

# Frame of Reference: Addressing the Challenges of Common Ground Representation in Situational Dialogs

Anonymous ACL submission

## Abstract

Common ground plays a critical role in situated spoken dialogs, where interlocutors must establish and maintain shared references to entities, events, and relations to sustain coherent interaction in a shared space and over time. With the increasing presence of embodied conversational agents and social robots, the ability to correctly ground this kind of conversational content in order to refer back later also becomes important for dialog systems. Prior studies have demonstrated that LLMs are capable of performing certain grounding acts like acknowledgments. However, relatively little work has investigated their capacity to leverage the grounded information, like in complex scenarios involving space and time (e.g., "let's go to that café near the park we went to yesterday"). To that end, in this work, we evaluate a model's ability to establish common ground by utilizing these "relational references" in the dynamic and shared environments of situated dialogs. We then test multiple methods for representing common ground and further propose approaches to improve their performance.

## 1 Introduction

In dialog, common ground refers to the shared knowledge, beliefs, and assumptions that accumulate between participants as a conversation progresses (Clark and Schaefer, 1989). Conversational Grounding is the process of building and maintaining this shared understanding. While recent advances in Large Language Models (LLMs) have demonstrated proficiency in performing certain grounding acts (Traum and Allen, 1994) such as acknowledging inputs for immediate context (Mohapatra et al., 2024b) - it remains unclear whether these behaviors represent genuine understanding or merely an "illusion of grounding". This distinction is critical: a system may mimic understanding through plausible acknowledgments, yet fail to utilize the information later effectively.



Figure 1: An example of people referring back to the common ground.

This challenge is amplified in dialogs that extend over longer periods. As dialog history grows, systems must move beyond context windows and employ memory management techniques to retrieve information from the established common ground. In order to give the systems a true grounding ability, the acknowledged information must therefore be represented and stored in a format that can be utilized later on, regardless of the size of the dialog. Its importance is particularly pronounced in spoken situational dialogs (dialogs that operate in the dynamic and shared environments) such as Embodied Conversational Agents (ECAs) (Cassell et al., 2001) and Social Robots (Chen et al., 2025) where the conversations can take place over multiple interactions. To become effective collaborators, these agents must be able to deploy information from their ongoing and previous interactions.

This gives rise to three central research questions that we address in this paper :

- RQ1 Benchmarking:** How do we measure the capability of a dialog system to establish a persistent, usable, and useful common ground?
- RQ2 Representation:** In real-world scenarios where dialog history exceeds context window size, how well do commonly used common ground representations perform in this

071	benchmark?	resource-constrained settings (mimicking real-life longer conversations exceeding context lengths of LLMs). Our experiments reveal that standard LLMs struggle with relational references in spoken situated dialogs, even with full context. We also test the performance of multiple commonly used common ground representation techniques for resource-constrained settings. However, they often perform worse than LLMs with full-context due to their failure to capture the relationships between entities which are necessary for accurate retrieval.	119
072	3. <b>RQ3 Improving Performance:</b> Given the		120
073	potential inadequacy of such representations,		121
074	how can we improve our system’s grounding		122
075	capabilities?		123
076	<b>Key Contributions and Findings:</b>		124
077	• <b>Benchmarking Grounding Capabilities:</b>		125
078	<a href="#">Kruijt and Vossen (2022)</a> demonstrated that		126
079	in situated dialogs, people often refer to es-		127
080	tablished common ground using references		128
081	containing relations to specific properties, en-		129
082	tities, or events. In this paper, we group such		130
083	references under the rubric of ‘relational ref-		131
084	erences’. To address RQ1, we thus propose		
085	that a system’s ability to resolve such com-		
086	plex relational references established through		
087	prior dialog can serve as a measure of its ca-		
088	pability to create persistent, useful, and usable		
089	common ground.		
090	For instance, in <a href="#">Figure 1</a> , the user refers to		
091	the phone via the shop where it was bought.		
092	Such recall of entities often lack a unique re-		
093	ference and could equally be described tempo-		
094	rally (e.g., first shop) instead of spatially (e.g.,		
095	shop next to the church). Without a ground-		
096	ing capability that establishes a persistent and		
097	informational common ground (whether inter-		
098	nal or external) that can be utilized later,		
099	a dialog system cannot successfully resolve		
100	these relational references.		
101	We validate <a href="#">Kruijt and Vossen (2022)</a> by ex-		
102	amining two natural dialog corpora to iden-		
103	tify how humans leverage common ground,		
104	finding multiple examples of such complex		
105	relational references. Informed by this,		
106	we introduce a new evaluation benchmark,		
107	IndiRef. It moves beyond the immediate		
108	simple grounding acts such as acknowledg-		
109	ments and creates a benchmark to evaluate		
110	whether dialog systems can resolve complex		
111	references to the established common ground		
112	(e.g., “that restaurant next to the museum that		
113	we went to in New York”), a critical require-		
114	ment for effective long-term conversation in		
115	spoken situated dialog.		
116	• <b>Evaluation of Current Techniques:</b> To an-		
117	swer RQ2, we evaluate LLMs of varying sizes		
118	in both full-dialog context (as a baseline) and		
		• <b>Synthetic Data &amp; RL Training:</b> To address RQ3, we identify a lack of multi-step reasoning capability for spoken dialogs in current LLMs along with a scarcity of high-quality training data. We develop a synthetic data generation pipeline and show that training models on these data using Reinforcement Learning (RL) improves their ability to handle complex references.	132
			133
			134
			135
			136
			137
			138
			139
			140
		By providing a benchmark to test the effective-	141
		ness of the common ground established by the di-	142
		alog system using relational references, exposing	143
		the limitations of current techniques, and demon-	144
		strating the efficacy of RL-based training using	145
		synthetic data, this work provides a pathway to-	146
		ward dialog systems capable of robust, situated	147
		conversational grounding over time.	148
		<b>2 Related Work</b>	149
		<a href="#">Traum and Allen (1994)</a> introduced the concept of	150
		Grounding Acts (GAs) that represent fundamental	151
		conversational moves like acknowledging or repair-	152
		ing. <a href="#">Mohapatra et al. (2024b)</a> showed that modern	153
		LLMs can successfully perform some of these GAs	154
		when the required context is limited to the last few	155
		utterances. The practical challenges of connecting	156
		these local grounding decisions to the long-term	157
		maintenance of a coherent common ground remain	158
		significant ( <a href="#">Mohapatra et al., 2024a</a> ).	159
		Earlier attempts to represent and manage the	160
		common ground, such as Domain Reference the-	161
		ory ( <a href="#">Denis, 2010</a> ), were limited in nature due to	162
		their reliance on rule-based approaches. Founda-	163
		tional work has largely focused on structured, sym-	164
		bolic approaches for task-oriented dialogs. Meth-	165
		ods like Dialog State Tracking (DST) ( <a href="#">Jacqmin</a>	166
		<a href="#">et al., 2022</a> ) and frame-based systems ( <a href="#">Blache and</a>	167

Houlès, 2021) represent conversational context as slots and values, capturing the user’s goals and constraints. While highly interpretable and effective for well-defined tasks, these approaches often lack the flexibility to handle the dynamic and unpredictable nature of long-term natural interactions.

More flexible representations such as dynamic knowledge graphs (KGs) (Walker et al., 2022; Bu et al., 2025) have been proposed to model entities and their relationships as they emerge in conversation. KGs allow for a richer, more interconnected model of context and can naturally handle entity-based references by linking mentions to nodes in the graph. However, a significant challenge arises in spoken dialogs where entities often lack a single, stable referring expression (e.g. ‘park next to the restaurant’ or ‘duck park’) or relation (e.g., ‘the room that has a TV’ and ‘TV is next to a table’ implies ‘room has a table’).

Neural memory architectures (Henaff et al., 2016; Rajendran et al., 2019; Wu and Yu, 2024) encode dialog history in distributed vector representations. However, this "black-box" nature, makes it difficult to verify whether their conversational behaviors, such as acknowledgments, stem from a truly grounded understanding or pattern matching, especially, when there is an absence of an evaluation metric.

Beyond tracking facts, a robust common ground representation requires modeling participant beliefs and perspectives (Traum and Larsson, 2003; Schlangen and Skantze, 2009), necessitating a shift from state tracking to belief tracking. Further frameworks (Di Maro et al., 2021; Khebour et al., 2024; Qiu et al., 2024) have proposed to improve reliability by verifying mutual understanding through the tracking of shared facts, evidence, and unresolved information. However, they still fail to address the core representational challenge: handling entities with ambiguous references and resolving their mentions over time. Moreover, most common ground tracking research (Shridhar et al., 2020; Williams et al., 2019; Madge et al., 2025) has been conducted on the immediate conversational frame as compared to grounding over multiple events. Wu et al. (2025) offers a dataset for multi-session dialogs but is devoid of complex relational references of situational dialogs.

With the increasing use of LLMs in spoken dialog systems, their finite context window poses a constraint on maintaining long-term conversational memory. A common workaround is dia-

log summarization, where models generate concise summaries of past exchanges to be injected into future contexts, enabling a personalized experience (Ramprasad et al., 2024). A more scalable solution involves Retrieval-Augmented Generation (RAG) (Lewis et al., 2020), which offloads the entire dialog history to external memory. During conversation, relevant excerpts are retrieved and fed into the LLM’s context, a technique used to maintain long-term persona memory in conversational agents (Shuster et al., 2021). However, they have not been tested on their ability to resolve complex relational references. Consequently, an effective common ground representation must maintain information that is efficiently retrievable, dynamically updatable, and capable of tracking participant perspectives.

