

# PyTOD: Programmable Task-Oriented Dialogue with Execution Feedback

Alexandru Coca<sup>1\*</sup>, Bo-Hsiang Tseng<sup>2</sup>, Pete Boothroyd<sup>2</sup>, Jianpeng Cheng<sup>3\*</sup>, Mark Gaynor<sup>2</sup>, Zhenxing Zhang<sup>2</sup>, Joe Stacey<sup>4\*</sup>, Tristan Guigue<sup>2</sup>, Héctor Martínez Alonso<sup>2</sup>, Diarmuid Ó Séaghdha<sup>2</sup>, Anders Johannsen<sup>2</sup>

<sup>1</sup>University of Cambridge, <sup>2</sup>Apple, <sup>3</sup>Meta, <sup>4</sup>Imperial College London  
ac2123@cam.ac.uk, ajohannsen@apple.com

## Abstract

Programmable task-oriented dialogue (TOD) agents enable language models to follow structured dialogue policies, but their effectiveness hinges on accurate state tracking. We present PyTOD, an agent that generates executable code to track dialogue state and uses policy and execution feedback for efficient error correction. To this end, PyTOD employs a simple constrained decoding approach, using a language model instead of grammar rules to follow API schemata. This leads to state-of-the-art state tracking performance on the challenging SGD benchmark. Our experiments show that PyTOD surpasses strong baselines in both accuracy and robust user goal estimation as the dialogue progresses, demonstrating the effectiveness of execution-aware state tracking.

## 1 Introduction

TOD agents provide natural language interfaces which enable users to control their digital environment to complete daily tasks. Such agents typically include a *dialogue state tracking* (DST) component, which maps the conversation history to a symbolic representation of the task-relevant information communicated during the exchange. At each turn, a *dialogue manager* (DM) uses this information to take *system actions*<sup>1</sup> necessary to help the user complete the task. The agent behaviour is controlled by a *dialogue policy* defined by application developers.

Domain adaptation has long been a challenge for state tracking (Jacqmin, 2022), dialogue management (Mosig et al., 2020) and end-to-end (E2E) (Zhao et al., 2023) agents, as it often requires developers to collect and annotate new datasets for re-training. To address this challenge, fine-tuning pre-trained language models (PLMs) (Radford et al., 2019; Raffel et al., 2020) within the schema-guided

\*Work done while at Apple.

<sup>1</sup>For example, retrieving information from a knowledge base or prompting the user to provide task constraints.

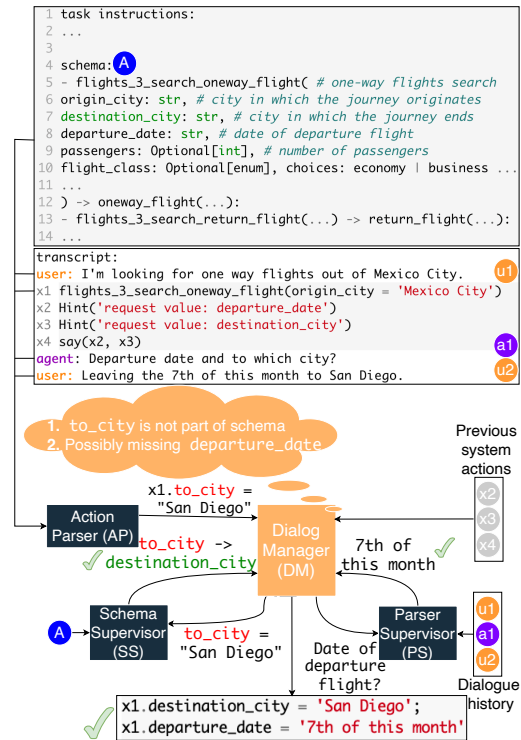


Figure 1: PyTOD overview. The action parser generates Python statements ( $x1.to\_city = "San\ Diego"$ ) representing the actions the user took at the current turn ( $u2$ ) given API schemata (A), dialogue history ( $u1$ ,  $a1$ ) and previous user actions ( $x1$ ). The dialogue manager (DM) executes the user action in a simulated environment. A schema supervisor is invoked by the DM to correct errors if predicted statements contain slot names that are not part of the API schema (e.g., the slot  $to\_city$  is mapped to  $destination\_city$ , a member of the schema). Given knowledge of previous system actions ( $x2$ ,  $x3$ ) the DM detects slot omissions and invokes a parser supervisor to correct them (e.g., flight date is recovered).

paradigm (Rastogi et al., 2019; Mosig et al., 2020) has emerged as a powerful approach. Schemata define the APIs accessible to the agent, including textual descriptions of their functions and parameters (referred to as *slots*). Mosig et al. (2020) extend schemata by including descriptions of the actions that TOD agents can perform. Zhao et al. (2023) build on these advances with AnyTOD, a state-of-the-art (SOTA) neuro-symbolic agent ca-

pable of following dialogue policies unseen during fine-tuning.

AnyTOD uses the schema and dialogue history to first generate a symbolic *state sequence*. This sequence identifies the API the user wishes to interact with and the slot values the user has mentioned. A second symbolic sequence - encoding the actions the user<sup>2</sup> and system have taken - is generated sequentially after the state sequence. Both sequences are subsequently interpreted by a *deterministic policy program* which recommends the next system action. State sequence is thus critical: prediction errors can prevent the system from taking the correct actions, leading to breakdowns in the interaction.

Despite its strengths, AnyTOD has some limitations. First, it re-estimates the state and action history at every dialogue turn solely based on the dialogue history and schema, which increases generation length and amplifies the risk of state-tracking errors. Second, it fails to exploit previous system actions to verify the correctness of the state sequence. Finally, its reliance on symbolic representations of state and action sequences requires additional system components to translate them into executable code, adding deployment complexity. We present PyTOD, a programmable dialogue system that addresses these challenges by directly communicating with its execution environment and policy programs to perform accurate state tracking for unseen APIs and domains, as discussed next.

**PyTOD overview** PyTOD incrementally generates dialogue states as code, using policy- and execution feedback for accurate state estimation (Figure 1). It operates as follows: (1) an action parser (AP) (§3.1) processes the user query to produce one or more Python statements; (2) the dialogue manager (DM) (§3.2) executes these statements, ensuring they valid according to the schema; (3) if invalid slots are found, the DM calls the schema supervisor (SS) (§3.2.2), a language model that constrains code generation according to the decoding schema; and (4) the DM evaluates the constrained output by comparing expected and current environment states, invoking a parser supervisor (PS) (§3.2.3) to recover from possible omissions or semantic errors. The dialogue state is derived by executing the program generated by PyTOD at each user turn.

**Contributions** Unlike PyTOD, TOD agents of-

<sup>2</sup>For example, providing a slot value or requesting information about a knowledge base item retrieved by the system.

ten optimise DST and DM independently, overlooking how policy information can enhance state generation accuracy. We further demonstrate that feedback from the execution environment enables language models to constrain decoding without requiring additional training data, with minimal developer effort and only a slight increase in system latency. PyTOD achieves SOTA performance on the challenging Schema-Guided Dialog (SGD) dataset (Rastogi et al., 2019).

As an additional contribution, we release `pytodlib`<sup>3</sup>, a Python library that simulates SGD APIs, including database responses and API behavior. This toolkit addresses the scarcity of conversational tool-use corpora, providing a valuable resource for benchmarking large language models’ ability to handle complex, multi-task, goal-oriented conversations while following predefined policies.

## 2 Simulated Environment

As a programmable dialogue system, PyTOD relies on a simulation of its operating environment and an execution engine, described in this section. Our simulation is based on SGD due to its large ontology, complex policy and high-quality annotations.

### 2.1 API simulation

We simulate the 58 APIs from the validation and test splits. Each API is implemented as a Python object that can be instantiated and updated by a dialogue agent.

Our simulation is based on a detailed analysis of the system and user dialogue acts in the SGD dataset. Accordingly, upon instantiation or update, the simulation returns structured *system action recommendations* or *hints*, that the agent can execute to continue the dialogue.

The specific action recommendations depend on the dialogue state. If the user has not specified all the required arguments for API execution, *slot filling hints* (x2, x3 in Figure 1) guide the agent towards requesting additional information from the user. When all task details are known, search queries return *entities* retrieved from databases, whereas for transactions, *confirmation hints* (x17 in Figure 2, turn 7) are suggested to confirm slot values. Upon execution, the APIs may return *notifications* if the task could not be completed as requested by the user and, in some cases, *alternative hints* to guide the agent to propose satisfiable

<sup>3</sup>Code: <https://github.com/apple/ml-pytod>.

task constraints (Figure 2, turn 8). Finally, agent initiative is simulated via related-task suggestions (Figure 2, turn 3) and *user prompt hints* (Figure 2, turn 7). See App. A.1 for implementation details.

External API and database responses are simulated using the semantic and entity annotations. The slot values extracted from the dialogue context are normalised prior to invoking APIs, allowing the simulation to return entities as Python objects that dialogue agents can inspect and reason over.

## 2.2 Execution engine

Code execution proceeds as follows:

1. A generated string is parsed to an Abstract Syntax Tree (AST) using Python’s `ast` library
2. The AST is converted into a Python callable
3. The callable is invoked, updating the state of the referenced APIs.

See App. A.3 for further details and examples.

## 3 PyTOD

### 3.1 Action parser

The AP parses user utterances into Python expressions, which are executed to carry out the user’s actions. The prompt consists of: (1) a header containing task-specific instructions and linearised schema, alongside a list of completed tasks and entities (e.g., flights) returned by them (§3.1.1); (2) a session transcript, where user and system turns are interleaved with Python code snippets representing user actions and execution outputs (§3.1.2); and (3) context-dependent instructions, dynamically rendered to provide additional task guidance and entity definitions as the conversation progresses (§3.1.3).

#### 3.1.1 Header

At the start of the conversation, the header instructions (Figure 1, top) prompt the PLM to identify the API that aligns with the user’s intent and extract any arguments specified by the user (Figure 7a, App. B.1.1). Schema API definitions are presented next, linearised as Python function signatures (Figure 1, lines 4–14). Each API name is followed by an *intent description* summarising its function (line 5), while arguments are annotated with their types and descriptions (lines 6–10). For *categorical slots*, which take *closed values* from a predefined set, the argument descriptions are prefixed with the list of valid options (line 10). Return types specify the

```
[1] user: Search a romantic chicken restaurant, in London.
x1 search_restaurant(where="London", cuisine="chicken")
x2 show(x1)
x3 next(x1) # type: restaurant; fields: live_band, ...
x4 say(x3)
agent: Chicken's House?
[2] user: Anything else?
x5 next(x1) # type: restaurant; fields: live_band, ...
agent: Nando's Oval?
[3] user: Love this one, do they have live music?
x6 select(x5, from_results=x1)
x7 suggest(task='book_table')
x8 say(x6.live_band, x7)
agent: Yes, they do! Want a table there?
[4] user: Maybe later. Weather tomorrow?
x9 suspend(x7)
x10 check_weather(city="London", date="tomorrow")
agent: 3 degrees, cloudy.
[5] user: Got it. Book that restaurant then!
x11 resume(x9)
x12 book_table(restaurant="Nando's Oval")
x13 Hint("request value: time")
agent: What time?
[6] user: 7.
x14 x12.time = "7"
x15 Hint("Confirm: people, restaurant, date")
agent: For two, 7pm at Nando's Oval, tomorrow?
[7] user: Hmm...
x16 conversation_pause()
x17 Hint("prompt for next action")
agent: Anything else I can do for you?
[8] user: Let's go for it!
x18 confirm(x12)
x19 Notification("Booking did not succeed")
x20 Hint("offer alternative: date")
agent: No luck for today. Tomorrow same time?
[9] user: Nah, leave it. Bye!
x21 decline_alternative()
x22 say()
```

Figure 2: Session transcript example. Natural Language Generation (NLG) calls shown only for turns 1 and 3.

entities produced by APIs, excluding their properties (line 12), which are displayed dynamically as entities are returned (§3.1.3).

