

Turning Conversations into Workflows: A Framework to Extract and Evaluate Dialog Workflows for Service AI Agents

Prafulla Kumar Choubey Xiangyu Peng Shilpa Bhagavath
Caiming Xiong Shiva Kumar Pentyala Chien-Sheng Wu

Salesforce AI Research

Correspondence: pchoubey@salesforce.com

Abstract

Automated service agents require well-structured workflows to deliver consistent and accurate responses to customer queries. However, such workflows are often undocumented, and their automatic extraction from conversations remains largely unexplored. In this work, we present a novel framework for extracting and evaluating dialog workflows from historical interactions. Our extraction process involves two key stages: (1) a retrieval step to select relevant conversations based on key procedural elements, and (2) a structured workflow generation step using question-answer-based chain-of-thought (QA-CoT) prompting. To comprehensively evaluate the quality of the extracted workflows, we introduce an automated simulation framework with agent and customer bots that measures their effectiveness in resolving customer issues. Extensive experiments on the ABCD and SynthABCD datasets show that our QA-CoT technique improves workflow extraction by 12.16% in average macro accuracy over the baseline. Moreover, our evaluation method closely aligns with human assessments, offering a reliable and scalable framework for future research.

1 Introduction

LLMs are transforming customer service by enabling automated bots to respond accurately and consistently to customer inquiries. These bots rely on well-structured workflows to guide interactions, ensure efficient issue resolution, and maintain adherence to company policies (Brown, 2023). However, high-quality workflows are scarce, making it difficult to ensure agent consistency and quality.

Traditional workflow creation methods typically depend on manual efforts (Bohus and Rudnick, 2009; Mosig et al., 2020; Mehri and Eskenazi, 2021; Chen et al., 2021b; Zhao et al., 2023; Robino, 2025), which are time-consuming, prone to becoming

User: My email address is salma.hussein@example.com, my order ID is 789012, and my username is salmah123.

Agent: Thank you for the information. Let me check the system for any errors related to your purchase. One moment, please.

Agent: I have verified the details, and it appears that the error is valid. I will process a refund for you. Is there anything else I can assist you with today?

Agent: It appears that the error is not valid. Could you please tell me your membership level?

Workflow snippet:

- If the error is valid, offer (or process) to refund.
- If the error is not valid (customer is wrong), ask the customer for their membership level.

Figure 1: An example showing the derivation of a workflow from historical conversations. Full workflow is shown in Figures 3 and 5.

ing outdated, and often fail to capture the full complexity of real-world interactions. Importantly, procedural knowledge is implicitly embedded in past customer-agent conversations. For example, Figure 1 illustrates a workflow segment derived from two conversations involving different customer scenarios. Yet, extracting and structuring this information remains an underexplored challenge. Automating workflow extraction from conversations offers a promising solution by transforming implicit knowledge into structured, reusable guidelines for both human agents and AI systems.

In this paper, we propose a novel two-step pipeline for automatically extracting dialog workflows from customer-agent interactions. These conversations often contain noise, such as agents combining multiple steps or sequencing them incorrectly. Moreover, customer interactions typically cover only a subset of the established rules, making it difficult to capture a complete workflow. To address these challenges, our pipeline first enhances the retrieval process by identifying key procedural components (Figure 2)—including intents, slot values, and resolution steps—and leveraging them to select the most relevant conversations while filtering out non-compliant ones.

In the second step, we introduce a structured QA-based chain-of-thought (QA-CoT) prompting

```

"intent": "Customer wants to refund an order before shipping.",
"slot_values": {
  "full_name": "Chloe Zhang",
  "username": "czhang94",
  "email": "czhang94@email.com",
  "order_id": "1553732700",
  "refund_method": "original credit card",
  "shipping_status": "not shipped"
},
"resolution_steps": [
  "Agent asked for the customer's full name.",
  "Agent requested the username, email, and order ID.",
  "Agent asked how the customer would like the refund processed.",
  "Agent inquired about the account ID.",
  "Agent asked about the shipping status of the order."
]

```

Figure 2: An example of procedural elements extracted from a conversation by the GPT-4o-mini LLM.

technique to systematically extract workflows from large conversation datasets. Extracting workflows at scale is challenging, as it requires capturing all critical decision points and ensuring their correct order within the workflow. Our QA-CoT approach addresses this by simulating an interactive exchange between a Guide and an Implementer agent. The Guide examines key aspects of the workflow, including preconditions, decision points, and step rationales, while the Implementer responds with insights derived from historical conversations. This structured QA process enables the LLM to focus on each decision point in isolation, ensuring that all essential procedural elements are accurately reflected in the final workflow.

Another key contribution of our work is a robust multi-step evaluation pipeline for dialog workflows, featuring an end-to-end methodology that closely aligns with human assessments. We simulate interactions between a service agent bot (equipped with the predicted workflow and relevant system data) and a customer bot (provided with the intent and user-specific details). After the conversation concludes, we evaluate task success by comparing the outcome produced using the predicted workflow with the expected outcome derived from the ground truth workflow (see Figure 5b for an example).

We evaluate our proposed QA-CoT method against standard prompting strategies using the ABCD (Chen et al., 2021b) dataset and an LLM-generated variant, SynthABCD. While ABCD originally grounds conversations in predefined dialog workflows, we repurpose it for the task of extracting workflows from conversations. To eliminate human-induced errors present in ABCD conversations, we generate SynthABCD using LLMs, ensuring error-free dialogues that adhere strictly to the same workflows. This enables controlled and

reliable evaluation under an idealized setting.

Our QA-CoT method outperforms other prompting strategies across multiple LLMs, including GPT-4o, Opus, Sonnet, and Gemini. For example, on the ABCD dataset, it achieves an 11.81% improvement over the baseline with GPT-4o. Moreover, incorporating explicit QA-CoT reasoning enhances the performance of state-of-the-art reasoning models such as OpenAI o1, o1-mini, o3-mini, and DeepSeek-R1 (Guo et al., 2025) across both datasets.

2 Related Work

Dialog workflows are a specialized form of procedural knowledge. While workflow extraction has received little attention, automatic procedural extraction has been widely studied, primarily focusing on “how-to” documents (Schumacher et al., 2012; Maeta et al., 2015; Chu et al., 2017; Park and Nezhad, 2018) and instructional videos (Ushiku et al., 2017; Zhou et al., 2018; Xu et al., 2020a). These works typically model linear sequences of explicitly stated actions, aiming to either predict procedural steps or generate summaries of task execution (Koupae and Wang, 2018). In contrast, our work tackles dialog workflows, where actions are often implicit and depend on previous steps, user inputs, and system responses. This introduces decision-dependent variability, making extraction significantly more challenging than predicting fixed procedural sequences.

Specific to dialog systems, there has been extensive research on studying structures in task-oriented dialogs (Jurafsky and Shriberg, 1997; Chotimongkol, 2008; Shi et al., 2019; Xu et al., 2020b; Chen et al., 2021a; Nath and Kubba, 2021; Wang et al., 2021; Rony et al., 2022; Lu et al., 2022; Qiu et al., 2022; Yin et al., 2023; Pryor et al., 2024; Burdisso et al., 2024), focusing on how dialogs evolve using dialog acts, intent-slot pairs, or turn-level dependencies. Our work parallels workflow discovery (Hattami et al., 2023; Raimondo et al., 2023; Min et al., 2023), which aims to predict the optimal next dialog action from the conversation’s current state and all available actions. However, unlike this, we focus on extracting global workflows applicable across conversations for a specific intent. This increases the complexity of the task, as the model must filter out noisy actions, and consolidate multiple potential actions sequences from different conversations. Additionally, we also propose a new

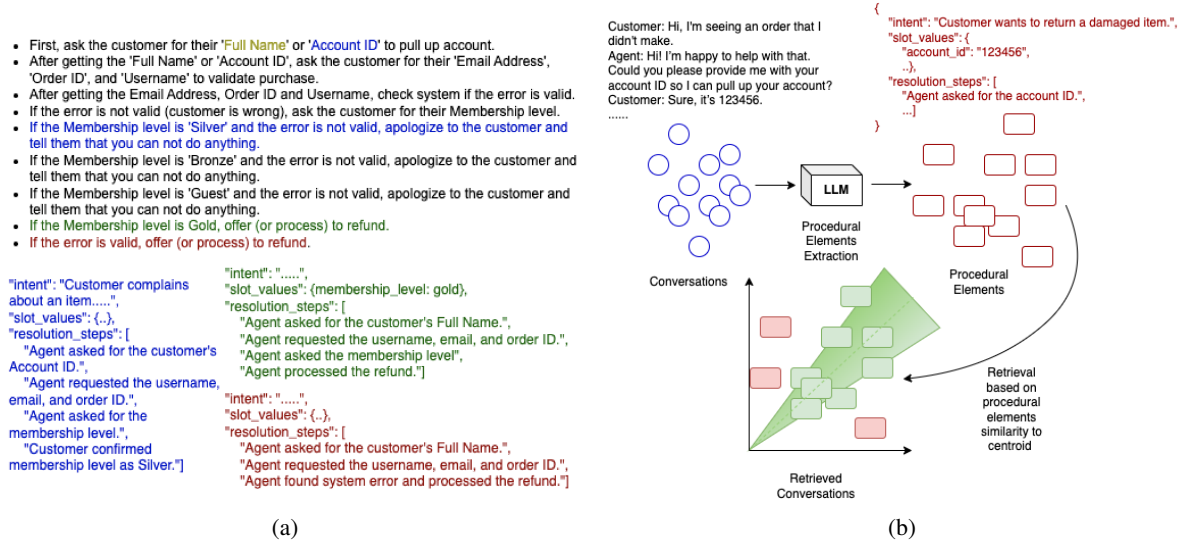


Figure 3: (a) An example workflow with procedural elements for 3 of the 10 possible distinct sub-flows. Sub-flows and their procedural elements are color-coded to match. (b) Steps in our proposed procedural element-based retriever. An example of complete procedural elements extracted from a conversation is shown in Figure 2.

```

..
Guide: What additional details are required to validate the customer's purchase?
Implementer: The agent must ask for the username, email address, and order ID to validate the purchase.
Guide: How does the customer's membership level impact the return process?
Implementer: The customer's membership level determines the return policy duration:
- Guest membership requires returns within 30 days with receipt.
- Bronze membership allows returns within 90 days or with original packaging.
- Silver membership permits returns within six months.
- Gold members have unlimited time for returns.
Guide: What should the agent do if the return is past the allowable period for the customer's membership level?
Implementer: If the return period has passed, the agent should check if the customer has the original receipt and packaging. If not, they politely inform the customer that a return cannot be processed.
..

