

# Grounded Misunderstandings in Asymmetric Dialogue: A Perspectivist Annotation Scheme for MapTask

Nan Li<sup>1</sup>, Albert Gatt<sup>1</sup>, Massimo Poesio<sup>1,2</sup>

<sup>1</sup>Utrecht University, Utrecht, the Netherlands

<sup>2</sup>Queen Mary University of London, London, the UK  
{n.li, a.gatt, m.poesio}@uu.nl

## Abstract

Collaborative dialogue relies on participants incrementally establishing common ground, yet in asymmetric settings they may believe they agree while referring to different entities. We introduce a perspectivist annotation scheme for the HCRC MapTask corpus (Anderson et al., 1991) that separately captures speaker and addressee grounded interpretations for each reference expression, enabling us to trace how understanding emerges, diverges, and repairs over time. Using a scheme-constrained LLM annotation pipeline, we obtain 13k annotated reference expressions with reliability estimates and analyze the resulting understanding states. The results show that full misunderstandings are rare once lexical variants are unified, but multiplicity discrepancies systematically induce divergences, revealing how apparent grounding can mask referential misalignment. Our framework provides both a resource and an analytic lens for studying grounded misunderstanding and for evaluating (V)LLMs' capacity to model perspective-dependent grounding in collaborative dialogue.

**Keywords:** reference expression, common ground, misunderstanding, collaborative task, annotation scheme

## 1. Introduction

In everyday collaborative dialogue, speakers and addressees continuously coordinate their understanding of what is being discussed (Clark and Brennan, 1991; Pickering and Garrod, 2004). A central mechanism for this coordination involves reference expressions (REs)—linguistic forms used to pick out entities in the shared context.

Utterances become part of the common ground only after being recognized and acknowledged by the addressee (Clark and Wilkes-Gibbs, 1986; Clark and Schaefer, 1989; Clark and Brennan, 1991), a process known as *grounding*. Traditional approaches to reference resolution assume that once grounding is achieved via explicit or tacit confirmation, both interlocutors successfully refer to the same entity. However, this assumption can fail in *asymmetric* settings where participants have access to different information.

Such asymmetric settings were the focus of the HCRC MapTask (Anderson et al., 1991; see Figure 1 for an example), where two participants navigate a route using slightly different maps. While participants in such settings show systematic adaptation behaviors to address information asymmetry (Bard et al., 2000; Viethen et al., 2011a,b; Healey et al., 2018), making full misunderstandings rare, few studies have examined the *personal interpretations* of REs by the two interlocutors and the subtle divergences that may persist even after apparent grounding. These limitations also exist in the MapTask annotations (discussed in Section 4.1), making them difficult to be used to identify misun-

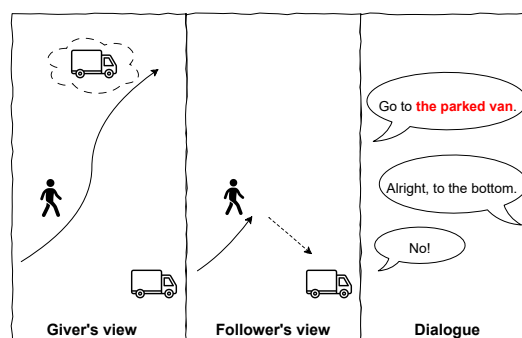


Figure 1: A simplified example of misunderstandings from MapTask dialogues. **Left:** the instruction *giver's* map contains two instances of *parked van* and the intended target is circled. **Middle:** the instruction *follower's* map shows only one *parked van* at the bottom that is shared with the *giver*. **Right:** the *giver* says “go to *the parked van*,” the *follower* acknowledges but grounds another instance, leading to correction (“No!”). Red highlights the reference expression.

derstandings.

To understand how interlocutors achieve and sometimes fail to achieve mutual understanding, we must first capture each participant's personal interpretation. Taking both perspectives allows us to quantify where apparent alignment hides subtle misalignments and to measure how frequent such misunderstandings actually are under different landmark discrepancy settings.

Prior frameworks for modeling dialogue states of-

fer rich tools to trace how shared understanding is established (Poesio and Traum, 1997; Cooper et al., 1999; Matheson et al., 2000; Lascarides and Asher, 2009; Schlangen and Skantze, 2011; Ginzburg, 2012). However, these symbolic, frame-based representations have not been widely integrated into reference-level annotation schemes, partly because they are difficult for large language models (LLMs) to reproduce in a constrained, schema-driven manner.

We introduce an annotation framework based on the MapTask annotations that explicitly captures the divergence between the speaker’s intended and the addressee’s interpreted referents during collaborative dialogue. Our scheme uses a hierarchy of attributes to represent personal interpretations and incremental grounding, making it both linguistically principled and operational for LLM annotation.

Applying this framework to MapTask dialogues with GPT-5 (OpenAI, 2025), we scale annotation across the entire MapTask corpus. While full misunderstandings remain uncommon, our analysis reveals systematic patterns of understanding state transitions. These findings not only shed light on how diverging interpretations emerge and resolve in collaborative tasks but also establish a benchmark for evaluating language-only and vision-language models’ grounding abilities.

Our research questions are:

**RQ1** *Is it possible to develop a personal-interpretation annotation scheme to capture nuanced understanding states for REs in collaborative tasks like MapTask?*

**RQ2** *Given conflicting personal interpretations, can we trace how understanding evolves across turns until the successful grounding?*

**RQ3** *Can LLMs, under a schema-constrained protocol, reliably annotate personal interpretations; and how does their output inform future evaluation of (V)LLMs on grounded dialogue?*

Our contributions include:

- We operationalize reference expression grounding as a graded, perspective-dependent alignment process, and propose an annotation scheme at the RE level.
- We present a scheme-constrained LLM-in-the-loop annotation pipeline of MapTask dialogues, producing 13k annotated REs with reliability estimates.
- We release analysis protocols that associate RE-level annotations with dual-perspective interpretations and incremental understanding states, enabling quantitative study of misalignment dynamics in MapTask dialogues and informing future evaluation of LLMs and VLMs on common grounding.

The remainder of this paper is structured as follows. Section 2 reviews related work. Section 3 introduces the MapTask corpus and its challenges for reference grounding. Section 4 details our annotation scheme and LLM-in-the-loop annotation pipeline. Section 5 analyzes understanding state distributions and transitions. Finally, Section 6 concludes with implications and future directions.

