

Repairs in a Block World: A New Benchmark for Handling User Corrections with Multi-Modal Language Models

Javier Chiyah-Garcia Alessandro Suglia Arash Eshghi

Heriot-Watt University, Edinburgh, United Kingdom

{fjc3, a.suglia, a.eshghi}@hw.ac.uk

Abstract

In dialogue, the addressee may initially misunderstand the speaker and respond erroneously, often prompting the speaker to correct the misunderstanding in the next turn with a Third Position Repair (TPR). The ability to process and respond appropriately to such repair sequences is thus crucial in conversational AI systems. In this paper, we first collect, analyse, and publicly release BLOCKWORLD-REPAIRS: a dataset of multi-modal TPR sequences in an instruction-following manipulation task that is, by design, rife with referential ambiguity. We employ this dataset to evaluate several state-of-the-art Vision and Language Models (VLM) across multiple settings, focusing on their capability to process and accurately respond to TPRs and thus recover from miscommunication. We find that, compared to humans, all models significantly underperform in this task. We then show that VLMs can benefit from specialised losses targeting relevant tokens during fine-tuning, achieving better performance and generalising better to new scenarios. Our results suggest that these models are not yet ready to be deployed in multi-modal collaborative settings where repairs are common, and highlight the need to design training regimes and objectives that facilitate *learning from interaction*. Our code and data are available at www.github.com/JChiyah/blockworld-repairs

1 Introduction

Unlike its formulation in much of the literature within NLP (see especially Wang et al., 2018, 2019), Natural Language Understanding (NLU) is not a unilateral, passive process, but an (inter)active one (see Schlangen, 2023, for expansive discussion): in everyday conversation, people continuously work together to *negotiate* shared understanding and coordination in order to move the conversation forward (Clark, 1996; Clark and Brennan, 1991; Goodwin, 1981; Healey et al., 2018; Mills,

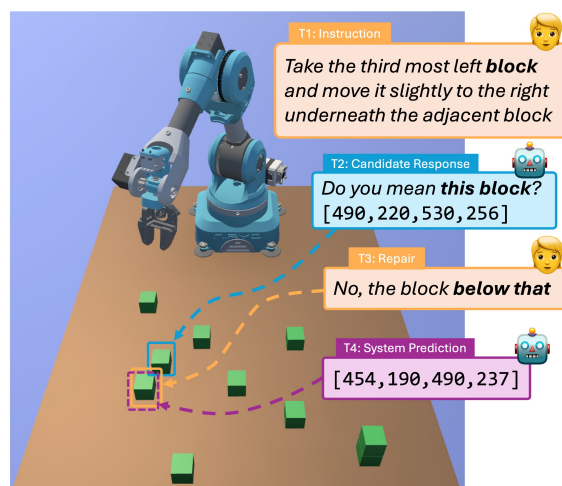


Figure 1: Example dialogue from BLOCKWORLD-REPAIRS: after predicting an incorrect response, VLMs must accurately interpret the repair to produce the correct bounding box prediction, a critical skill for human-robot collaboration tasks.

2007). One of the key interactional processes that enables this is called *repair* (Schegloff et al., 1977; Schegloff, 1992) – see Figure 1: a set of universal and highly systematised corrective feedback mechanisms for dealing with *miscommunication* when it arises in conversation (Enfield et al., 2013; Dingemans et al., 2015).

Therefore, the ability to interpret and generate effective repair sequences is crucial to *robust* and *faithful* Conversational AI technology. This need is especially acute in settings where more fine-grained levels of understanding are required, for example in embodied human-machine collaboration. Indeed the frequency of miscommunication in human dialogue is known to vary with both the overarching task and the medium of communication (Colman and Healey, 2011).

In this paper, we focus on a particular class of corrective feedback, so called, *Third Position Repairs* (henceforth TPR; Schegloff, 1992) in multi-

modal settings. These occur when the addressee initially misunderstands the speaker (see Figure 1 at T1, the *trouble source* turn), responds based on this misunderstanding (at T2), which in turn reveals the misunderstanding to the addressee who then goes on to correct the misunderstanding (at T3). The addressee then provides a new response (at T4) to the original request (at T1) based on the correction. TPRs have been largely neglected in the NLP community, likely due to underestimation of the importance of repair phenomena in dialogue, and thereby lack of appropriate data and frameworks for model training and evaluation. Repair is of particular interest in (multi-modal) language modelling since it almost invariably involves highly focused / specific revisions to the representations already built by the model (e.g. revisions to the referents of referring expressions, like in our case here), based on the context-dependent structure of the corrective feedback. Repairs can thus be used as probes into how models represent meaning and the structure of the overarching task (see Chiyah-Garcia et al. (2023); Madureira and Schlangen (2023); Benotti and Blackburn (2021) for different implementations of this idea). And as we show here, this is also consequential for how language models should be trained to process TPRs including the specific objective functions involved.