### 3 RQ1 : Benchmarking

To measure grounding beyond superficial acknowledgment patterns, we turned to how humans leverage common ground. Inspired by Kruijt and Vossen (2022), who showed that people frequently use relational references to access shared knowledge, we propose that a model’s ability to resolve such references can serve as a robust metric for a model’s grounding capabilities, one that can deal with dialogs where entities may not have a single, unambiguous referring expression.

To better understand the mechanisms of relational reference, we analyzed two dialog datasets previously used in grounding research. The Meetup dataset (Ilinykh et al., 2019) places two participants on a 2D grid and tasks them with finding each other despite only seeing their own locations. Success in this setting depends on describing rooms, maintaining a record of shared knowledge, modeling the partner’s perspective, and clarifying ambiguities, all of which are central challenges for grounding. In contrast, the Spot the Difference (STD) (Lopes et al., 2018) dataset involves two speakers, each given slightly different images, who must coordinate over an audio channel to identify discrepancies. While Meetup provides a scenario closer to real world situations for building common ground and referring back to it, we also used STD as it provides a different context to help test the robustness of our representation technique. Meetup corpus contains 5131 utterances, and Spot the difference contains 54 dialogs with 4934 utterances.

We manually analyzed each utterance in the dialog datasets to identify references to previously

Phenomena	Dataset	
	Meetup	Spot the difference
Temporal Reference	5	4
Spatial Reference	34	72
Comparative Reference	0	6
Attributive Reference	33	8
Inferred grounding	177	67

Table 1: Frequency of different reference types

mentioned entities (details in Appendix A). This analysis yielded a taxonomy of four primary categories: 1) **Temporal**, referring to an entity by its order of appearance (e.g., "*the Thai restaurant that we went to after watching Spiderman*"); 2) **Spatial**, locating an entity relative to another (e.g., "*the bottle on the table*"); 3) **Attributive**, describing an entity by its intrinsic properties (is-a: "*yellow house*" or has-a: "*sofa room*"); and 4) **Comparative**, using a relative property (e.g., "*bigger chair*"). Furthermore, we observed implicit grounding mechanisms akin to the "indirect acknowledgments" identified by Mohapatra et al. (2024a). We termed them **inferred grounding**, where participants ground information that was implied rather than explicitly stated. For instance, when a navigator is asked by their partner whether they see a "*pink towel*" in their bathroom, that navigator also grounds the fact that their partner has visited a bathroom with a pink towel, later referring to it as "*the bathroom with the pink towel that you had seen*". The frequencies found from our analysis are presented in Table 1. Given the general difficulty of isolating specific phenomena within a limited number of spoken dialog corpora (Allison et al., 2006; Crible, 2022), the identification of multiple occurrences in these two small datasets (though limited in number) offers evidence for the use of ‘relational references’ to refer back to established common ground.

### 3.1 IndiRef data creation

Our goal here is to extend beyond prior work to evaluate a model’s ability to establish a persistent, useful and usable common ground. This capability becomes critical given the finite context windows of LLMs, which necessitate the efficient storage and retrieval of shared information. To this end, we introduce IndiRef, a benchmark consisting of Question/Answer (Q/A) pairs where each question seeks information regarding an entity that is identified via a relational reference.

We focus on three primary reference types based on their conversational frequency in Table 1: **Temporal**, **Spatial**, and **Attributive**. We also used the fourth category **inferred grounding**, to specifi-

cally probe the model’s understanding of implicit grounding. As noted previously (Table 1), the natural density of these references in the Meetup and STD datasets is low; therefore, we manually constructed 400 Q/A pairs (100 per category) to ensure sufficient coverage and unambiguous answers.

**Robustness:** IndiRef is designed to be adversarial to superficial heuristics. By taking dialogs containing multiple entities of the same type (e.g., containing several ‘sofas’), we force the model to use the provided contextual reference (e.g., "*color of the sofa in the tv room*" vs "*color of the sofa in the room in front of the bathroom*") to disambiguate the target entity rather than relying on simple keyword matching. A correct answer thus serves as strong evidence of successful reference resolution, indicating that the model was able to ground the information properly in the first place.

**Perspective:** Additionally, the benchmark evaluates perspective-taking through the use of deictics (Shopen, 1985). Questions are framed as the final utterance of a participant, requiring the model to respond as the interlocutor. This tests the model’s ability to distinguish between “self” and “other” within the shared context. For instance, correctly answering "*What was the color of the sofa in the second house that I visited?*" requires the model to recall the information shared by the questioner about their sequence of actions. However, this changes fundamentally if the question replaces ‘I’ with ‘you’, which would require the model to access shared information pertaining to its own character’s experiences, thereby testing its ability to maintain distinct perspectives within the common ground.

The multi-dimensional nature of IndiRef provides a crucial metric to guide the development of models capable of creating more robust and usable common ground. For further details on benchmark construction, illustrative examples, and our human validation study that yielded an accuracy of 87.5%, please refer to Appendix B.

## 4 RQ2 : Representation

### 4.1 Baseline - Full Dialog Context

To evaluate common ground representation capabilities, we adopt a framework inspired by (Zhang et al., 2025) utilizing three operators: Writer ( $W$ ), Reader ( $R$ ), and Generator ( $G$ ). As a baseline, we assess the LLM’s internal ability to ground context by providing the full dialog history ( $h_t$ ) (contain-

LLMs	Meetup				Spot the difference			
	Temporal	Spatial	Attributive	Inferred	Temporal	Spatial	Attributive	Inferred
Gemma2 - 2B	0.20 / 0.18	0.18 / 0.16	0.24 / 0.26	0.26 / 0.16	0.18 / 0.12	0.18 / 0.10	0.20 / 0.12	0.16 / 0.10
Llama3.1 - 8B	0.38 / 0.32	0.46 / 0.38	0.46 / <b>0.44</b>	0.20 / 0.20	0.40 / <b>0.38</b>	<b>0.30 / 0.32</b>	<b>0.36 / 0.36</b>	0.26 / 0.16
Gemma2 - 9B	0.48 / 0.38	0.50 / 0.46	0.38 / 0.40	0.26 / 0.18	0.34 / 0.20	0.24 / 0.16	0.28 / 0.16	0.20 / 0.10
Gemma2 - 27B	<b>0.50 / 0.44</b>	<b>0.58 / 0.56</b>	<b>0.48 / 0.44</b>	0.28 / 0.26	0.40 / 0.30	0.28 / 0.20	0.32 / 0.26	0.26 / 0.12
Qwen QwQ - 32B	0.38 / 0.32	0.52 / 0.38	0.44 / 0.40	<b>0.40 / 0.40</b>	<b>0.42 / 0.34</b>	0.26 / 0.22	0.26 / 0.20	<b>0.28 / 0.20</b>

Table 2: FEM ( $\uparrow$ ) / LLM-as-Judge ( $\uparrow$ ) accuracy for baselines where the entire dialog history is used as the context to answer the question. Best results in **bold**.

LLMs	Meetup				Spot the difference			
	Temporal	Spatial	Attributive	Inferred	Temporal	Spatial	Attributive	Inferred
Full Dialog [from Table 2]	0.38 / 0.32	<b>0.46 / 0.38</b>	<b>0.46 / 0.44</b>	0.20 / 0.20	<b>0.40 / 0.38</b>	<b>0.30 / 0.32</b>	<b>0.36 / 0.36</b>	<b>0.26 / 0.16</b>
Summarisation	0.32 / 0.28	0.34 / 0.26	0.30 / 0.25	<b>0.28 / 0.18</b>	0.22 / 0.16	0.24 / 0.16	0.34 / 0.26	0.20 / 0.12
Chunking (NV-Embed-V2)	0.24 / 0.20	0.08 / 0.06	0.16 / 0.08	0.22 / 0.24	0.02 / 0.02	0.04 / 0	0.06 / 0.02	0.12 / 0.08
Chunking (BM25)	0.26 / 0.24	0.20 / 0.16	0.20 / 0.18	0.24 / <b>0.26</b>	0.04 / 0	0 / 0	0.10 / 0.08	0.12 / 0.10
Agentic Ontology-based	<b>0.40 / 0.36</b>	0.38 / 0.34	0.38 / 0.30	0.24 / 0.22	0.20 / 0.14	0.12 / 0.14	0.20 / 0.24	0.18 / 0.14

Table 3: FEM ( $\uparrow$ ) / LLM-as-Judge ( $\uparrow$ ) accuracy for various representations using Llama 3.1 8B. Best in **bold**.

ing the speaker information, utterances, and the time stamps) and tasking them with generating a response ( $u_{t+1}$ ) to a question in the preceding utterance. We test a range of open-weight instruction-tuned models: Gemma2 (2B, 9B, 27B)(Gemma, 2024), Llama 3.1 8B(Llama, 2024), and reasoning-focused Qwen-QWQ-32B (Qwen, 2025). In this setup,  $W$  appends each new utterance to the history,  $R$  provides the entire history to the model, and  $G$  produces the subsequent utterance.

$$W(h_{t-1}, u_{t-1}) = h_t; R(u_t) = h_t; G(h_t, u_t) = u_{t+1}$$

## 4.2 Resource Constrained Setting

As dialogs extend over time, retaining the complete interaction history becomes prohibitive in terms of context window, computation and inference latency. Although the IndiRef dataset fits within modern LLM context windows, we artificially restrict context access to simulate the ‘longer context’ scenarios required for persistent agents. Despite lacking extreme temporal length, the underlying corpora of IndiRef, particularly the Meetup dataset, offer a critical testing ground because of their fragmented context i.e., speakers discuss and recall information exchanged across different contexts(rooms in this case) as the conversation progresses. Under our restricted context settings, this dynamic effectively simulates the episodic transitions of long-term interactions, forcing the model to resolve references that lie outside the immediate context window.