As tasks are completed, the header is updated with descriptions of completed tasks and relevant entity definitions (Figure 7b, App. B.1.1).

#### 3.1.2 Session transcript

The session transcript records the dialogue history, interleaved with *PyTOD program statements* that capture user and system actions (Figure 2).

PyTOD program lines start with *intermediate variables* (Stacey et al., 2024) to enable incremental updates to the dialogue state based on user actions. This design minimises generation length and eschews re-estimating the entire dialogue state after every turn. Moreover, unlike SOTA DST approaches which extract slot values from both agent and user utterances, agent-mentioned slot values are extracted from the database results and API

1	<b>search API call</b> [e.g., <code>search_restaurant</code> ]	INFORM_INTENT	[1]
	<b>Generated when:</b> a user initiates a task which entails retrieving entities from a DB the agent can access (e.g., browsing). <b>Syntax:</b> keywords-only function call. Function name is a concatenation of service name and user intent. <b>Side effect:</b> updates API state with arguments provided by the user		
2	<b>iteration</b> [next]	OFFER, REQ_ALTS	[1], [2]
	<b>Generated when:</b> search queries are fully specified or user browses through search results. <b>Syntax:</b> <code>evar next(sapivar)</code> where <code>sapivar</code> references a search API call. <b>Side effect:</b> fetches top search result from the database and stores it in <code>evar</code> .		
3	<b>selection</b> [select]	SELECT	[3]
	<b>Generated when:</b> user selects an entity suggested by the agent (e.g., a restaurant). <b>Syntax:</b> <code>select(evar, from=api)</code> where <code>evar</code> is a search result and <code>api</code> references a search API call. <b>Side effect:</b> API state updated with selected entity properties (e.g., <code>restaurant="Nando's Oval"</code> at turn [3] in Figure 2).		
4	<b>suspend task</b> [suspend]	NEGATE_INTENT	[3], [4]
	<b>Generated when:</b> a user switches context, suspending a follow-up task suggested by the agent. <b>Syntax:</b> <code>svar suspend(sugvar)</code> where <code>sugvar</code> is a variable bound to an agent suggested task (e.g., x7 in Figure 2). <b>Side effect:</b> <code>svar</code> stores state relevant to the follow-up task (e.g., <code>restaurant="Nando's Oval"</code> in Figure 2).		
5	<b>resume task</b> [resume]	INFORM_INTENT	[5]
	<b>Generated when:</b> a user continues a previously suspended task. <b>Syntax:</b> <code>resume(svar)</code> where <code>svar</code> is a variable referencing a suspended task. <b>Side effect:</b> Ensures resumed tasks (e.g., x12 in Figure 2) are updated with state stored in <code>svar</code> .		
6	<b>transaction API call</b> [e.g., <code>book_restaurant</code> ]	INFORM_INTENT	[5]
	<b>Generated when:</b> a user initiates a task that changes their device state or that of an external service (e.g., bookings). <b>Syntax:</b> See #1. <b>Side effect:</b> See #1.		
7	<b>provide task constraints</b> [-]	INFORM	[6]
	<b>Generated when:</b> a user provides task constraints or corrects slot values. <b>Syntax:</b> dot assignment: <code>variable.field = value</code> , where <code>variable</code> references a prior call and <code>field</code> is an API slot. <b>Side effect:</b> updates referenced API state with slot values provided by the user.		
8	<b>pause</b> [conversation_pause]	—	[7]
	<b>Generated when:</b> detecting an out-of-scope utterance. <b>Syntax:</b> zero-arguments function call to <code>conversation_pause()</code> . <b>Side effect:</b> carries over the current API state to the next turn.		
9	<b>confirm</b> [confirm]	AFFIRM	[8]
	<b>Generated when:</b> a user confirms a transaction for execution. <b>Syntax:</b> <code>confirm(apivar)</code> where <code>apivar</code> references a transaction API call. <b>Side effect:</b> updates API state agent values (e.g., <code>people=2</code> , <code>date="tomorrow"</code> , <code>restaurant="Nando's Oval"</code> in Figure 2).		
10	<b>decline suggestions</b> [decline_alternative]	NEGATE	[9]
	<b>Generated when:</b> a user declines agent's suggestion to update task constraints following a transaction failure. <b>Syntax:</b> zero-arguments function call to <code>decline_alternative()</code> . <b>Side effect:</b> carries over the current API state to the next turn.		

Table 1: PyTOD user action sub-language summary. Shaded rows show equivalent DSTC8 dialogue acts (Rastogi et al., 2020) and [digits] point to syntax examples and utterances grounded in these statements in Figure 2.

responses so are guaranteed to be correct. These design choices mitigate the impact of copy errors, which typically affect the accuracy of generative DST models (Cho et al., 2022).

**User actions** The user actions PyTOD can parse are listed Table 1, with syntax examples in Figure 2. While the actions are closely related to the dialogue acts commonly defined in TOD corpora, App. B.1.2 shows the proposed sub-language can represent more complex task-oriented interactions.

**Communicative system actions** Executing user actions updates the dialogue state. Furthermore, as discussed in §2.1, the dialogue policy simulation returns system action recommendations based on the dialogue state. The agent selects appropriate actions by passing them as positional arguments to the `say` routine (e.g., turns 1 and 3 in Figure 2). `say`

is an interface to the NLG module, which generates the next agent utterance given the selected action recommendations. For efficiency, the `say` calls and action recommendations are suppressed from the transcript after agent utterance generation – only the agent utterances are included in the session transcript history.

**Non-communicative system actions** are persistent in session history and provide cues for state tracking and language generation. These include database call markers which signal entity retrieval and guide the model to predict subsequent iteration calls (x2 in Figure 2) and the `perform` statement (x32 in Figure 9, App. B.1.1) which marks successful execution of transactional API calls.

We provide further session transcripts with detailed explanations of PyTOD operation in App. E.



### 3.1.3 Context-dependent instructions

Context-dependent instructions are rendered dynamically based on generated statements. They list properties of retrieved entities which allow PyTOD to answer questions about them (e.g., turn 3 Figure 2; see Figure 8a in App. B.1 for prompt) and may encode policy (see Figure 8b, App. B.1).

## 3.2 Dialogue manager

The DM coordinates interaction with the user by taking actions recommended by the dialogue policy upon AP output execution. To ensure successful execution, it constrains and validates the AP outputs as described in the remainder of this section.

### 3.2.1 AP output constraints

AP-generated statements must be valid Python expressions with declared variables to execute. The DM enforces these constraints, inserting `parse_errors` into the transcript if parsing fails. Additionally, it restricts generated API names to those listed in the AP header, correcting errors by minimizing edit distance. Once constrained, statements are executed. If the AP predicts a slot name outside the decoding schema, the DM invokes its SS component to constrain it accordingly.

### 3.2.2 Schema supervisor

The SS generates a prompt using the schema and the AP output, based on three generic templates (App. B.2.1). It follows a multiple-choice question answering (MQA) format (Figure 3), where the question and options depend on the AP error. If the AP predicts an unknown slot, the prompt lists the all the slot names in the schema with descriptions<sup>4</sup> as answer choices, instructing the model to select the option corresponding to the slot which best matches the AP output (Figure 3a). If the slot name is unknown but its value is listed in the schema, the prompt includes only categorical slot definitions and possible values (Figure 3b). For cases where the AP outputs a slot from a training schema but not the current task schema, the model is presented with slot descriptions from the task schema and instructed to select the closest paraphrase (Figure 3c).

### 3.2.3 Parser supervisor

The DM invokes the PS when SS-constrained AP assignment expressions fail to provide values for slots the system previously requested from the user.

<sup>4</sup>Descriptions are replaced with data type for integer-valued slots and value enumeration for categorical slots.

Given the definitions, which keyword matches `'to_city'`?

- a) `origin_city`: flight departure city
- **b) `destination_city`: city in which the journey ends**
- c) `number_of_bags`: `int`
- d) `ticket_fare_class`: `economy/business`
- e) none: the definitions do not describe `'to_city'`

Answer:

(a) Unknown slot name. The `SS output` is converted to `destination_city`, replacing `to_city` in the AP output.

Here are some definitions:

- `subtitle_language`: language to use for subtitles (or none)

Given these, `'subtitle_free = true'` is a synonym of:

- **a) `subtitle_language = none`**
- b) `subtitle_language = english`
- c) options do not describe `'subtitle_free = true'`

Answer:

(b) Unknown slot name (closed value). The `SS output` is converted to `subtitle_language = none`, replacing `subtitle_free = true` in the AP output.

Which sentence paraphrases `'date of car pickup (pickup)'`?

- **a) the first date to start using the rental car**
- b) the date to return the car
- c) no option above paraphrases `'pickup'`

Answer:

(c) Memorised slot name. The `SS output` is converted to `pickup_date`, replacing `pickup` in the AP output.

Figure 3: Illustration of prompts generated by the SS to **constrain AP generations** to the decoding schema.

Q: Answer the following questions. Output "unanswerable" if the question cannot be answered given the conversation.

In the conversation:

...

**user**: Great, now search one way flights out of Mexico City.

**agent**: Departure date and to which city?

**user**: Leaving the `7th of this month` to `San Diego`.

1) city in which the journey ends?  
2) date of departure flight?

Answer:

Figure 4: Illustration of the prompt generated by the PS to handle AP omissions and semantic errors. The expected SS output is 1) `San Diego` 2) `7th of this month`. The second answer is used to correct an omission by concatenating the constrained AP output in Figure 1 (`x1.destination_city = "San Diego"`) with `x1.departure_date = "7th of this month"`.

In response, the PS generates a prompt using the schema and dialogue history, based on a simple template (App. B.3). The prompt follows an extractive QA format (Figure 4), where the questions correspond to the schema descriptions of the omitted slots (e.g., `Departure date of the flight?` in Figure 1); the context for answering them is limited to the dialogue history of the current task.

The PS corrects slot omissions and semantic errors. First, it extracts values for all requested slots missing from the AP output. Then, for open-value slots, if a predicted answer exactly matches a value already assigned to a slot in the constrained AP output, the system assumes a semantic error and replaces the predicted slot name with the corresponding omitted one. If no such match is found, the PS appends assignment expressions that bind the predicted answers to the omitted slots and updates the current task (*viz departure\_date*, Figure 1).

## 4 Experimental Setup

We evaluate PyTOD on schema-guided DST. While Zhao et al. (2023) also evaluate their system on next-action prediction (NAP), PyTOD is designed such that, given the correct dialogue state, it is guaranteed to take a correct action according to the dialogue policy. Hence, NAP is highly correlated with DST performance and does not provide additional insight into system behavior. Such turn-based evaluations often fail to reflect a system’s ability to satisfy complex user goals in real interactions (Takanobu et al., 2020; Elizabeth et al., 2024).

### 4.1 Datasets and metrics

**Datasets** We experiment on SGD, which consists of 16,142 training dialogues spanning 26 *service schemata*<sup>5</sup>. The test set comprises 4,201 dialogues covering 21 service schemata. SGD is designed to measure TOD generalization: among the 90 distinct task sequences in the test set, 85.6% involve a task grounded in a schema *unseen* during training, which corresponds to 77% of the dialogues.