Our contributions are thus as follows: 1) we collect and release BLOCKWORLD-REPAIRS (BW-R), a dataset of collaborative dialogues in a tabletop manipulation task that, by design, has a high potential for referential ambiguity and is focused on complex multi-modal task instructions as well as TPRs; 2) we establish a human baseline for the whole task through an in-person human study where participants try to solve the task using these dialogues; 3) we show that state-of-the-art Vision Language Models (VLM) can learn to process TPR sequences through a specialised training regime, with some opportunity for generalisation; and 4) considering the substantial performance gap between the best models and humans, we present an in-depth error analysis comparing the two, revealing that models struggle with references that humans find easy.

2 Background

Communicative Grounding, Miscommunication and Repair What is often ignored or glossed over in today’s research on Conversational AI is that conversation involves collaborative effort from

speakers and addressees to ensure that what is said is understood sufficiently for the task at hand before the conversation can move forward: this almost continuous exchange of both positive and negative feedback is called *communicative grounding*¹ (see Clark, 1996, and many others following). This exchange of feedback enables interactants to coordinate their (linguistic and non-linguistic) actions in conversation (Eshghi et al., 2022; Mills, 2014), and enables shared languages to be established and sustained (Healey et al., 2018). *Miscommunication* occurs when one of the interactants detects a problem in their own, or another’s understanding, and is usually dealt with immediately using different forms of corrective feedback, or, as they are collectively called, *repair* (Schegloff, 1992). These are classified in the literature based on who initiates the repair; who performs it; and where the actual correction takes place. For example, Clarificational Exchanges (CE) whereby the addressee produces a request for clarification (CR) and the original speaker provides a response / correction in the next turn, constitute other-initiated, self-repairs. On the other hand, TPRs, our focus in this paper, are self-initiated, self-repairs on the third turn.

Computational Models of Repair Considerable attention has been paid to computational models for the interpretation and generation of *self-repair* (see Hough and Schlangen, 2015; Hough, 2015; Shalymov et al., 2017; Eshghi and Ashrafzadeh, 2023; Buß and Schlangen, 2011; Hough and Purver, 2012, among others): a class of repairs whereby the speaker corrects themselves on the fly within the same conversational turn (e.g. “User: I want to go to London uhm sorry Paris”). Similarly, the crucial role of generating and responding to CRs (e.g. “Pardon/what/who?”) in conversational models has long been recognised (see San-Segundo et al., 2001; Purver, 2004; Purver and Ginzburg, 2004; Rieser and Moore, 2005; Rodríguez and Schlangen, 2004; Rieser and Lemon, 2006, among others), but existing systems either remain limited (e.g. Curry et al. (2018)) or do not support this at all – see Purver et al. (2018) for an overview. Recent work tries to identify when to pose a CR (Addlesee et al., 2024; Madureira and Schlangen, 2024, 2023; Zhu et al., 2021; Shi et al., 2022; Addlesee and Eshghi, 2021), but few evaluate the ability of models to process their responses (Gervits et al.,

¹Not to be confused with *symbol grounding* but see (Larsson, 2018) for how the two are related.

2021; Aliannejadi et al., 2021). Chiyah-Garcia et al. (2023) show that the ability of Vision and Language Models (VLM) to effectively process clarificational exchanges depends on the level of granularity of the model’s cross-modal representation, and therefore also the specific objectives used to train the models.

Our work here is closest in spirit to that of Balaraman et al. (2023) who provide the first large dataset of Third Position Repairs, and use it to evaluate LLMs. However, they focus only on unimodal repairs in the context of Conversational Question Answering, and do not perform any fine-tuning of the models they evaluate, like we do here.