## 2. Related work

**Collaborative Reference and Grounding** Reference expressions are central to establishing and maintaining common ground in interactive dialogue (Clark and Wilkes-Gibbs, 1986; Clark and Schaefer, 1989). Successful mutual understanding is constrained by many factors, e.g., communication medium, interactive alignment at different linguistic levels (Clark and Brennan, 1991; Pickering and Garrod, 2004). Previous work reveals that speakers do not always fully accommodate addressee knowledge, exhibiting egocentric biases even when explicitly aware of information asymmetry (Horton and Keysar, 1996; Keysar et al., 2000; Lane et al., 2006)

This theoretical foundation has motivated NLP tasks to study grounded reference in collaborative settings with asymmetric information. The HCRC MapTask (Anderson et al., 1991) and iMap corpus (Louwerse et al., 2007) both instantiate navigation tasks where participants have slightly different maps, eliciting rich adaptation behavior (Bard et al., 2000; Viethen et al., 2011a,b). More recent work includes the OneCommon Corpus (Udagawa and Aizawa, 2019), a partially-observable environment for both participants. These approaches typically treat grounding as a single end state rather than tracking divergent personal interpretations throughout the dialogue.

Current VLMs have excelled at some object identification tasks (Yu et al., 2016), where referents (objects in images) are assumed to be deterministically grounded once specified. Some recent work shows that VLMs struggle to adopt different frames of reference and visual perspectives (Zhang et al., 2024; Lee et al., 2025). Moreover, existing benchmarks rarely evaluate whether models can track evolving, asymmetric interpretations as dialogue progresses. Our work addresses this gap by annotating personal interpretations at the RE level, enabling future evaluation of whether (V)LLMs understand incremental grounding processes in (asymmetric) dialogue.

**Understanding State Tracking** Meaning coordination operates through continuous “running repairs” rather than discrete error correction episodes (Healey et al., 2018). This repair-driven

coordination directs the participants' evolving understanding states, and motivates the need to track divergent interpretations as they emerge and resolve through integration.

Formal dialogue theories have modeled interlocutors' intentions and repairs through structured information state models. Some proposals (Poesio and Traum, 1997; Matheson et al., 2000) extended Discourse Representation Theory to integrate semantic and pragmatic information, showing how turn-taking, discourse segmentation, and grounding are represented through conversation acts in common ground. Ginzburg (2012) proposed the KoS framework<sup>1</sup>, where unresolved Questions Under Discussion (QUDs) persist until explicitly addressed. Lascarides and Asher (2009) modeled agreement and correction dynamics but assumed a highly idealized scenario which excludes misunderstandings. Schlangen and Skantze (2011) treated dialogue processing as an incremental hypothesis-formation and revision process operating at a sub-utterance level. Some recent work leveraged annotations from other meaning representations, e.g., Abstract Meaning Representation, to represent common ground structure and information states (Khebour et al., 2024; Lai et al., 2025). We are informed by this line of work, but focus on RE-level perspectivist interpretations and inferring understanding states through the alignment and misalignment of these interpretations.

**LLM as an Annotator** LLMs have demonstrated strong performance on various semantic and pragmatic annotation tasks (Eichin et al., 2025; Chen et al., 2025; Qamar et al., 2025). They also achieve comparable results to human annotators in subjective tasks such as sentiment analysis (Bojić et al., 2025). However, challenges remain for scheme-constrained annotation tasks requiring structured outputs (Ettinger et al., 2023). Recent advances in grammar-constrained decoding and JSON schema-based generation offer promising solutions, enabling LLMs to generate outputs that adhere to predefined structures while maintaining semantic accuracy (Geng et al., 2023; Park et al., 2025). Our work builds on these developments by employing a multi-layer attribute resolution pipeline to guide LLMs in capturing personal interpretations in collaborative dialogue.

### 3. MapTask: Annotation and Beyond

#### 3.1. Development and Follow-up Studies

HCRC MapTask is a corpus documenting unscripted dialogues in a collaborative route-following

---

<sup>1</sup>A term loosely connected to *conversation-oriented semantics* (Ginzburg, 2012).

task with asymmetric maps. Participants must collaborate to reproduce a route on one participant's map (the instruction *giver's*) on the other's (the instruction *follower's*). They cannot see each other's maps and can only communicate verbally, with eye contact possible in half of the dialogues. The intentional differences between maps on routes, landmark names, and landmark placements often lead to ambiguous or misleading instructions. This makes the MapTask corpus a valuable resource and testbed for studying grounding, reference, and repair theories in dialogue.

Subsequent research analyzed reference expressions in MapTask from multiple perspectives. Bard et al. (2000) show that speakers reduce articulatory clarity when mentioning landmarks for the second time, even when the addressee has not heard the first mention or cannot see the referent, suggesting that intelligibility control is driven primarily by the speaker's own knowledge rather than careful modeling of the listener's information state. Varges (2005) treated spatial descriptions in MapTask as reference expressions that distinguish particular map points from distractors, observing that speakers frequently use spatial relations (e.g., "above the west lake") and vague quantifiers (e.g., "about two inches southwest") when referring to featureless route points around named landmarks. Marge and Rudnicky (2011) analyzed miscommunication patterns in MapTask, documenting how instruction followers detect ambiguity (e.g., multiple possible referents) and misunderstanding through clarification questions, repair strategies, and explicit feedback when requested actions appear impossible or unclear. Poesio et al. (2004) applied the MATE annotation scheme (Poesio et al., 1999) to an Italian adaptation of MapTask to make anaphoric relations explicit.

#### 3.2. Task Design and Landmark Discrepancies

The HCRC MapTask corpus consists of 16 pairs of maps and 128 dialogues, with 8 dialogues recorded for each map. In each dialogue, two participants are given visually similar but non-identical maps. One participant, the *giver*, has a route marked on the giver's map and must guide the other participant, the *follower*, to reproduce this route on the follower's map through verbal communication alone. Both maps share the same starting point and most landmarks, but only the giver's map displays the complete route and endpoint. Critically, participants cannot see each other's maps, creating an asymmetric information setting.

The corpus contains 267 unique landmarks (differentiated by names) across all maps. While both participants' maps share many identical landmarks

(132 landmarks, 49.4%), the remaining 135 landmarks (50.6%) exhibit systematic discrepancies. We classify these discrepancies into three categories<sup>2</sup> and present the distribution by discrepancy type in Table 1.

Discrepancy Type	Count	%	Ref.	%
Lexical	20	7.5	914	7.0
Existence	99	37.1	2878	22.0
Multiplicity	16	6.0	956	7.3
Discrepant	135	50.6	4748	36.3
Identical	132	49.4	8333	63.7
Total	267	100.0	13081	100.0

Table 1: Distribution of landmark types and corresponding reference expressions in the MapTask corpus.

**Lexical Discrepancy** (20 landmarks forming 10 pairs, 7.5%): Landmarks appear at the same location on both maps with identical icons but slightly different name labels. For instance, one map may label a landmark “cliffs” while the other reads “sandstone cliffs”; or “carved wooden pole” versus “totem pole”.

**Existence Discrepancy** (100 landmarks, 37.5%): Landmarks appear on one participant’s map but are entirely absent from the other’s. The participant viewing the map without the landmark must either signal the absence or imagine its location based on the partner’s description.