```

Figure 4: A snippet of the QA chain-of-thought generated by the GPT-4o model for the *return_color* intent using conversations from the ABCD data. The extracted QA pairs highlight key preconditions based on membership level.

QA-CoT prompting method for workflow extraction and introduce a robust end-to-end evaluation framework to assess the accuracy of the extracted workflows.

3 Dialog Workflow Extraction

The dialog workflow extraction task aims to identify and organize all essential steps required to resolve customer issues across various scenarios into structured guidelines. For example, Figure 3a illustrates a workflow for handling customer complaints about a bill they never purchased, while Figure 5a shows the corresponding state machine diagram, highlighting different sub-flows. These sub-flows account for scenarios such as assisting gold, silver, guest, or bronze members; addressing cases with or without system errors; and accommodating customers who provide either their full name or account details. In this example, the goal is to generate a workflow that comprehensively captures all variables, including membership levels, the pres-

ence of system errors, and the type of customer information available.

We focus on two key aspects to improve dialog workflows extraction: retrieval, which involves selecting the most relevant historical conversations for the task, and extraction, which uses a structured QA-CoT prompting approach to derive workflows from the selected conversations.

3.1 Procedural Element-based Retrieval

Our goal is to extract all valid step sequences for resolving customer issues, ensuring comprehensive scenario coverage while filtering out incorrect or noisy conversations. Since our focus is on the underlying process instead of surface-level dialogs, we propose a two-step selection strategy. First, we extract key procedural elements—such as intent, relevant slot values, and the resolution steps taken by the agent—from each conversation using GPT-4o-mini (example in Figures 2 and 3a, prompt in Figure 10). This ensures selection is based on

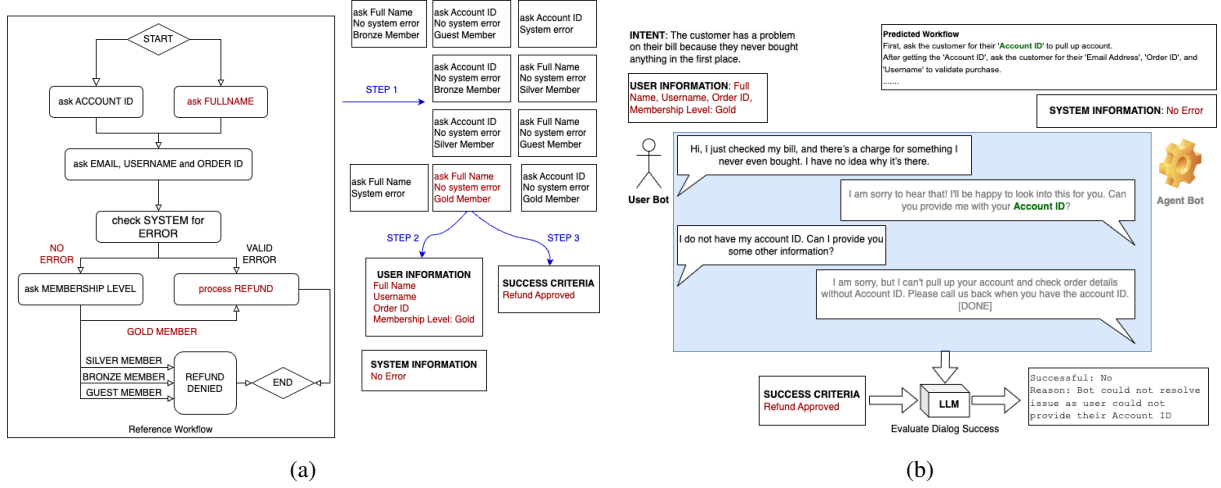


Figure 5: (a) Flowchart of the workflow in Figure 3a, illustrating 10 possible customer scenarios [Step 1, E2E pipeline], along with an example of user information, system information, and the success criteria for one scenario [Steps 2 and 3]. (b) Simulation of the user (system) bot based on intent and user information (system information and the predicted workflow) [Step 4], followed by a final evaluation of dialogue success [Step 5].

procedural content rather than surface-level dialog.

Next, we embed each conversation using its extracted key procedural elements with the OpenAI ‘text-embedding-3-small’ model, then we cluster and compute the centroid for each intent. We then select the top-K conversations closest to each centroid using cosine similarity, ensuring the chosen conversations represent the most commonly followed resolution steps (Figure 3b). By focusing on key procedural elements rather than full conversations, we retain meaningful process information while filtering out noise. As we demonstrate in Section 6.1, directly clustering full conversation embeddings captures extraneous details and conversational variability, leading to significantly lower performance compared to procedural element embeddings.

While this approach effectively captures frequent resolution steps, it may reduce diversity. Given the inherent noise in conversations, distinguishing between true procedural variations and spurious deviations remains a challenge for LLMs. Our strategy prioritizes common procedural patterns while minimizing noise, enabling more reliable workflow extraction.

3.2 QA-CoT for Workflows Extraction

A naive approach to extracting workflows is to use simple prompting with retrieved conversations. However, this often results in workflows that lack fine-grained details. For instance, in the example workflow shown in Figure 3a, the generated workflow might rigidly follow a sequence by checking

for errors first and then asking the membership level even when a system error occurs.

To address these limitations, we propose QA-CoT, a structured framework inspired by multi-agent debate and collaboration systems (Du et al., 2023; Zhang et al., 2024; Musumeci et al., 2024; Hu et al., 2025). In QA-CoT, a Guide-Implementer interaction generates QA pairs (example in Figure 4, prompt in Figure 18) to capture fine-grained procedural knowledge. The Guide formulates targeted questions based on past conversations, focusing on clarifying preconditions, decision points, and logical dependencies. For instance, in the example workflow shown in Figure 3a, the Guide would explicitly ask: “What is the next step if there is a system error?” and “What is the next step if there is no system error?” The Implementer, using the same past conversations, provides detailed responses, such as “asking for the membership level” or “processing a refund” based on the scenario. To improve workflow extraction, we augment historical conversations with the Guide-Implementer exchange and prompt the LLM (Figure 19) to generate workflows that incorporate fine-grained procedural details.

We experimented with two approaches for simulating this interaction. The first is multi-turn prompting, where the Guide and Implementer take turns. The Guide asks a question, and the Implementer provides a response. This process continues until the Guide has no further clarification questions or a maximum of 25 turns is reached. The

second approach generates the entire exchange in a single pass. Empirically, the single-pass method yields superior results with GPT-4o (see Table 4), while also being computationally efficient. Given our large-scale experiments across 8 LLMs and two datasets, we adopt the single-pass strategy.

4 Dialog Workflow Evaluation

Evaluating workflow requires careful attention to the relationships and dependencies between constituent steps to ensure the predicted workflow aligns with the expected decision-making process. We encountered significant challenges in performing this evaluation at the workflow level, primarily due to the complex dependencies between steps that are difficult to assess consistently. In an initial study, we asked human annotators to evaluate whether each step in the ground-truth workflows was accurately represented in the predicted workflows, using a set of 10 LLM-extracted workflows. The results showed moderate inter-annotator agreement, with a Cohen’s κ score of 0.45, highlighting the difficulties even humans face in identifying nuanced procedural differences.

To address these challenges, we propose an end-to-end (E2E, Figure 5) evaluation framework that works in five stages¹:

- *Decomposing Workflows into Sub-flows*: We first decompose each workflow into its constituent sub-flows. As shown in Figure 5a (Step 1), this process involves extracting all 10 possible paths from START to END, which result from variations in membership levels, the type of user input (full name or account ID), and the occurrence of system errors. This approach ensures comprehensive coverage of all decision branches in our evaluation (prompt shown in Figure 11).
- *Generating User and Agent Bot Information*: Next, we map each sub-flow to the relevant user and system information required for bot simulation. In the subflow highlighted in red in Figure 5a (Step 2), the user bot should provide only their full name, username, order ID, and membership level (e.g., silver), while the agent bot has access to internal systems that confirm no error occurred on the company’s end (prompt shown in Figure 13).

¹We also explore alternative evaluation strategies such as QA-based evaluation (Fabbri et al., 2021). These strategies proved less effective than our proposed E2E approach and are discussed in detail in Appendix D.

- *Defining Success Criteria*: We further link each sub-flow to a success criterion, which represents the expected outcome of the conversation if the agent follows the ground-truth workflow. In the sub-flow highlighted in red in Figure 5a (Step 3), the agent is expected to approve the user’s refund request (prompt shown in Figure 13).
- *Dialog Simulation*: Next, we simulate interactions between a customer bot and a service agent bot, where the customer bot conveys intent and provides user-specific details, while the agent bot utilizes system data and acts according to the predicted workflow. The customer bot responds strictly to explicit agent’s requests or seeks alternatives when unable to provide the requested information. Meanwhile, the agent bot executes the predicted workflow by requesting data (e.g., asking for a username), performing actions (e.g., issuing refunds, checking system information), or terminating the conversation if no further steps are possible. The interaction ends when the agent either completes the workflow successfully or cannot proceed due to missing information (see Figure 5b, prompts in Figures 15 and 16).
- *Success Evaluation*: Once the conversation ends, we evaluate whether the simulated interaction meets the success criteria. For instance, in Figure 5b, the interaction fails because the user bot was unable to provide the account ID (prompt shown in Figure 14).

We evaluated the accuracy of our E2E simulation pipeline and found that the overall accuracy across steps 1 to 4 was 94.81%. Moreover, step 5 exhibited strong alignment with human evaluation, achieving a Cohen’s Kappa score of 0.92. These results suggest that our E2E evaluation closely mirrors human judgments (for details, see Appendix B).

5 Experimental Settings

5.1 Dataset

We use two datasets in our evaluations: the human-developed ABCD and the LLM-synthesized SynthABCD. The ABCD dataset is inherently noisy (Appendix C.3), as conversations were generated through human interactions between a customer and a support agent. This process introduced inconsistencies, such as customers providing information before being prompted and agents deviating from

workflows by reordering, skipping, or merging steps. While these variations reflect real-world interactions, they complicate workflow extraction. To address these issues, we create SynthABCD, a synthetic dataset derived from ABCD workflows using an LLM. SynthABCD reduces noise by strictly following predefined workflows, thereby enabling more accurate evaluation of LLMs' workflow extraction capabilities.