VLMs for Situated Collaborative Tasks

Thanks to recent advances in the development of LLMs, there have been many attempts to derive Vision+Language models that use pretrained visual encoders to solve vision-language tasks (e.g., (Liu et al., 2024; Laurençon et al., 2024), *inter alia*). It is important to note that most of these models are trained to follow instructions that can be specified in a single turn. Many tasks can be included in this category such as visual question answering (Antol et al., 2015), image captioning (Lin et al., 2014), and referential expression resolution (Yu et al., 2016). To solve these tasks, models are trained to maximise the likelihood of the responses in the training data using supervised learning methods or using reinforcement learning from human feedback (Ouyang et al., 2022), ignoring the fact that a conversation is a process that unfolds over multiple turns where each of which matters to make correct decisions. For instance, in tasks such as Visual Dialogue (Das et al., 2017), it is well-known that a very limited number of responses are dependent on the dialogue history (Agarwal et al., 2020), making them not a suitable benchmark for truly collaborative tasks requiring the ability to establish common ground. As highlighted by Suglia et al. (2024), many Embodied AI benchmarks (e.g., ALFRED (Shridhar et al., 2020), Simbot Arena (Gao et al., 2023), etc.) also have similar issues considering that the level of ambiguity is reduced to the minimum, minimising the need for any form of correction. For this reason, BW-R represents the first benchmark that is aimed at assessing the ability of current VLMs to resolve TPRs—an important capability which is essential for establishing common ground.

3 The BlockWorld-Repair Dataset

We build on top of the Block World dataset by (Bisk et al., 2016), an instruction-following task where a robot manipulator has to move blocks on a virtual board. Humans provided instructions in natural language referring to a block to pick up and a new location to drop it with complex visual and spatial descriptions (i.e., “*the block on the top-right corner moves behind the middle-top block*”).

Prior works have struggled to handle the most ambiguous setting with blank blocks (Tan and Bansal, 2018; Mehta and Goldwasser, 2019; Dan et al., 2021). However, here we argue that single-turn instructions are insufficient to solve highly ambiguous environments or real-world situations. Even the most detailed referring expressions in the Block World may fail to uniquely identify a referent (see Figure 1), which repairs could alleviate by narrowing down the candidate pool and introducing new information to identify the correct referent (T3 in Figure 1). Similar to how humans misinterpret and subsequently repair referential ambiguities with TPRs, VLMs must also be capable of handling multi-modal TPRs in dialogues, as they are essential for mutual coordination.

Differences from the original Block World The original dataset consists of pairs of images and single-turn instructions, requiring agents to resolve inherently ambiguous commands. However, it does not account for dialogue or clarificational exchanges as strategies to overcome environmental ambiguity, which humans commonly use to repair miscommunications. To address this, we extend the dataset by incorporating brief dialogues where humans interact with the agent and produce Third Position Repairs. We preserve the original Block World instructions and test set while adding our dataset as an extension to explore how humans resolve ambiguity through clarification and repair strategies.

Ecological validity In contrast to more recent photo-realistic simulated environments (e.g., AI2Thor (Kolve et al., 2017)), we build BW-R using the Unity engine to simulate a realistic tabletop manipulation task that gives us the ability to specifically control for the conditions that require collaboration between agents to correctly complete the task. This is also similar to previous works that sacrifice photo-realistic vision to assess the systematic generalisation ability of VLMs for tabletop

manipulation tasks (e.g., VIMA-Bench (Jiang et al., 2023)), or works that focus on dialogue coordination in highly ambiguous scenarios (e.g., The Cups dataset (Dobnik et al., 2020)). Appendix A provides further comparisons to related datasets.

3.1 Dataset Collection

We built a dialogue interface on Amazon Mechanical Turk (AMT) where workers chat with a robot agent and provide instructions to move blocks in a virtual world, similar to the original dataset. The robot agent is prone to mistakes and verifies its actions by pointing to blocks or locations. AMT workers can then collaborate with the robot to correct misunderstandings via a short dialogue through complex, context-dependent TPRs, based on the robot’s indicated position. Figure 1 shows an example dialogue.

The robot performs the correct action 70% of the time² and selects the appropriate block or location after the corresponding repair, irrespective of the quality of the TPR. Consequently, humans provided at most one initial instruction specifying the block to pick and where to place it, one TPR indicating the **source block** to move, and one TPR specifying the **target position** for the block. The initial instructions are thus equivalent to the original commands.