**Multiplicity Discrepancy** (16 landmarks, 6.0%): Landmarks appear twice on one map (always the giver’s) but only once on the other (the follower’s). Shared instances of this type occupy the same location on both maps but are positioned farther from the route than the unique, non-shared instances.

### 3.3. Annotations Leveraged in This Work

The dataset of the MapTask corpus is publicly available online<sup>3</sup> with rich multi-level annotations. For our experiment and analysis, we leverage the following annotation layers: (1) **Timed Units**, word-level transcriptions aligned with audio timestamps, providing the basic segmentation of speech and text. (2) **Reference Expressions**, annotations of reference expressions linked to landmark IDs,

<sup>2</sup>In the original MapTask paper (Anderson et al., 1991), the authors claim that they also add some phonetic contrast features to landmark names to create discrepancies; we assign them to the following three categories according to our definition and observation.

<sup>3</sup><https://groups.inf.ed.ac.uk/maptask/>

recorded based on timed units. (3) **Move Tags and Transactions**, move-level annotations following the coding scheme of Carletta et al. (1997), capturing the communicative function of utterances (e.g., instruct, acknowledge, query), recorded based on timed units. Transactions are sequences of moves that move from one point to another on maps.

We conducted preprocessing to integrate these annotations into a unified format suitable for LLM annotation. First, we reassigned landmark identifiers (IDs) to distinguish between giver-side and follower-side interpretations (in Section 4.1), as the original landmark ID scheme fails to highlight the discrepancies in some cases. Second, we segmented the dialogue transcripts into transactions to control the context window when providing context to the model. Finally, we generated structured prompts combining reference expressions, new unified landmark ID candidates, and context dialogue to support the annotation workflow described in Section 4.2.

## 4. Modeling Personal Interpretations

### 4.1. Approach and General Principles

We introduce a perspectivist, RE-level annotation scheme that records the speaker-intended and addressee-interpreted landmark IDs, plus five binary attributes in a fixed decision order.

**Landmark ID Design** The original MapTask annotations assign landmark IDs by landmark names, which treat landmarks of multiplicity discrepancies (e.g., two “parked van” at the giver’s map, one at the follower’s map) as the same entity, but treat lexical variants (e.g., “cliffs” versus “sandstone cliffs”) as different entities. This conflation obscures whether participants successfully align their interpretations. To address this limitation, we introduce a new unified landmark ID assigning strategy that distinguishes giver-side and follower-side landmark instances. In this paper, we use the acronyms `mtlm` and `umlm` to refer to the original MapTask landmark ID and our unified landmark ID, respectively.

Our landmark IDs follow the format `<map-id>_<landmark-name>#<ordinal>@<side>`, where:

- `map-id`: The original Map ID (`m0`, `m1`, ..., `m15`)
- `landmark-name`: The landmark name as it appears on the map (e.g., `parked_van`, `cliff`)
- `ordinal`: Positional index (0, 1, ...) for landmarks of the same type occurring multiple times on maps, ordered from bottom to top.

Landmarks of this type have the same ordinal number if they are at the same location on both maps, which means it is possible that the only “parked van” on the follower’s map is `m0_parked_van#1@f` if the giver’s map has two “parked van”, while on the giver’s map `m0_parked_van#0@g` denotes the unique instance that is at a lower place.

- `side`: Map side (`g` for giver, `f` for follower)

When a reference expression refers to multiple landmarks simultaneously, we concatenate IDs with the operator `+`, e.g., `m0_parked_van#0@g+m0_parked_van#1@g`.

Beyond landmark IDs, we annotate five binary attributes that jointly characterize the nature of the expression and the addressee’s understanding state. These attributes are designed in a hierarchical decision cascade that models the incremental resolution of reference and embeds Chain-of-Thought (Wei et al., 2022) into LLMs’ reasoning process.

**`is_quantificational`** Indicates whether the reference expression functions as an existence query from the speaker’s perspective rather than a definite reference. For instance, “a parked van” in “Do you have a parked van?” is quantificational, whereas “the parked van” in “Go past the parked van” is not. Such quantificational expressions do not typically presuppose a specific referent in Map-Task contexts.

**`is_specified`** Indicates whether the dialogue context provides sufficient evidence to infer the addressee’s interpretation. If the addressee fails to respond or produces an utterance unrelated to the reference expression, the interpretation remains unspecified.

**`is_accommodated`** Indicates whether the addressee successfully address the reference expression without signaling comprehension failure. Accommodation fails when the addressee explicitly requests repetition, clarification, or indicates mishearing (e.g., “What?” or “Pardon?”).

**`is_grounded`** Indicates whether the addressee links the reference expression to a specific landmark on the maps. An accommodated expression may fail to ground if the addressee cannot identify a corresponding landmark (e.g., “I don’t have that” or “Where is it?”). Grounding is a prerequisite for assigning a landmark ID to the addressee’s interpretation.

**`is_imagined`** Indicates whether the addressee’s interpretation refers to a landmark

absent from his/her own map. This occurs when the addressee, having recognized an existence discrepancy, mentally projects the landmark’s location based on the speaker’s description. In such cases, we assign the speaker’s landmark ID to the addressee’s interpretation and mark `is_imagined = true`, signaling that alignment was achieved.

## 4.2. Annotation Pipeline

Informed by recent advances in Prompt Engineering (Brown et al., 2020; Wei et al., 2022; Sahoo et al., 2024), we develop a heuristic prompting annotation pipeline leveraging LLMs to assist with the perspective-taking annotation task. To seek a balance between information density and context length (Liu et al., 2023; Kuratov et al., 2024), we provide the dialogue text from the start of the conversation until the end of the current transaction as context. We use a scheme-constrained prompt with GPT-5 via the OpenAI Batch API<sup>4</sup> and enforce a JSON-schema output to scale annotation with quality control. All parameters are set to default.

For each reference expression, we construct a structured prompt containing:

- **Background** Overview of the MapTask task design and landmark discrepancy types, and a general introduction of the annotation task.
- **Task description**: Detailed description of the annotation task, emphasizing the approach of using landmark IDs and 5-layer attribute hierarchy to annotate.
- **Landmark ID explanation**: Explanation of the unified landmark ID format.
- **Annotation rule**: Step-by-step workflow instructions operationalizing the five-attribute decision and landmark ID resolution cascade.
- **Output format**: JSON schema with required fields.
- **Target reference expressions**: All reference expressions in the current transaction to annotate with their IDs.
- **Dialogue context**: Dialogue transcript from the start of the conversation until the end of the current transaction, with reference expressions highlighted by brackets.
- **Dialogue acts**: Move-level annotations indicating communicative function in time order, e.g. “[g utt:1 move:instruct] starting off we are above a caravan park [f utt:2 move:acknowledge] mmhmm”.