In the first step of the SynthABCD synthesis, we use the same step 1 from E2E evaluation framework to generate all possible sub-flows. Next, we pair each subflow with the corresponding workflow to guide the LLM in simulating several conversations between a customer and an agent for each subflow. Additional data details, such as the intents used for ABCD and SynthABCD and human evaluation for SynthABCD, are discussed in Appendix C.

5.2 Conversations Retrievers

We evaluate our procedural element similarity-based retrieval strategy (Proc-Sim) against three baselines: random selection, procedural element diversity-based retrieval (Proc-Div), and conversation similarity-based retrieval (Conv-Sim). For Conv-Sim, we follow the same procedure as Proc-Sim but use full conversation embeddings instead of procedural element embeddings. For Proc-Div, our goal is to enhance the diversity of the selected conversations while filtering out noise. We begin by removing noisy conversations, excluding those in the lowest 10th percentile based on cosine similarity to the centroid of the procedural element embeddings. From the remaining set, we select one conversation that is furthest from the centroid as the initial candidate. Subsequently, we iteratively select additional conversations that are maximally distant from the centroid of the procedural elements of the already selected set (Gonzalez, 1985).

5.3 Workflows Generation Method

We compare the performance of our proposed QA-CoT model against several prompting techniques.

- **Basic:** We use a single-step prompt over selected conversations to generate workflows directly (prompt in Figure 20).
- **Reflection (Shinn et al., 2023):** We build upon Basic by identifying and addressing coverage gaps in the workflow. After generating the initial workflow using basic prompting, a reflection

prompt is applied to identify missing steps or inconsistencies, which are then refined to improve the workflow iteratively (prompts in Figures 21 and 22).

- **Plan (Wang et al., 2023):** This approach first prompts the LLM to devise a structured plan for extracting workflows given the conversations and then applies a second LLM prompt to execute the plan and generate workflows from the same set of conversations (prompts in Figures 23 and 24).
- **Ensemble:** We start by generating four workflows using basic prompting, each based on a different random order of conversations. The LLM then evaluates these workflows, checking step-by-step correctness based on cross-consistency among them, and produces a final consolidated workflow by selecting the most consistent steps across all versions (Chen et al., 2023) (prompt in Figure 25).
- **QA-CoT+Reflect:** We apply a reflection prompt to the QA chain-of-thought generated through the Guide-Implementer interaction, identifying and refining missing questions before generating the final workflow (prompt in Figure 26).

For consistency, we generate two workflows per intent by randomly varying the order of input conversations, and report the average performance across these two runs in all experiments. We evaluate performance using three metrics: macro accuracy (the arithmetic mean of the percentage of correct sub-flows per intent), micro accuracy (the percentage of total correct sub-flows across all intents), and #utt, which denotes the average number of utterances in the simulated conversations.

6 Results and Analysis

6.1 Conversations Retrieval

For each conversation retrieval approach, we select subsets containing 25, 50, 75, 100, and all conversations and prompt four LLMs (GPT-4o, gemini-1.5-pro, opus-3, and sonnet-3.5) to extract workflows. We plot macro accuracy and micro accuracy for each strategy, as shown in Figure 6.

Our results show that **selecting a subset of conversations consistently yields better performance than using all available conversations**, highlighting that simply applying an LLM to the entire dataset is not the most effective approach. GPT-4o and Sonnet-3.5 achieve peak performance

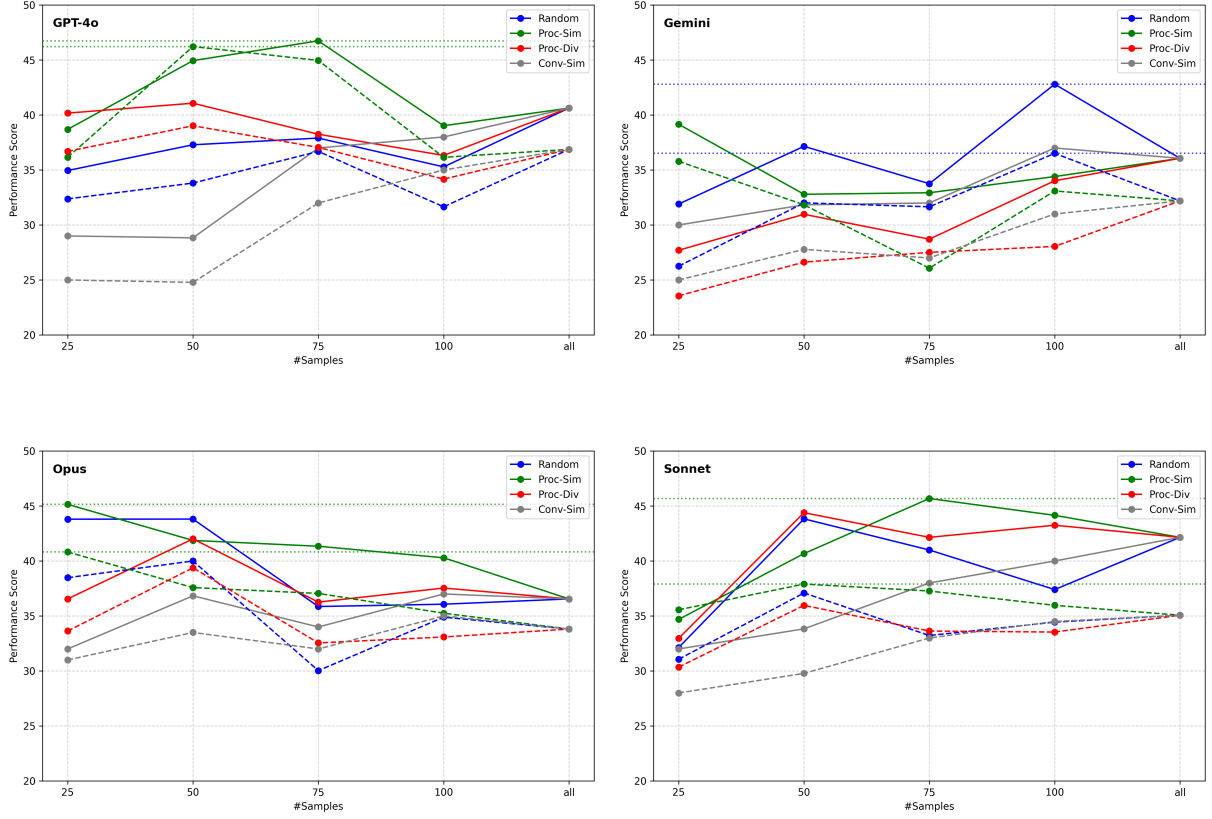


Figure 6: Performance comparison of conversation selection strategies (Proc-Sim, Conv-Sim, Random, and Proc-Div) with varying numbers of selected conversations. Results are shown for four LLMs: gpt-4o, gemini-1.5-pro, opus-3, and sonnet-3.5, using E2E macro (solid lines) and micro (dashed lines) accuracies.

with 75 conversations, Opus-3 peaks at 25, and Gemini-1.5-Pro performs best with 100.

Among the retrieval strategies, **Proc-Sim is the most effective for three of the four LLMs, and consistently outperforms Conv-Sim**, suggesting that retrieval based on explicit procedural elements is more reliable than relying on full conversation embeddings. Gemini-1.5-Pro is the only model that achieves its best performance with the random selection strategy, indicating that it is less sensitive to noise in the input conversations. In contrast, Proc-Div consistently performs worse, highlighting that prioritizing diversity can introduce noise from real-world conversations. When exposed to a more diverse but noisier set, the model may fail to identify the correct procedure and instead incorporate outlier steps, resulting in reduced performance.

To evaluate different prompting techniques and LLMs for workflow extraction, we adopt the best-performing retrieval strategy and number of conversations for each LLM individually (e.g., for GPT-4o, we use Proc-Sim with 75 conversations). We use macro accuracy to select the best configuration,

as it accounts for differences in workflow complexity by treating each workflow equally, rather than favoring those with more sub-flows.

6.2 Workflows Generation

As shown in Table 1, all baselines (Reflection, Plan, and Ensemble) fail to consistently outperform basic prompting across LLMs. Notably, both Reflection and Ensemble incorporate additional refinement steps to improve workflow extraction; however, these refinements do not consistently enhance performance and may introduce noise. Similarly, Reflection prompting on QA-CoT fails to yield performance gains, further reinforcing that refinement steps may not always benefit workflow extraction.

On the other hand, **QA-CoT outperforms almost all baselines on both the ABCD and SynthABCD datasets**, demonstrating that augmented context with Guide and Implementer interaction is beneficial for workflow generation. On average, QA-CoT improves macro accuracy by 8.73% on ABCD and 15.59% on SynthABCD. The most significant improvement is observed on SynthABCD

Method	GPT-4o			gemini			opus			sonnet		
	Macro	Micro	#utt	Macro	Micro	#utt	Macro	Micro	#utt	Macro	Micro	#utt
ABCD												
Basic	46.74	44.96	8.69	42.80	36.51	7.89	45.15	40.82	9.99	45.68	41.0	8.86
Reflection	44.43	42.09	8.51	37.14	35.61	8.51	50.89	45.86	10.47	51.80	49.46	10.57
Plan	37.29	38.30	7.71	41.0	38.12	6.75	35.77	32.91	7.90	37.56	34.71	7.39
Ensemble	46.13	45.32	8.66	38.71	38.67	8.42	41.80	34.35	8.51	46.47	42.44	9.29
QA-CoT	58.55	56.29	10.60	48.09	46.40	8.40	52.36	55.75	11.92	56.28	53.44	11.49
+ Reflection	55.45	55.39	10.73	43.10	46.22	8.36	56.84	56.65	11.54	54.44	54.67	10.56
SynthABCD												
Basic	68.91	61.95	8.71	74.73	75.65	8.59	91.43	91.52	12.29	51.23	45.43	7.49
Reflection	57.42	52.61	8.24	61.90	60.43	9.08	91.85	92.82	12.47	49.14	41.08	7.68
Plan	78.40	68.91	8.55	72.77	70.65	7.97	74.85	70.21	9.07	73.02	70.65	9.32
Ensemble	84.31	81.30	9.62	81.13	79.78	10.41	78.70	73.26	9.21	55.97	45.43	7.99
QA-CoT	86.53	86.08	11.96	84.04	81.96	9.70	89.73	90.21	12.27	88.38	86.73	10.46
+ Reflection	87.87	89.13	12.28	76.59	73.91	9.89	87.12	85.21	11.83	82.59	85.65	11.08

Table 1: Performance comparison of different workflow extraction systems.