As a result of this data collection, we obtained 795 dialogues after filtering out low-quality and problematic ones, which include 795 initial instructions, 629 source block TPRs and 635 final target position TPRs, resulting in 2059 total entries. We provide further details about the AMT setup, task and dataset statistics in Appendix B.1.

3.2 Quality Control and Human Baselines

Since we used an automated agent to collect the dialogues and AMT is prone to provide low-quality data (Saravanos et al., 2021), we validated a data subset through an in-person study. We recruited 22 participants to take on the robot agent’s role, following the instructions and repairs previously given by human workers.

Participants interacted with a similar interface to the data collection, with a randomly selected dialogue displayed in the chat window. Their task was to choose the source block and a target position they believed the dialogue referenced within a large image of the virtual table. They did not

²For comparison, the best models only pick the correct block 54% of the time (Tan and Bansal, 2018).

have access to the true block configurations and received no feedback on whether they selected the correct block/location. Participants did not see repeated entries (same dialogue), and at least two participants annotated each entry. After filtering out incorrect or problematic entries, we obtained 991 action annotations.

Participants performed well overall and rated the AMT instructions highly, achieving 68% accuracy in selecting the correct block to move, rising to 75% after repairs, with similar success in target position prediction. These results emphasise the need for models to process TPRs accurately. We analyse human performance in §6.

4 Experiments

This section evaluates the ability of VLMs to process instructions and repairs in an instruction-following task within a situated context. The repair turns in the BW-R are meaningful only when interpreted alongside the robot’s pointing position at the time of the repair, based on the first instruction; thus, they are incomplete on their own. Consequently, we have dialogue triplets that are intrinsically connected and can only be comprehended as a whole: the *initial instruction*, the *incorrect candidate prediction*, and the *repair*. We argue that more general models must be able to process both initial instructions and any subsequent repairs.

4.1 Experimental Setup

4.1.1 Dataset

We combine the Block World and BW-Rdata (70/30 train/test), which contains two types of entries: *single-instructions* (1 turn) and *TPRs* (3 turns: instruction, candidate response, repair). We ablate the data during our experiments to analyse fine-grained information about the models’ capabilities. Refer to Appendix B for further dataset details.

4.1.2 Tasks

Inspired by Bisk et al. (2016), we evaluate models on two prediction tasks: 1) *Source block prediction*, which consists of predicting the block to pick up out of the 10 candidates in the virtual world; and 2) *Target position prediction*, which consists of predicting where to drop the block on the table.

Previous works have formulated these tasks to aid the prediction, such as finding a reference block and an offset distance from it (Bisk et al., 2016) or quadrants in a board (Mehta and Goldwasser, 2019). For our experiments, and follow-

ing common visual grounding objectives in VLMs (Wang et al., 2022; Liu et al., 2024; Chen et al., 2023), we prompt models to generate the coordinates of the bounding box for the block or location ($[x_{min}, y_{min}, x_{max}, y_{max}]$) in the range of 0-1, which we then scale to the image size. Predictions not matching this format during parsing are marked as out of distribution and removed. We use specific prompts for each model to match their training prompting strategy (see Appendix C).

4.1.3 Metrics

For evaluation, we use the metrics originally proposed in the Block World (Bisk et al., 2016). For **source block accuracy**, we use Intersection over Union and select blocks above a threshold. We compute **block distances** by converting bounding boxes to their respective XYZ coordinates in the virtual world and then calculating the distance in blocks between the predicted and true locations.

4.1.4 Models

We experiment with two state-of-the-art open-source models: Idefics2 8B (Laurençon et al., 2024), and LLaVA 1.5 7B (Liu et al., 2023). Both models have strong visual-textual cross-modality grounding capabilities and have been pre-trained on a mix of relevant vision and language objectives which facilitate bounding box generation. After preliminary tests with the “chatty” version of Idefics2, we decided to use the instruction-tuned version in all our experiments. We also provide zero-shot evaluations for GPT-4o (OpenAI, 2024), a frontier VLM model³.