**Metrics** We evaluate performance using *joint goal accuracy (JGA)*, which measures the percentage of dialogue turns where all slot-value pairs are predicted correctly. To assess generalization, we report JGA separately for seen and unseen services. This distinction highlights each model’s ability to generalize to unseen slots and values and to correctly interpret API descriptions that were not encountered during training. We compute JGA using the official evaluator<sup>6</sup> but extend the SGD state annotations with canonical slot values (App. C.1).

JGA evaluates context-to-state accuracy but if a slot correctly predicted in an earlier turn is later omitted when the state is re-estimated, the agent may prompt the user to needlessly repeat informa-

tion, degrading interaction quality. We introduce C-JGA, a stricter metric that enforces state consistency: under C-JGA, a turn contributes to the JGA only if the state at all previous turns in the same task is jointly correct.

### 4.2 Baselines

*T5DST* (Lee et al., 2022) jointly encodes the dialogue history and a slot description to generate the corresponding slot value. Each slot is processed independently, so one forward pass is required for each slot in the schema at every turn. *D3ST* (Zhao et al., 2022) predicts all active slot values in a single pass. *SDT-Seq* (Gupta et al., 2022) encode the dialogue history alongside a sample conversation and its state sequence to learn DST via in-context finetuning. *PROMPTER* (Aksu et al., 2023) employ prefix-finetuning for zero-shot domain adaptation. *LLaMA* (Feng et al., 2023) implements schema-guided DST with LLMs, and *LDST* (*ibid.*) is a LLaMA finetuned with diverse, handcrafted prompts. *AnyTOD* is the SOTA E2E agent on SGD.

### 4.3 State tracking with PyTOD

**API retrieval** Consistent with all prior work on SGD, we assume service knowledge at each turn. Hence, PyTOD retrieves APIs from the AP header (§3.1.1) and not the entire assistant schema.

**Action parser** The AP is fine-tuned starting from FlanT5 (Chung et al., 2024), as described in App. C.2. Upon receiving a user utterance, the AP generates program statements which are validated by DM as described in §3.1.

**Schema supervisor** The SS constrains the AP output based on the decoding schema (Figure 3). We use MQA prompts and PLM knowledge for *zero-shot schema supervision* with FlanT5 (3B) (Chung et al., 2024), eliminating the need for slot paraphrase collection. This makes LLM-guided constrained decoding simple to implement.

**Parser supervisor** The PS corrects slot omissions and semantic errors (Figure 4), with finetuning prompts constructed from the same dialogues as those used for AP finetuning. We perform *multitask learning* for action parsing and parsing supervision. This allows PyTOD to function as a single model which performs corrections on DM request, with schema supervision handled by an off-the-shelf LLM. See fine-tuning details in App. C.2.

**Dialogue state prediction** The validated transcripts are executed line-by-line using `pytodlib`’s execution engine (Section 2.2 and App. A.3), which

<sup>5</sup>A service exposes multiple APIs representing user intents.

<sup>6</sup>Available at <https://bit.ly/3B7jD1c>.

Size	Model	JGA	C-JGA	JGA (Seen)	JGA (Unseen)	#
220M	PROMPTER	–	–	–	49.4	1
	T5DST	72.6	–	89.7	66.9	2
	D3ST †	71.2	62.2	93.2	63.8	3
	AnyTOD	–	–	89.9	62.4	4
	SDT-Seq †	77.5	68.7	93.5	72.2	5
780M	D3ST †	76.5	67.7	93.8	70.8	6
	SDT-Seq †	82.7	74.2	94.1	78.9	7
7B	LLaMA*	75.3	–	–	–	8
	LDST*	84.5	–	–	–	9
220M	PyTOD (B)	76.8	72.7	91.0	71.8	10
780M	PyTOD (L)	82.2	78.4	92.1	78.9	11

Table 2: PyTOD DST performance. Rows marked with † report the results of our replication study in App. C.3. \* indicates training with prompt augmentation and evaluation with custom scripts as opposed to official evaluator.

changes the simulation state as described in Table 1. At every user turn, the state of the dialogue is the state of the current API. The API state can be directly accessed from provided evaluation scripts and shares format with the SGD state annotations.

## 5 Main Result

TOD agents are deployed in resource-constrained settings, so we implement the AP and PS components using a small PLM and compare PyTOD against SOTA models of similar size and LLMs. Table 2 shows that PyTOD closely matches or outperforms all baselines, while also achieving higher consistency across turns. In particular, PyTOD (B) surpasses D3ST and T5DST by absolute margins of 5.6% and 4.2%, respectively, with stronger performance on unseen services. PyTOD (L) achieves a similar improvement over D3ST (#6&11).

D3ST and T5DST re-estimate the state from scratch at each turn, meaning their JGA increases if early errors are later corrected. In reality, however, these errors would alter the dialogue flow, so the JGA increases due to error recovery overestimate real-world performance. When adjusted for consistency, D3ST JGA drops sharply: 9% (220M, #3) and 8.8% (780M, #6). In contrast, PyTOD incrementally predicts user actions based on its *predicted* past actions, resulting in smaller JGA drops of 4.1% (B, #10) and 3.8% (L, #11). These smaller decreases stem from two key factors. First, when users change goals mid-task, incorrect states may still be corrected. Second, parameters incorrectly copied from previous tasks at the start of a new task may be later overridden by their correct values.

PyTOD retrieves APIs from the prompt while the JGA of all other models is invariant to intent parsing errors, making the figures less sensitive to anno-

Size	PS	SS	Multitask	JGA	JGA Seen	JGA Unseen	#
220M	✓	✓	✓	76.8	91.0	72.1	1
	✗	✓	✗	75.8	90.1	71.0	2
	✗	✗	✓	64.2	88.2	56.2	3
	✗	✗	✗	64.4	88.8	56.3	4
780M	✓	✓	✓	82.2	92.1	78.9	5
	✗	✓	✗	80.6	91.5	77.0	6
	✗	✗	✓	74.6	90.8	69.2	7
	✗	✗	✗	74.1	90.4	68.6	8

Table 3: Contribution of parser supervisor (PS) and schema supervisor (SS) to PyTOD performance. **Multitask** indicates joint training of the action parsing and parsing supervision tasks. Rows 10&11 from Table 2 are repeated in rows 1&5 to facilitate comparisons.

tation errors (§6.3) but less reflective of real-world TOD performance. PyTOD’s JGA therefore more accurately reflects practical deployment scenarios where intent errors impact interaction quality.

**Descriptions vs demonstrations** Beyond surpassing schema-guided baselines, PyTOD performs competitively with approaches requiring additional developer effort, such as SDT-Seq. The JGA gap between PyTOD and SDT-Seq is minimal: 0.7% (B, #5&10) and 0.5% (L, #7&11), with SDT-Seq performing better on seen domains.

However, SDT-Seq requires developers to manually craft example demonstrations for each intent in addition to service schemas, increasing deployment overhead. Furthermore, SDT-Seq does not perform API retrieval, requiring an external intent detection model. Finally, PyTOD demonstrates substantially higher consistency, outperforming SDT-Seq by 5.0% (#5&10) and 4.2% (#7&11) in C-JGA, reinforcing its stability in multi-turn interactions.

**LLM comparison** PyTOD compares favourably with LLMs. Compared to LLaMA, PyTOD (B) improves JGA by 1.5% (#8&10) while being 32 times smaller. PyTOD (L) improves it by 6.9% while being 10 times smaller (#8&11). These results show that automatic prompt generation given execution errors and policy feedback enables small language models to perform competitively with LLMs while operating at a fraction of a cost. LDST is trained with data generated with multiple prompt templates, outperforming our best model by 2.3% (#9&11). We do not perform augmentation.

## 6 Analysis and Discussion

### 6.1 Ablation study

**Parser supervisor** Table 3 shows the PS improves PyTOD JGA by an average of 1.0% (Base, #1&2) and 1.6% (Large, #5&6). Analysing 210 errors

AP Size	SS Size	PS	JGA	JGA Seen	JGA Unseen	#
220M	3B	✓	76.8 (+5.6%)	91.0	72.1	1
		✗	75.8 (+4.6%)	90.1	71.0	2
	780M	✓	75.1 (+3.9%)	89.9	70.2	3
		✗	71.3 (+0.1%)	88.8	65.4	4
	220M	✓	74.1 (+2.9%)	89.7	68.8	5
		✗	70.8 (-0.4%)	89.0	64.7	6
780M	3B	✓	82.2 (+5.7%)	92.1	78.9	7
		✗	80.6 (+4.1%)	91.5	77.0	8
	780M	✓	81.2 (+4.7%)	91.2	77.9	9
		✗	78.6 (+2.1%)	90.3	74.7	10
	220M	✓	80.5 (+4.0%)	91.2	76.9	11
		✗	77.9 (+1.4%)	90.3	73.7	12

Table 4: PyTOD performance as function of SS size. Checkmarks (✓) and crosses (✗) indicate PS presence or absence, respectively. Rows 10 and 11 from Table 2 are repeated in rows 1 and 7, for easy comparison. Bracketed numbers represent absolute improvements with respect to our implementation of D3ST (Table 2, #3&6)

AP Size	SS Size	Samples per sec	Runtime (sec)	Relative Latency	#
220M	-	9.110	5875.12	1.00	1
	220M	6.816	7919.79	1.35	2
	780M	6.766	7946.26	1.35	3
	3B	6.323	8588.64	1.46	4
780M	-	3.331	16269.70	1.00	5
	220M	2.971	18300.93	1.12	6
	780M	2.964	18341.15	1.13	7
	3B	2.903	18771.32	1.15	8

Table 5: Latency analysis. Numbers reported are for models without PS (marked with ✗ in Table 4). Each figure is an average of three runs of making predictions for  $\approx 53k$  test set turns on an NVIDIA A100 GPU.

corrected by the PS in the best-performing PyTOD (Large) run (82.6% JGA) reveals that slot omissions occur more frequently than semantic errors (61% vs. 39%). The most common semantic error stems from confusion between similar slots (e.g., start date vs. end date). Sharing AP and PS parameters has a negligible effect on parsing accuracy (#3&4, #7&8), simplifying PyTOD deployment.

**Schema supervisor** While the AP generates slot names absent from the decoding schema for unseen services, the predicted slots often retain the correct semantics. For instance, the AP produces new slot names such as *travel\_starts\_from* instead of *journey\_starts\_from*. Other times, the AP outputs slots seen in a training schema implementing the same domain as an (unseen) decoding schema (e.g., outputs *hotel\_name* instead of *place\_name*), reflecting real-world challenges where TOD agents must support integration of services similar to the ones they have been trained on without further fine-tuning. The SS effectively mitigates both of these challenges, improving JGA by 11.4% for PyTOD (Base, Table 3, #2&4) and 6.5% for PyTOD (Large, #6&8). Notably, the SS reduces system latency while improving performance: the 220M AP+SS system (#2) runs 1.89 times faster compared to the

780M AP (#8) while being 1.7% more accurate.

Beyond performance gains, LLM-constrained decoding simplifies deployment compared to grammar-based approaches. Unlike the latter, which require re-engineering to align with new backbone PLM tokenization schemes, LLM-based constrained decoding allows seamless AP updates.

## 6.2 Does SS need to be an LLM?

Table 4 shows that PLM-constrained decoding remains effective even as the SS size decreases: JGA drops by only 2.7% (#1&5) for PyTOD (B) and 1.7% (#7&11) for PyTOD (L) when reducing SS from 3B to 220M parameters, so LLM-guided constraint is feasible with small PLMs. The PS recovers some of the errors, improving JGA by an average of 3.55% (PyTOD B) and 2.6% (PyTOD L). Regardless of the AP, SS and PS configurations PyTOD consistently outperforms D3ST.