	Macro	Micro	#utt
Ground-truth Workflow	96.93	96.97	10.58

Table 2: Performance when ground-truth workflows are used to simulate the system bot.

with the Sonnet model, where QA-CoT boosts macro accuracy by 37.15%. On ABCD, it achieves the highest gain of 11.81% with the GPT-4o model. These improvements are consistent across open-access LLMs (Appendix A.1), highlighting the robustness of QA-CoT in enhancing workflow extraction. Additionally, as discussed in Appendix A.2, QA-CoT remains effective even when applied to noisy, unfiltered conversation sets, yielding substantial gains over basic prompting.

There is one outlier: the Opus model on SynthABCD, where Reflection prompting performs better. Notably, Opus achieves near-optimal performance in this setting, with a macro accuracy of 91.43%, while the upper bound—attained when the agent bot follows the ground-truth workflow—is 96.93% (Table 2). Unlike predicted workflows, which may introduce errors, the ground-truth workflow guarantees strictly correct actions, representing the theoretical performance ceiling.

We also observe that LLMs perform substantially better on SynthABCD, underscoring their ability to extract accurate workflows from clean, error-free conversations. This finding highlights the importance of accurate conversation selection in enhancing workflow extraction performance.

Lastly, we observe that ground truth workflows have an average of 10.58 utterances per conversation (Table 2), whereas model-predicted workflows tend to produce longer interactions despite lower accuracy. For example, the best-performing

predicted workflows, generated by Opus on SynthABCD, average 12.47 utterances (1.89 more than the ground truth) but achieve a macro accuracy that is 5.08 points lower than that of the ground-truth workflows.

Our results also reveal a positive correlation between macro accuracy and utterance count, suggesting that higher-quality predicted workflows often result in longer conversations. However, this trend may also reflect inefficiencies, for example, when a model unnecessarily splits a correct step into multiple substeps. Since the E2E metrics focus on goal completion rather than interaction length, such decompositions are not penalized. Therefore, we recommend using utterance count as a complementary efficiency metric, favoring workflows that achieve task completion in fewer steps when multiple candidates exhibit comparable success rates.

6.3 Performance of Reasoning LLMs

In Table 3, we present the performance of four reasoning LLMs (DeepSeek-R1, o1, o1-mini, and o3-mini) on both the ABCD and SynthABCD datasets. We use 75 conversations selected using the Proc-Sim retriever. Our results show that **simply prompting a reasoning LLM does not consistently improve performance**. Among the evaluated models, only o3-mini demonstrates a significant improvement over non-reasoning LLMs on ABCD, achieving performance levels close to QA-CoT prompting with non-reasoning models. Furthermore, applying **QA-CoT with reasoning LLMs consistently improves performance over basic prompting across both datasets**. This finding highlights that explicit QA-based chain-of-thought can benefit even reasoning LLMs in extracting dialog workflows.

Method	o1			o1-mini			o3-mini			DeepSeek R1		
	Macro	Micro	#utt	Macro	Micro	#utt	Macro	Micro	#utt	Macro	Micro	#utt
ABCD												
Basic	36.72	35.97	7.70	44.20	40.46	9.73	57.06	57.19	9.52	46.10	40.10	8.42
QA-CoT	39.0	33.09	8.50	60.74	64.74	12.05	61.20	61.51	11.26	47.24	44.06	7.89
SynthABCD												
Basic	59.57	52.39	7.65	69.57	67.39	10.56	79.50	78.91	9.14	65.26	51.52	7.50
QA-CoT	80.60	76.95	9.62	87.16	87.17	12.43	89.49	90.21	10.31	66.95	55.65	8.41

Table 3: End-to-End performance comparison of different reasoning LLMs.

To better understand the limitations in performance improvement, we analyze both the thinking process and final outputs of DeepSeek-R1 on the ABCD and SynthABCD datasets (Figure 7). Our analysis reveals that while the thinking tokens capture the general reasoning process, they often lack fine-grained preconditions. For example, instead of explicitly detailing the relationship between the number of days and membership type when determining return eligibility, the model generates a generic statement such as “determine return eligibility based on purchase date.” This omission can lead to incomplete workflow extraction, as critical decision criteria such as a 6-month return window for silver members versus a 30-day window for guest members are not captured. In contrast, QA-CoT with GPT-4o includes these fine-grained details, as shown in Figure 4.

7 Error Analysis

We manually reviewed all eight QA-CoT systems to identify cases where they failed to handle all scenarios correctly (example in Figure 8). One common mistake was **confusing system-available information with details that should be requested from the user**. This issue was particularly frequent in shipping-related tasks, where models incorrectly assumed they should check the shipping status in the system instead of asking the customer. o1 and DeepSeek-R1 made this mistake in five intents, while o3-mini was the only model that avoided it. This error did not stem from the conversation data but rather from the LLM’s parametric knowledge, as it appeared in both the ABCD and SynthABCD datasets. Another common issue was **models tend to ignore alternative options**. Instead of allowing a choice, they often predicted only one option. For example, when a workflow required asking for either the “full name or the account ID”, models typically picked just one (e.g., ask “full name”) instead of keeping both possibilities. A third major issue was **failing to follow conditional logic at**

branching points. For example, a workflow might require asking for the membership level only if there is no system error. However, predicted workflows often ignored this condition and included a fixed sequence, combining steps incorrectly instead of following the intended logic.

We observed two additional patterns that did not affect accuracy but impacted efficiency. First, models often **merged multiple steps into a single step**, particularly during the early steps in workflows. For example, a workflow might specify sequential steps such as verifying the account ID or full name, followed by email ID, order ID, and username. Rather than maintaining these as distinct steps, models frequently combined them, thereby reducing the average number of utterances required to resolve an issue. Second, when the original workflow allowed multiple conditions to proceed, such as purchase date within 90 days, original receipt, or original packaging, models tended to convert this flexibility into a strict sequence. For instance, instead of allowing any of the conditions to be sufficient, models often predicted checking the purchase date first, then asking about the receipt if the date exceeded 90 days, and so on. Although these behaviors do not reduce accuracy, they negatively impact efficiency by increasing the number of utterances compared to the ground-truth workflow.

8 Conclusion

We introduced a novel framework for extracting and evaluating dialog workflows from historical conversations. Our proposed QA-CoT prompting approach outperforms standard prompting strategies, yielding a 12.16% improvement in average macro accuracy. Furthermore, the strong alignment between our end-to-end evaluation and human assessments highlights the robustness and reliability of our evaluation framework. We believe this work provides a robust foundation for structured workflow extraction and evaluation, and will encourage future research to develop more effective methods.

9 Limitations

Our workflow extraction and evaluation methods are designed for Service AI agents. Although they work well for this domain, applying them to more complex domains may require some adjustments. Additionally, our conversation selection method groups conversations by intent. An extra step for intent classification may be needed when intent labels are unavailable.

Our E2E evaluation aligns closely with the human judgment, making it a reliable way to assess the accuracy of workflows. However, for more complex workflows in other domains, human evaluation can provide additional insights to further refine the evaluation framework.

SynthABCD offers a controlled testing environment, ensuring consistent benchmarking. Although it does not capture the full complexity of real-world interactions, models that perform well on SynthABCD provide a good starting point. To improve robustness, we recommend additional testing on real-world datasets such as ABCD to ensure that the models perform well in diverse scenarios.

References

- Dan Bohus and Alexander I. Rudnicky. 2009. [The ravenclaw dialog management framework: Architecture and systems](#). *Computer Speech & Language*, 23(3):332–361.
- Kristi Brown. 2023. [Automate this! — power up your einstein bot with flow](#).
- Sergio Burdisso, Srikanth Madikeri, and Petr Motlicek. 2024. [Dialog2flow: Pre-training soft-contrastive action-driven sentence embeddings for automatic dialog flow extraction](#). *arXiv preprint arXiv:2410.18481*.
- Bingkun Chen, Shaobing Dai, Shenghua Zheng, Lei Liao, and Yang Li. 2021a. [Dsbert:unsupervised dialogue structure learning with bert](#). *Preprint*, arXiv:2111.04933.
- Derek Chen, Howard Chen, Yi Yang, Alexander Lin, and Zhou Yu. 2021b. [Action-based conversations dataset: A corpus for building more in-depth task-oriented dialogue systems](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3002–3017, Online. Association for Computational Linguistics.
- Xinyun Chen, Renat Aksitov, Uri Alon, Jie Ren, Kefan Xiao, Pengcheng Yin, Sushant Prakash, Charles Sutton, Xuezhi Wang, and Denny Zhou. 2023. [Universal self-consistency for large language model generation](#). *Preprint*, arXiv:2311.17311.
- Ananlada Chotimongkol. 2008. [Learning the structure of task-oriented conversations from the corpus of in-domain dialogs](#).
- Cuong Xuan Chu, Niket Tandon, and Gerhard Weikum. 2017. [Distilling task knowledge from how-to communities](#). *Proceedings of the 26th International Conference on World Wide Web*.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2023. [Improving factuality and reasoning in language models through multiagent debate](#). *Preprint*, arXiv:2305.14325.
- Alexander R Fabbri, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2021. [Qafacteval: Improved qa-based factual consistency evaluation for summarization](#). *arXiv preprint arXiv:2112.08542*.
- Teofilo F. Gonzalez. 1985. [Clustering to minimize the maximum intercluster distance](#). *Theoretical Computer Science*, 38:293–306.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint arXiv:2501.12948*.
- Amine El Hattami, Stefania Raimondo, Issam Laradji, David Vazquez, Pau Rodriguez, and Chris Pal. 2023. [Workflow discovery from dialogues in the low data regime](#). *Preprint*, arXiv:2205.11690.
- Zhe Hu, Hou Pong Chan, Jing Li, and Yu Yin. 2025. [Debate-to-write: A persona-driven multi-agent framework for diverse argument generation](#). *Preprint*, arXiv:2406.19643.
- Dan Jurafsky and Elizabeth Shriberg. 1997. [Switchboard swbd-damsl shallow-discourse-function annotation coders manual](#).
- Mahnaz Koupaee and William Yang Wang. 2018. [Wikihow: A large scale text summarization dataset](#). *ArXiv*, abs/1810.09305.
- Bo-Ru Lu, Yushi Hu, Hao Cheng, Noah A. Smith, and Mari Ostendorf. 2022. [Unsupervised learning of hierarchical conversation structure](#). *Preprint*, arXiv:2205.12244.
- Hirokuni Maeta, Tetsuro Sasada, and Shinsuke Mori. 2015. [A framework for procedural text understanding](#). In *Proceedings of the 14th International Conference on Parsing Technologies*, pages 50–60, Bilbao, Spain. Association for Computational Linguistics.
- Shikib Mehri and Maxine Eskenazi. 2021. [Schema-guided paradigm for zero-shot dialog](#). In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 499–508, Singapore and Online. Association for Computational Linguistics.