4.2 Evaluating Models on TPRs

VLMs should inherently be able to process TPRs as they are fundamental components of human dialogues. We initially assess the VLMs’ out-of-the-box capabilities in a zero-shot setting, and then proceed to distil these capabilities through fine-tuning. In both cases, we prompt the models with the task instruction, followed by the conversation and provide the relevant image of the blocks. We fine-tune for 1 epoch using parameter-efficient methods (e.g., LoRA (Hu et al., 2021)) following the recommended hyperparameters for each model. We train a different model for each sub-task so models can fully leverage the input prompt and avoid interference between the two tasks (i.e., source block and

³To facilitate further analysis, we will release its predictions as part of our code release.

target position prediction).

Test Data	Train Data	Model	Source		Target
			Acc ↑	Mean (SD) ↓	Mean (SD) ↓
Instructions		Idefics2-zs	0.18	5.06 (±3.5)	9.55 (±2.5)
		LLaVA-zs	0.24	4.40 (±3.3)	5.74 (±3.0)
	IO	Idefics2-ft	0.33	3.33 (±3.0)	4.93 (±2.6)
		LLaVA-ft	0.25	3.77 (±2.9)	5.06 (±2.9)
	RO	Idefics2-ft	0.21	4.36 (±3.1)	5.38 (±2.4)
		LLaVA-ft	0.12	5.05 (±3.1)	6.32 (±2.9)
	Full	Idefics2-ft	0.30	3.39 (±2.9)	4.75 (±2.6)
		LLaVA-ft	0.22	3.82 (±2.9)	4.46 (±2.4)
		GPT-4o-zs	0.30	3.83 (±3.3)	4.22 (±2.6)
Repairs		Idefics2-zs	0.00	3.74 (±1.9)	4.13 (±2.1)
		LLaVA-zs	0.13	3.29 (±2.1)	3.45 (±1.2)
	IO	Idefics2-ft	0.21	4.19 (±2.7)	5.11 (±2.6)
		LLaVA-ft	0.18	3.43 (±2.4)	3.49 (±1.3)
	RO	Idefics2-ft	0.58	2.19 (±3.1)	7.07 (±2.7)
		LLaVA-ft	0.07	3.20 (±1.7)	3.67 (±1.3)
	Full	Idefics2-ft	0.26	3.43 (±2.4)	5.81 (±2.3)
		LLaVA-ft	0.44	2.66 (±2.8)	3.69 (±2.2)
		GPT-4o-zs	0.41	2.75 (±3.1)	2.95 (±2.2)
Random Baseline			0.10	6.50 (±3.0)	6.26 (±2.9)

Table 1: Model performance on source selection and target position prediction tasks on zero-shot (zs) or fine-tuned (ft) on data subsets: Instructions-Only (IO), Repairs-Only (RO) and Full. We compare source block accuracy (↑) and mean block distances (↓). Lower distances indicate predictions closer to the correct location.

Results Table 1 shows the performance of testing the VLMs on either instruction or repair-only data subsets with a mix of training BW-R data. Zero-shot models typically perform above the random baseline, initially demonstrating their visual-grounding capabilities. Fine-tuned models exhibit more nuanced results: they partially learn the task through training, particularly when trained and tested on the same type of entries (instructions or repairs). However, employing the full dataset does not consistently enhance performance: Idefics2 appears to adapt well to the training data or shows greater data efficiency, whereas LLaVA achieves better results with the full data and therefore suggests stronger generalisation capabilities.

Overall, the models struggle to exploit TPRs, especially in zero-shot settings. Intuitively, repairs should facilitate the task by introducing new information that refines the candidate pool. However, because repairs are often context-dependent (i.e., based on the candidate response), models require strong object-centric representations (Bengio et al., 2013; Seitzer et al., 2023) and a deeper understanding of relational dynamics between objects to effectively leverage TPRs (e.g., “the block to

the right”) (Ilinykh and Dobnik, 2021). This is a particular issue of current vision transformers, which frequently struggle with positional concepts such as spatial understanding (Ilinykh and Dobnik, 2022; Pantazopoulos et al., 2024), and tend to rely on language biases (Salin et al., 2022). GPT-4o is unique in its ability to process these repairs effectively and improve its performance with instructions. We see that models can learn to process TPRs with some training, and the following section explores methods to further distil this capability.

5 Learning to Process TPRs

During training, models encountered two types of entries: single-turn instructions or dialogues up to the TPR. In human conversations, we can effortlessly manage both scenarios; therefore it is reasonable to expect that models trained to predict bounding boxes from a single instruction should also be capable of handling repairs to their predictions. However, our previous results indicate that achieving both capabilities simultaneously is challenging, and that generalization incurs a cost, even when models address the same task.

