and Harsh Jhamtani. 2023. [Ontologically faithful generation of non-player character dialogues](#).
- Xi Ye, Qiaochu Chen, Isil Dillig, and Greg Durrett. 2023. [SatLM: Satisfiability-aided language models using declarative prompting](#). In *The 3rd Workshop on Mathematical Reasoning and AI at NeurIPS’23*.
- Ori Yoran, Tomer Wolfson, Ben Bogin, Uri Katz, Daniel Deutch, and Jonathan Berant. 2023. [Answering questions by meta-reasoning over multiple chains of thought](#).
- Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. [Personalizing dialogue agents: I have a dog, do you have pets too?](#) In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

A Dataset Construction Details

Forming examples for different reasoning types:

We do not need to extract any standing instructions z for examples in NONEAPPLICABLE. For examples in PLAIN, each (domain, slot, value) triple was extracted and written in natural language via an if-else template. Since each slot is independent of each other, this set of instructions form z . MULTIHOP examples were formed by creating a hierarchy of slots associated with the same domain (like *seating_class* is dependent on *airlines*). If the subsequent dialogue states contained the same dependent slots, then that example was categorized as a MULTIHOP example, where the primary slot value was obtained from the dialogue or one of

the standing instructions. MULTIDOMAIN examples were dialogues from SGD that were inherently multi-domain because they required API calls from different domains. These reasoning types were created through a deterministic process based on the existing SGD data.

MULTIPREFERENCE examples were formed by duplicating one of the ground truth standing instructions from PLAIN, MULTIHOP and MULTIDOMAIN, and substituting a value with another relevant entity. Meanwhile, CONFLICT examples were formed with examples from PLAIN or MULTIHOP. We added information that conflicts with the gold standing instruction like asking for *Mexican* restaurants when the standing instruction is about preference for Italian restaurants.

Post-processing: We unified domains such as *Restaurant_1*, *Restaurant_3* as *Restaurants*. *Restaurant_2* was renamed as *HouseStays*. We also deduplicated the slot names under these domains like *location* and *area* was converted to *area*. Similarly, the *Services* domain was expanded as *Salons*, *Doctors*, and *Dentists* instead. All the examples were constructed only from the domains and examples available in the training set of SGD. In addition to removing domains whose combination doesn't make sense in the MULTIDOMAIN reasoning type, we also remove MULTIDOMAIN examples which do not have any attributes for the second domain.

For paraphrasing the templated instructions, we prompted GPT-3 to generate paraphrases with three distinct prompts to promote diversity.

Prompt 1: Write a colloquial paraphrase for the given sentences. Refrain from using if then format

Prompt 2: Reword the following in your own words. Keep the same meaning. Change the sentence structure to exclude if then format:

Prompt 3: Reword the following in your own words. Keep the same meaning. Make the sentences sound like instructions or commands.

Change the sentence structure to exclude if-then format. If the sentence starts with "If I ask for xyz", also reword that xyz part.

We replace the templated standing instruction randomly with one of the paraphrases leading to 4097 unique instructions across the dataset.

B Experiment Details

For the selection experiments involving BM25 and Contriever, N was varied from 1 to 10 and

chosen according to the best exact match on the dev set ($N=4$ for BM25, $N=2$ for CONTRIEVER). For LLMs, the K for demonstration exemplars was varied from 1 to 5 (with $K=5$ being best for ICL-DYNAMIC and other interpretation tasks). For the MULTI-PASS experiments, we varied K for three additional rounds and found that providing one additional pass had the best results on the development set. The temperature for all the LLM-based experiments was 0. We provide the prompt templates for the different experiments at <https://github.com/nikitacs16/nlsi>

For evaluation, all the outputs were converted to lowercase and double quotes were unified to a fixed unicode. Using "vs" and "versus" was unified to "versus". The models were not penalised if they produced *subcategory* instead of *event_type* arising due to the noise in the base dataset. For the interpretation evaluation, the API calls were converted to function_name-slot-value triples per slot-value per API call. In the case of examples multiple API calls, the models had a tendency to include every attribute in a single API call instead of separate API calls. To penalise this in the exact match, if the number of predicted API calls was not equal to ground truth API calls the model received an exact match of 0.

C Additional Results

We report the results by reasoning type for experiments using base LLM as GPT-4 in Table 5. The trends are similar to the trends discussed in Section 5.3. As we could not obtain the interpretation results in time, we look at the per reasoning type for the selection task while comparing ICL and MULTI-PASS. We see that there is a clear improvement over PLAIN (86.7 v/s 92.0), MULTIPREFERENCE (46.7 v/s 61.4) and MULTIDOMAIN(30.5 v/s 40.2) confirming that a second pass over the predictions improves the results.

Type	ORACLE	DIRECT	JOINT	ICL-D	ICL
NONEAPPLICABLE	68.2	57.3	48.8	61.4	62.6
PLAIN	77.9	67.6	70.5	69.7	65.0
MULTIHOP	65.5	56.4	47.3	59.1	57.9
MULTIPREFERENCE	55.8	24.1	32.6	42.6	38.2
MULTIDOMAIN	30.9	16.1	12.6	12.0	07.6
CONFLICT	70.2	35.0	32.0	33.5	22.3

Table 5: Per reasoning type exact match on the interpretation task (GPT-4). ICL-D is ICL-DYNAMIC. All the methods find PLAIN easiest while struggling at MULTIDOMAIN. Different methods show different trends without a consistent winner