**Latency analysis** Table 5 confirms that SS size has a minimal impact on PyTOD latency. Expanding SS from 220M to 3B increases latency by only 11% for PyTOD (Base) and just 2% for PyTOD (Large). This is expected since the SS prompts are short and the MQA formatting enables SS to constrain decoding with a single token. Most of the latency increase stems from on-demand SS model loading, which can be optimized by keeping SS in memory at the cost of a higher memory footprint.

## 6.3 Error analysis

**Seen services** An analysis of 20 dialogues from each of the 3 services with JGA below SDT-Seq average, reveals that *RideSharing\_2* (*RS\_2*) and *Movies\_1* (*MOV\_1*) contribute most to the discrepancy. For *RS\_2*, 70% of the errors involve incorrect slot values: the system consistently misinterprets requests like *cheapest ride* as *regular* instead of *pool*, suggesting that PyTOD could benefit by implementing the AP and PS with a PLM that has stronger world knowledge. For *MOV\_1*, value errors are primarily *span errors*, where the model fails to capture the full movie title or crosses span boundaries. Notably, annotation errors where the annotated slot values are absent from the user utterance predominate in *Travel\_1* samples.

**Unseen services** Annotation errors significantly impact unseen service performance, particularly in *Trains\_1*, where intent paraphrase errors (App. D.1) prevent PyTOD from tracking state due to its reliance on retrieved train schedules. When search tasks succeed, *copy errors* - where PyTOD



Error Type	RideSharing_2	Movies_1	Travel_1	RentalCars_3	Trains_1	Music_3	Messaging_1	Total
	78.2 (-13.9%) ✓	80.7 (-11.4%) ✓	92.5 (+0.4%) ✓	55.3 (-23.6%) ✗	60.2 (-18.7%) ✗	64.0 (-14.9%) ✗	65.7 (-13.2%) ✗	-
Missed	15.0(3)	5.9(2)	9.5(2)	3.7(1)	9.4(3)	20.0(5)	-	16(9.2)
Value	70.0(14)	32.4(11)	-	22.2(6)	9.4(3)	36.0(9)	-	43(24.9)
Copy	15.0(3)	29.4(10)	4.8(1)	48.1(13)	21.9(7)	28.0(7)	78.6(11)	52(30.1)
Annot.	-	17.6(6)	71.4(15)	14.8(4)	40.6(13)	16.0(4)	21.4(3)	45(26.0)
Intent	-	-	-	-	9.4(3)	-	-	3(1.7)
Other	-	14.7(5)	14.3(3)	11.1(3)	9.4(3)	-	-	14(8.1)
<b>Total</b>	<b>20</b>	<b>34</b>	<b>21</b>	<b>27</b>	<b>32</b>	<b>25</b>	<b>14</b>	<b>173(100)</b>

Table 6: Errors of the best PyTOD (L) run, for seen (✓) and unseen (✗) services. Second row shows service JGA and its absolute deviation with respect to the average seen/unseen JGA in Table 2 (row 11). Percentages indicate the fraction of total service errors and (raw counts) are shown. **Missed**=omitted slot, **Value**=incorrect value, **Copy**=copied incorrect value/did not copy value, **Annot.**=annotation err., **Intent**=incorrect intent.

fails to propagate query parameters (e.g., *album* from *LookupMusic* to *PlayMedia*) reduce JGA in *Movie\_3* and *RentalCars\_3*. For example, while *GetCarsAvailable* succeeds 90% of the time, *ReserveCar* succeeds in just 26.23% of cases. This highlights systematic failures which could be addressed through targeted improvements.

In *Messaging\_1*, copy errors stem from incorrect co-reference resolution: the *location* slot is resolved to cities instead of addresses due to training schema bias: *location* always describes cities in training. TOD agents that support multiple services are prone to similar errors when concept names differ across training schemas, and fine-tuning alone can lead to overfitting to schema-specific naming conventions. Future work will explore enhanced supervision mechanisms to mitigate these biases.

## 7 Related Work

**Few- and zero-shot DST** Our approach and baselines (§4.2) extend prior transfer learning work on cross-domain DST generalization via schema descriptions (Lin et al., 2021b) and QA tasks (Lin et al., 2021a). Recently, large-scale proprietary LLMs (e.g., ChatGPT (OpenAI, 2022)) have shown strong DST performance with no (Heck et al., 2023) or few (Li et al., 2024; Hu et al., 2022; Feng et al., 2023; Wu et al., 2023, *inter alia*) training dialogues. Like PyTOD, these approaches predict state updates, represented in JSON format (Wu et al., 2023), as SQL queries (Hu et al., 2022), function calls (Stricker and Paroubek, 2024; Li et al., 2024) or code (King and Flanigan, 2023). Unlike PyTOD, they are not policy-guided nor do they operate in a simulated environment; instead they track slots from user and agent utterances - a limitation that degrades performance (Wu et al., 2023).

These methods focus on data-efficient DST, relying on LLMs at inference. Recent works (Kulkarni et al., 2024; Finch and Choi, 2024) use LLMs for data generation, improving generalization without costly per-turn inference and addressing concerns over cost, resource availability, and privacy (Heck et al., 2023; Feng et al., 2023). Distillation (Lee et al., 2024c; Dong et al., 2024) and data augmentation (Feng et al., 2023; Li et al., 2024) provide alternatives, though they still depend on billion-scale models. To improve efficiency, subsequent works (Lee et al., 2024a,b) employ self-correction (Xie et al., 2022; Tian et al., 2021) to maintain accuracy. PyTOD targets stricter resource constraints, achieving accurate and consistent DST with models an order of magnitude smaller. It eschews corrector fine-tuning or feedback generation, using readily available execution feedback and simple QA prompts for zero-shot correction with one token. Its policy- and environment-driven corrections avoid per-turn supervision, improving efficiency.

**TOD agents** Transfer learning (Zhao et al., 2023) and prompting (Shu et al., 2022; Zhang et al., 2023; Li et al., 2023; Saley et al., 2024, *inter alia*) also apply to TOD agent development. `pytodlib` provides challenges beyond MultiWOZ through policy and ontology complexity (Hudeček and Dusek, 2023; Lee et al., 2024a; Gao et al., 2024), providing a testbed for advancing TOD agents while addressing the dearth of conversational tool-use corpora (Lu et al., 2024).

## 8 Conclusion

We introduced PyTOD, a model that generates code incrementally, tracks state through execution, and uses policy and execution feedback to automatically generate error-correcting prompts. By coupling state tracking with execution, PyTOD shows enhanced cross-turn consistency and thus improved real-world reliability. We release `pytodlib`, a simulation grounding SGD, to advance research on zero-shot TOD agents and conversational tool use. Future work will focus on enhancing PyTOD’s robustness to copy errors and applying it to LLM-based zero- and few-shot E2E dialogue modeling.

### Limitations

**API retrieval** Like most state-of-the-art approaches on the SGD and MultiWOZ datasets, PyTOD assumes knowledge of the service schema at each turn (but not the API). In real-world sce-

narios, however, virtual assistants must first infer the schema before tracking dialogue state. This is particularly challenging in SGD, where multiple services within the same domain exhibit fine-grained differences. For instance, Buses\_1 and Buses\_2 both implement the FindBus APIs, yet disambiguation between the two services is only possible if users mention fare\_type, an *optional* slot. Similarly, the Buses\_3 service (from the test set) can only be distinguished if the category slot is provided. As discussed in §6.3, PyTOD relies on accurate intent parsing, making it susceptible to service disambiguation errors that will degrade performance.

**Schema robustness** We have not explicitly evaluated PyTOD’s robustness to linguistic schema variations, which are known to affect transfer learning-based DST systems. While AP accuracy will degrade, PyTOD’s SS - the primary contributor to its performance (§6.1) - is expected to mitigate this impact since it performs zero-shot corrections using MQA prompts. Future work could enhance AP and PS robustness using transfer learning from QA task (Lin et al., 2021a; Cho et al., 2023), knowledge-seeking turn grounding (Coca et al., 2023a) or synthetic schemas (Coca et al., 2023b), none of which require additional annotation.

**Prompt optimisation** While PS and AP latency are optimized through dialogue history truncation and minimal generation length, the AP prompt itself remains unoptimised. The AP header already maintains a stack of completed tasks and retrieved entities, effectively summarising the dialogue history. However, our prompts contain both the header and transcript, introducing redundancy and increasing system latency.

Additionally, the API header presents an opportunity for personalised conversational intelligence. Embedding and storing it in a vector database could enable retrieval across multi-session conversations, improving continuity and personalisation. We will explore this in future work.

**Grammar-constrained decoding** We opted for a deep learning approach to constraining AP output to the schema due to the complexity of working with dynamically generated grammar rules needed to constrain decoding to a set of valid tokens. We considered this advanced optimisation a research topic for future work.

**Interactive evaluation** Our results demonstrate that state-of-the-art DST models exhibit stability errors when evaluated using C-JGA. As DST systems improve, we believe evaluating models in real-world user interactions will be essential for assessing their practical viability.

## Acknowledgments

Alexandru Coca’s doctoral studies are supported by the EPSRC grant EP/R513180/1. He would like to acknowledge generous financial support from the Cambridge Philosophical Society through the Research Studentship. The authors thank Ankit Bhattarai from the University of Cambridge for his contribution to the SDT-Seq replication study and evaluation. Xiaofeng Wu and Stephen Pulman from Apple alongside anonymous reviewers are acknowledged for their improvement suggestions.

## References

- Ibrahim Taha Aksu, Min-Yen Kan, and Nancy Chen. 2023. [Prompter: Zero-shot adaptive prefixes for dialogue state tracking domain adaptation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4588–4603, Toronto, Canada. Association for Computational Linguistics.
- Hyundong Cho, Andrea Madotto, Zhaojiang Lin, Khyathi Chandu, Satwik Kottur, Jing Xu, Jonathan May, and Chinnadhurai Sankar. 2023. [Continual dialogue state tracking via example-guided question answering](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 3873–3886, Singapore. Association for Computational Linguistics.
- Hyundong Cho, Chinnadhurai Sankar, Christopher Lin, Kaushik Sadagopan, Shahin Shayandeh, Asli Celikyilmaz, Jonathan May, and Ahmad Beirami. 2022. [Know thy strengths: Comprehensive dialogue state tracking diagnostics](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 5345–5359, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. [Scaling instruction-finetuned language models](#). *J. Mach. Learn. Res.*, 25:70:1–70:53.