- Do June Min, Paloma Sodhi, and Ramya Ramakrishnan. 2023. [Workflow-guided response generation for task-oriented dialogue](#). *Preprint*, arXiv:2311.08300.
- Johannes E. M. Mosig, Shikib Mehri, and Thomas Kober. 2020. [STAR: A Schema-Guided Dialog Dataset for Transfer Learning](#). *arXiv e-prints*, arXiv:2010.11853.
- Emanuele Musumeci, Michele Brienza, Vincenzo Suriani, Daniele Nardi, and Domenico Daniele Bloisi. 2024. [Llm based multi-agent generation of semi-structured documents from semantic templates in the public administration domain](#). *Preprint*, arXiv:2402.14871.
- Apurba Nath and Aayush Kubba. 2021. [Tscan : Dialog structure discovery using scan](#). *ArXiv*, abs/2107.06426.
- Hogun Park and Hamid R. Motahari Nezhad. 2018. [Learning procedures from text: Codifying how-to procedures in deep neural networks](#). *Companion Proceedings of the The Web Conference 2018*.
- Connor Pryor, Quan Yuan, Jeremiah Liu, Mehran Kazemi, Deepak Ramachandran, Tania Bedrax-Weiss, and Lise Getoor. 2024. [Using domain knowledge to guide dialog structure induction via neural probabilistic soft logic](#). *Preprint*, arXiv:2403.17853.
- Liang Qiu, Chien-Sheng Wu, Wenhao Liu, and Caiming Xiong. 2022. [Structure extraction in task-oriented dialogues with slot clustering](#). *Preprint*, arXiv:2203.00073.
- Stefania Raimondo, Christopher Pal, Xiaotian Liu, David Vazquez, and Hector Palacios. 2023. [Improving generalization in task-oriented dialogues with workflows and action plans](#). *Preprint*, arXiv:2306.01729.
- Giorgio Robino. 2025. [Conversation routines: A prompt engineering framework for task-oriented dialog systems](#). *Preprint*, arXiv:2501.11613.
- Md Rashad Al Hasan Rony, Ricardo Usbeck, and Jens Lehmann. 2022. [Dialogk: Knowledge-structure aware task-oriented dialogue generation](#). *Preprint*, arXiv:2204.09149.
- Pol Schumacher, Mirjam Minor, Kirstin Walter, and Ralph Bergmann. 2012. [Extraction of procedural knowledge from the web: a comparison of two workflow extraction approaches](#). In *Proceedings of the 21st International Conference on World Wide Web, WWW '12 Companion*, page 739–747, New York, NY, USA. Association for Computing Machinery.
- Weiyang Shi, Tiancheng Zhao, and Zhou Yu. 2019. [Unsupervised dialog structure learning](#). *Preprint*, arXiv:1904.03736.
- Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflexion: Language agents with verbal reinforcement learning](#). *Preprint*, arXiv:2303.11366.
- Atsushi Ushiku, Hayato Hashimoto, Atsushi Hashimoto, and Shinsuke Mori. 2017. [Procedural text generation from an execution video](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 326–335, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Jianhong Wang, Yuan Zhang, Tae-Kyun Kim, and Yunjie Gu. 2021. [Modelling hierarchical structure between dialogue policy and natural language generator with option framework for task-oriented dialogue system](#). *Preprint*, arXiv:2006.06814.
- Lei Wang, Wanyu Xu, Yihuai Lan, Zhiqiang Hu, Yunshi Lan, Roy Ka-Wei Lee, and Ee-Peng Lim. 2023. [Plan-and-solve prompting: Improving zero-shot chain-of-thought reasoning by large language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2609–2634, Toronto, Canada. Association for Computational Linguistics.
- Frank F. Xu, Lei Ji, Botian Shi, Junyi Du, Graham Neubig, Yonatan Bisk, and Nan Duan. 2020a. [A benchmark for structured procedural knowledge extraction from cooking videos](#). In *Proceedings of the First International Workshop on Natural Language Processing Beyond Text*, pages 30–40, Online. Association for Computational Linguistics.
- Jun Xu, Zeyang Lei, Haifeng Wang, Zheng-Yu Niu, Hua Wu, Wanxiang Che, and Ting Liu. 2020b. [Discovering dialog structure graph for open-domain dialog generation](#). *Preprint*, arXiv:2012.15543.
- Congchi Yin, Piji Li, and Zhaochun Ren. 2023. [Ctrlstruct: Dialogue structure learning for open-domain response generation](#). In *Proceedings of the ACM Web Conference 2023, WWW '23*, page 1539–1550. ACM.
- Yusen Zhang, Ruoxi Sun, Yanfei Chen, Tomas Pfister, Rui Zhang, and Serkan Ö. Arik. 2024. [Chain of agents: Large language models collaborating on long-context tasks](#). *Preprint*, arXiv:2406.02818.
- 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.
- Luowei Zhou, Chenliang Xu, and Jason J. Corso. 2018. [Towards automatic learning of procedures from web instructional videos](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI'18/IAAI'18/EAAI'18*. AAAI Press.

```

..
<think>
Okay, let's tackle this problem. The user wants me to identify the steps necessary to resolve a customer's issue based on the provided
conversations. They specified that the steps should be essential, consolidated, and formatted as a numbered list with pre-conditions,
required inputs, and specific conditions.

First, I need to go through all the conversations to understand the common steps agents take to resolve the return requests. Since there
are 75 conversations, I'll look for patterns and repeated actions.

Looking at the conversations, the first thing agents usually do is ask for the customer's name or account ID to pull up their account.
Then, they request the username, email address, and order ID to validate the purchase. Next, they check the membership level (guest,
bronze, silver, gold) because the return policy varies based on that. For example, guests might have a shorter return window,
while gold members have unlimited returns.

Agents then verify the purchase date. If the purchase is beyond the standard return period for the customer's membership level, they
ask if the customer has the receipt or original packaging. If they do, the return is processed. Without the receipt, returns might
be denied, especially for guests or lower-tier members.

After confirming eligibility, the agent collects the customer's address to generate a shipping label. Finally, they ask how the
customer wants to process the return: by mail, in-store, or at a drop-off center.

I need to make sure all these steps are included, but without breaking them down by individual conversation. Also, pre-conditions
like membership level affecting return eligibility and required inputs like receipt or address should be mentioned.

Let me list out the steps:

1. Collect customer's full name or account ID to access their account.
2. Obtain username, email address, and order ID to validate the purchase.
3. Determine the customer's membership level (guest, bronze, silver, gold) to apply the appropriate return policy.
4. Verify the purchase date to check if it's within the allowed return period for their membership level.
5. If the purchase is outside the standard return period, check if the customer has the receipt or original packaging (pre-condition
for returns beyond standard period).
6. Collect the customer's shipping address to generate a return label.
7. Confirm the return method (mail, in-store, drop-off center) and provide instructions.

I should check if any steps are missing. For example, in some cases, if the customer doesn't have a receipt and it's past the return
period, the agent can't process the return. That's covered in step 5. Also, the return method is always asked, so step 7 is necessary.
Each step includes required inputs (like address) and conditions (membership level affecting return window).

Yes, this covers all the essential steps without unnecessary details. The list is consolidated and formatted as per the instructions.
</think>
..

```

Figure 7: A snippet of <think> tokens generated by the DeepSeek-R1 model for the *return_color* intent using conversations from the ABCD data. Compared to the QA-CoT (Figure 4), think tokens lack fine-grained conditions based on membership level.

	Macro	Micro	#utt
Multi-turn QA	52.50	49.28	9.28
Single-turn QA	58.55	56.29	10.60

Table 4: Performance of multi-turn prompting and single-pass prompting for simulating the Guide-Implementer interaction. The single-pass approach achieves better accuracy.

A Ablation Studies

A.1 Results for Open-Access LLMs

Table 5 shows results for three open-access LLMs on both the ABCD and SynthABCD datasets. Across all models and settings, QA-CoT consistently outperforms the Basic method, demonstrating its robustness in both standard and synthetic task environments. For example, on ABCD, QA-CoT improves macro accuracy by more than 10 points for Qwen2.5-72B-Instruct and by 6 points for Qwen2.5-7B-Instruct, while Llama-3.1-405B-Instruct shows a more modest gain of approximately 2.7 points. On SynthABCD, where conversations strictly follow workflows, Qwen2.5-72B-Instruct shows substantial improvement, achieving a macro accuracy of 90.88 with QA-CoT compared

to 65.67 with the Basic method.

A.2 Effect of Retrieval on QA-CoT prompting

Table 6 presents an ablation study comparing QA-CoT prompting with and without retrieval across four closed LLMs. QA-CoT consistently outperforms the Basic baseline, even without retrieval, highlighting its robustness to outlier or noisy conversations. When paired with the Proc-Sim retriever, QA-CoT yields additional gains for GPT-4o, Gemini, and Opus, suggesting that relevant conversational context can further enhance reasoning. For Sonnet, however, retrieval does not lead to improvement, indicating that the model is already robust to noise when using QA-CoT.

B Human Evaluations of E2E Framework

We evaluate the effectiveness of our E2E framework using 18 distinct intents from the ABCD dataset². First, we manually assess the automatic scenario construction process against ground-truth

²All annotators hold graduate degrees in computer science and have over eight years of experience in NLP within industry or academia.

	Llama-3.1-405B-Instruct			Qwen2.5-72B-Instruct			Qwen2.5-7B-Instruct		
Method	Macro	Micro	#utt	Macro	Micro	#utt	Macro	Micro	#utt
ABCD									
Basic	51.82	52.33	10.45	43.24	38.66	8.90	34.29	33.99	8.73
QA-CoT	54.49	55.03	11.05	53.31	52.87	11.69	40.39	37.58	9.72
SynthABCD									
Basic	88.06	89.35	10.74	65.67	57.61	8.73	57.30	56.30	9.96
QA-CoT	91.01	90.65	11.26	90.88	91.73	12.04	64.35	62.17	10.33

Table 5: End-to-End performance comparison of different open-access LLMs.