To handle repairs, models cannot rely only on the last turn but must consider the context from the initial user instruction and the following candidate output. The intermediate predicted bounding box, even if inaccurate, also provides crucial information within the context of the TPR, as repairs rarely repeat a fully-formed instruction (i.e., “*The right one!*”). In the case of BW-R, repairs are relative to where the robot is currently pointing and often provide partial information to correct the robot’s action. During fine-tuning, following common approaches in VLM training (Laurençon et al., 2024), VLMs calculate the cross-entropy loss for all the tokens in the input including both the intermediate (wrong) and last (correct) generation. Thus, the models from §4.2 are learning from incorrect tokens, and therefore these tokens are somehow affecting the quality of the VLMs generations.

5.1 Masking Token Loss

To encourage VLMs to process and correctly handle TPRs, we experiment with different training regimes during fine-tuning. Specifically, as shown in Figure 2, we define different masking criteria for the cross-entropy loss and assess their effect on downstream performance.

Default loss, VLM’s default cross-entropy applied

to all tokens (Laurençon et al., 2024).

User-turn loss, we only calculate the loss for the user turns and the completion target. In this case, the intermediary prediction would not influence the overall loss but models could still benefit from learning instructions and the task format.

Completion-only loss to focus solely on the generated tokens and ignore all previous user and assistant turns in the entry. We expect these models to have the strongest generalisation capabilities, at the cost of not learning about the input instructions or corrections.

We follow the same procedure for fine-tuning as in Section 4.2, training one model per sub-task with the default hyperparameters.




Figure 2: Masking criteria for the cross-entropy loss.

5.2 Results

Table 2 shows the fine-tuned VLMs on the full data across different masking criteria⁴. We first observe that applying a mask on all tokens but completions has large benefits for both VLMs for repairs in source prediction, with strong improvements particularly when fine-tuning with repairs-only or the full dataset. This loss allows models to learn from TPRs and transfer some task knowledge from instructions to repairs when using the full data. In other settings, this loss has small negative effects on the VLM’s performance, especially when models are trained and tested on the same data. This is somewhat expected as calculating the loss for the complete input has benefits, such as using many more tokens during training, learning the input format and quickly adapting to the task. However, these models do not seem to learn the task well and thus do not generalise.

One of our assumptions is that the intermediate predictions hurt models, as these would calculate losses for incorrect tokens. The results of masking these tokens in the loss (user-turns only) suggest that this loss does not help the VLMs as much as

⁴All model results are provided in Appendix E.

	Model	Loss	Source		Target
			Acc \uparrow	Mean (SD) \downarrow	Mean (SD) \downarrow
Instructions	Idefics2	Default	0.30	3.39 (± 2.9)	4.75 (± 2.6)
		User-turn	0.36	3.12 (± 3.0)	4.59 (± 2.6)
		Completion	0.33	3.33 (± 3.0)	4.52 (± 2.6)
	LLaVA	Default	0.22	3.82 (± 2.9)	4.46 (± 2.5)
		User-turn	0.24	3.65 (± 2.9)	4.60 (± 2.5)
		Completion	0.19	3.93 (± 2.7)	4.51 (± 2.8)
Repairs	Idefics2	Default	0.26	3.43 (± 2.4)	5.81 (± 2.4)
		User-turn	0.32	2.94 (± 2.6)	6.00 (± 2.0)
		Completion	0.47	2.29 (± 2.4)	5.90 (± 2.7)
	LLaVA	Default	0.44	2.66 (± 2.8)	3.69 (± 2.2)
		User-turn	0.37	2.69 (± 2.5)	4.04 (± 2.4)
		Completion	0.54	2.01 (± 2.5)	4.56 (± 2.7)

Table 2: Fine-tuned VLMs with different loss criteria.

completion loss, and has a mixed bag of effects (positive or negative).

Most of the improvements with either of the losses only apply to source prediction. The models struggle with target position predictions and do not show clear benefits from these losses. Although these sub-tasks expect a similar output, predicting a target location is much harder, as it does not depend on a particular salient object in the image (i.e., a block) but on spaces between objects. It is thus unsurprising that VLMs pre-trained with visual grounding objectives (e.g., RefCOCO (Yu et al., 2016)) do not perform well when we are not referring to an object. In these cases, it seems that the additional prompt tokens are helpful for models to learn the task and masking them hurts their performance.