<sup>4</sup><https://platform.openai.com/docs/guides/batch>

### Prompt Template (Abbreviated)

```
<background> Background on MapTask, landmark discrepancies, and reference resolution </background>
<task_description> Annotate personal interpretations for reference expressions </task_description>
<landmark_id_explanation> Unified landmark ID format and semantics </landmark_id_explanation>
<annotation_rule>
Follow the annotation workflow for each RE step by step:
Step 0: Identify the speaker and the addressee, and match them with the giver and follower based on the actual situation; fill in the landmark ID
for the speaker's interpretation (intended landmark) first.
Steps 1-5: Check and fill in attributes is_quantificational, is_specified, is_accommodated, is_grounded, and is_imagined in order.
Step 6: If grounded, fill the landmark ID for the addressee's interpretation.
Step 7: Provide a concise (<50 words) evidence-based reason summarizing the judgement in the reason field for each RE.
</annotation_rule>
<output_format> JSON schema with required fields </output_format>
<target_ref_ids> IDs of reference expressions to annotate </target_ref_ids>
<context_dialogue> Dialogue transcript with surrounding context </context_dialogue>
<context_dialogue_acts> Move-level annotations for context </context_dialogue_acts>
<landmark_candidates> All landmarks on both maps with IDs </landmark_candidates>
```

Figure 2: Prompt template for annotation, abbreviated due to space limits. Figure 4 in Appendix A.1 provides a more detailed version of the prompt template.

- **Landmark candidates:** List of landmarks on maps with their `umlm` IDs, based on the scheme defined in the previous section.

The prompt (see Figure 2 and 4) guides the model through a sequential and heuristic annotation workflow. First, the model identifies the speaker and addressee roles (giver/follower) and assigns the speaker's intended landmark ID. Second, it evaluates `is_quantificational` to determine whether the expression presupposes a referent. If not quantificational, the model proceeds through the remaining four attributes in order. Finally, if `is_grounded=true`, the model assigns the addressee's interpreted landmark ID. The model is asked to give a concise textual reason citing dialogue evidence for our further analysis.

### 4.3. Quality Evaluation

We apply the annotation pipeline to all 128 dialogues in the MapTask corpus, resulting in 13,077 annotated reference expression instances<sup>5</sup>.

To evaluate the reliability of the annotation pipeline, we selected 3 dialogues (dialogue IDs: `q1ec2`, `q1nc3`, `q1nc7`) for complete human verification, comprising 504 reference expressions. We compute micro-averaged accuracy and f1-scores for each binary attribute by comparing LLM predictions against the human gold standard. The human gold standard was collected by one experienced

<sup>5</sup>The LLM omitted annotations for 4 REs (from 3 different dialogues), despite their presence in both the dialogue context and the RE list, likely due to uncertainty during the generation process.

annotator (one of the paper authors) following the same annotation guidelines, with ambiguous cases discussed and resolved with other authors.

Attribute	N	Acc.	F1	Errors
<code>is_quantificational</code>	504	100.00%	1.00	0
<code>is_specified</code>	466	97.85%	0.990	10
<code>is_accommodated</code>	455	97.58%	0.988	11
<code>is_grounded</code>	447	96.20%	0.980	17
<code>is_imagined</code>	422	97.63%	0.891	10
<b>RE-level: 28/504 REs (5.6%) with <math>\geq 1</math> error</b>				<b>48</b>

Table 2: Evaluation metrics by attribute across 3 dialogues (`q1ec2`, `q1nc3`, `q1nc7`; micro-averaged). N indicates the number of REs evaluated at each attribute (determined by gold annotations). Bottom line reports the RE-level error count (unique REs with  $\geq 1$  attribute errors), which reflects an accuracy of 94.4%.

The LLM annotations exhibit relatively high validity across multiple evaluation dimensions when compared to human gold standard. For all grounded reference expressions, the LLM achieves 95.5% accuracy and 99.5% micro-F1, demonstrating reliable recovery of each participant's intended and interpreted referents. Table 2 then presents the micro-averaged evaluation results for the five binary attributes across all 3 dialogues. The LLM achieves high performance across all attributes, with accuracy ranging from around 97% to 100% and F1 scores from 0.89 to 1.00. At the reference expression level, 28 of 504 REs (5.6%) contained at least one attribute error, totaling 48 attribute-level mismatches. Attribute-level error breakdown

shows that `is_grounded` has the highest error count (17), largely stemming from *false negatives* where the RE is grounded in gold but missed by the model, which typically occur when grounding is inferred from implicit confirmation or reuse rather than explicit echoing.

## 5. Discussion

### 5.1. Classification of Understanding States

Our annotation framework derives understanding states from hierarchical attribute decisions at the RE level. We define three primary categories that capture the alignment or divergence between speaker and addressee interpretations:

**Aligned** The addressee successfully accommodates and grounds the reference expression to the speaker’s intended landmark, and the landmark `mTlm` IDs match between speaker and addressee interpretations. This represents successful reference resolution.

**Misunderstood** The addressee grounds the reference expression to a different landmark from that the speaker intended. Both participants believe successful reference has occurred, yet their interpretations diverge. Such silent misalignments may persist undetected unless subsequent dialogue reveals the discrepancy.

**Pending** The interlocutors await further information or signal comprehension failure. This category includes several communication breakdown scenarios corresponding to the attributes discussed in Section 4.1, except for `is_imagined`.

For each RE, we infer the interlocutors’ understanding state from our annotated data, and Table 3 presents the distribution. The “misunderstood” type reaches 7.07%. However, as discussed in Section 3.2, 7.5% of landmarks exhibit *lexical discrepancies* where the same landmarks are labeled with different names on maps. For instance, when the giver says “the old mill” (referring to `m12_old_mill@g`) and the follower only has `m12_mill_wheel@f` on his/her map, they may initially correct each other but usually quickly reach an agreement that they are referring to the “same” landmark at the same location. In such cases, despite the initial lexical mismatch, we assume they have reached an aligned understanding state.

Verified by our observations, we unified the 10 pairs of lexical variants and adjust the “misunderstood” cases caused by them into “Aligned” ones.

Table 3 demonstrates the understanding state distribution after this unification as well. Now the misunderstanding rate drops to 1.82%, a significant reduction from the original 7.07%.

State	Before		After	
	Count	%	Count	%
Aligned	8,750	66.9	9,435	72.1
Pending	3,403	26.0	3,403	26.0
Misunderstood	924	7.07	239	1.82
Total REs	13,077	100.0	13,077	100.0

Table 3: Distribution of understanding states before and after lexical variant unification in the LLM-annotated MapTask reference expression dataset.

The relatively low misunderstanding rate aligns with Healey et al. (2018)’s analysis of rarity of sustained misunderstandings in collaborative dialogue: participants actively monitor alignment through continuous “running repairs” and employ clarification strategies to prevent persistent divergence.

Discrepancy Type	Total REs	Misunderstanding	Rate (%)
Lexical	914	13	1.4
Multiplicity	956	115	12.0
Existence	2,878	93	3.2
Identical	8,333	18	0.2
Total	13,081	239	1.8

Table 4: Misunderstanding counts and rates by landmark discrepancy type.