	GPT-4o			gemini			opus			sonnet		
Method	Macro	Micro	#utt	Macro	Micro	#utt	Macro	Micro	#utt	Macro	Micro	#utt
Basic _{all}	40.63	36.87	7.87	36.06	32.19	7.44	36.55	33.81	8.28	42.15	37.41	8.85
QA-CoT _{all}	55.05	51.07	10.04	48.65	45.32	8.96	57.88	54.31	11.02	57.46	56.11	10.86
QA-CoT _{retrieve}	58.55	56.29	10.60	52.18	55.03	8.67	59.26	57.37	11.51	55.67	54.85	11.25

Table 6: Ablation results comparing the effectiveness of QA-CoT prompting with and without retrieval across four closed LLMs. Basic_{all} denotes the baseline using all conversations with no QA-CoT prompting; QA-CoT_{all} applies QA-CoT prompting to all conversations without retrieval; QA-CoT_{retrieve} applies QA-CoT prompting to the top-k conversations retrieved by our Proc-Sim retriever, where the number of retrieved conversations is tuned individually for each model based on validation performance.

1. Gather customer information:
 - Full name
 - Username
 - Email address
 - Order ID
2. Verify membership level and purchase date:
 - Confirm membership level (Guest, Bronze, Silver, Gold)
 - Confirm purchase date to ensure it falls within the return window (30 days for Guest, 90 days for Bronze, 6 months for Silver, unlimited for Gold)
3. Check for receipt or original packaging if outside the return window:
 - If the purchase is outside the return window, confirm if the customer has the receipt or original packaging.
4. Confirm return eligibility:
 - If eligible, proceed with the return process.
 - If not eligible, inform the customer and provide any alternative solutions if available.
5. Collect shipping address for the return label:
 - Full address
6. Determine return method:
 - Ask the customer if they prefer to return by mail, in store, or at a drop-off center.
7. Generate and send the return label:
 - Email the return label to the customer.
8. Provide instructions for the return:
 - Inform the customer to print the label and attach it to the package.
 - Provide any additional instructions based on the chosen return method.
9. Confirm completion and offer further assistance:
 - Confirm that the return process has been initiated.
 - Ask if there is anything else the customer needs help with.

Figure 8: Example showing errors in the extracted workflow. Instead of two steps ‘First, ask the customer for their Full Name or Account ID to pull up account.’ and ‘After getting the Full Name or Account ID, ask the customer for their Email Address, Order ID, and Username to validate purchase.’, it combines them into one and also exclude **Account ID**. It also merge the steps ask ‘membership level’ and ask ‘purchase date’ while the ground-truth workflow requires asking membership level first and then asking ‘purchase date’ or ‘original receipt’ or ‘original packaging’ or ‘address’ depending on the membership level.

workflows (Step 1, Figure 5a). Across these intents, human annotators identified 231 possible scenarios, while GPT-4o generated 230, of which 224 were correct. Next, we evaluate the mapping of

scenarios to customer information, system information, and success criteria (Steps 2 and 3, Figure 5a). Among the 224 correct scenarios, GPT-4o successfully generated all required details in 219 cases. Across all three steps, GPT-4o achieved an overall accuracy of 94.81% (219 out of 231 scenarios), demonstrating its effectiveness in capturing and mapping essential workflow details for simulating user and agent interactions.

Finally, to assess the LLM’s ability to evaluate conversation success, two annotators manually labeled successful completions for 105 simulated conversations. Inter-annotator agreement was high, with a Cohen’s κ of 0.966. Agreement between GPT-4o’s evaluations and human judgments was also strong, with a Cohen’s κ of 0.922. These results confirm the robustness of our E2E framework for workflow evaluation.

C Dataset

C.1 ABCD Dataset

The original ABCD dataset contains 55 intents, but we focus on 21 complex intents that involve multi-step workflows and exclude the remaining 34. The excluded intents consist of non-procedural issues—such as FAQs about products, pricing, timings, membership, or features—as well as simpler, linear workflows. The selected 21 complex intents include multiple sub-flows to handle diverse scenarios, such as varying membership levels.

Intent	Type	#Scenarios	ABCD	SynthABCD
refund_initiate	simple	1		
refund_update	simple	1		
refund_status	simple	1		
return_stain	complex	18	✓	✓
return_size	complex	18	✓	
return_color	complex	18	✓	
status_mystery_fee	complex	10	✓	✓
status_delivery_time	complex	10	✓	✓
status_payment_method	complex	8	✓	✓
status_quantity	complex	8	✓	✓
manage_upgrade	complex	14	✓	✓
manage_downgrade	complex	26	✓	✓
manage_create	complex	26	✓	✓
manage_cancel	complex	26	✓	✓
recover_username	simple	1		
recover_password	simple	1		
reset_2fa	simple	1		
status_service_added	complex	10	✓	✓
status_service_removed	complex	8	✓	✓
status_credit_missing	complex	6	✓	✓
manage_change_address	simple	1		
manage_change_name	simple	1		
manage_change_phone	simple	1		
manage_payment_method	simple	1		
bad_price_competitor	simple	1		
bad_price_yesterday	simple	1		
out_of_stock_general	simple	1		
out_of_stock_one_item	simple	1		
promo_code_out_of_date	complex	12	✓	✓
promo_code_invalid	complex	12	✓	
mistimed_billing_already_returned	complex	10	✓	✓
mistimed_billing_never_bought	complex	10	✓	✓
missing	simple	1		
cost	complex	8	✓	✓
status_due_amount	simple	1		
status_due_date	simple	1		
manage_pay_bill	simple	1		
manage_extension	complex	8	✓	✓
manage_dispute_bill	complex	12	✓	✓
credit_card	simple	1		
shopping_cart	simple	1		
search_results	simple	1		
slow_speed	simple	1		

Table 7: Intents used in E2E evaluation framework for the ABCD and SynthABCD datasets. *return_stain*, *return_size* and *return_color* share the same workflow, as do *promo_code_out_of_date* and *promo_code_invalid*. Therefore, we only use one from each group in SynthABCD.

C.2 SynthABCD Dataset

Among the 21 intents used in our E2E study, three workflows are duplicates of each other, and two others are also duplicates, resulting in a total of 18 distinct workflows. We focus on these 18 intents. To introduce variability, we create 50 random user profiles with varying attributes such as name, profession, and city of residence, generating two distinct synthetic conversations for each scenario. The prompts used for converting workflows into scenarios and for simulating conversations are provided in Figures 11 and 12, respectively.

Our SynthABCD pipeline employs the same automatic scenario construction process as the E2E

framework. Across the 18 intents, human annotation identified 231 possible scenarios, while the GPT-4o model generated 230 scenarios. Of these, 224 were correctly extracted, yielding an accuracy of 96.97%. With only one missed scenario and six incorrect extractions, these results demonstrate the effectiveness of our conversation synthesis process and its applicability to new workflows. All intents used in our evaluation are reported in Table 7.

C.3 Data Quality

We evaluate the compliance of conversations in both the ABCD validation set and SynthABCD with their corresponding ground-truth workflows

Dataset	F	NA	NF	NC Conv.
ABCD	39.0%	57.07%	3.93%	28.57%
SynthABCD	43.56%	56.21%	0.22%	2.71%

Table 8: Data quality evaluation of ABCD and SynthABCD datasets. F: followed, NA: not applicable, NF: not followed, NC Conv.: non-compliant conversation

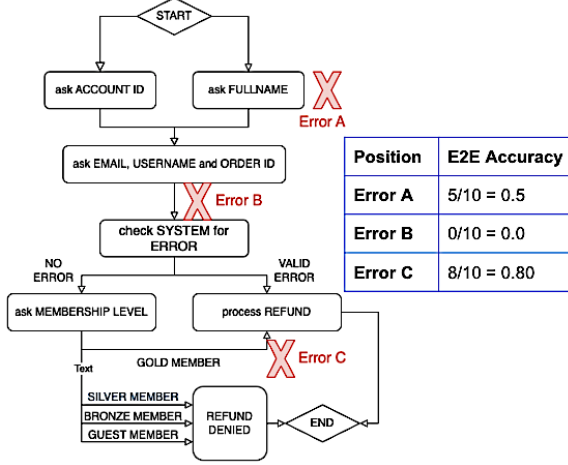


Figure 9: Illustration of End-to-End (E2E) accuracy evaluation. For this workflow, there are 10 scenarios based on membership level, system error, and whether the user can provide their full name or account ID. In the case of error B, none of the scenarios can be completed successfully, resulting in an accuracy of 0. For error A, only 50% of scenarios where the user provides an account ID can be completed, yielding an accuracy of 0.5. For error C, 2 scenarios fail, including those for gold members with either account ID or full name, resulting in an accuracy of 0.8.

using the GPT-4o model (prompt shown in Figure 17). For each conversation, we assess whether each workflow step is followed, deemed not applicable, or not followed. We also compute the percentage of non-compliant conversations, defined as those that fail to follow at least one step in the ground-truth workflow.

As shown in Table 8, the results reveal notable differences between the two datasets. In ABCD, 39.0% of workflow steps are followed, compared to 43.56% in SynthABCD. More importantly, the percentage of unfollowed steps (NF) is significantly lower in SynthABCD (0.22%) than in ABCD (3.93%), demonstrating the effectiveness of controlled conversation synthesis. This reduction in noise aligns with SynthABCD’s goal of enhancing data consistency. As a result, SynthABCD achieves 97.29% compliant conversations, a substantial improvement over ABCD’s 71.43%, reflecting its higher overall quality.

D Alternative Evaluation Frameworks for Dialog Workflows

In addition to our primary E2E evaluation, we explored several alternative frameworks to assess step-level correctness in workflow reconstruction. While these methods offer fine-grained insights, they fail to capture the overall effectiveness of a workflow and show low agreement with human evaluations. In particular, approaches that evaluate steps in isolation overlook dependencies across the workflow, often leading to over- or underestimation of practical performance. As illustrated in Figure 9, not all errors carry equal weight—some, like error B, may block all scenarios and result in 0% accuracy, while others, like error C, affect only a subset of cases and reduce accuracy to 0.8. Step-level evaluations treat all errors uniformly, failing to account for their actual impact on workflow execution. In contrast, E2E evaluation offers a more accurate and reliable measure, closely aligning with human judgments and directly reflecting real-world utility. Since workflows are ultimately judged by their ability to guide complete processes, we adopt E2E evaluation as our primary metric.