Furthermore, we see that models fine-tuned on instructions alone, with or without masking, acquire some task capabilities, but do not generalise well to repairs, and vice-versa. When trained with the full data and completion-only loss, VLMs learn what is important regardless of the conversation, and reach their best results, learning what is important particularly for repairs.

During these fine-tuning experiments, we identify two main factors at play: 1) the size of the training data; and 2) the incorrect candidate responses. When training on repairs, the size (i) becomes the bigger issue as masking considerably reduces the available tokens during training, which is crucial when data is smaller. Models do better when there is more training data, as happens with instructions, but they train on incorrect bounding boxes from the intermediate system turn (ii), hurting their performance. Handling repairs allows us to evaluate the models’ ability to interpret the initial instruction,

	Model	Source		Target
		Acc \uparrow	Mean (SD) \downarrow	Mean (SD) \downarrow
Inst.	Idefics2 - ft	0.21	3.97 (± 4.0)	5.52 (± 3.1)
	LLaVA - ft	0.16	4.08 (± 4.1)	4.63 (± 3.4)
	GPT-4o	0.26	4.60 (± 4.6)	4.30 (± 3.4)
	Human Participants	0.68	1.59 (± 1.6)	3.64 (± 3.4)
	Repairs	Idefics2 - ft	0.43	2.46 (± 2.5)
LLaVA - ft		0.60	1.82 (± 1.8)	4.82 (± 2.7)
GPT-4o		0.50	2.78 (± 2.8)	2.58 (± 1.8)
Human Participants		0.75	1.41 (± 1.4)	2.77 (± 2.8)

Table 3: Humans compared to the best models for the BW-R test subset with human annotations.

candidate response, and repair turns. Calculating losses on completions only ensures that models learn a more robust prompt representation that is useful when generalising across both types of data.

6 Error Analysis

We want models that can collaborate with humans even when misunderstandings arise, and thus we need to understand the differences in how models and humans solve the same visual-grounding task. This section attempts to analyse models beyond performance metrics and provides insight into the models’ behaviour, such as where mismatches arise compared to how humans process TPRs. We use the best models from §5 trained on the whole data using completions-only loss.

6.1 Human Comparison

To compare models with humans on this task, we use the data samples that we collected from our in-person human study, in which participants saw the same data that models are evaluated on. The results in Table 3 show clear differences, with models struggling to reach human performance but demonstrating significant benefits from TPRs after training. We can see that, despite Idefics2 and LLaVA having lower source accuracy for instructions, they both beat GPT-4o on distance for source. In this case, accuracy fails as a metric for prediction quality, since lower distances better reflect true proximity. We also see impressive results for LLaVA processing TPRs after our training for source prediction. Regarding target position, this task is considerably more challenging for humans and models, reflected in worse scores overall. LLaVA and Idefics2 struggle to process repairs related to target locations, in contrast to GPT-4o which outperforms the human participants.

Difficulty	Model	Source		Target
		Acc \uparrow	Mean (SD) \downarrow	Mean (SD) \downarrow
Easy	Idefics2	0.35	3.09 (± 3.1)	5.72 (± 3.1)
	LLaVA	0.41	2.62 (± 2.6)	3.98 (± 2.5)
	GPT-4o	0.45	3.50 (± 3.5)	2.18 (± 1.3)
	Humans	1.00	.00 (± 0)	.81 (± 0.5)
	Humans	1.00	.00 (± 0)	.81 (± 0.5)
Medium	Idefics2	0.37	2.32 (± 2.3)	5.30 (± 2.3)
	LLaVA	0.42	2.19 (± 2.2)	4.54 (± 3.0)
	GPT-4o	0.42	3.05 (± 3.1)	2.86 (± 2.0)
	Humans	0.50	2.24 (± 2.2)	2.28 (± 1.3)
	Humans	0.50	2.24 (± 2.2)	2.28 (± 1.3)
Hard	Idefics2	0.22	3.82 (± 3.8)	6.00 (± 3.2)
	LLaVA	0.25	4.30 (± 4.3)	5.77 (± 3.3)
	GPT-4o	0.16	4.46 (± 4.5)	5.30 (± 3.7)
	Humans	0.00	5.62 (± 5.6)	6.25 (± 3.3)
	Humans	0.00	5.62 (± 5.6)	6.25 (± 3.3)

Table 4: Model performances across subsets with increasing levels of complexity (*easy*, *medium*, *hard*).