- Alexandru Coca, Bo-Hsiang Tseng, Jinghong Chen, Weizhe Lin, Weixuan Zhang, Tisha Anders, and Bill Byrne. 2023a. [Grounding description-driven dialogue state trackers with knowledge-seeking turns](#). In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 444–456, Prague, Czechia. Association for Computational Linguistics.
- Alexandru Coca, Bo-Hsiang Tseng, Weizhe Lin, and Bill Byrne. 2023b. [More robust schema-guided dialogue state tracking via tree-based paraphrase ranking](#). In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1443–1454, Dubrovnik, Croatia. Association for Computational Linguistics.
- Xiaoyu Dong, Yujie Feng, Zexin Lu, Guangyuan Shi, and Xiao-Ming Wu. 2024. [Zero-shot cross-domain dialogue state tracking via context-aware auto-prompting and instruction-following contrastive decoding](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 8527–8540, Miami, Florida, USA. Association for Computational Linguistics.
- Michelle Elizabeth, Morgan Veyret, Miguel Couceiro, Ondrej Dusek, and Lina Maria Rojas-Barahona. 2024. [Do large language models with reasoning and acting meet the needs of task-oriented dialogue?](#) *CoRR*, abs/2412.01262.
- Yujie Feng, Zexin Lu, Bo Liu, Liming Zhan, and Xiao-Ming Wu. 2023. [Towards llm-driven dialogue state tracking](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 739–755. Association for Computational Linguistics.
- James D. Finch and Jinho D. Choi. 2024. [Diverse and effective synthetic data generation for adaptable zero-shot dialogue state tracking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 12527–12544, Miami, Florida, USA. Association for Computational Linguistics.
- Haoyu Gao, Ting-En Lin, Hangyu Li, Min Yang, Yuchuan Wu, Wentao Ma, Fei Huang, and Yongbin Li. 2024. [Self-explanation prompting improves dialogue understanding in large language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC/COLING 2024, 20-25 May, 2024, Torino, Italy*, pages 14567–14578. ELRA and ICCL.
- Raghav Gupta, Harrison Lee, Jeffrey Zhao, Yuan Cao, Abhinav Rastogi, and Yonghui Wu. 2022. [Show, don't tell: Demonstrations outperform descriptions for schema-guided task-oriented dialogue](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 4541–4549. Association for Computational Linguistics.
- Michael Heck, Nurul Lubis, Benjamin Ruppik, Renato Vukovic, Shutong Feng, Christian Geishauer, 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.
- Yushi Hu, Chia-Hsuan Lee, Tianbao Xie, Tao Yu, Noah A. Smith, and Mari Ostendorf. 2022. [In-context learning for few-shot dialogue state tracking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 2627–2643. Association for Computational Linguistics.
- Vojtěch Hudeček and Ondrej Dusek. 2023. [Are large language models all you need for task-oriented dialogue?](#) In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 216–228, Prague, Czechia. Association for Computational Linguistics.
- Léo Jacqmin. 2022. [« est-ce que tu me suis ? » : une revue du suivi de l'état du dialogue \(“do you follow me ?” : a review of dialogue state tracking\)](#). In *Actes de la 29e Conférence sur le Traitement Automatique des Langues Naturelles. Volume 2 : 24e Rencontres Etudiants Chercheurs en Informatique pour le TAL (RECITAL)*, pages 1–19, Avignon, France. ATALA.
- Brendan King and Jeffrey Flanigan. 2023. [Diverse retrieval-augmented in-context learning for dialogue state tracking](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 5570–5585, Toronto, Canada. Association for Computational Linguistics.
- Atharva Kulkarni, Bo-Hsiang Tseng, Joel Ruben Antony Moniz, Dhivya Piraviperumal, Hong Yu, and Shruti Bhargava. 2024. [SynthDST: Synthetic data is all you need for few-shot dialog state tracking](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1988–2001, St. Julian's, Malta. Association for Computational Linguistics.
- Chia-Hsuan Lee, Hao Cheng, and Mari Ostendorf. 2024a. [Correctionlm: Self-corrections with SLM for dialogue state tracking](#). *CoRR*, abs/2410.18209.
- Chia-Hsuan Lee, Hao Cheng, and Mari Ostendorf. 2024b. [OrchestraLLM: Efficient orchestration of language models for dialogue state tracking](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1434–1445, Mexico City, Mexico. Association for Computational Linguistics.



- Harrison Lee, Raghav Gupta, Abhinav Rastogi, Yuan Cao, Bin Zhang, and Yonghui Wu. 2022. [SGD-X: A benchmark for robust generalization in schema-guided dialogue systems](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 10938–10946. AAAI Press.
- Seanie Lee, Jianpeng Cheng, Joris Driesen, Alexandru Coca, and Anders Johannsen. 2024c. [Effective and efficient conversation retrieval for dialogue state tracking with implicit text summaries](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 96–111, Mexico City, Mexico. Association for Computational Linguistics.
- Zekun Li, Zhiyu Chen, Mike Ross, Patrick Huber, Seungwhan Moon, Zhaojiang Lin, Xin Dong, Adithya Sagar, Xifeng Yan, and Paul A. Crook. 2024. [Large language models as zero-shot dialogue state tracker through function calling](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 8688–8704. Association for Computational Linguistics.
- Zekun Li, Baolin Peng, Pengcheng He, Michel Galley, Jianfeng Gao, and Xifeng Yan. 2023. [Guiding large language models via directional stimulus prompting](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Zhaojiang Lin, Bing Liu, Andrea Madotto, Seungwhan Moon, Zhenpeng Zhou, Paul Crook, Zhiguang Wang, Zhou Yu, Eunjoon Cho, Rajen Subba, and Pascale Fung. 2021a. [Zero-shot dialogue state tracking via cross-task transfer](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7890–7900, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, Zhou Yu, Andrea Madotto, Eunjoon Cho, and Rajen Subba. 2021b. [Leveraging slot descriptions for zero-shot cross-domain dialogue StateTracking](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5640–5648, Online. Association for Computational Linguistics.
- Jiarui Lu, Thomas Holleis, Yizhe Zhang, Bernhard Aumayer, Feng Nan, Felix Bai, Shuang Ma, Shen Ma, Mengyu Li, Guoli Yin, Zirui Wang, and Ruoming Pang. 2024. [Toolsandbox: A stateful, conversational, interactive evaluation benchmark for LLM tool use capabilities](#). *CoRR*, abs/2408.04682.
- Johannes E. M. Mosig, Shikib Mehri, and Thomas Kober. 2020. [STAR: A schema-guided dialog dataset for transfer learning](#). *CoRR*, abs/2010.11853.
- OpenAI. 2022. [Introducing ChatGPT](https://openai.com/blog/chatgpt). <https://openai.com/blog/chatgpt>.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2019. [Towards scalable multi-domain conversational agents: The schema-guided dialogue dataset](#). *CoRR*, abs/1909.05855.
- Abhinav Rastogi, Xiaoxue Zang, Srinivas Sunkara, Raghav Gupta, and Pranav Khaitan. 2020. [Schema-guided dialogue state tracking task at DSTC8](#). *CoRR*, abs/2002.01359.
- Vishal Vivek Saley, Rocktim Jyoti Das, Dinesh Raghu, and Mausam . 2024. [Synergizing in-context learning with hints for end-to-end task-oriented dialog systems](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 5596–5612, Miami, Florida, USA. Association for Computational Linguistics.
- Raphael Shu, Elman Mansimov, Tamer Alkhouli, Nikolaos Pappas, Salvatore Romeo, Arshit Gupta, Saab Mansour, Yi Zhang, and Dan Roth. 2022. [Dialog2api: Task-oriented dialogue with API description and example programs](#). *CoRR*, abs/2212.09946.
- Joe Stacey, Jianpeng Cheng, John Torr, Tristan Guigue, Joris Driesen, Alexandru Coca, Mark Gaynor, and Anders Johannsen. 2024. [LUCID: LLM-generated utterances for complex and interesting dialogues](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 4: Student Research Workshop)*, pages 56–74, Mexico City, Mexico. Association for Computational Linguistics.
- Armand Stricker and Patrick Paroubek. 2024. [A few-shot approach to task-oriented dialogue enhanced with chitchat](#). In *Proceedings of the 25th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 590–602, Kyoto, Japan. Association for Computational Linguistics.



Ryuichi Takanobu, Qi Zhu, Jinchao Li, Baolin Peng, Jianfeng Gao, and Minlie Huang. 2020. [Is your goal-oriented dialog model performing really well? empirical analysis of system-wise evaluation](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue, SIGdial 2020, 1st virtual meeting, July 1-3, 2020*, pages 297–310. Association for Computational Linguistics.

Xin Tian, Liankai Huang, Yingzhan Lin, Siqi Bao, Huang He, Yunyi Yang, Hua Wu, Fan Wang, and Shuqi Sun. 2021. [Amendable generation for dialogue state tracking](#). In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 80–92, Online. Association for Computational Linguistics.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *CoRR*, abs/1910.03771.

Yuxiang Wu, Guanting Dong, and Weiran Xu. 2023. [Semantic parsing by large language models for intricate updating strategies of zero-shot dialogue state tracking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023, Singapore, December 6-10, 2023*, pages 11093–11099. Association for Computational Linguistics.

Hongyan Xie, Haoxiang Su, Shuangyong Song, Hao Huang, Bo Zou, Kun Deng, Jianghua Lin, Zhihui Zhang, and Xiaodong He. 2022. [Correctable-DST: Mitigating historical context mismatch between training and inference for improved dialogue state tracking](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 876–889, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Xiaoying Zhang, Baolin Peng, Kun Li, Jingyan Zhou, and Helen Meng. 2023. [SGP-TOD: building task bots effortlessly via schema-guided LLM prompting](#). *CoRR*, abs/2305.09067.

Jeffrey Zhao, Yuan Cao, Raghav Gupta, Harrison Lee, Abhinav Rastogi, Mingqiu Wang, Hagen Soltau, Izhak Shafran, and Yonghui Wu. 2023. [AnyTOD: A programmable task-oriented dialog system](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 16189–16204, Singapore. Association for Computational Linguistics.

Jeffrey Zhao, Raghav Gupta, Yuan Cao, Dian Yu, Mingqiu Wang, Harrison Lee, Abhinav Rastogi, Izhak Shafran, and Yonghui Wu. 2022. [Description-driven task-oriented dialog modeling](#). *CoRR*, abs/2201.08904.

## A pytodlib

We describe key pytodlib features, encouraging interested readers to consult our software release.

### A.1 Service APIs

Figure 5 shows a sample implementation of the Buses\_3 service from SGD. It provides two APIs: (1) FindBus, a *search* or *query* API (Figure 5a) that enables users to query a bus schedule database using natural language, and (2) BuyBusTicket, a *transactional* API (Figure 5b) that allows users to purchase a ticket for an itinerary proposed by the TOD agent based on FindBus search results or by specifying ticket details directly.

```
@register_command(service="Buses_3")
class FindBus(SearchCommand):
    from_city: SearchCommandArgument[str] = SearchCommandArgument()
    to_city: SearchCommandArgument[str] = SearchCommandArgument()
    departure_date: SearchCommandArgument[str] = SearchCommandArgument()
    num_passengers: SearchCommandArgument[str] = SearchCommandArgument()
    category: SearchCommandArgument[str] = SearchCommandArgument()

    def __init__(self, dialogue_id: DialogueID):
        super().__init__(dialogue_id)
```

(a) Search API.

```
@register_command(service="Buses_3")
class BuyBusTicket(ConfirmedCommand):
    from_city: ConfirmedCommandArgument[str] = ConfirmedCommandArgument()
    to_city: ConfirmedCommandArgument[str] = ConfirmedCommandArgument()
    departure_time: ConfirmedCommandArgument[str] = ConfirmedCommandArgument()
    departure_date: ConfirmedCommandArgument[str] = ConfirmedCommandArgument()
    num_passengers: ConfirmedCommandArgument[str] = ConfirmedCommandArgument()
    additional_luggage: ConfirmedCommandArgument[str] = ConfirmedCommandArgument()

    def __init__(self, dialogue_id: DialogueID):
        super().__init__(dialogue_id)
```

(b) Transactional API.

Figure 5: Implementation of the SGD Buses\_3 service

The APIs define slots as class variables implementing the *descriptor protocol*<sup>7</sup>. This enables advanced functionality such as *execution error feedback*: setting an undefined attribute on an API could return a string describing the error instead of raising an exception. Another example is *type coercion*: descriptors cast slot values to the data types specified in the schema, and may be configured to provide natural language feedback if the conversion fails, providing yet another opportunity for generating error-correction prompts. We leave these experiments to future work.

### A.2 Policy simulation

The SGD conversations are generated by sampling from a *policy graph* (Rastogi et al., 2019; Mosig et al., 2020), which outlines the intended flow of a dialogue. Both search and transactional APIs require zero or more specific slots to function. To provide them, the system processes the initial user turn and takes actions to elicit missing slot values. Once all required slots are filled, search APIs can construct a valid database query, while transactional APIs execute an external service call (e.g., to a ticket booking service).