D.1 Evaluation Methods

D.1.1 Automatic QA Evaluator

We explored a QA-based evaluation framework that measures the accuracy of each step in a workflow by verifying its correct dependence on pre-conditions, such as the outcome of a previous step (Fabbri et al., 2021). For example, if the ground truth workflow includes a step to ask shipment status’ followed by either ask membership level if the product has shipped’ or ‘issue refund if the product has not shipped,’ our goal is to verify that the predicted workflow reflects this conditional logic accurately. To achieve this, we decompose ground-truth workflows into atomic questions, such as: Q: What is the next step after asking shipment status if the product has not shipped? A: issue refund. Q: What is the next step after asking shipment status if the product has shipped? A: ask membership level. By structuring evaluation in this way, we ensure a step-wise correctness check while maintaining consistency. We use the GPT-4o model both for generating QA pairs from ground-truth workflows and for evaluating the predicted workflows.

D.1.2 Embedding Similarity

Embedding similarity method quantifies structural and semantic similarity using cosine distance between workflow representations. Specifically, we represent both the ground-truth and predicted workflows as embedding vectors and define the similarity score as:

$$S_{\text{embedding}} = \frac{1}{\text{cosine_distance}(\mathbf{w}_{\text{ref}}, \mathbf{w}_{\text{pred}})} \quad (1)$$

where \mathbf{w}_{ref} and \mathbf{w}_{pred} are the vector representations of the ground-truth and predicted workflows, respectively. A higher score indicates greater alignment between the two workflows. To obtain these embeddings, we use OpenAI text-embedding-3-small model.

D.1.3 Edit Distance

Edit-distance-based evaluation quantifies structural differences by computing the minimum number of operations required to transform the predicted workflow into the ground-truth workflow. We define three types of operations:

- **Insertion:** Adding a missing step to align with the ground-truth workflow.
- **Deletion:** Removing an extraneous step that does not appear in the ground-truth workflow.
- **Reordering:** Adjusting the sequence of steps to match the ground-truth workflow.

To compute edit distance, we use the GPT-4o that aligns the predicted workflow to the ground-truth workflow using these operations. The final score is defined as:

$$S_{\text{edit}} = \frac{1}{\text{number_of_edits}} \quad (2)$$

where a higher score indicates fewer modifications needed for alignment, implying better prediction accuracy.

D.1.4 Step-Accuracy Evaluation

We use the GPT-4o model to label each step in the ground-truth workflow as either correctly or incorrectly covered with respect to the predicted workflow. The final score is computed as:

$$S_{\text{step}} = \frac{\text{number_of_correct_steps}}{\text{total_number_of_steps}} \quad (3)$$

where a higher score indicates that a greater proportion of steps in the ground-truth workflow are correctly captured by the predicted workflow.

D.1.5 Likert-Scale Evaluation

We use the GPT-4o model to assign a similarity score between 1 and 100 based on structural and semantic alignment between predicted and ground-truth workflows.

D.2 Human Evaluation Setup

To evaluate the alignment of alternative evaluation protocols, we conduct a human study on workflows for 11 complex intents, each accompanied by 4 predicted workflows. The evaluation process began with decomposing each workflow into a series of QA pairs. For the gold ground-truth workflows, one author wrote the QA pairs, while a second author verified their quality. This dual-step validation ensures that the questions were clear and accurately represented the underlying steps and conditional logic inherent in the workflows. From the 11 workflows, we obtain a total of 141 QA pairs. Additionally, the 4 predicted workflows for each intent collectively yield 564 QA pairs, providing a comprehensive dataset for evaluation.

Next, two authors independently labeled the correctness of each QA pair based on the corresponding predicted workflow. This independent labeling process resulted in a Cohen’s κ score of 0.742, indicating substantial agreement between the annotators. For QA pairs labeled as correct and incorrect, we achieve an F1 score of 0.937 for correct labels, demonstrating high agreement on identifying correct steps, and 0.82 for incorrect labels, indicating slightly lower agreement in identifying errors. This further indicates that detecting errors by reasoning on previous steps and their outcomes is challenging, even for humans.

To assess the overall alignment between the predicted and ground-truth workflows, we calculate a final score for each predicted workflow based on the percentage of correctly answered questions. This score shows a correlation score of 0.835 between human annotators, indicating strong consistency in their evaluations. We further analyze correlations between step-level evaluation methods and human judgments, as reported in Table 9. Among automated evaluation methods, the QA-based metric achieves the highest correlation (0.708) with human annotations, indicating its relative effectiveness in step-level evaluation. However, these correlation scores remain lower than those observed with End-to-End (E2E) evaluation, reinforcing our decision to prioritize E2E evaluation as the primary

Automatic Evaluation	Correlation
embedding	0.412
edit-distance	0.587
step-accuracy	0.631
likert-scale	0.678
QA-based (Proposed)	0.708

Table 9: Correlation between different automatic evaluation metrics and human judgements.

metric.

E Prompts

Extract intent, slot values and resolution steps from the customer service chat provided under <CONVERSATION> section below.

STEP-BY-STEP INSTRUCTIONS

1. Exclude any non-essential statements such as greetings, apologies, and expressions of gratitude.
2. Be concise, complete, and truthful to the facts mentioned in the conversation.
3. Include any all the numerical and categorical details such as personal names, addresses, phone numbers etc.
4. Follow the exact sequential order in which the agent took different actions in "resolution_step".
5. Respond using the JSON format:

```
{
  "intent": "What is the customer issue? Use max 50 characters.",
  "slot_values": {
    "dictionary of slots and values collected from the customer or provided to the customer by the agent formatted as {slot: value}"
  },
  "resolution_steps": ["List of steps/ actions taken by the agent to resolve the customer's issue."]
}
```

Example Conversation:

Customer: Hi, I'd like to return an item I purchased.

Agent: Hi! I'm happy to help with that. Could you please provide me with your order ID so I can pull up your account?

Customer: Sure, it's 123456.

Agent: Thank you. Can you confirm the reason for your return?

Customer: The item arrived damaged.

Agent: I'm sorry to hear that. We'll get this sorted for you right away. Would you prefer a replacement or a refund?

Customer: I'd like a refund, please.

Agent: Got it. I've initiated the return process for a refund. You'll receive an email with a prepaid return label. Once we receive the item, we'll process your refund within 3-5 business days. Can I help you with anything else?

Customer: No, that's all. Thanks!

Agent: You're welcome! If you need further assistance, feel free to reach out. Have a great day!

Example Output:

```
{
  "intent": "Customer wants to return a damaged item.",
  "slot_values": {
    "order_id": "123456",
    "return_reason": "damaged item",
    "refund_preference": "refund"
  },
  "resolution_steps": [
    "Agent asked for the order ID.",
    "Agent asked the reason for the return.",
    "Agent asked whether the customer wants a replacement or refund.",
    "Agent initiated the return process for a refund.",
    "Agent informed the customer about receiving a prepaid return label via email.",
    "Agent explained that the refund will be processed within 3-5 business days after receiving the item."
  ]
}
```

Figure 10: The prompt for extracting key elements: issue description, slots-values and resolution steps from ABCD conversations [Conversation Selection].

Identify all branching conditions in the provided policy. Do not include common steps. Generate one branch per line.

Example Policy 1:

```
{policy_1}
```

Example Output 1:

```
{output_1}
```

Example Policy 2:

```
{policy_2}
```

Example Output 2:

```
{output_2}
```

Figure 11: The prompt for mapping workflows to scenarios [SynthABCD and E2E Evaluation].

You are given below a dialog policy that a customer service agent uses to resolve a customer's issue.

Policy

policy

Read the policy carefully and simulate a conversation between an user and an agent that follows the below sub-policy.

```
{subflow}
```

Instruction

- Generate conversation that strictly follow the provided sub-policy.
- Make sure that the conversation is started by user.

User's details

Name: {user_name}

Profession: {user_profession}

Address: {city}

- Make sure that agent only asks the information in subflow or common steps in the dialog policy. Avoid any unnecessary information.

Example Policy 1:

```
{example_policy}
```

Example Subflow 1:

```
{example_subflow}
```

Example Output 1:

```
{example_conv}
```

Figure 12: The prompt for generating conversations given a workflow and a scenario [SynthABCD].

You are given a dialog workflow and a specific user scenario. Generate detailed list of user information, system information and the corresponding expected outcome for the given scenario based on the workflow. Respond following the below json format.

```
{
  "user information": ["detailed list of information that the user provide to the agent to complete the conversation in the given scenario"],
  "system information": ["detailed list of information that the agent check in system to complete the conversation in the given scenario"],
  "outcome": "expected outcome reflecting the resolution in the given scenario"
}
```

Example Policy 1:

```
{policy_1}
```

Example Scenario 1:

```
{scenario_1}
```

Example Output 1:

```
{output_1}
```

Example Policy 2:

```
{policy_2}
```

Example Scenario 2:

```
{scenario_2}
```

Example Output 2:

```
{output_2}
```

Example Policy 3:

```
{policy_3}
```

Example Scenario 3:

```
{scenario_3}
```

Example Output 3:

```
{output_3}
```

Example Policy 4:

```
{policy_4}
```

Example Scenario 4:

```
{scenario_4}
```

Example Output 4:

```
{output_4}
```

Figure 13: The prompt for extracting user information, system information and success criteria given a scenario and a workflow [E2E evaluation].

You are given a dialog policy and corresponding criteria for successful completion of conversation with a customer.

Given a conversation between an agent and a customer, check if the conversation ends successfully or not.

Dialog Policy:

```
{policy}
```

Success Criteria:

```
{outcome}
```

Conversation:

```
{conv}
```

Respond in json following the format below.

```
{
  "successful": "yes/no"
  "explanation": "explain how the conversation went?"
}
```

Figure 14: The prompt for evaluating successful completion of a conversation between the customer and service bot [E2E evaluation].

You are a customer talking to an agent to have your issue resolved. You are given the issue description and the slots-values that you share with the agent. Respond to the agent and provide them the requested information if possible.