We compare three commonly used methods for representing common ground under these constraints: summarization, utterance chunking, and ontology creation.

**Summarization:** A method shown by Padmakumar et al. (2023) to be a highly effective representation of open-dialog context for transformer-based models. Its utility is further underscored by its adoption in modern LLM chatbots for personalization (Zhang et al., 2024) by summarizing each session separately. In this scheme,  $W$  creates a summary ( $s_t$ ) of the whole dialog, which is stored, retrieved by  $R$ , and used by  $G$  to form the next utterance.

$$W(s_{t-1}, u_{t-1}) = s_t; R(u_t) = s_t; G(s_t, u_t) = u_{t+1}$$

**Utterance Chunking:** Our first method to explicitly group information to enhance retrieval. Here,  $W$  segments the dialog into overlapping chunks of utterances ( $c_i$ ). For our experiments, we define a chunk as seven utterances with a stride of three to ensure informational continuity.  $R$  then retrieves the top three chunks most relevant to the user’s utterance by matching its embedding against the embeddings of each chunk. To isolate the performance of this structural approach, our Reader only uses the embeddings of the concatenated utterances within a chunk, and does not use additional embeddings of chunk summaries as done by (Wu et al., 2025). To better understand how embeddings affect the retrieval of important context, we tested two types of embedders: a sparse model (BM25 (Robertson et al., 1994)) and a dense model (NV-Embed-V2 (Lee et al., 2024)).

$$W(h_{t-1}, u_t) = [c_1, \dots, c_i]$$

$$R(u_{t+1}) = \text{top}_k(\text{sim}(e(u_{t+1}), [e(c_1), \dots, e(c_i)]))$$

$$G(u_t, [c_1, \dots, c_k]) = u_{t+1}$$

**Ontology:** We use an **agentic** approach which investigates the use of an **ontology** ( $O$ ) to create a

structured common ground. The Writer ( $W$ ) processes dialog chunks to extract key entities, assigning each a unique id and mapping its properties, inter-entity relations and source speaker as a proxy for estimating each speaker’s belief. This rich structure enables complex relational reasoning (e.g., if ‘*place A has object b*’ and ‘*b is kept on top of c*’, then ‘*A also has c*’). To preserve temporal context,  $W$  also maintains a chronological log of events, which includes the involved entity IDs and speakers, serving as a primary index for querying. Given this highly structured knowledge,  $R$  uses a multi-step query process to retrieve and process information before sending it to  $G$ . It first decomposes the user’s request into a series of sub-queries ( $q$ ), each prefixed with a command:  $RAG[n]$ ,  $Process$ , or  $Final$ . A  $RAG[n]$  query retrieves the top- $n$  relevant contexts ( $c$ ) from the extracted common ground. Here  $n$  is also generated by the LLM. A  $Process$  query then performs operations on this retrieved data, such as filtering or selecting specific items. The final retrieved context and the  $Final$  query are passed to  $G$  to synthesize the final answer. All operations are implemented using langchain (Chase, 2022) for the vectorstore and NV-Embed-V2 for embedding generation.

$$\begin{aligned}
 W(O_{t-1}, u_{t-1}) &= O_t \\
 R(u_t) &= [q_1 \dots q_m] \\
 R(q_i, O_t, c_{i-1}) &= c_i \\
 G(q_m, c_{m-1}) &= u_{t+1}
 \end{aligned}$$

### 4.3 Evaluation

To comprehensively evaluate performance, we employ a Fuzzy Exact Matching (FEM) metric. For each question, an answer is considered correct if it contains a predefined keyword or phrase. Recognizing that keyword presence alone may not guarantee accuracy, we also incorporate an LLM-as-a-judge approach (Zheng et al., 2023). A Qwen2.5-32B-Instruct (Yang et al., 2024) model, provided with the question and both the correct and generated answers, determines correctness based on semantic equivalence. Crucially, dialog context was intentionally excluded from the LLM’s input to ensure its judgment focuses solely on the answers’ semantic similarities. Since the relational references point to only one correct referent, the overall performance for both fuzzy matching and the LLM-as-a-judge approach was quantified using accuracy. To validate the reliability of our automated judge, we conducted a human evaluation on a subset of answer pairs. We observed a 90% agreement rate,

demonstrating a strong alignment between human and LLM judgments (details in Appendix C).

Table 2 presents the performance of various LLMs when provided with the full dialog context. For Meetup, results indicate that model performance improves with increasing size, with Gemma2-2B consistently exhibiting the lowest accuracies across all categories. The relative performance of Fuzzy Matching and LLM-as-a-judge remains consistent across models. Slight variations occur when Fuzzy Matching scores differ due to longer answers with multiple entities or alternate representations (e.g., digits vs. words). Surprisingly, Gemma2-27B emerges as the top performer across all evaluated models, outperforming even Qwen-QWQ which is a slightly larger model specifically trained for enhanced reasoning. Further analysis of Qwen-QWQ’s responses revealed difficulties in logical reasoning and frequent hallucinations (Appendix; Figure 12), potentially attributable to a scarcity of spoken dialogs in its training data. Despite this, Qwen-QWQ excelled in inferred grounding questions, a category where other models performed poorly, underscoring the critical role of reasoning for grounding implicit information. It also had better fuzzy exact match scores due to its verbose nature. However, in STD, Gemma models underperformed compared to Llama. We hypothesize that the longer dialog context in STD as compared to Meetup helped Llama which uses RoPE over Gemma which uses sliding window attention. Furthermore, the models demonstrated significant limitations in differentiating between speaker perspectives (e.g., ‘your’ vs ‘mine’) and in isolating information conveyed by individual speakers, often resulting in the erroneous merging of conversational content. It is critical to note that **none of the baselines achieved an accuracy exceeding 50% across all the categories.**

Table 3 details the performance of various resource-constrained representation techniques. Due to space constraints, we primarily report results for the Llama3.1-8B model; results for the Gemma2 2B and 9B models are provided in Appendix F. The general trend across all models remained consistent. We did not test larger models due to resource limitations. All the representation techniques yielded poorer performance compared to the baseline, a likely consequence of information loss during the extraction of relevant information. We noticed that sparse embedding (BM25) demonstrated a slight advan-

LLMs	Meetup				Spot the difference			
	Temporal	Spatial	Attributive	Inferred	Temporal	Spatial	Attributive	Inferred
Original Setting [Table 2]	0.38 / 0.32	0.46 / 0.38	0.46 / <b>0.44</b>	0.20 / 0.20	0.40 / 0.38	0.30 / 0.32	0.36 / 0.36	0.26 / 0.16
In-Context Learning	<b>0.60 / 0.56</b>	<b>0.58 / 0.54</b>	<b>0.62 / 0.58</b>	0.42 / 0.34	0.48 / 0.52	<b>0.48 / 0.40</b>	0.52 / 0.54	<b>0.36 / 0.44</b>
GRPO Trained	0.58 / 0.52	<b>0.66 / 0.54</b>	<b>0.62 / 0.60</b>	<b>0.46 / 0.42</b>	<b>0.52 / 0.54</b>	<b>0.48 / 0.42</b>	<b>0.56 / 0.58</b>	<b>0.38 / 0.42</b>

Table 4: FEM ( $\uparrow$ ) / LLM-as-Judge ( $\uparrow$ ) accuracy on entire dialog history for enhanced Llama 3.1-8B. Best in **bold**.

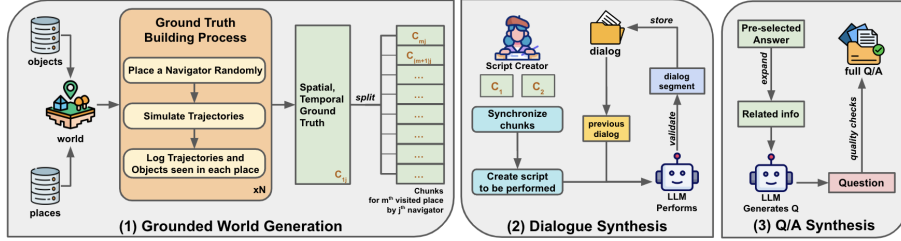


Figure 2: Pipeline for the synthetic data generation involving three main phases: (1) Grounded World Generation, creates the factual basis of the scenario; (2) Dialogue Synthesis, where LLM generates conversation constrained by the script; and (3) Q/A Synthesis, which extracts and formats question-answer pairs for training.

tage over dense embedding, which may be attributed to its enhanced capability in named-entity retrieval. The agentic ontology-based system outperformed other resource-constrained representation techniques. While it performed relatively well on the Meetup dataset, it struggled with the STD dataset, often merging the representations of images of both participants due to their similarity. However, its better performance suggests that a multi-step retrieval process, particularly one that explicitly extracts speaker-specific information regarding entities and events, helps the LLM understand the context better.

## 5 RQ3 : Improving Performance

Error Analysis of the thinking tokens from baseline models indicated that they primarily struggled to link references to entities requiring multi-hop reasoning. Despite general advancements in reasoning capabilities, our findings showed that LLMs continue to face significant difficulties within situated dialog settings. Further, a significant challenge lies in their understanding of overlapping information from multiple speakers, preventing the construction of a reliable grounding representation. Thus, our initial research focused on enhancing the performance of baseline models leveraging the entire dialog history. We implemented In-Context Learning (Dong et al., 2022) by augmenting the input prompts with examples for each reference type to instruct the LLMs on how to accurately ground the context within the dialogs. The prompts (see Appendix G) were further enriched with explicit guidance on decoupling the information contributed by each speaker, unless a common entity was being

discussed. As shown in Table 4, this resulted in a significant improvement in LLM performance across all categories in both datasets. These findings indicate that enhanced prompt design can substantially improve the models' comprehension of spoken conversations, a critical prerequisite for effective grounding.