Table 4 demonstrates how these misunderstandings distribute across different types of landmark discrepancies. The distribution shows a pattern: multiplicity discrepancies account for 48.5% of all misunderstandings (115/239 cases) despite representing only 7.3% of reference expressions in the corpus. The misunderstanding rate of this type is more than 6× higher than the corpus average (12.0% versus 1.8%). This reveals the success of the MapTask asymmetric design in eliciting misunderstandings, as multiplicity discrepancies inherently create ambiguity about which specific landmark instance is being referred to: when a landmark type appears twice on the giver’s map but only once on the follower’s map, participants often assume referential uniqueness (Lane et al., 2006; Bard et al., 2000). For example, the giver refers to instance #0, the follower accommodates to their visible instance #1, and both believe successful reference has occurred. This silent divergence persists unless subsequent spatial descriptions reveal the mismatch.

For “misunderstood” cases within the existence discrepancy and identical landmarks, we find that the most frequently occurring cause is still lexical similarity. For example, “white mountain” on Map

m9 with the highest misunderstanding occurrences in the identical type is usually confused with “slate mountain” on the same map, which has the highest misunderstanding occurrences in the existence discrepancy type. “Rock fall” with the second highest misunderstanding occurrences in the existence discrepancy type is often misunderstood as “stones” on the same map. These landmarks all have different names, icons, and locations.

## 5.2. Reference Chain

To analyze how understanding evolves across repeated references to the same landmark, we extracted *reference chains*—sequences of reference expressions targeting the same landmark (same `mt_lm`) within a dialogue. Across 128 dialogues, we identified 1,665 reference chains with a mean length of 6.89 REs per landmark (range: 1–70).

Discrepancy Type	# Chains	Mean Length	Max Length	Median Length
Identical	939	8.16	70	6
Lexical	79	10.70	61	9
Existence	539	4.09	31	3
Multiplicity	108	7.00	26	5

Table 5: Reference chain length by discrepancy type.

The length of these chains partly reflects how persistently participants negotiate understanding of landmarks and use these landmarks to generate spatial descriptions. Table 5 summarizes chain lengths by discrepancy type. Lexical variants yield the longest chains on average (10.7 references), as participants often need multiple turns to resolve naming mismatches and confirm shared understanding. Landmarks with no discrepancy generate the second longest chains on average, while existence discrepancies yield the shortest. This suggests that existence discrepancies are usually quickly ruled out as reference anchors when participants realize the landmark is absent from one map, whereas references to non-discrepant landmarks persist longer because they support iterative spatial description and navigation coordination, aligning with previous findings (Viethen et al., 2011a; Reitter and Moore, 2014).

To illustrate how understanding states evolve through a reference chain, we present a simple multiplicity discrepancy case from dialogue `q7ec2`. The landmark *vast meadow* appears twice on the giver’s map (`m11_vast_meadow#0@g`, lower/unique; `m11_vast_meadow#1@g`, upper/shared) but only once on the follower’s (`m11_vast_meadow#1@f`, upper/shared).

Table 6 traces five reference expressions to this landmark (see Figure 5 in Appendix A.2 for the dialogue excerpt). The giver initiates with “a meadow” (`ref.112`), intending the lower instance (`#0@g`).

The follower queries “a vast meadow?” (`ref.114`) and the giver confirms (`ref.115`). However, the follower identifies his/her own, the upper shared, vast meadow (`#1@f`), creating a **misunderstanding**: both have landmarks in mind, yet they are different ones.

Ref	Spk	Expression	Giver ID	Follower ID	State
112	G	“a meadow”	<code>#0@g</code>	<code>ungrounded</code>	Pending
114	F	“a meadow”	<code>#0@g</code>	<code>ungrounded</code>	Pending
115	G	“a vast meadow”	<code>#0@g</code>	<code>#1@f</code>	<b>Misunderstood</b>
116	F	“a vast meadow”	<code>#1@g</code>	<code>#1@f</code>	Aligned
119	F	“vast meadow”	<code>#1@g</code>	<code>#1@f</code>	Aligned
127	G	“a vast meadow”	<code>#0@g</code>	<code>ungrounded</code>	Pending

Table 6: Reference chain for *vast meadow* in dialogue `q7ec2`. Landmark IDs (interpretations of the giver and the follower) are abbreviated (full prefix: `m11_vast_meadow`).

In the next turn (`ref.116`), the follower describes the location as “just below the blacksmiths and carpenters,” which corresponds to the upper meadow shared on both maps. The giver accepts this reframing via confirmation, and both participants align on the upper instance (`#1`). This alignment persists through `ref.119`. However, later in the dialogue (`ref.127`), the giver again mentions “a vast meadow” in a different spatial context, referring back to the lower instance (`#0@g`), but the follower cannot ground it, reverting the state to *pending*.

This chain exemplifies how multiplicity discrepancies create recurring alignment challenges: the same landmark name maps to different physical locations on the maps for participants, producing misunderstandings that may resolve through spatial negotiation but can resurface later in the dialogue.

## 5.3. Turns to Ground

To further disentangle the efforts participants invest in coordinating understanding of landmarks from using grounded landmarks for navigation, we measure the number of turns (counted per utterance) between a chain’s first reference and its first aligned state, and term this measure *turns-to-ground*. Figure 3 shows the cumulative distribution of turns-to-ground by discrepancy type across 1,334 chains that reached alignment. The longer negotiation for multiplicity likely reflects the cognitive complexity of distinguishing between multiple same-named landmarks versus simply confirming/denying existence. Participants must coordinate not only on the landmark name but also on which specific instance is intended, requiring more iterative spatial descriptions and confirmations.

There are 331 chains that never reach an aligned state at all, indicating that participants sometimes stop early and abandon negotiation on some land-

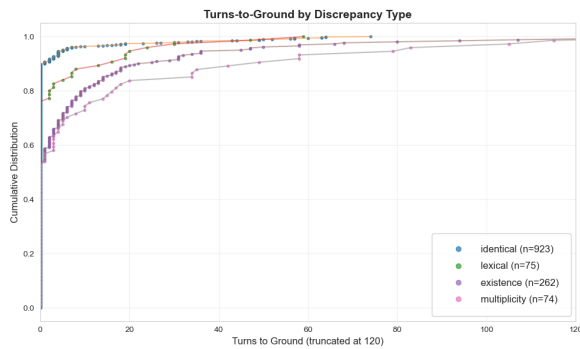


Figure 3: Cumulative distribution of turns-to-ground by discrepancy type.

marks. This aligns with task-oriented dialogue efficiency.

#### 5.4. Special Case

In dialogue `q8ec4`, the giver mistakenly takes the finish as the start point, leading to a persistent spatial misalignment that lasts for nearly half the conversation. The speaker insists on using absolute directions and distances (e.g., “thirty degrees up the way” and “two inches to your left”) to give instructions rather than landmark-relative spatial descriptions, extending the misunderstanding further.