STEP-BY-STEP INSTRUCTIONS:

- Identify the information requested by the agent.
- Check the provided issue description and the slots-values and identify if you can provide the requested information to the agent.
- If you can, respond to the agent with the requested information.
- If you don't have the requested information, respond that "you don't have the requested information and ask if there is any other information you can provide".
- Only provide the information that is requested by the user.

Issue:

```
{issue}
```

Slots-Values that you can provide to the agent:

```
info
```

Make up free form slot-values, e.g., user name, full name, order ID, account ID, email and address.

Conversation History:

```
{history}
```

Figure 15: The prompt for simulating a customer bot [E2E Evaluation].

You are a customer service agent trying to solve a customer's issue. You are given the dialog policy, conversation history and the system information. Respond to the customer by strictly following the provided dialog policy.

STEP-BY-STEP INSTRUCTIONS:

- Read the provided conversation history and identify the current dialog state.
- Match dialog state with the dialog policy, and identify the next step to address the customer's issue.
- If the next step involves checking system information (e.g., checking system for an error), check the available system information and inform that to the customer. If the required system information is not available, reply "DONE".
- If the next step involves requesting information from customer (e.g., username), ask customer for the required information.
- If the next step involves taking some action (e.g., issuing refund), inform the customer that you have successfully taken that action.
- Respond to the customer based on the identified next step.
- If the policy does not describe the next step based on the current dialog state, conclude the conversation by generating "DONE".
- If you can not take the next step for any reason, conclude the conversation by generating "DONE".
- If the issue has been successfully resolved, conclude the conversation by generating "DONE".
- Avoid repeating yourself or requesting information that has already been mentioned in the conversation history.

Dialog Policy
...
{policy}
...
System Information
{info}
Conversation History
{history}

Figure 16: The prompt for simulating a service agent bot [E2E Evaluation].

You are a quality assurance manager tasked to assess whether an agent followed the established agent guidelines for resolving a customer's issue. The guidelines offer detailed, rule-based instructions for agents to be followed in a step-by-step manner. However, agents may not consistently adhere to these instructions. Your task is to evaluate the agent's compliance by comparing the steps taken in the provided conversation against the established guidelines.

****Important: Approach your task step-by-step. Carefully evaluate each rule in the guidelines.****

Step-by-Step Instruction

- For each rule in guidelines, determine whether that rule is applicable to conversation or not.
- Some rules are context-specific, they define actions contingent on preceding action outcomes, making them applicable only in specific conversational situations.
- Other rules are universally applicable, dictating actions that consistently follow a preceding action, regardless of the outcome.
- If the rule is applicable, check whether the agent followed the prescribed steps accurately or not.
- Evaluate the accuracy of both the action, and the order of the action.
- Here are some rule for handling specific scenarios
- If the customer has already provided certain details (such as Full Name, Account ID, refund amount, item details, etc.) either voluntarily or as requested by the agent in a prior step: Do not penalize the agent for not following guidelines by not requesting the same details again.
- Do not penalize the agent for not following guidelines even if the agent request for the same details again.
- If a rule includes multiple actions, the order of those actions is irrelevant and should not be penalized. For instance, [ask 'Email Address', 'Order ID', and 'Username'] or [ask 'Order ID', 'Email Address', and 'Username'] are equivalent.
- Generate output in the JSON format given below, offering both the response and an explanation for each rule.

```
{
  "Rule_1":{
    "response": 'followed'/'not applicable'/'not followed',
    "explanation": ''
  },
  "Rule_2":{
    "response": 'followed'/'not applicable'/'not followed',
    "explanation": ''
  },
  ..
}
```

Figure 17: The prompt for conversation's compliance evaluation given the ground-truth workflow [Data Quality].

You are a QA simulator consisting of two agents: a Guide and an Implementer. Both agents are tasked with collaboratively reconstructing the process of resolving customer issues by analyzing historical conversations.

Below are the profiles for Guide and Implementer:

Guide:

- Asks targeted questions to understand the steps taken in a given scenario.
- Focuses on clarifying the preconditions, decision points, and the logic behind the steps.
- Ensures that all possible customer contexts and edge cases are addressed.

Implementer:

- Review past conversations to answer the questions.
- Provides detailed explanations about the actions taken at each step and their rationale.

Below is an example discussion between a guide and an implementer agent:

Guide: What is the first step an agent should take when booking a restaurant?

Implementer: ask name and phone number

Guide: What is the next step after asking the name and phone number?

Implementer: ask preferred date and time for the reservation

Guide: What is the next step after asking the preferred date and time for the reservation?

Implementer: ask party size

Guide: What is the next step after asking the party size?

Implementer: ask for any special requests

Guide: What is the next step after asking about any special request?

Implementer: check availability

Guide: What is the next step if restaurant is booked at requested date and time?

Implementer: offer alternative time and date

Guide: What is the next step if restaurant is available at requested date and time?

Implementer: book and inform the customer

Let's think step-by-step and generate a discussion between the Guide and the Implementer based on below conversations.

Figure 18: The prompt for generating question-answer-based chain-of-thoughts [Workflow Extraction: QA-CoT].

Identify the steps necessary to resolve the customer's issue based on the provided conversations and the discussion between Guide and Implementer Agents.

- Include only the essential actions needed for issue resolution, excluding unnecessary steps.
- Use the discussion between Guide and Implementer Agents to identify important steps in the workflow.
- Create a unified, consolidated list of steps without breaking them down by individual conversation.
- Formatting Instructions:
 1. Use simple and concise language for each step, mentioning any pre-conditions where applicable.
 2. Organize the steps in a numbered list for clarity.
 3. Include relevant details, such as required inputs and specific conditions, for each step.

Figure 19: The prompt for extracting workflows given historical conversations and the discussion between the guide and implementer agents [Workflow Extraction: QA-CoT].

Identify the steps necessary to resolve the customer's issue based on the provided conversations.

- Include only the essential actions needed for issue resolution, excluding unnecessary steps.
- Create a unified, consolidated list of steps without breaking them down by individual conversation.
- Formatting Instructions:
 1. Use simple and concise language for each step, mentioning any pre-conditions where applicable.
 2. Organize the steps in a numbered list for clarity.
 3. Include relevant details, such as required inputs and specific conditions, for each step.

Figure 20: The prompt for basic prompting strategy. [Workflow Extraction Baseline: Basic].

Reflect on the following past chats and assess how well the dialog workflow covers the different scenarios mentioned. Your goal is to focus on the most representative scenarios discussed in the chats, excluding outlier conversations where agent may have taken wrong steps. For each valid scenario, consider whether the workflow addresses it comprehensively, and if not, identify the gaps.

Evaluate Coverage Gaps

- Where does the workflow fail to provide guidance or actions for specific scenarios presented in the past chats?
- Are there any special cases (e.g., different membership tiers, exceptions, or advanced user queries) that the workflow misses entirely?
- Ignore outlier conversations where agent actions are inconsistent with most other conversations from the same scenario.
- For each gap identified, explain how the current workflow might be extended or modified to address these uncovered cases.

Suggestions for Enhancing Coverage

- What additions or modifications could be made to the workflow to cover the uncovered scenarios more effectively?

Figure 21: The prompt for identifying coverage gap in a workflow given historical conversations. [Workflow Extraction Baseline: Reflect].

Improve the below dialog workflow based on the provided conversations and the LLM feedback.

- Include only the essential actions needed for issue resolution, excluding unnecessary steps.
- Create a unified, consolidated list of steps without breaking them down by individual conversation.
- Ignore feedback that is specific to an individual customer and does not address a general dialog scenario.
- Formatting Instructions:
 1. Use simple and concise language for each step, mentioning any pre-conditions where applicable.
 2. Organize the steps in a numbered list for clarity.
 3. Include relevant details, such as required inputs and specific conditions, for each step.

Figure 22: The prompt for improving coverage given conversations and LLM-feedback. [Workflow Extraction Baseline: Reflect].

Your task is to identify the steps necessary to resolve the customer's issue based on the provided conversations.

{convs}

Let's first understand the problem and devise a plan to solve the problem.

Figure 23: The prompt for generating plan from historical conversations. [Workflow Extraction Baseline: Plan].

Identify the steps necessary to resolve the customer's issue based on the provided conversations and the plan.

- Include only the essential actions needed for issue resolution, excluding unnecessary steps.
- Use the plan to identify important steps in the workflow.
- Create a unified, consolidated list of steps without breaking them down by individual conversation.
- Formatting Instructions:
 1. Use simple and concise language for each step, mentioning any pre-conditions where applicable.
 2. Organize the steps in a numbered list for clarity.
 3. Include relevant details, such as required inputs and specific conditions, for each step.
- Only generate the final guideline, outlining important steps required for solving customer's issue.

Figure 24: The prompt for generating workflows from plan and historical conversations. [Workflow Extraction Baseline: Plan].

Aggregate and improve upon the following four dialog workflows used to resolve the customer's issue.

Workflow 1:
{wf1}

Workflow 2:
{wf2}

Workflow 3:
{wf3}

Workflow 4:
{wf4}

Use the provided example conversations to decide which of the mentioned workflows is correct at each step or to identify any steps that are missing from all workflows.
{convs}

STEP-BY-STEP INSTRUCTIONS:

- Identify all the steps necessary to resolve the customer's issue.
- Include only the essential actions needed for issue resolution, excluding unnecessary steps.
- Create a unified, consolidated list of steps without breaking them down by individual conversation.

Formatting Instructions:

1. Use simple and concise language for each step, mentioning any pre-conditions where applicable.
2. Organize the steps in a numbered list for clarity.
3. Include relevant details, such as required inputs and specific conditions, for each step.

Figure 25: The prompt for ensemble prompting strategy. [Workflow Extraction Baseline: Ensemble].

Correct the given LLM-generated discussion between two collaborative agents, the Guide and the Implementer.

Reflect on the given discussion:

- Evaluate whether the discussion cover all the steps comprehensively.
- Look for discussion parts that may belong to outlier conversations (e.g., rare scenarios or edge cases) and determine if they should be part of the main workflow.

Correct and Complete:

- Add any missing discussions to bridge gaps and ensure logical continuity.
- Reorganize or exclude discussion about steps in outlier conversations if they do not fit the main workflow.
- Clearly specify the conditions under which any alternative steps or edge cases arise.

Provide the corrected and complete discussion following the format of original discussion between the guide and implementer agents. Do not add any justifications.

Figure 26: The prompt used for refining the QA chain-of-thought [Workflow Extraction Baseline: QA-CoT+Reflect].