<sup>7</sup>See <https://bit.ly/4hLB4aU>.

```

1 def maybe_parse_api_call_or_action(
2     program_statement: str
3 ) -> Command | str | None:
4     """Parse an action parser output to a python API or action name."""
5     def string_to_py_cmd(ast_expr: ast.Expr) -> Command | str:
6         match ast_expr:
7             # matches a fcn call e.g.,
8             # buses_3_find_buses(Location="Cambridge")
9             case ast.Call(ast.Name() as func, args, keywords):
10                command_name = func.id # "buses_3_find_buses"
11                match command_name:
12                    # if the command is not a user/system action
13                    # it must be an API call
14                    case tool if tool not in USR_SYS_ACTIONS:
15                        # build FindBus command from the schema
16                        command = _build_command(command_name)
17                        # set slot values
18                        _set_command_properties(command, args, keywords)
19                        return command
20                    case _:
21                        # return the action name, parsed downstream
22                        return command_name
23
24 match ast.parse(program_statement).body:
25     # matches function calls e.g. find_bus(), next(x1), etc
26     case [ast.Expr(value=ast.Call() as expr)]:
27         api_py_obj = string_to_py_cmd(expr)
28         return api_py_obj
29     case _:
30         return

```

Figure 6: Sample execution engine code, showing how a program statement is interpreted as a python object.

Slot-filling is abstracted in the Command interface<sup>8</sup>, and the SearchCommand<sup>9</sup> interface (implementing search use cases) and the ConfirmedCommand<sup>10</sup> interface (implementing transaction-specific policies) inherit. All concrete service implementations (e.g., Buses\_3) extend the search and/or transaction interfaces. Upon execution, the interfaces return system actions, such as show, perform, Hint and Notification. For example, if a user says, *I need a bus from London to Manchester.*, executing FindBus returns a system action Hint(request value: departure\_date)<sup>11</sup>, which an agent can verbalise to ask for the missing constraints. We encourage the interested reader to consult the detailed examples in Table 12 and revisit Figure 1 in §1, Figure 2 in §3.1.2 and Figure 9 in App. B.1 to deepen their understanding of the agent-environment interaction protocol pytodlib implements.

Unlike a majority of examples in the STARv2 (Zhao et al., 2023) dataset, where users provide constraints only when prompted and the system follows a purely deterministic behaviour, SGD conversations follow a complex policy graph, that frequently includes user-initiated constraint specification and extends beyond the standard slot-filling approach in MultiWOZ. We encourage readers to explore the documentation of the simulation package in our software supplementary.

<sup>8</sup>See src/simulation/command.py

<sup>9</sup>See src/simulation/search\_command.py

<sup>10</sup>src/simulation/confirmed\_command.py

<sup>11</sup>The other arguments in Figure 5 are optional, so are not requested by the system.

your task is to identify which of the following tasks the user wants to complete, and update it accordingly as the user provides more information during the conversation. you should also carefully follow any developer instructions to answer user questions and complete follow-up tasks.

(a) Task instruction

```

1. considered one-way air travel options for their itinerary.
the 'flights_3_search_oneway_flight' api returned a
'oneway_flight' object (x11).

class oneway_flight:
    """object returned by the 'flights_3_search_oneway_flight' api.
    in subsequent api calls, you should pass, to compatible
    arguments, the last value mentioned for the following entity
    properties:
    properties
    -----
    origin_city: str
        city in which the journey originates
    destination_city: str
        city in which the journey ends
    departure_date: str
        date of departure flight
    airlines: enum
        name of airline operating the flight
    passengers: int
        number of passengers to find flight seats for
    flight_class: enum
        fare class of flight booking
    number_checked_bags: int
        number of bags to check in
    """

your task is to identify ...

```

(b) Task stack

Figure 7: Action parser header prompt components.

### A.3 Execution engine

The execution engine converts the DM’s output into a Python object. Figure 6 illustrates how a program statement is interpreted as an API object implemented by a service (e.g., FindBus in Figure 5a). The execution process begins by parsing the statement into an AST (line 24). If the tree matches a function call signature (line 26), a conversion function first extracts the function name and arguments (lines 6–10). If the function corresponds to an API implemented by the schema (line 14), the engine retrieves the command from a registry, instantiates the appropriate object (e.g., FindBus), and assigns slot values (lines 15–18). For user or system actions, the function name is returned (lines 20–22), and object instantiation is handled downstream. Interested readers should consult the documentation of the execution package in our software release for further details.

## B Prompts

### B.1 Action parser

#### B.1.1 Header

Figure 7 presents a sample AP prompt header. If no task has been completed, the prompt begins with task instructions (Figure 7a). Once a task is completed, a summary of the task and the returned

```

prompt:
user: weather on 11th of this month in london, england.
0 weather_1_get_weather(date = '11th of this month', city = 'london, england')
1 show(x0)
2 next(x0) # type: forecast
developer: the user may request specific information about the 'forecast' object (x2)
properties listed below. pass the relevant property or properties to 'say' to answer
(eg, say(x2.precipitation, x2.humidity)).
- precipitation: the possibility of rain or snow in percentage
- humidity: percentage humidity
- wind: wind speed in miles per hour
- temperature: temperature in fahrenheit
- date: date for the weather
agent: 71 degrees 15 percent chance
user: how humid though? windy?
3

target:
say(x2.humidity, x2.wind)

```

(a) Post-iteration context-dependent instructions.

```

1 ...
2 agent: alright, so you'd like to send $1,470 from your checking account
3 to the checking account belonging to yumi?
4 user: yes. how many days will that take?
5 confirm(x7)
6 developer: in the event of a 'banks_2_transfer_money' failure,
7 copy the values of the keywords below the agent last mentioned
8 to subsequent 'banks_2_transfer_money' calls:
9 - recipient_account_type
10 developer: unless a signal indicates a 'banks_2_transfer_money' calling
11 error, the properties
12 - transfer_time: number of days for the transfer to go through
13 may be communicated to the user upon their request by referencing
14 'x13' while calling 'say' (eg, say(x13.transfer_time)).

```

(b) Post-confirmation context-dependent instructions.

Figure 8: Sample context-dependent instructions.

entity is prepended before the task instructions (Figure 7b). The entity definitions include a docstring instructing the model to copy relevant argument values to the parameters of subsequent API calls.

## B.1.2 Session transcript

This section provides additional insight regarding the generality of PyTOD programs and system actions.

**Handling compositional utterances** Figure 9 demonstrates how PyTOD can handle compositional user utterances requiring nested function calls. At turn 9, the user declines alternative restaurant booking arrangements proposed by the agent, requesting directions to their next meeting. To resolve the meeting location, the agent must retrieve user’s schedule from their calendar, sort it and then read the location of the next event (x29 - x31).

Such utterances cannot be represented by AnyTOD or other state-of-the-art TOD agents and DST models, which are limited to parsing slot-value pairs from user and agent utterances.

**System actions** §3.1.2 introduced the perform statement as a marker of successful API execution. An example of its syntax and use is shown in Figure 9 at turn 9. By passing the reference of the variable bound to this statement to the subsequent say call, the NLG module can generate an utterance which informs the user their task executed successfully. This is equivalent to a NOTIFY\_SUCCESS in

```

agent: For two, 7pm at Nando's Oval, tomorrow?
[7] user: Lemme think about it ...
x15 conversation_pause()
x16 Hint("prompt for next action")
x17 say(x17)
agent: Take your time, I'll be here!
[8] user: Go ahead actually!
x18 confirm(x12)
x19 Notification("Booking did not succeed")
x20 Hint("offer alternative: date")
x26 say(x19, x20)
agent: No luck with the booking for today. Tomorrow same time?
[9] user: Nah, leave it. Directions to my next meeting, quick,
just noticed the time!
x27 decline_alternative()
x28 Hint("prompt for next action")
x29 get_meeting(sort_by="date")
x30 next(x29)
x31 get_directions(to=x30.location)
x32 perform(x31)
x33 say(x31, x32)
agent: Sure! It's in 3 minutes in the ARR building, room April.
[10] user: K, gotta dash! Bye!
x34 say()
agent: Have a productive one!

```

Figure 9: An alternative continuation of the conversation in Figure 2 after turn 6. The perform system action an example of how PyTOD parses compositional utterances (turn 9) are depicted. Unlike in Figure 2 where some of the say calls were omitted for clarity, all calls to the NLG module are shown.

the DSTC8 dialogue act tagset.

See Table 12 for a summary of all PyTOD system and user actions along with further examples.

## B.1.3 Context-dependent instructions

Context dependent instructions are formatted as **developer:** turns, appear in the prompt after an iteration (i.e., next) or confirmation instructions. In the former case (Figure 8a), they provide an itemized list of entity properties that the user may request along with their natural language descriptions. A brief instruction precedes this list, prompting the language model to invoke the say routine to communicate the requested information to the user. In the latter case (Figure 8b), additional system policy instructions relevant to state tracking may be included. For instance, line 6-9 in Figure 8b illustrate how PyTOD can be guided to correctly parse API parameters following an API calling error.

## B.2 Dialogue manager

### B.2.1 Schema supervisor

Figure 10 shows the SS prompt generation templates. A common input to these is slot\_schemas, a list of dictionaries containing the names, descriptions, data type and possible values (for categorical slots) for the active service. These are formatted by developer-defined filters (e.g., slot\_definition\_formatter). The template for constraining the value of a categorical slot to one of

```

1 Given the definitions, which keyword below best matches
2 {{ predicted_argument_with_definition }}
3
4 {% for slot in slot_schemas -%}
5 - {{ loop.index | int2alpha }} {{ slot | slot_definition_formatter }}
6 {%if loop.last -%}
7 - {{ (loop.index + 1) | int2alpha }} none:
8 the definitions do not describe {{ predicted_argument }}
9 {% endif -%}
10 {%- endfor %}
11 Answer:

```

(a) Unknown slot name. Sample prompt in Figure 3a.

```

1 Here are some definitions:
2 {% for slot in slot_schemas %}
3 - {{ slot | slot_definition_formatter }}
4 {%- endfor %}
5 { set cnt = [0] %}
6 Given these, {{ predicted_keyword }} is a synonym of:
7 {% for slot in slot_schemas -%}
8 {% for value in slot.possible_values -%}
9 {% if cnt.append(cnt.pop() + 1) %}{{ endif %}
10 - {{ cnt[0] | int2alpha }} {{ slot.name }} {{ "=" }} {{ value | lower }}
11 {%- endfor -%}
12 {%if loop.last %}
13 - {{ (cnt[0] + 1) | int | int2alpha }} \
14 options do not describe {{ predicted_keyword }}
15 {% endif -%}
16 {%- endfor %}
17 Answer:

```

(b) Unknown slot name (closed value). Sample prompt in Figure 3b.

```

1 Here are some definitions:
2 - {{ arg_def.name }}: {{ arg_def.description | lower }}
3
4
5 Given these, {{ predicted_keyword }} is a synonym of:
6 {% for value in arg_def.possible_values %}
7 - {{ loop.index | int2alpha }} \
8 {{ arg_def.name }} {{ "=" }} {{ value | lower }}
9 {%- endfor %}
10
11 Answer:

```

(c) Unknown categorical slot value

```

1 Which sentence below paraphrases {{ memorised_argument_description_with_name }}?
2
3 {% for slot in slot_schemas -%}
4 - {{ loop.index | int2alpha }} {{ slot | slot_definition_formatter }}
5 {%if loop.last -%}
6 - {{ (loop.index + 1) | int2alpha }} \
7 none of the sentences above paraphrases {{ memorised_argument_name }}
8 {% endif -%}
9 {%- endfor %}
10 Answer:

```

(d) Memorised slot name. Sample prompt in Figure 3c.