To enable the model to develop inherent grounding capabilities applicable to diverse settings, we used reinforcement learning to improve the model's multi-hop reasoning capability. We used Group Relative Policy Optimization (GRPO) (Shao et al., 2024) to train the model to identify the important context by providing positive rewards for each accurate response. A primary challenge was the absence of suitable training data for situational spoken dialogs. Thus, we generated synthetic data to resolve data scarcity. This approach not only resolved data scarcity but also facilitated arbitrary scaling, fine-grained control over domains, and improved coverage of edge cases.

### 5.1 Synthetic Data

Given that off-the-shelf LLMs lack proficiency in spoken dialog, we designed a systematic, multi-phase synthetic data generation framework with an "environment-first, dialog-later" approach. Inspired by procedural content generation (Smith, 2015) and recent advances in synthetic dialog generation (Suresh et al., 2025), our framework is engineered to produce a grounded situated dialog corpus. The process is structured into three phases as shown in Figure 2:

**Grounded World Generation:** First, a rich and verifiable simulated 'world' is programmatically constructed, starting from a list of places and re-

LLMs	Meetup				Spot the difference			
	Temporal	Spatial	Attributive	Inferred	Temporal	Spatial	Attributive	Inferred
Without GRPO training [Table 3]	0.40 / 0.36	0.38 / 0.34	0.38 / 0.30	0.24 / 0.22	0.20 / 0.14	0.12 / 0.14	0.20 / 0.24	0.18 / 0.14
With GRPO training	0.48 / 0.46	0.44 / 0.42	0.52 / 0.44	0.36 / 0.38	0.20 / 0.24	0.22 / 0.28	0.28 / 0.32	0.16 / 0.20

Table 5: FEM ( $\uparrow$ ) / LLM-as-Judge ( $\uparrow$ ) accuracy for agentic ontology-based representation after GRPO training.

lated objects previously generated (see Appendix D). Two algorithmic navigators explore this world, and their experiences are logged to establish the spatio-temporal ground truth (i.e., what they saw, where they went, etc.) that we can then split into smaller chunks, where  $C_{m,j}$  is the chunk for the  $m^{th}$  visited place by the  $j^{th}$  navigator.

**Dialog Synthesis:** We then synthesize dialogs where the two navigators describe their locations to find each other. Dialogs are constrained by the rules and facts of the created world. A programmatic Script Creator synchronizes the chunks of the two Navigators and generates the script for the LLM to follow, thus generating dialog segments based on the event logs. In this way, we rely on the Script Creator for the ‘reasoning’ part of the task (i.e., noticing similarities, addressing ambiguities, resolving references, etc), while the LLM focuses on constrained Utterance Generation. This generation process is iterative and considers the previously generated utterances for cohesiveness.

**Q/A Synthesis:** The final stage automatically identifies scenarios to generate accurate Question-Answer (Q/A) pairs requiring complex reference resolution. The process is deterministic: a fact is preselected from the Ground Truth as the ‘expected answer’, and the LLM is prompted to generate a question, constrained to the context, that elicits this specific answer. These are designed to probe various aspects of common ground, including spatial, temporal, and relational facts, and undergo a strict validation process. We generated  $\sim 600$  Q/A pairs from various synthetic dialog scenarios.

Table 4 demonstrates that reinforcement learning on the synthetic data improved model accuracy by 15-20% not only on the Meetup dataset (which closely resembles the training data), but also on the Spot-the-Difference dataset. This indicates that the models successfully acquired the underlying concepts needed to resolve such relational references to the common ground. We note that for the incorrect responses, the models would often reason in the right direction but then hallucinate while providing the final answer. We expect this to improve with larger models but, due to lack of computational resources, we did not train on them.

To assess whether the enhanced contextual understanding of GRPO-trained models improves their representational capabilities, we applied the agentic ontology-based representation using the same GRPO-trained model trained on synthetic Q/A pairs, given its promising results in earlier experiments. As shown in Table 5, the GRPO-trained model outperformed its non-GRPO counterpart on the Meetup dataset but continued to struggle with the Spot-the-Difference dataset. The main issue concerned representing the two images as distinct within the common ground, as the models frequently merged them due to their high degree of similarity. While more carefully designed task-specific prompts could improve performance, these would limit generalizability. Thus, our findings fall short of our broader goal: developing a universal model for a persistent and useful common ground representation capable of resolving complex relational references to it while relying only on spoken utterances and minimal contextual descriptions.

## 6 Conclusion and Future Work

In this work, we used relational references to measure the capability of a dialog system to establish a persistent and useful common ground by developing the IndiRef benchmark. While our baselines with full dialog context struggled to consistently achieve an accuracy of 50%, models mimicking resource-constrained settings performed even worse. Our experiments showed that both in-context learning and GRPO training with synthetic data improved an LLM’s ability to resolve relational references from the full dialog history. The same GRPO-trained models moderately also enhanced common ground extraction for long-term use in agentic setups due to better multi-step reasoning. However, inconsistent performance highlight the need for more robust representation methods. For future work, while our study focused on interpretable textual representations of the common ground, exploring more efficient parametric methods is a promising direction. We believe that our benchmark and proposed methods provide a solid foundation for advancing research into more robust common ground representation in dialog agents.

## 694 Limitations

695 Certain limitations offer avenues for future re-  
696 search. First, the dataset used for our test cases,  
697 while selected to approximate spontaneous spoken  
698 dialog, has inherent constraints. The ideal dataset  
699 would comprise spontaneous, multi-session spoken  
700 dialogs. The logistical challenges in acquiring such  
701 data necessitated the use of two specialized dia-  
702 log corpora, Meetup and Spot the Difference that,  
703 while effective in providing a range of grounding  
704 phenomena through a game-based framework, may  
705 not entirely capture the nuances of spontaneous  
706 multi-topic and multi-session human interaction.  
707 Although the corpora may lack the full breadth  
708 of really long conversation, the dataset used for  
709 our corpora, especially Meetup, provides a critical  
710 testing ground by featuring fragmented context i.e.  
711 speakers discuss and recall information exchanged  
712 in previously visited simulated environments (dif-  
713 ferent rooms), effectively mimicking the transition  
714 between distinct contexts. A resource-constrained  
715 setting using such dialogs thus allows us to rigor-  
716 ously test how different grounding representations  
717 sustain performance under such fragmented context  
718 constraints which are also inherent to long-term,  
719 multi-session interactions.

720 Second, due to computational resource limita-  
721 tions, our empirical evaluation was restricted to the  
722 8-billion parameter Llama 3.1 instruction-tuned  
723 model especially for the GRPO-based training. A  
724 central hypothesis in current NLP research is that  
725 increasing model scale can reduce issues like hal-  
726 lucination. While we conjecture that our findings  
727 would benefit from this effect, validating this hy-  
728 pothesis was beyond the scope of our resources.  
729 A valuable direction for future work would be to  
730 replicate and extend our experiments using larger  
731 models.

732 Lastly, while our test cases provide a solid founda-  
733 tion for testing the model’s ability to use estab-  
734 lished common ground, they do not comprehen-  
735 sively test it’s capacity to handle ‘assumptions’ and  
736 ‘beliefs’. These deeper cognitive elements, while  
737 tilting towards theory of mind research, are still  
738 essential for robust conversational grounding and  
739 present a clear direction for future investigation.

## 740 References

741 Ben Allison, David Guthrie, and Louise Guthrie. 2006.  
742 Another look at the data sparsity problem. In *Text,*

*Speech and Dialogue*, pages 327–334, Berlin, Hei-  
delberg. Springer Berlin Heidelberg.

Philippe Blache and Matthis Houlès. 2021. [Common ground, frames and slots for comprehension in dialogue systems](#). pages 33–43.

Chenyang Bu, Guojie Chang, Zihao Chen, CunYuan Dang, Zhize Wu, Yi He, and Xindong Wu. 2025. [Query-driven multimodal GraphRAG: Dynamic local knowledge graph construction for online reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 21360–21380, Vienna, Austria. Association for Computational Linguistics.

J Cassell, T Bickmore, L Campbell, H Vilhjálms-  
son, and H Yan. 2001. [More than just a pretty face: conversational protocols and the affordances of embodiment](#). *Knowledge-Based Systems*, 14(1):55–64.

Harrison Chase. 2022. Langchain. <https://github.com/langchain-ai/langchain>.

Huili Chen, Yubin Kim, Kejia Patterson, Cynthia Breazeal, and Hae Won Park. 2025. [Social robots as conversational catalysts: Enhancing long-term human-human interaction at home](#). *Science Robotics*, 10(100):eadk3307.

Herbert H. Clark and Ed Schaefer. 1989. Contributing to discourse. *Cogn. Sci.*, 13:259–294.

Ludivine Crible. 2022. [Studying discourse from corpus and experimental data: Bridging the methodological gap](#). *Discours*, 30.

Alexandre Denis. 2010. [Generating referring expressions with reference domain theory](#). In *Proceedings of the 6th International Natural Language Generation Conference*. Association for Computational Linguistics.

Maria Di Maro, Antonio Origlia, and Francesco Cutugno. 2021. [Cutting melted butter? common ground inconsistencies management in dialogue systems using graph databases](#). In *IJCoL, 7-1, 2*, pages 157–190.

Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. 2022. [A survey on in-context learning](#). In *Conference on Empirical Methods in Natural Language Processing*.