In the annotation, the LLM fails to capture the misalignment since the beginning, partly because of an omniscient bias: the model can “see” both maps at once, so it underestimates asymmetric knowledge and treats contradictions as noise rather than evidence of a wrong starting point. Prior studies show that human addressees can achieve more efficient communication than overhearers, while VLMs perform poorly at overhearing (Schober and Clark, 1989; Wang et al., 2025). Future evaluations could explicitly test omniscience and perspective (analyzer versus participant versus overhearer conditions) to test models’ perspective-taking ability in grounded dialogue.

## 6. Conclusion

Understanding when interlocutors truly ground reference expressions versus only appear to do so requires capturing both sides’ interpretations. This work establishes such a perspectivist view by mapping reference expressions to understanding states in asymmetric collaborative dialogue. Our annotation scheme disentangles the speaker’s intended and the addressee’s interpreted referents through a five-attribute hierarchy and unified landmark IDs, operationalized via a scheme-constrained prompting pipeline.

Our quantitative analyses reveal that given the rarity of full misunderstandings, multiplicity discrep-

ancies act as a consistent hazard for eliciting misalignment. Our statistics also support prior studies on interlocutors’ preferences for reusing grounding-established REs (Healey et al., 2018; Viethen et al., 2011a; Reitter and Moore, 2014).

We plan to establish evaluations based on these annotated dialogues for evaluating (V)LLMs’ incremental grounding abilities under different information settings, moving beyond object identification toward perspective-taking in collaborative dialogue.

## 7. Limitations

**Annotation Reliability** Our annotation pipeline relies on a single LLM (GPT-5), and the human gold standard was produced by one annotator (with ambiguous cases discussed among co-authors). This constrains claims about cross-model generalizability and inter-annotator agreement. Although our annotation scheme is designed as a structured and heuristic decision cascade with JSON-schema validation and explicit explanations — achieving 94.4% RE-level accuracy across 504 reference expressions (see Section 4.3) — these results reflect one model under one set of conditions. Comprehensive evaluation across multiple LLMs and human annotators remains an important direction for future work.

**Spatial Understanding Precision** Our perspectivist scheme tracks landmark-level interpretations but does not capture finer-grained spatial reasoning. The MapTask involves route navigation between landmarks, yet our annotations focus solely on landmark reference resolution. Path choices, direction ambiguities, and spatial relations between landmarks remain unannotated.

**Multimodal information** The original MapTask corpus was collected in conversational settings where participants could perceive prosody, and even have eye contact in half of the cases. Our annotation pipeline relies exclusively on dialogue transcripts and maps, potentially missing non-verbal signals that influence grounding. For instance, backchannel responses like “mmhmm” may carry different conversational functions (acknowledgment versus query) depending on intonation, which our text-only approach cannot distinguish. While we demonstrate that text-based LLM annotation achieves reasonable accuracy, incorporating multimodal information could further improve annotation quality and theoretical coverage.

## 8. Ethics Statement

This research involves annotated dialogue data from the HCRC MapTask corpus, which consists

of recordings of human participants engaged in a collaborative task. The corpus is publicly available for research purposes and doesn't contain any sensitive personal information.

## 9. Code and Data Availability

We release our full code, annotation prompt templates, and the complete annotated dataset to facilitate reproducibility<sup>6</sup>.

## 10. Acknowledgements

We appreciate the helpful comments and suggestions from the anonymous reviewers. We would also like to thank Amy Isard, Matthew Aylett, and Martín Villalba for their valuable input on the analysis of misunderstanding in MapTask during the early stages of this project. This work is funded by the Dutch Research Council (NWO) through the AiNed Fellowship Grant NGF.1607.22.002, *Dealing with Meaning Variation in NLP*.

## 11. References

- Anne H Anderson, Miles Bader, Ellen Gurman Bard, Elizabeth Boyle, Gwyneth Doherty, Simon Garrod, Stephen Isard, Jacqueline Kowtko, Jan McAllister, Jim Miller, et al. 1991. The HCRC Map Task corpus. *Language and speech*, 34(4):351–366.
- Ellen Gurman Bard, Anne H Anderson, Catherine Sotillo, Matthew Aylett, Gwyneth Doherty-Sneddon, and Alison Newlands. 2000. [Controlling the Intelligibility of Referring Expressions in Dialogue](#). *Journal of Memory and Language*, 42(1):1–22.
- Ljubiša Bojić, Olga Zagovora, Asta Zelenkauskaite, Vuk Vuković, Milan Čabarkapa, Selma Veseljević Jerković, and Ana Jovančević. 2025. Comparing large language models and human annotators in latent content analysis of sentiment, political leaning, emotional intensity and sarcasm. *Scientific reports*, 15(1):11477.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Jean Carletta, Amy Isard, Stephen Isard, Jacqueline C Kowtko, Gwyneth Doherty-Sneddon, and Anne H Anderson. 1997. The reliability of a dialogue structure coding scheme. *Computational linguistics*, 23(1):13–31.
- Huiyao Chen, Meishan Zhang, Jing Li, Min Zhang, Lilja Øvrelid, Jan Hajič, and Hao Fei. 2025. Semantic role labeling: A systematical survey. *arXiv preprint arXiv:2502.08660*.
- Herbert H. Clark and Susan E. Brennan. 1991. [Grounding in communication](#). In Lauren B. Resnick, John M. Levine, and Stephanie D. Teasley, editors, *Perspectives on Socially Shared Cognition*, pages 127–149. American Psychological Association.
- Herbert H. Clark and Edward F. Schaefer. 1989. [Contributing to discourse](#). *Cognitive Science*, 13(2):259–294.
- Herbert H. Clark and Deanna Wilkes-Gibbs. 1986. [Referring as a collaborative process](#). *Cognition*, 22(1):1–39.
- Robin Cooper, Staffan Larsson, Colin Matheson, Massimo Poesio, and David Traum. 1999. Coding instructional dialogue for information states. *Deliverable D1*, 1.
- Florian Eichen, Yang Janet Liu, Barbara Plank, and Michael A Hedderich. 2025. Probing llms for multilingual discourse generalization through a unified label set. *arXiv preprint arXiv:2503.10515*.
- Allyson Ettinger, Jena Hwang, Valentina Pyatkin, Chandra Bhagavatula, and Yejin Choi. 2023. [“you are an expert linguistic annotator”: Limits of LLMs as analyzers of Abstract Meaning Representation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8250–8263, Singapore. Association for Computational Linguistics.
- Saibo Geng, Martin Josifoski, Maxime Peyrard, and Robert West. 2023. [Grammar-constrained decoding for structured NLP tasks without finetuning](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10932–10952, Singapore. Association for Computational Linguistics.
- Jonathan Ginzburg. 2012. *The Interactive Stance*. Oxford University Press.
- Patrick G. T. Healey, Gregory J. Mills, Arash Eshghi, and Christine Howes. 2018. [Running Repairs: Coordinating Meaning in Dialogue](#). *Topics in Cognitive Science*, 10(2):367–388.