Figure 10: Schema supervisor prompt templates

the values listed in the schema (Figure 10c) is a special case of the template for constraining unknown slot names with categorical values (Figure 10b): only one slot definition corresponding to the one predicted by the AP is displayed and the none option is removed to ensure the output is one of the values enumerated in the schema.

### B.3 Parser supervisor

Figure 11 shows the PS prompt generation template. The DM filters the transcript to extract **user** and **agent** turns relevant to the current task, ignoring previous tasks which are not relevant for predicting the current slot value. These turns are processed by the `conversation_formatter` filter, which prepends the conversation role to the utterance. The schemata of the slots requested at the previous turned are passed to the template as `slot_list`. The `question_formatter` filter the formats the slot descriptions by lowercasing them and appending a question mark.

```

1 Q: Answer the following questions. Output "unanswerable" \
2 if the question cannot be answered given the conversation.
3
4 In the conversation:
5
6 {{source_turns | conversation_formatter }}
7
8 {%- for item in slot_list %}
9 {{ loop.index }}) {{ item | question_formatter }}
10 {%- endfor %}
11
12 Answer:

```

Figure 11: Parser supervisor template

## C Experimental Details

### C.1 Slot values normalisation

In SGD, the dialogue state updates when the user either explicitly provides a slot value or accepts a system-proposed value. Traditional TOD agents track the latter by extracting slot-value pairs from agent utterances. In contrast, as described in rows 3 and 9 of Table 1 (in §3.1.2), PyTOD updates the state by executing `select` and `confirm` commands. These commands read relevant slot values from entities retrieved via database queries (for search-based interactions) or API responses (for transactional interactions). Since SGD system call annotations<sup>12</sup> are canonicalised, PyTOD must normalise open-valued parameters extracted from the dialogue history before making API calls and ensure that system-proposed slot values are de-normalised for evaluation.

In practice, normalisation is performed by looking up the surface form of a predicted value in a mapping that links surface forms to their canonical counterparts. This table is easily constructed from SGD semantic annotations, as illustrated in Figure 12. Instead of de-normalising slot values copied from entities or API responses, we equivalently extend the state annotations to include their canonical forms. This ensures that slots tracked via execution are directly comparable to the reference values used by the DSTC8 evaluator.

### C.2 PyTOD implementation

We finetune PyTOD with the hyperparameters in Table 7 until the development set JGA for *unseen* services is maximised. The learning rate is constant, with no scaling. We follow the same protocol when finetuning `google/flan-t5-large`, except that we allocate a training budget of just one epoch. All reported results are an average of three runs

<sup>12</sup>These include call parameters, entity properties, and API responses.



```

1 "frames": [
2   {
3     "actions": [
4       {
5         "act": "INFORM",
6         "canonical_values": [
7           "11:30"
8         ],
9         "slot": "time",
10        "values": [
11          "half past 11 in the morning"
12        ]
13      }
14    ],
15    "service": "Restaurants_1",
16    "slots": [
17      {
18        "exclusive_end": 49,
19        "slot": "time",
20        "start": 22
21      }
22    ],
23    "state": {
24      "active_intent": "ReserveRestaurant"
25      "requested_slots": [],
26      "slot_values": {
27        "time": [
28          "half past 11 in the morning",
29          "11:30",
30        ],
31        ...
32      }
33    }
34  }
35 ]

```

Figure 12: Semantic frame for the utterance *I would like it to be half past 11 in the morning*. The action annotations (line 3 - 14) are processed to extend slot-value annotation with the corresponding canonical value (e.g., line 29). The ellipsis in line 31 marks slot-value pairs which were omitted for clarity.

with different random seeds.

**Action Parser (AP)** The AP is fine-tuned to learn user action prediction and communicative system action prediction based on a dataset derived from the 16, 142 dialogues in the SGD training set. After filtering examples that exceed the maximum source length of 2048 tokens, the final training set contains 296, 013 examples.

**Parser Supervisor (PS)** In contrast to the SS, the PS is fine-tuned to handle slot omissions and semantic errors. The fine-tuning dataset for the PS is therefore constructed by creating 34, 105 extractive QA prompts similar to the one in Figure 4 from the SGD training split. Several strategies are employed during data creation.

First, the dialogue history in each prompt is filtered to include only turns relevant to the active service, focusing the model on the immediate context. Second, the order of questions is randomized

Hyperparameter	Value
Pretrained model	google/flan-t5-base
Optimizer	Adafactor
Batch size	32
Learning rate	0.0001
Warm-up steps	1500
Number of epochs	2
Evaluation frequency	1500 steps

Table 7: PyTOD training hyperparameters.

to prevent the model from learning any positional biases.

Finally, the prompts are augmented with negative examples to train the model to avoid hallucination. With a probability of 0.5, each prompt includes questions for slots whose values are not mentioned in the context. To train the model to make fine-grained distinctions, there is a subsequent 50% chance that one of these unanswerable questions will probe for a slot that is semantically similar to one already present (e.g., asking for a check-out date when the user has only provided a check-in date).

The PS shares its parameters with the AP. This allows PyTOD’s core logic to be handled by a single, small pre-trained language model, which requests corrections from an LLM only when validation fails. This design, mediated by the Dialogue Manager, balances performance with the inference costs associated with larger models.

The data pre-processing scripts are included in the linked code release.

### C.3 Replication study

We finetune D3ST and SDT-Seq using the transformers (Wolf et al., 2019) library (v4.35.2). With the exception of FlanT5 (Chung et al., 2024) D3ST, which was trained on two NVIDIA A100 GPUs (80GB), all models were trained on a single NVIDIA RTX 3090 GPU (24GB).

#### C.3.1 D3ST

We set all training parameters to match those reported by Zhao et al. (2022) and pre-process the data using their official script<sup>13</sup>. As the original work does not specify a model selection metric or evaluation frequency, we evaluate every 5, 000 steps and select the best checkpoint based on overall JGA on the development set. Training is terminated early if accuracy does not improve within 15, 000 steps (approximately 3 epochs).

<sup>13</sup>Available at <https://bit.ly/4aKe9KL>.

Size	Model	JGA	JGA Seen	JGA Unseen	Setting	#
220M	D3ST (Zhao et al., 2022)	72.9	92.5	66.4	-	1
	D3ST (Flan-T5, ours)	71.2	93.2	63.8 (65.0/64.8/61.7)	A	2
	D3ST (Flan-T5, ours)	70.7	92.9	63.3 (61.9/65.3/62.7)	B	3
780M	D3ST (Zhao et al., 2022)	80.0	93.8	75.4	-	4
	D3ST (Flan-T5, ours)	76.5	93.8	70.8 (69.9/71.6/70.9)	A	5

Table 8: D3ST replication results. Numbers in brackets show the metric values for each experiment run, three-runs averages are shown otherwise.

Our results differ by an absolute 1.7% from the published JGA for the base model (Table 8, rows 1 & 2). We observe a 0.8% improvement on seen services, but a 2.6% drop on unseen services.

To rule out overfitting, we increase the evaluation frequency to 900 steps and select the model maximising the *unseen services* JGA, stopping the training after 1 epoch of no improvement. However, this leads to a slight performance drop (Table 8, rows 2 & 3). Hence, finetune Flan-T5 (780M) with the best settings, achieving seen services performance on par with the published results but a 4.6% discrepancy on unseen services (rows 4 & 5).

### C.3.2 SDT-Seq

We set training hyperparameters to the values reported by Gupta et al. (2022), using the data processing scripts from the official SDT code release<sup>14</sup>. Since the evaluation frequency and model selection metric are unspecified, we evaluate every 1600 steps, selecting the model with the highest development set overall JGA. We closely replicate the reported results (Table 9).

Size	Model	JGA	JGA Seen	JGA Unseen	#
220M	SDT-Seq (Gupta et al., 2022)	76.3	-	-	1
	SDT-Seq (Flan-T5, ours)	77.5	93.5	72.2	2
780M	SDT-Seq (Gupta et al., 2022)	83.3	-	-	3
	SDT-Seq (Flan-T5, ours)	82.7	94.1	78.9	4

Table 9: SDT-Seq replication results. The reported numbers are averaged over five runs, each using a distinct set of demonstrations to construct the fine-tuning prompts.

## D Extended Analysis and Discussion

### D.1 Annotation errors

Table 10 presents sample errors identified in our analysis in §6.3, showing intent paraphrase errors in #2&3. While our DM performs argument-based disambiguation to identify intent, *GetTrainTickets* and *FindTrains*, the two *Trains\_1* intents, share all their arguments. Consequently, misparaphrased intent annotations prevent PyTOD from retrieving

<sup>14</sup>Available at <https://bit.ly/4aKe9KL>.

#	Utterance	Annotation	Explanation	Service
1	Today at 2 in the afternoon.	<i>pickup_time='2 in the afternoon'</i>	<i>start_date='today'</i> in utterance	RC_3
2	I'm in the mood for some music and would like to play some songs.	<i>intent=LookUpMusic</i>	The utterance semantics is better represented as <i>intent=PlayMedia</i>	MUS_3
3	I need a train ticket with a fully refundable feature.	<i>intent=FindTrains class='Flexible'</i>	The utterance semantics is better represented as <i>intent=GetTrainTickets</i> .	TR_1
4	Okay, what about attractions there. I need Place of Worship, and something with no entry fee.	<i>intent=FindAttractions category="Place of Worship" free_entry=True good_for_kids=True</i>	No <i>good_for_kids=True</i> mention in utterance.	TRA_1

Table 10: Sample annotation errors identified during the error analysis. *RC\_3=RentalCars\_3*, *MUS\_3=Music\_3*, *TR\_1=Trains\_1*, *TRA\_1=Travel\_1*.

train schedules, leading to degraded DST performance. Analysis of 30 additional dialogues from this domain, we found that intent confusion caused state errors in 20 out of 50 cases.

Table 10 highlights some utterances contain slots mentions without corresponding user action annotations (#1) while others sometimes fail to paraphrase actions parametrised by boolean or categorical slots (#4). These issues were difficult to identify with the methods available to Rastogi et al. (2019). During PyTOD development we, however, identified 451 dialogues across the train, development and test releases<sup>15</sup> while experimenting with fault-tolerant execution. This suggests that programmable dialogue systems, while challenging to develop, can serve as valuable tools for dataset quality improvements.

### D.2 C-JGA breakdown

To support future comparisons, we provide a breakdown of C-JGA metrics for our models and replicated baselines across seen and unseen services.

Size	Model	C-JGA	C-JGA (Seen)	C-JGA (Unseen)	#
220M	D3ST (Flan-T5, ours)	62.2	86.0	54.3	1
	SDT-Seq (Flan-T5, ours)	68.7	86.6	62.8	2
780M	D3ST (Flan-T5, ours)	66.5	87.9	61.0	3
	SDT-Seq (Flan-T5, ours)	74.2	88.0	69.6	4
220M	PyTOD (Base)	72.7	87.3	67.8	5
780M	PyTOD (Large)	78.4	89.1	74.9	6

Table 11: Breakdown of C-JGA from Table 2.

<sup>15</sup>We include dialogue IDs in our code release.