Team Gemma. 2024. [Gemma 2: Improving open language models at a practical size](#). *ArXiv*, abs/2408.00118.

Mikael Henaff, Jason Weston, Arthur Szlam, Antoine Bordes, and Yann LeCun. 2016. [Tracking the world state with recurrent entity networks](#). *ArXiv*, abs/1612.03969.

Nikolai Ilinykh, Sina Zarriß, and David Schlangen. 2019. [Meet up! a corpus of joint activity dialogues in a visual environment](#). In *Proceedings of the 23rd*

743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755  
756  
757  
758  
759  
760  
761  
762  
763  
764  
765  
766  
767  
768  
769  
770  
771  
772  
773  
774  
775  
776  
777  
778  
779  
780  
781  
782  
783  
784  
785  
786  
787  
788  
789  
790  
791  
792  
793  
794  
795



907	Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. <a href="#">Retrieval augmentation reduces hallucination in conversation</a> . In <i>Findings of the Association for Computational Linguistics: EMNLP 2021</i> , pages 3784–3803, Punta Cana, Dominican Republic. Association for Computational Linguistics.	962
908		963
909		964
910		
911		
912		
913		
914	Gillian Smith. 2015. An analog history of procedural content generation. In <i>FDG</i> , pages 1–6. Boston, MA.	
915		
916	Sathya Krishnan Suresh, Wu Mengjun, Tushar Pranav, and EngSiong Chng. 2025. <a href="#">DiaSynth: Synthetic dialogue generation framework for low resource dialogue applications</a> . In <i>Findings of the Association for Computational Linguistics: NAACL 2025</i> , pages 673–690, Albuquerque, New Mexico. Association for Computational Linguistics.	965
917		966
918		967
919		968
920		969
921		
922		
923	David Traum and James Allen. 1994. A "speech acts" approach to grounding in conversation.	
924		
925	David R. Traum and Staffan Larsson. 2003. <i>The Information State Approach to Dialogue Management</i> , pages 325–353. Springer Netherlands, Dordrecht.	
926		
927		
928	Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Galouédec. 2020. Trl: Transformer reinforcement learning. <a href="https://github.com/huggingface/trl">https://github.com/huggingface/trl</a> .	
929		
930		
931		
932		
933	Nicholas Thomas Walker, Stefan Ultes, and Pierre Lison. 2022. <a href="#">Graphwoz: Dialogue management with conversational knowledge graphs</a> . <i>ArXiv</i> , abs/2211.12852.	
934		
935		
936		
937	Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: deep self-attention distillation for task-agnostic compression of pre-trained transformers. In <i>Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20</i> , Red Hook, NY, USA. Curran Associates Inc.	
938		
939		
940		
941		
942		
943		
944	Tom Williams, Fereshta Yazdani, Prasanth Suresh, Matthias Scheutz, and Michael Beetz. 2019. <a href="#">Dempster-shafer theoretic resolution of referential ambiguity</a> . <i>Auton. Robots</i> , 43(2):389–414.	
945		
946		
947		
948	Di Wu, Hongwei Wang, Wenhao Yu, Yuwei Zhang, Kai-Wei Chang, and Dong Yu. 2025. <a href="#">Longmemeval: Benchmarking chat assistants on long-term interactive memory</a> . <i>ICLR</i> , abs/2410.10813.	
949		
950		
951		
952	Qingyang Wu and Zhou Yu. 2024. <a href="#">Stateful memory-augmented transformers for efficient dialogue modeling</a> . In <i>Findings of the Association for Computational Linguistics: EAACL 2024</i> , pages 853–867, St. Julian's, Malta. Association for Computational Linguistics.	
953		
954		
955		
956		
957		
958	An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 23 others. 2024. Qwen2.5 technical report. <i>arXiv preprint arXiv:2412.15115</i> .	962
959		963
960		964
961		
	Zeyu Zhang, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. 2025. <a href="#">A survey on the memory mechanism of large language model-based agents</a> . <i>ACM Trans. Inf. Syst.</i> , 43(6).	965
		966
		967
		968
		969
	Zhehao Zhang, Ryan A. Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, Ruiyi Zhang, Jiuxiang Gu, Tyler Derr, Hongjie Chen, Ju-Ying Wu, Xiang Chen, Zichao Wang, Subrata Mitra, Nedim Lipka, and 2 others. 2024. <a href="#">Personalization of large language models: A survey</a> . <i>ArXiv</i> , abs/2411.00027.	970
		971
		972
		973
		974
		975
		976
		977
	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena. In <i>Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS '23</i> , Red Hook, NY, USA. Curran Associates Inc.	978
		979
		980
		981
		982
		983
		984
		985
	<b>A Conversational Data Analysis</b>	986
	The reference analysis in our conversational data was conducted through an iterative annotation process involving three independent raters with a background in Natural Language Processing. As we wanted to inductively discover the reference typology, we did not start with a pre-defined schema. Hence, the methodology was executed in multiple phases: We first took a few dialogs and each person marked the phrases of 'relational' references independently where we used the term 'relational reference' to group a reference to an entity or an event using associated contextual information, such as time, space, interlocutor identity etc. After this process, we did an Inter-rater reliability on the reference phrases discovered by the three annotators. On achieving a Krippendorff's Alpha of 0.84, we annotated such reference phrases for the rest of the corpora. Once completed, we grouped the reference phrases manually and named each of the phrases. This led to the specific relational reference groups discussed in the paper i.e. Spatial, Temporal, Comparative, Attributive and Indirect.	987
		988
		989
		990
		991
		992
		993
		994
		995
		996
		997
		998
		999
		1000
		1001
		1002
		1003
		1004
		1005
		1006
		1007
		1008
	<b>B Benchmark Creation</b>	1009
	The creation of our test cases involved an initial selection of 50 dialogs per category from both	1010
		1011

1012 datasets. The authors of this paper then under- 1064  
 1013 took the task of creating question/answer pairs that 1065  
 1014 contained relational references to entities within 1066  
 1015 the established common ground of each dialog. A 1067  
 1016 key methodological consideration was the exclu- 1068  
 1017 sive use of conversational data to formulate these 1069  
 1018 questions. This approach ensures that the com- 1070  
 1019 mon ground is derived solely from the information 1071  
 1020 discussed, rather than any external knowledge in- 1072  
 1021 dividual participants may have from the images  
 1022 they viewed. For validation, the benchmark was  
 1023 evaluated by 10 volunteers with expertise in Nat-  
 1024 ural Language Processing. These volunteers were  
 1025 provided with the dialogs and asked to answer the  
 1026 corresponding questions. We randomly selected  
 1027 10% of test cases from each category for this vali-  
 1028 dation process and each test case was provided to  
 1029 two volunteers. We then checked the accuracy of  
 1030 the answers from the human evaluators using the  
 1031 gold answers of the benchmark. The accuracy rate  
 1032 of 87.5% affirmed the correctness of our question-  
 1033 answer pairs in the benchmark. Since no human  
 1034 evaluator pair received more than one question, we  
 1035 didn't calculate the inter-rater reliability between  
 1036 them. However, we would like to mention that we  
 1037 didn't find a single case where both the annota-  
 1038 tors got the answer wrong(i.e. different from gold  
 1039 answer). Thus the 87.5% accuracy refers to the  
 1040 fact that the pairs agreed with each other on 87.5%  
 1041 of the instances. A subsequent error analysis in-  
 1042 dicated that the mistakes were a result of human  
 1043 error where they missed the mention of certain in-  
 1044 formation.

1045 Illustrative examples for each test case category  
 1046 are provided in Figures 3, 4, 5, and 6. Specifically,  
 1047 we present two examples (Spatial and Temporal)  
 1048 from the Spot the Difference dialogs and two ex-  
 1049 amples (Attributive and Inferred) from the Meetup  
 1050 dialogs. In these figures, we have highlighted the  
 1051 essential segment of the dialog containing the an-  
 1052 swer (in purple) to the question (in violet). We also  
 1053 provide the keywords needed for the exact match  
 1054 analysis.

1055 As can be seen in the examples, our questions  
 1056 are always regarding an object that has multiple  
 1057 instances with different properties. For example,  
 1058 Figure 3 provides an example of inferred ground-  
 1059 ing, where, while A asks B if his child's room has a  
 1060 teddy bear, B says no. This should ground two facts  
 1061 - A's room has a teddy bear while B's room doesn't.  
 1062 However, in the question, A asked about his child's  
 1063 room. This should have made B answer about A's

room but since in the original dialog, the answer  
 to the question is 'nope', our models generally pro-  
 vide a negative answer. '...' in our figures indicate a  
 time gap between the chunk of utterances as can be  
 seen in Figure 6 where we have a gap from the time  
 [02:21] till [10:45] and again till [12:04]. Figure 6  
 references a white cloud mentioned in preceding  
 images, yet the question is specifically directed at  
 the final image.

[00:15] **B:** hello  
 [00:29] **B:** I've found a child's room.  
 [00:31] **A:** i am in a child  
 [00:34] **A:** 's room too  
 [00:42] **A:** blue comforter and teddy bear?  
 [00:43] **B:** nope.  
 [00:48] **B:** There are two beds. with a pur-  
 plish mattress, and a biege comforter across  
 each  
 [00:57] **A:** oh i can look for that one  
 [01:03] **B:** okay.  
 [01:14] **B:** there's also a set of pink and  
 white towels hanging by the foot of the bed  
 [02:01] **B:** there's a small white nightstand  
 between them, with a little lamp on it.  
 [02:01] **B:** I'm heading to a new location.  
 [02:01] **A:** i am in the bathroom  
 [02:12] **A:** where are you at?  
 [02:16] **B:** I'm outside. Let's move around  
 a bit.  
 [02:21] **A:** k  
 ...  
 [08:23] **A:** Was there a teddy bear in that  
 child's room that I visited?  
**B:** Yes there was  
**Exact Match:** Yes

Figure 3: Test example for inferred grounding from Meetup dataset

...