---

<sup>6</sup><https://github.com/chnln/grounded-misunderstandings-in-maptask>

- William S Horton and Boaz Keysar. 1996. When do speakers take into account common ground? *Cognition*, 59(1):91–117.
- Boaz Keysar, Dale J Barr, Jennifer A Balin, and Jason S Brauner. 2000. Taking perspective in conversation: The role of mutual knowledge in comprehension. *Psychological science*, 11(1):32–38.
- Ibrahim Khalil Khebour, Kenneth Lai, Mariah Bradford, Yifan Zhu, Richard A. Brutti, Christopher Tam, Jingxuan Tu, Benjamin A. Ibarra, Nathaniel Blanchard, Nikhil Krishnaswamy, and James Pustejovsky. 2024. [Common Ground Tracking in Multimodal Dialogue](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3587–3602. ELRA and ICCL.
- Yury Kuratov, Aydar Bulatov, Petr Anokhin, Ivan Rodkin, Dmitry Sorokin, Artyom Sorokin, and Mikhail Burtsev. 2024. Babilong: Testing the limits of llms with long context reasoning-in-a-haystack. *Advances in Neural Information Processing Systems*, 37:106519–106554.
- Kenneth Lai, Lucia Donatelli, Richard Brutti, and James Pustejovsky. 2025. A model of information state in situated multimodal dialogue. In *16th International Conference on Computational Semantics*, page 292.
- Liane Wardlow Lane, Michelle Groisman, and Victor S. Ferreira. 2006. [Don't Talk about Pink Elephants! Speakers' Control over Leaking Private Information during Language Production](#). *Psychological Science*, 17(4):273–277.
- Alex Lascarides and Nicholas Asher. 2009. Agreement, disputes and commitments in dialogue. *Journal of semantics*, 26(2):109–158.
- Phillip Y Lee, Jihyeon Je, Chanho Park, Mikaela Angelina Uy, Leonidas Guibas, and Minhyuk Sung. 2025. Perspective-aware reasoning in vision-language models via mental imagery simulation. *arXiv preprint arXiv:2504.17207*.
- Nelson F Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. Lost in the middle: How language models use long contexts. *arXiv preprint arXiv:2307.03172*.
- Max M Louwerse, Nick Benesh, Mohammed E Hoque, Patrick Jeuniaux, Gwyneth Lewis, Jie Wu, and Megan Zirnstein. 2007. Multimodal communication in face-to-face computer-mediated conversations. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 29.
- Matthew Marge and Alexander I Rudnicky. 2011. Towards overcoming miscommunication in situated dialogue by asking questions. In *AAAI Fall Symposium: Building Representations of Common Ground with Intelligent Agents*.
- Colin Matheson, Massimo Poesio, and David Traum. 2000. [Modelling Grounding and Discourse Obligations Using Update Rules](#). In *1st Meeting of the North American Chapter of the Association for Computational Linguistics*.
- OpenAI. 2025. [Gpt-5](#). Large Language Model.
- Kanghee Park, Timothy Zhou, and Loris D'Antoni. 2025. Flexible and efficient grammar-constrained decoding. *arXiv preprint arXiv:2502.05111*.
- Martin J. Pickering and Simon Garrod. 2004. [Toward a mechanistic psychology of dialogue](#). *Behavioral and Brain Sciences*, 27(02).
- Massimo Poesio, Florence Bruneseaux, and Laurent Romary. 1999. The mate meta-scheme for coreference in dialogues in multiple languages. In *ACL'99 Workshop Towards Standards and Tools for Discourse Tagging*, pages 65–74.
- Massimo Poesio, Rodolfo Delmonte, Antonella Bristot, Chiran Luminata, and Sara Tonelli. 2004. [The venex corpus of anaphora and deixis in spoken and written italian](#).
- Massimo Poesio and David R. Traum. 1997. [Conversational Actions and Discourse Situations](#). *Computational Intelligence*, 13(3):309–347.
- Ayesha Qamar, Jonathan Tong, and Ruihong Huang. 2025. Do llms understand dialogues? a case study on dialogue acts. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 26219–26237.
- David Reitter and Johanna D Moore. 2014. Alignment and task success in spoken dialogue. *Journal of Memory and Language*, 76:29–46.
- Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. A systematic survey of prompt engineering in large language models: Techniques and applications. *arXiv preprint arXiv:2402.07927*.
- David Schlangen and Gabriel Skantze. 2011. A general, abstract model of incremental dialogue processing. *Dialogue & Discourse*, 2(1):83–111.
- Michael F Schober and Herbert H Clark. 1989. [Understanding by addressees and overhearers](#). *Cognitive Psychology*, 21(2):211–232.

Takuma Udagawa and Akiko Aizawa. 2019. [A Natural Language Corpus of Common Grounding under Continuous and Partially-Observable Context](#).

Sebastian Varges. 2005. Spatial descriptions as referring expressions in the maptask domain. In *Proceedings of the Tenth European Workshop on Natural Language Generation (ENLG-05)*.

HAE Viethen, Robert Dale, and Markus Guhe. 2011a. Generating subsequent reference in shared visual scenes: Computation vs. re-use. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP) Edinburgh, Scotland*, pages 1158–1167. Unknown Publisher.

Henriette Viethen, Robert Dale, and Markus Guhe. 2011b. [The Impact of Visual Context on the Content of Referring Expressions](#). In *Proceedings of the 13th European Workshop on Natural Language Generation*, pages 44–52. Association for Computational Linguistics.

Zhengxiang Wang, Weiling Li, Panagiotis Kaliosis, Owen Rambow, and Susan Brennan. 2025. [LVLMs are bad at overhearing human referential communication](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 16758–16782, Suzhou, China. Association for Computational Linguistics.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. 2016. Modeling context in referring expressions. In *European conference on computer vision*, pages 69–85. Springer.

Zheyuan Zhang, Fengyuan Hu, Jayjun Lee, Freda Shi, Parisa Kordjamshidi, Joyce Chai, and Ziqiao Ma. 2024. Do vision-language models represent space and how? evaluating spatial frame of reference under ambiguities. *arXiv preprint arXiv:2410.17385*.

## A. Appendices

### A.1. Annotation Prompt Template

Figure 2 in the main text shows an abbreviated version of the annotation prompt template. Below

we provide the full content of each section in the prompt template<sup>7</sup>. Each section is wrapped in the XML style in the actual prompt, for example, the background section is enclosed by `<background>` and `</background>`.