## E Annotated Session Transcript Examples

Context	Transcript
<p><b>U<sub>1</sub></b>: Please make a payment to Mary.</p> <p><b>Transcript explanation:</b></p> <ul style="list-style-type: none"> <li>- The AP parses turns where the user introduces a new task to <i>draft</i> API calls (<code>make_payment</code>).</li> <li>- The SS is triggered if API call parameters are not in the schema (e.g., <i>recipient</i> is generated instead of <i>receiver</i>).</li> <li>- All transcript entries are bound to a variable (e.g., <code>x1</code>) so that they can be referenced by other instructions.</li> <li>- Unless all required arguments are known, API execution returns possible system actions, termed <i>slot filling hints</i>. To make a payment, the user must provide a recipient, amount and specify if they wish to make a private transaction or not. No value is known for the last two slots so executing the API recommends these two actions.</li> <li>- PyTOD executes a recommended action by passing its corresponding variable to the NLG call. Here it chooses to request the payment amount so <code>x2</code> is passed to <code>say</code>. Equivalently, it could have chosen to ask the user if they wish the transaction to be private (in which case <code>say(x3)</code> would have been generated) or requested both slots at once (case in which <code>say(x2, x3)</code> would have been generated).</li> </ul>	<pre>x1 make_payment(receiver="Mary") x2 Hint("Ask user to provide: amount") x3 Hint("Ask user to provide: private_transaction") x4 say(x2)</pre>
<p><b>S<sub>2</sub></b>: How much?  <b>U<sub>2</sub></b>: Pay her sixty dollars.</p> <p><b>Transcript explanation:</b></p> <ul style="list-style-type: none"> <li>- The system utterance <b>S<sub>2</sub></b> is the output of the <code>say(x2)</code> call in the previous turn. We do not consider end-to-end evaluation so in our evaluation we use the ground truth system and user utterances from the corpus.</li> <li>- The AP parses the user utterance to an assignment instruction, which is PyTOD's slot filling mechanism.</li> <li>- Since <code>x1</code> is bound to <code>make_payment</code>, the assignment updates the properties of the draft command.</li> <li>- For efficiency, the action recommendations and <code>say</code> are not shown in the prompt for user action parsing so the dialogue context for assignment prediction is <b>U<sub>1</sub></b>, <b>x1</b>, <b>S<sub>2</sub></b> and <b>U<sub>2</sub></b> (along with the AP header (§3.1.1), as shown in Figure 1).</li> <li>- Given the <code>say(x2)</code> call, the DM anticipates the user will provide the amount in the next turn and uses this information for verification. Missing information would trigger a PS call.</li> <li>- After the update, the <code>private_transaction</code> required argument is still not known, so when the assignment is executed, the API recommends a system action (<code>x6</code>) to find the value of this parameter.</li> </ul>	<pre>x5 x1.amount="sixty dollars" x6 Hint("Ask user to provide: private_transaction") x7 say(x6)</pre>
<i>New conversation (Restaurant Search &amp; Booking)</i>	
<p><b>U<sub>1</sub></b>: I'm after a Chinese in central London.</p> <p><b>Transcript explanation:</b></p> <ul style="list-style-type: none"> <li>- When all the required arguments to the <code>find_restaurant</code> search query are known, execution returns a list of restaurants matching the user request, which is bound to <code>x1</code>.</li> <li>- The dialogue manager signals that search returned some results by inserting <code>show(x1)</code> in the transcript. The only positional argument of <code>show</code> is a reference to the API call that returned the results.</li> <li>- Given <b>U<sub>1</sub></b>, and the transcript lines <code>x1</code> and <code>x2</code>, PyTOD generates the instructions bound to <code>x3</code> and <code>x4</code>.</li> <li>- <code>next(x1)</code> fetches the top returned entity (a <code>Restaurant</code> object) and binds it to a variable (<code>x3</code>).</li> <li>- Entities can be passed to NLG calls as positional arguments (<code>x4</code>).</li> </ul>	<pre>x1 find_restaurant(area="centre", city="London", cuisine="Chinese") x2 show(x1) x3 next(x1) x4 say(x3)</pre>
<p><b>S<sub>2</sub></b>: There are many options, the first is Golden Wok, 500 ft away.  <b>U<sub>2</sub></b>: What else?</p> <p><b>Transcript explanation:</b></p> <ul style="list-style-type: none"> <li>- When called with an entity variable, the NLG generates an utterance that provides the user with information relevant in context.</li> <li>- When then user requests an alternative, the results list is iterated using a <code>next</code> instruction (<code>x5</code>) and the result is communicated to the user (<code>x6</code>).</li> </ul>	<pre>x5 next(x1) x6 say(x5)</pre>
<p><b>S<sub>3</sub></b>: The Burnt Kidney?  <b>U<sub>3</sub></b>: The first one sounds like a winner.</p> <p><b>Transcript explanation:</b></p> <ul style="list-style-type: none"> <li>- The AP generates the <code>select</code> statement when the user selects one of the alternatives proposed by the agent</li> <li>- The <code>select</code> positional argument is bound to a search result fetched with <code>next</code>.</li> <li>- <code>from_results</code> is a reference to the query that returned the list of results.</li> <li>- <code>select</code> execution returns two system action recommendations: a <i>suggestion</i> (<code>x8</code>) and a <i>task prompting hint</i> (<code>x9</code>).</li> <li>- <code>x8</code> binds a <code>book_restaurant</code> draft command updated with relevant details from history (e.g., <code>restaurant_name</code>, <code>location</code>).</li> <li>- The agent chooses to pass the suggestion to the NLG module (<code>x10</code>).</li> </ul>	<pre>x7 select(x3, from_results=x1) x8 suggest(task="book_restaurant") x9 Hint("Prompt user for next task") x10 say(x8)</pre>
<p><b>S<sub>5</sub></b>: Book it now?  <b>U<sub>5</sub></b>: Not now. Tell me my checking balance.</p> <p><b>Transcript explanation:</b></p> <ul style="list-style-type: none"> <li>- The AP generates a <code>suspend</code> instruction when the user declines a suggestion.</li> <li>- <code>suspend</code> takes variables bound to a <code>suggest</code> statement as input.</li> <li>- In this case, the AP generates the <code>suspend</code> instruction and <code>check_balance</code> API call at once.</li> <li>- <code>check_balance</code> only requires the <code>account</code> argument, so executing it returns a list of accounts (<code>x12</code>), which are signalled to the agent (<code>x13</code>).</li> <li>- Like before, PyTOD generates the instructions <code>x14</code> and <code>x15</code> at once.</li> <li>- The top result, an <code>Account</code> entity, is fetched from results (<code>x14</code>).</li> <li>- The agent passes the <code>Account</code> entity to the NLG module to provide the user with the relevant information (<code>x15</code>).</li> </ul>	<pre>x11 suspend(x6) x12 check_balance(account="checking") x13 show(x12) x14 next(x12) x15 say(x14)</pre>

Continued on next page

Table 12 – continued from previous page

Context	Transcript
<p><b>S<sub>6</sub>:</b> There are \$10,535.55.  <b>U<sub>6</sub>:</b> Okay, let's book a table there! Next Monday, at 7.</p> <p><b>Transcript explanation:</b>  - The AP parses the last user utterance to two actions, bound to x16 and x17.  - When the user refers to a task they previously suspended, a resume statement whose only positional argument is a variable bound to a suspend statement is generated.  - Executing a resume statement returns a draft API call updated with relevant parameters from the dialogue history: a book_restaurant instance with known restaurant_name and location in this case (which had been bound to x11 when the task was suspended).  - Resumed tasks may be updated through assignments (x17).  - When all the arguments for a transaction are known, draft API call execution returns a confirmation hint (x18).  - The agent requests the user to confirm by passing the confirmation hints to the NLG module (x19).</p>	<pre>x16 resume(x11) x17 x16.date="Next Monday"; x16.time="7" x18 Hint("Confirm: restaurant_name, date, time") x19 say(x18)</pre>
<p><b>S<sub>7</sub>:</b> Please confirm your Golden Wok booking, 7th April at 7 pm.  <b>U<sub>7</sub>:</b> You got it!</p> <p><b>Transcript explanation:</b>  - When the user confirms a transaction, the AP generates a confirm statement whose positional argument is a variable bound to the draft API.  - confirm calls execute the API. Task completion failures are indicated by Notification messages (x21).  - If transactions can be completed given an alternative set of constraints, an alternative_hint system action indicating possible alternatives is returned by the command (x22).  - Task status and alternatives are communicated to the user (x23).</p>	<pre>x20 confirm(x16) x21 Notification("Booking did not succeed") x22 Hint("Alternative: time=9:15pm") x23 say(x21, x22)</pre>
<p><b>S<sub>8</sub>:</b> Only 9:15 is available. Happy with that?  <b>U<sub>8</sub>:</b> I'm afraid not.</p> <p><b>Transcript explanation:</b>  - The decline_alternative instruction is generated by the AP to mark negation.  - Executing decline_alternative returns task prompting hints (x25) which are communicated to the user (x26).</p>	<pre>x24 decline_alternative() x25 Hint("Prompt user for next task") x26 say(x25)</pre>
<p><b>S<sub>9</sub>:</b> Anything else I can do for you?  <b>U<sub>9</sub>:</b> Let me think...</p> <p><b>Transcript explanation:</b>  - The AP generates conversation_pause instructions to mark out-of-domain utterances or situations when the user does not take an action.</p>	<pre>x27 conversation_pause() x28 Hint("Prompt user for next task") x29 say(x28)</pre>
<i>New conversation (Search &amp; Play a Movie)</i>	
<p><b>U<sub>1</sub>:</b> I want to watch a Disney movie.</p> <p><b>Transcript explanation:</b>  - slice is analogous to show (see above), but binds a small number of entities the agent should preset to the user.  - when called with a slice reference, say produces an utterance where multiple entities are mentioned.</p>	<pre>x1 find_movie(producer="Disney") x2 slice(x1) x3 say(x2)</pre>
<p><b>S<sub>2</sub>:</b> Let's see ... Snow White, Frozen 1, or Frozen 2?  <b>U<sub>2</sub>:</b> Good taste, let's go with Frozen 2.</p> <p><b>Transcript explanation:</b>  - select can take an entity name (e.g., movie) as a keyword argument instead of a positional argument when multiple entities are mentioned in a single utterance.</p>	<pre>x4 select(movie="Frozen 2", from_results=x1) x5 suggest(task="play_movie") x6 Hint("Prompt user for next task") x7 say(x5)</pre>
<p><b>S<sub>3</sub>:</b> Play it?  <b>U<sub>3</sub>:</b> On the living room TV!</p>	<pre>x8 play_movie(movie="Frozen 2", device="TV", area="living room") x9 Hint("Confirm: area, living room") x10 say(x9)</pre>
<p><b>S<sub>4</sub>:</b> Confirm Frozen 2, in the living room?  <b>U<sub>4</sub>:</b> Yes! When did it come out?</p> <p><b>Transcript explanation:</b>  - perform takes as a sole argument a variable bound to a transactional intent, and is inserted in the transcript by the DM to mark successful transaction completion.  - An entity (e.g., SelectedMovie) is bound to the transaction variable (x8) upon successful execution.  - Passing variables bound to perform (x12) to the NLG enables the agent to communicate task success to the user.  - Information requested by the user is communicated by passing the entity to the NLG call and specifying the attribute that should be communicated (x13). This is possible since the entity definitions are displayed in the prompt after confirmations, as shown in Figure 8b.</p>	<pre>x11 confirm(x8) x12 perform(x8) x13 say(x12, x8.release_year)</pre>
<p><b>S<sub>5</sub>:</b> Playing your movie, released in 2019.  <b>U<sub>5</sub>:</b> Thanks, have a good one!</p> <p><b>Transcript explanation:</b>  - The AP generates say with no arguments to mark conversation end.</p>	<pre>x14 say()</pre>

Table 12: Further session transcript examples, listing the entire set of PyTOD system and user actions along with their semantics. Explanations provide detailed, step-by-step explanations of transcript generation. As the conversation advances, the user and system utterances and the transcript updates of all previous turns are concatenated to form the context for predicting the transcript updates given the current turn (formed of the last system and user utterances).