[02:37] **A:** I am in the basement

[02:37] **B:** I'm in a basement.

[02:49] **B:** Mine has a white staircase

[02:54] **A:** No

[03:03] **A:** mine has wooden stair case

[03:21] **B:** Mine has a leather chair. Should I try to move towards you?

[03:35] **A:** Sure

[03:37] **B:** Wooden? What else?

[03:53] **A:** water heater and washer and dryer

[04:08] **A:** a plastic chair and a screen door

[04:08] **A:** I'm heading to a new location.

[04:08] **A:** I have found a staircase

[04:34] **A:** The staircase is on my right and there is a hallway in front of me

[04:58] **A:** Find it?

[05:01] **B:** found the staircase

[05:15] **A:** is it the one with the hallway in front of you?

[05:27] **B:** it is to my east

[05:30] **A:** and the staircase on the right?

[05:38] **B:** its on my left

[05:53] **A:** We're not at the same staircase

...

[07:12] **B:** What was the material of the chair present in the **white staircase basement**?

**A:** It was a leather chair.

**Exact Match:** leather

Figure 4: Test example for attributive grounding from Meetup dataset.

...

[09:45] **A:** Okay I see.

[09:46] **A:** Two sharks on the right top

[09:49] **A:** of the image.

[09:50] **A:** They are blue

[09:51] **A:** and gray.

[09:52] **A:** Or

[09:52] **A:** another kind of blue.

[09:53] **B:** The one below.

[09:55] **A:** Yes.

[09:56] **B:** In my

[09:57] **B:** picture it's green, the one on the top is blue. Which color did you say it was yours?

[10:01] **A:** They are both blue and

[10:03] **A:** gray.

[10:03] **B:** Okay.

[10:04] **B:** The down one in mine is green.

[10:07] **A:** Okay.

[10:08] **A:** They are facing the right.

[10:10] **B:** Correct.

[10:10] **A:** Okay and they have

[10:12] **A:** two fins.

[10:15] **B:** Two.

[10:16] **A:** Fins.

[10:16] **B:** Yes.

[10:17] **A:** aa so below the sharks I have three yellow

[10:21] **A:** octopuses.

[10:23] **B:** I have two.

...

[19:02] **A:** What were the colors of the animals right above the octopuses in my image?

**B:** There were blue and gray sharks above the octopuses in your image.

**Exact Match:** blue and gray

Figure 5: Test example for spatial grounding from Spot the difference dataset

...

[02:11] **B:** Ahan.

[02:11] **A:** <ahm>

[02:12] **A:** To the left there is a cloud a gray cloud with raindrops underneath.

[02:18] **B:** Oh so it's a white cloud.

[02:21] **A:** Yeah.

[02:21] **B:** And there's no rain drops.

[02:21] **A:** Okay.

...

[10:45] **A:** Let's switch to the next image now.

...

[12:04] **A:** Okay and up in the sky to the right there's <aa> three clouds.

[12:09] **B:** I have four.

[12:10] **B:** <ahm>

[12:10] **A:** Okay.

[12:11] **B:** Two big ones on the right and two little ones <on the> on the left and upper side.

[12:17] **A:** Okay I have three ones and

[12:21] **A:** I have <aa> the the small one is <aa> above the two big ones.

[12:27] **B:** <ahm>

[12:28] **B:** Okay <ahm>.

...

[14:07] **B:** Do you remember the number of clouds in my last image?

**A:** You had four clouds.

**Exact Match:** four

Figure 6: Test example for temporal grounding from Spot the difference dataset

**C LLM-As-A-Judge Effectiveness** 1073

To validate the reliability of our automated evaluator (LLM-as-a-Judge), we conducted a human evaluation study. We recruited 10 participants with a mixed background in Natural Language Processing, Cognitive Science and Neuro-Science. Each of them annotated a distinct set of 4 samples, resulting in a total evaluation set of  $N = 40$  instances. Participants were presented with a question, the gold reference answer, and the model-generated answer. They were instructed to determine whether the generated answer possessed the semantic equivalence of the gold answer. Treating the human annotations as ground truth, we calculated the accuracy of the LLM-Judge which achieved an accuracy of 90% (36/40 matches with human judgment). It demonstrated a strong alignment between human and LLM judgements. Upon conducting a qualitative error analysis on the discrepancies, we observed that the LLM-Judge occasionally exhibits excessive leniency. It tends to overlook specific attribute mismatches if the core object class aligns. For example, given the question "What is on the table?", if the gold answer is *blue bottle* and the generated answer is *red bottle*, the LLM may classify them as equivalent because both refer to a bottle, thereby ignoring the color hallucination. However, the high accuracy of LLMs as judges indicated their good overall performance that made us use them as automatic evaluators.

**D Synthetic Data Generation** 1103

We expand upon the Synthetic Data generation process to offer a comprehensive understanding of the framework's design and operation. The core objective of this generation process extends beyond mere factual accuracy to the creation of dialogs that are inherently spoken-like and situational. We specifically design the framework to model the complex dynamics of conversational grounding, where participants actively establish and maintain mutual understanding of entities and facts. This is achieved by programmatically incentivizing and modeling real-world conversational phenomena, including high volumes of back-and-forth exchanges, frequent backchanneling (e.g., "uh-huh," "yeah"), and explicit affirmations or clarifications of shared knowledge. By ensuring the agents are constantly listening to and reliably building upon each other's words, the resulting corpus aims to a deeper, more reliable form of collaborative communication.

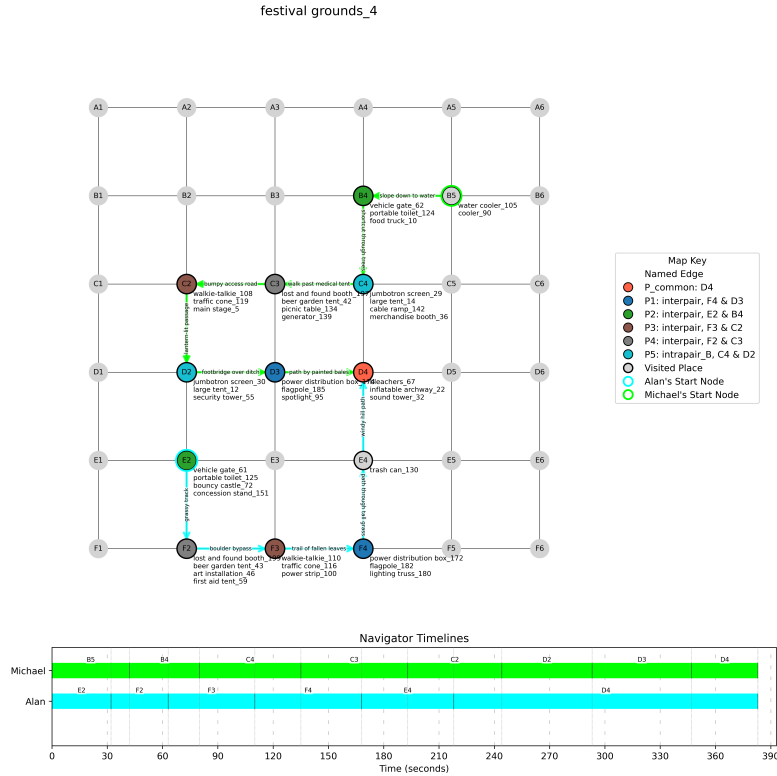


Figure 7: A visualization of a full\_scenario. Unique objects, with their IDs, are crafted to have similar characteristics to induce ambiguity that has to be resolved. The ambiguity can arise from a navigator’s experience alone (intrapair) or for their experience combined (interpair). Each chunks ( $\mathcal{C}_{m,j}$ ) is delimited by the densely dotted lines.

### D.1 Grounded World Generation

The environment is not a simple random collection of items; it’s built from a structured taxonomy (see Figure 7). The process starts with a curated list of abstract places, like festival grounds or historic museum. For each place, we have generated inventories of objects and edges (paths). Objects are coherent with the place and highly detailed with properties and relations to other objects. This granular data provides the raw materials for a rich, fact-based world.

The simulation is a time-series event logger. Once a place is selected and populated with objects, two non-sentient virtual navigators are placed within it. The system then simulates their movement, creating a full\_scenario file that logs their ‘trajectories’ over a set time span. This file is the definitive ground truth, recording what each navigator ‘sees’, their paths, and their arrival/departure times. This detailed record ensures that every piece of information later referenced in the dialog is verifiable. The continuous Ground Truth is then partitioned into smaller chunks ( $\mathcal{C}_{m,j}$ ), where  $j$  denotes the navigator and  $m$  denotes the index of the visited location  $\mathcal{P}$  (or time segment).

### D.2 Dialog Synthesis

A programmatic "Script Writer" acts as the link between the simulated world and the LLM. It takes the structured information from a time segment, algorithmically ‘reasons’ over it, and converts this into a "script" prompt for the LLM to “perform” like an actor. This pre-processing is key to off-loading the reasoning burden from the LLM. For example, the model will directly receive something like Listing 1.