<b>Background</b>
MapTask dialogues involve a route giver guiding a follower using two similar but non-identical maps, and they cannot see each other’s map. The giver has the route; some landmarks differ by name (lexical), existence (present on only one map), or multiplicity (appearing twice on one map). You are an expert annotator for reference expressions (REs). Now for each RE, you should annotate the speaker’s intended and the addressee’s interpreted landmarks using provided landmark IDs. Note that the speaker and addressee are determined by who produces and who receives the RE, not by the fixed giver/follower roles.
<b>Task Description</b>
Annotate every RE listed in the <code>target_ref_ids</code> section using landmark IDs from the candidate inventory. Also fill in the boolean attributes: <code>is_quantificational</code> (from the speaker’s perspective, whether the expression is used to check or deny the existence of a landmark rather than to refer to a specific one; the following four other attributes are from the addressee’s perspective), <code>is_specified</code> (whether the addressee provides any response that shows they have noticed the RE rather than skipping or ignoring it), <code>is_accommodated</code> (whether the RE is successfully received and understood by the addressee, false if the addressee mishears, misunderstands, or asks for repetition), <code>is_grounded</code> (whether the RE is linked to any concrete landmark on the maps by the addressee), and <code>is_imagined</code> (optional, whether the grounded landmark exists only on the speaker’s map). Base every decision strictly on the provided dialogue context and dialogue-act tags. Provided information includes: (1) dialogue context: dialogue transcript from the beginning of the dialogue until the end of the current transaction (a transaction is a series of moves from one point to another on the maps); (2) maps: both the giver’s and follower’s maps with map names at the bottom; (3) dialogue act tags: dialogue context annotated with dialogue acts; (4) landmark candidates: a list of all landmarks on both maps. Pick up IDs from the <code>landmark_candidates</code> section, and do not invent new IDs.
<b>Landmark ID Explanation</b>

<sup>7</sup>To enforce the model’s output to adhere to the required format, we not only include the output format section in the prompt, but also attach a JSON-schema file as output format parameter when calling the OpenAI Response API.

Landmark ID syntax: `<map-id>_<concept>#<ordinal>@<side>`. `<concept>` is the landmark name from the original dataset; `side`  $\in$  {g, f} indicates which map the landmark is located on; `ordinal` (0 or 1) is used only when a concept appears twice on the same map; `ordinal` numbers are ordered from 0 to 1 by their occurrences from bottom to top on the same map, and the number is the same across maps at matching locations. For simultaneous multi-target references, join their IDs with +.

#### Annotation Rule

Follow the annotation workflow for each RE step by step:

**Step 0:** Identify the speaker and the addressee from the sentence header; fill in the speaker's interpretation (intended landmark ID) first. Exception: if the speaker is querying or denying a landmark previously introduced by the addressee (e.g., "where is the running river?" / "I don't have the running river"), set the speaker's interpretation to empty ("").

**Step 1** (`is_quantificational`): Set to true when the RE is used only to check or deny existence (e.g., "Do you have a parked van?" / "I don't have any roman baths"); otherwise false.

**Step 2** (`is_specified`): When not quantificational, determine if the addressee shows any uptake of the RE in the context, rather than ignoring or skipping it without reaction; if so, set `is_specified=true`, otherwise false. Exception—implicit reuse: if this expression was previously grounded between the interlocutors with no subsequent repair/denial, the current RE may be treated tacitly grounded; set `is_specified=true`, and at step (4) mark it as grounded.

**Step 3** (`is_accommodated`): If specified, set to false when the addressee fails to process the RE, e.g., mishearing, confusion, or requesting repetition; otherwise true.

**Step 4** (`is_grounded`): When accommodated, set to true only if the addressee links the RE to a specific landmark on the maps; otherwise false.

**Step 5** (`is_imagined`): When grounded, set to true if the addressee's interpretation is only on the speaker's map; otherwise false.

**Step 6:** If grounded, fill in the landmark ID for the addressee's interpretation.

**Step 7:** Provide a concise (<50 words) evidence-based reason summarizing the judgement in the reason field for each RE.

#### Output Format

Return strict JSON only:

```
{
  "dialogue_id": str,
  "landmark_reference_expressions": [
    {
      "ref_id_unif": str,
      "is_quantificational": bool,
      "is_specified": bool,
```

```
      "interpretations": {
        "is_accommodated": bool,
        "is_grounded": bool,
        "is_imagined": bool,
        "giver": str,
        "follower": str
      },
      "reason": str
    },
    ...
  ]
}
```

#### Dynamic Sections

The remaining prompt sections are instantiated based on all referenced expressions for each request:

- `<target_ref_ids>` lists the RE IDs to annotate in the current request;
- `<context_dialogue>` provides the dialogue transcript from the start of the conversation through the current transaction;
- `<context_dialogue_acts>` provides the same context annotated with dialogue acts;
- `<landmark_candidates>` lists all landmarks on both maps with their `umlm` IDs.

Figure 4: Full annotation prompt template.

## A.2. Reference Chain Dialogue Excerpt

Table 6 in the main text traces the understanding state transitions for the landmark *vast meadow* in dialogue *q7ec2*. Below is the relevant dialogue excerpt providing context for those reference expressions (`refs` 112, 114, 115, 116, 119, 127). Reference expressions are enclosed in `<...>`; the six reference expressions are **bolded**.

#### Dialogue Excerpt: q7ec2 (utterances 131–161)

```
[giver utt:131] i've got «a
meadow»ref.112 to the left of «that»
due south pee porashin
[follower utt:132] «a vast
meadow»ref.114
[giver utt:133] «a vast meadow»ref.115
[follower utt:134] i've got «a vast
meadow»ref.116 up just below «the
blacksmiths» and «carpenters»
[giver utt:135] well is that huge
big and there's nothing at all in
that space
[follower utt:136] uh-huh except
[giver utt:137] about halfway up the
page
[follower utt:138] «vast meadow»ref.119
about halfway up the page it's empty
```

```
on mine
[giver utt:139] well you want to
well anyway follow «crane bay» round
«the curve»
... (utterances 140–160: participants navigate
along crane bay and attractive cliffs) ...
[giver utt:161] just below i've got
«a vast meadow»ref.127 there
[follower utt:162] right where will
i stop that line across the top
```

**Figure 5:** Dialogue excerpt from q7ec2  
(utterances 131–161).

**Context:** The giver's map contains two *vast meadow* instances: #0@g (lower, unique to giver) and #1@g (upper, shared with follower). The follower has only the upper instance (#1@f). At ref.112, the giver refers to the lower meadow. The follower's "a vast meadow?" (utt:132) is a clarification query. When the giver confirms (ref.115), the follower identifies the upper shared meadow instead, producing a misunderstanding. At ref.116, the follower's spatial description ("below the blacksmiths and carpenters") shifts the giver's grounding target to the upper/shared meadow, and both align. After a 28-utterance gap discussing other landmarks, the giver returns to "a vast meadow" (ref.127) in a new spatial context, but the follower is not able to ground it — the state reverts to pending.