Listing 1: Example of a  $\mathcal{C}_{m,j}$  chunk for the LLM

```

"place_type": "bakery production facility",
"time_span_s": [175, 212],
"navigator_A": {
  "name": "Michael",
  "zone_type": "packaging and storage zone",
  "status": "arrived & departing",
  "arrived_from": "office area connection",
  "objects_to_mention": [
    {"type": "packaging machine"},
    {
      "type": "label applicator",
      "property": "wipe-on",
      "relation": "holding",
      "reference": "label roll holder"
    }
  ],
  {"type": "delivery cart"},
  {"type": "refrigeration unit"}
]

```

```

1177 }
1178 ... //continues

```

The Script Creator systematically introduces complexity to the dialog (i.e., noticing similarities, addressing ambiguities, etc.) This includes generating sentences that reference objects seen by the other navigator (indirect reference) or that require a deeper understanding of the shared context.

Listing 2: ‘Reasoned’ part by the Script Writer

```

1186 ... //continues
1187
1188 "current_observer": {
1189   "name": "Michael",
1190   "full_details_of_what_Michael_is_seeing": {
1191     "type": "delivery cart",
1192     "property": "fixed axle",
1193     "relation": "secured by",
1194     "reference": "securing strap"
1195   },
1196   "ask_question_about": "Michael asks about a
1197   delivery cart seen before by Kenneth",
1198   "to_avoid_in_question": "**Michael does NOT
1199   ask about specific properties that only
1200   Kenneth has seen.**"
1201 },
1202 "to_discuss_with": {
1203   "partner_name": "Kenneth",
1204   "full_details_of_what_Kenneth_saw": {
1205     "type": "delivery cart",
1206     "property": "solid rubber",
1207     "relation": "near",
1208     "reference": "loading ramp"
1209   },
1210   "seen_after_taking": "pallet movement lane",
1211   "that_place_also_had": "waste compactor"
1212 }

```

The LLM then outputs the dialog:

Listing 3: Example of LLM-generated dialog segment

```

1215 Michael: I just arrived in the packaging and
1216 storage zone via the office area connection.
1217 Kenneth: Uh-huh, I'm still in the mixing and
1218 preparation zone.
1219 Michael: There's a packaging machine here
1220 with an adhesive seal protected by a safety
1221 interlock.
1222 Kenneth: Got it. My packaging machine here
1223 uses an ultrasonic seal and receives from a
1224 product feed conveyor.
1225 Michael: Noted. There's also a label
1226 applicator with a wipe-on method holding a
1227 label roll holder.
1228 Kenneth: Mine's a tamp-apply applicator
1229 triggered by a product sensor. Any delivery
1230 carts around?
1231 Michael: Yes, there's a delivery cart with a
1232 fixed axle secured by a securing strap.
1233 Kenneth: Mine's near a loading ramp with
1234 solid rubber wheels.
1235 ... //continues
1236

```

### D.3 Q/A Synthesis 1238

First, we automatically identify "questionable moments" - where there are multiple entities of similar nature and a relational reference to them would only be resolved with a good grounding mechanism. We do this using both the ground truth ( $\mathcal{C}_{all}$ ) and the dialog transcript ( $\mathcal{D}$ ), algorithmically analyzing the conversation and flagging instances in a structured list. 1239-1246

A fact is preselected from the Ground Truth as the “expected answer”, and the LLM is prompted to generate a question that elicits this specific answer. The complex context for a moment is captured in a JSON fact: 1247-1251

Listing 4: Example of a fact for Q/A generation

```

1252 "fact_id": 17,
1253 "chunk_id": "chunk_007",
1254 "question_type": "inferred_comparison_result",
1255 "questioner": "Kenneth",
1256 "answerer": "Michael",
1257 "context": {
1258   "property_of_obj_seen_by_Michael": "fixed
1259   axle",
1260   "property_of_obj_seen_by_Kenneth": "solid
1261   rubber",
1262   "ambiguous_object_type": "delivery cart",
1263   "disambiguating_fact": "the one that was
1264   where there was also a waste compactor"
1265 },
1266 "expected_answer": "solid rubber",
1267

```

Given the whole context, the LLM generates the following questions in both a long and a short form giving us two variety of the same question: 1269-1271

Listing 5: Example of the generated and validated Question Pair

```

1272 "question_full": "Of the two delivery carts we
1273 discussed-yours with a fixed axle and mine
1274 with solid rubber wheels-which one was
1275 located where there was also a waste
1276 compactor?"
1277 "question_short": "Which type of delivery cart
1278 was near the waste compactor?",
1279

```

### D.4 Validation and Quality Control 1281

To ensure structural integrity and semantic accuracy, we employ an algorithmic validation pipeline active *during generation*. The system parses the LLM output in real-time; if a check fails, a specific error message is fed back to the model to prompt self-correction. If the model fails repeatedly, the instance is discarded. **Dialogs:** We enforce conversational pacing by limiting consecutive utterances, we verify that movement intentions are stated, and we ensure all mandatory objects seen in  $\mathcal{C}_{m,j}$  are 1282-1291

mentioned. To account for linguistic variability, we employ a sliding-window semantic match (using all-MiniLM-L6-v2 of the MiniLM family of models (Wang et al., 2020)) rather than strict keyword matching, verifying that specific properties and relations defined in the Ground Truth are contextually present. **Questions:** We verify that generated questions contain mandatory scenario keywords, ensuring they are tied to the Ground Truth. Furthermore, we enforce strict constraints to prevent answer leakage and validate point-of-view consistency (e.g., ensuring correct pronoun usage for the observer vs. questioner).

Of all the questionable facts identified, only  $\sim 600$  ( $\sim 36\%$ ) of them were successfully converted into validated Q/A pairs. This high rejection rate indicates the effective pruning of logical inconsistencies and factual errors, significantly enhancing corpus quality.

## E GRPO Training

For our GRPO training, we used 600 Question/Answer pairs from 60 artificially generated dialogs. While we can generate much more synthetic data, we were constrained by the computing resources and decided to use only 600 instances which showed to be enough for the model as the reward stabilized towards the end of the training as seen in Figure 8. The huggingface TRL (von Werra et al., 2020) library was used for GRPO training.

### E.1 Hyperparameters

We used a batch size of 16 for the RL training. The training used 8 A100 GPUs in total out of which 4 A100s were used by VLLM for the inference. Here are our other hyperparameters : learning rate -  $1e-6$ , GRPO beta - 0.025, weight\_decay - 0.1, adam\_beta1 - 0.9, adam\_beta2 - 0.99, warm-up ratio - 0.1, and a cosine learning scheduler.

## F Gemma Results

Tables 6 and 7 contain all the results for Gemma2 2B and 9B models. As seen from the Tables, the overall pattern remains the same as compared to Llama 3.1-8B. The 2 Billion model performs relatively poor compared to the 9 Billion model across all the categories. Chunking in general provides the worst results among all the representation techniques suggesting the importance of good abstraction of information that preserves the links between the information across different segments of the

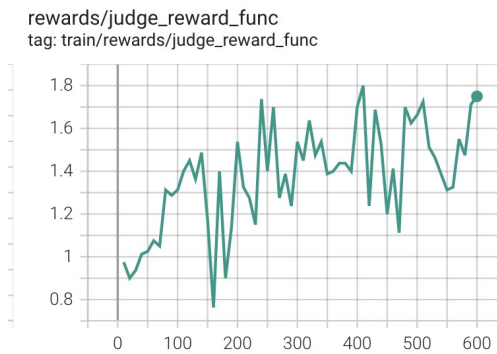


Figure 8: Training reward over time while using the synthetic data. Here the y-axis is the reward while x-axis is the step.

conversation. Surprisingly, for the gemma2 models that struggled with longer context of STD, the agentic system performs better than when provided with the full dialog for Meetup test cases. This suggests that providing a structure might help it reason better for longer contexts.

## G Prompts

During testing, we utilized the system prompts shown in Figures 9, 10 and 11. Our approach leverages chain-of-thought prompting (Figure 10), which we found to be more effective than direct questioning. This prompt outlines the game’s rules to help the LLM understand the context. However, it differs from the in-context prompt (Figure 11) by not specifying how to formulate the answer. We have omitted the prompts for our Agentic method from this paper due to the numerous large prompts used in its different phases, but they will be included in the open-source code release. Furthermore, we deliberately excluded textual descriptions of the images from the prompts. This ensures that the questions can only be answered using the conversational context, thereby testing the model’s ability to form an accurate common ground representation based on what has been discussed.

## H License

All tools and models used in this research were employed under licenses that permit academic use. Specifically, we utilized the Huggingface TRL library and the Sentence Transformer model all-MiniLM-L6-v2 under the Apache-2.0 license. The Llama 3.1, Gemma 2, and Qwen models were used in accordance with the Llama 3.1 Community License Agreement, Google’s usage license, and the Apache-2.0 license respectively.

LLMs	Meetup				Spot the difference			
	Temporal	Spatial	Attributive	Inferred	Temporal	Spatial	Attributive	Inferred
Full Dialog [from Table 2]	0.20 / 0.18	0.18 / 0.16	0.24 / 0.26	0.26 / 0.16	0.18 / 0.12	0.18 / 0.10	0.20 / 0.12	0.16 / 0.10
Summarisation	0.16 / 0.14	0.12 / 0.12	0.20 / 0.14	0.18 / 0.14	0.12 / 0.10	0.14 / 0.12	0.12 / 0.10	0.08 / 0.02
Chunking (NV-Embed-V2)	0.16 / 0.14	0.04 / 0	0.12 / 0.10	0.22 / 0.16	0 / 0	0.02 / 0	0.02 / 0	0.02 / 0
Chunking (BM25)	0.24 / 0.20	0.24 / 0.20	0.20 / 0.18	0.16 / 0.12	0 / 0	0 / 0	0.06 / 0.02	0.04 / 0
Agentic Ontology-based	0.40 / 0.36	0.32 / 0.26	0.26 / 0.22	0.20 / 0.16	0.12 / 0.06	0.10 / 0.04	0.14 / 0.08	0.06 / 0.10

Table 6: FEM ( $\uparrow$ ) / LLM-as-Judge ( $\uparrow$ ) accuracy for various representations using Gemma2-2B.

LLMs	Meetup				Spot the difference			
	Temporal	Spatial	Attributive	Inferred	Temporal	Spatial	Attributive	Inferred
Full Dialog [from Table 2]	0.48 / 0.38	0.50 / 0.46	0.38 / 0.40	0.26 / 0.18	0.34 / 0.20	0.24 / 0.16	0.28 / 0.16	0.20 / 0.10
Summarisation	0.28 / 0.22	0.24 / 0.18	0.28 / 0.24	0.16 / 0.14	0.12 / 0.10	0.10 / 0.10	0.12 / 0.12	0.16 / 0.14
Chunking (NV-Embed-V2)	0.22 / 0.16	0.10 / 0.10	0.18 / 0.16	0.22 / 0.22	0 / 0	0 / 0.02	0 / 0	0.04 / 0
Chunking (BM25)	0.34 / 0.26	0.22 / 0.24	0.24 / 0.26	0.26 / 0.24	0.02 / 0	0.02 / 0	0.02 / 0	0.06 / 0.02
Agentic Ontology-based	0.44 / 0.46	0.40 / 0.38	0.42 / 0.36	0.26 / 0.24	0.32 / 0.32	0.36 / 0.32	0.30 / 0.26	0.28 / 0.16

Table 7: FEM ( $\uparrow$ ) / LLM-as-Judge ( $\uparrow$ ) accuracy for various representations using Gemma2-9B.

## I Use of AI

We declare that we have used LLMs for helping us paraphrase some of the sentences in this paper in order to better convey our findings to the community. However, we would like to mention that it was not used for writing big chunks of the paper. Every single line written here has been thoroughly read and revised by the authors.

You are a strict evaluator. Given the Question, LLM response and the correct response, judge whether the LLM response and the correct response both have the same meaning provided the question.

**\*KEEP IN MIND TO ALWAYS FOLLOW THESE RULES\* -**

- 1) You have to give 'DIFFERENT' if they are not having the same meaning and 'SAME' if they mean the same.
- 2) If the words are synonyms then they shall be considered the same.
- 3) If the response is a negation of the other response then it does not have the same meaning and hence should be 'DIFFERENT'.
- 4) If correct answer is kitchen but the response is dining room, then it should be 'DIFFERENT' because they are different. Do not reason that they have similar purpose. Just because the place might have similar function does not make it the same.
- 5) Sometimes the answers from LLM might contain reasoning, but you should only provide the score based on the final answer provided which is at the end of the answer.
- 6) If the two responses are negations but in different styles, they should still be considered as the 'SAME'. eg. 'do not recall presence of a cat' is the same as 'no cats were there' are the same, the first one is more indirect form of saying it.
- 7) If the main content of the answers are the same then answer 'SAME'. The style might be different. For example, 'yes there was a car there' and 'a car was there' are the 'SAME' in meaning even though one is more affirmative than the other.

You have to give the answer in the following format:

<reasoning>

(your reasoning here in one or two sentences where you concretely mention the reason why they mean the same or not.)

</reasoning>

<answer>

(your final verdict here)

</answer>

Figure 9: The system prompt for LLM as a judge.

You are participant {current\_participant} who is navigating a shared environment with another participant {other\_participant}. You are in a virtual space that has lots of different areas. The participants are initially located in different areas.

The goal is for the two participants to locate themselves in the same area. In order to achieve this goal, the participants communicate with one another and describe the area they find themselves in. On the basis of those descriptions, they move to different areas and describe their new areas to the other participant.

The two participants need to find themselves in the same area. In the last utterance, participant {other\_participant} asks a question. Think about the question, get all the information required from the dialog as participant {current\_participant} and then answer that question as your response. First reason and interpret the question. Identify whose experience you have to refer to {other\_participant} or {current\_participant} or both. Fully break it down and then give the answer. Reason step by step and then answer.

The final answer should be in the format -

<reasoning>

(your reasoning here which requires analysing the dialog and understanding what are the important information to look at before answering. Take all the important information into consideration step by step in your reasoning.)

</reasoning>

<answer>

(your final answer here. DO NOT REASON HERE!!)

</answer>.

Figure 10: The system prompt for the LLM when provided with the full dialog for Meetup test cases. It was also used for the GRPO trained models.

[FIRST PARAGRAPH SAME AS Figure 10]

The two participants need to find themselves in the same area. In the last utterance, participant {other\_participant} asks a question. Answer the question. First understand what it means and then answer as the other user. In other words, if A asked the question then answer as if you are B and vice versa. Remember they start at different locations and try and meet and are may be in different rooms and may have parallel observations which are valid from their perspective. Identify whose experience and observations you have to refer to. First reason and interpret the question. Identify whose experince you have to refer to, then A or B or both fully break it down and then give the answer.

Remember most importantly -

1) Reason step by step and then answer. Example:

Question from A: do you remember the park I visited?

Although B has to answer here, the reference is to A's experience, so check what A observed.

Similarly, For, B: Do you remember the garden i visited? -> A should answer based on what B experienced.

2) Most Crucial: If the question is about "my x" (e.g., "my dining room", "my office"), always refer to the one who asked the question. So:

- "My dining room" (asked by A) → Refer to A's observation of their dining room.

- "My chandelier" (asked by B) → Refer to B's observation of their chandelier.

3) Also, prioritize semantic ownership and pronoun resolution carefully. Don't let "described by B" override the ownership stated in the question. Always prioritize the experience of the person who owns the object/location. If A says "My chandelier," then even if B is answering, the correct grounding is in A's description — because A is the owner. Use the other speaker's description if that's whose experience is referenced in the question.

4) If observations conflict strongly, assume two different entities till it has been resolved. Example conflict: A: "I see one white fridge", B: "I have five yellow fridges" → this means they are in two different kitchens.

Example grounding from dialog:

00:02:01 — A: I'm in a kitchen and see a fridge

00:02:10 — B: yellow?

00:02:15 — A: No not yellow but white.

00:03:10 - A: Found another kitchen. White walls.

Spatial Grounding: A is in Kitchen A, B is in Kitchen B. A then moves to kitchen C.

Object Grounding: Fridge A (white) in Kitchen A — seen by A, Fridge B (yellow) in Kitchen B — seen by B, Kitchen C has white walls - seen by A.

Temporal Grounding: A first goes to kitchen A, B simultaneously goes to kitchen B. A then move and finds another kitchen C which is the second kitchen that A saw.

The final answer should be in the format -

<reasoning> ... [Like in Figure 10] ... </answer>

Figure 11: The system prompt for in-context learning with full dialog history for Meetup test cases.

**Question :** Can you please remind me which room is positioned east of the dining room that I visited?

**Correct Answer :** Attic

**Generated Answer :** Okay, let's see what the user needs here. They're asking for the next utterance in the conversation based on the provided dialog history. The last message from A is a question: "Can you remind me which room is positioned east of the dining room that I visited?" So, I need to figure out how B would respond. Looking through the conversation history, we need to track where each participant has been. Let's start by recalling the key points: First, Participant B started in a dining room with a red chandelier (timestamp [00:00:27]). Then, A moved to the east and mentioned being in an attic ([00:01:04]), but then later described being in a room with a curved ceiling, ceiling fan, and triangular shape, possibly an attic again. At [00:02:43], A mentions finding a jacuzzi with beach chairs, which seems like a bathroom or maybe a spa room. But then moves to a pool area ([00:03:31] mentions brown tiles and pool). Meanwhile, B went to a street ([00:03:48]) and then an attic with a bed and closet, which was different from A's attic without a bed. Eventually, B ended up in the room A was referring to, which had a window and a wooden table on the right, plus a 2007 date ([00:05:04]). The question now is about the room east of the original dining room. From earlier messages, when A first said they were in the dining room with a red chandelier, then moved east to what they thought was an attic. However, after moving, the next image B received was a library/museum building ([00:01:16 Image B]). Wait, actually, the timestamps show that after A says they're heading east, the next Image B is the library building. Hmm, but maybe the movement directions correspond to changing rooms. Since in virtual spaces, moving east could lead them to adjacent rooms. Wait, the initial dining room was B's starting point. When A was in the dining room too? At [00:01:04], A says in dining room with red chandelier, so they were both in the same dining room at some point. But since they needed to get to the same room, perhaps moving east from there led A to a different room. The Image B after A's move to east showed the library/museum building. So maybe moving east from the dining room leads to the library? Because the next image after A moving east was the library. So when A asks which room is east of the dining room, B would refer to that library or museum. But looking at the latest interactions, B is currently in a room with a window, wooden table on the right, and a 2007 date. That room's description ([00:05:04]) is an empty room with a dining table, chairs, desk, laptop, etc. Maybe that's a different area. Alternatively, B might remember previous movements. Alternatively, maybe the library was the room east of dining because after A went east, the next image was the library. Therefore, B would say something like "the library or museum building." Also, considering that in the final messages, they ended up in a room with the 2007 date, which might not be the library, but another spot. So the response should be B telling A that the room east of the dining room is the library or museum. Let me check the exact description again. The Image B after A moved east was: "large, unique building with a distinctive dome...library or museum." So B should mention that. Hence, the answer is probably pointing out the library/museum as the eastern room. </think> [00:05:42] B: The room east of the dining room looked like a library or museum with a big dome. check the building with cars parked out front and traffic lights

Figure 12: Failure case from Qwen-QWQ where it produces a long answer despite reasoning in its <reasoning> tags separately