6.2 Impact of Task Difficulty

To better pinpoint the strengths and weaknesses of the different models, we categorise the BW-R data into difficulty levels according to human performance⁵. We use lower human performance as an indirect proxy for the complexity of the example, and define three levels as follows: 1) **Easy**: both human annotators for the same entry correctly predicted the block (100% source accuracy) or their target position was within 1 block unit distance away from the true location; 2) **Medium**: at least one out of the two annotators correctly predicted the block (50% source accuracy) or their target position was below the human’s mean distance away from the location (between 1 and 3.22 units away); 3) **Hard**: neither of the annotators correctly predicted the block (0% source accuracy) or their target position was further than the human’s mean distance (above 3.22 units).

The results of evaluating VLMs on these difficulty subsets (Table 4) indicate that models and humans do not struggle on the same entries. Models show a more uniform distribution of performance, with improvements over humans in the most difficult subset but significantly underperforming with easier references. This behaviour seems to follow the trend with the average word number in the subsets⁶, suggesting that easy examples (most words) may employ longer, more complex terms to describe references, whereas hard examples (fewest words) may have easier terms for models, but are too underspecified for humans.

⁵We provide this analysis according to GPT-4o performance in Appendix D.

⁶See Appendix D.

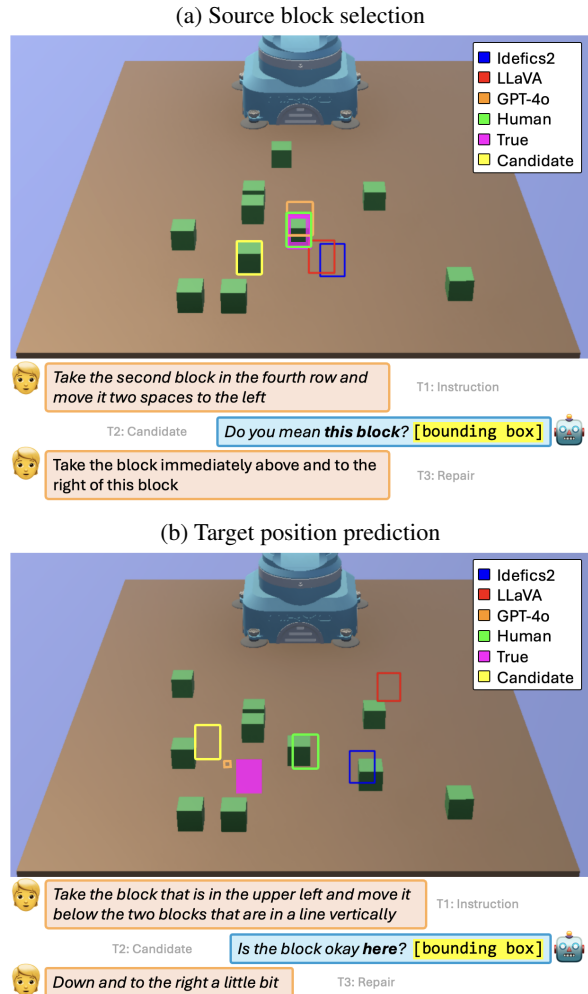


Figure 3: Two medium-difficulty dialogues with the bounding boxes predicted by the VLMs and humans.

6.3 Qualitative Analysis

We further analyse 50 BW-R dialogues to understand the proficiency differences of VLMs with that of human participants (see Figure 3 for some examples). We find that humans and GPT-4o tend to have similar predictions when processing candidate responses and their repairs. Models are usually able to process simple spatial repairs (i.e., left, right, above) (Chiyah-Garcia et al., 2023), however, they particularly struggle with more abstract concepts (i.e., rows, columns) (Ilinykh et al., 2022). In that regard, both Idefics2 and LLaVA exhibit comprehension of the simpler referring expressions, but have a challenging time with longer or more complex sequences. Idefics2 also seems prone to over-correct (e.g., *right* pushes predictions to the absolute right). We also notice that VLMs sometimes predict bounding boxes that do not align with blocks in the source prediction task, which is un-

expected given their pre-training (and fine-tuning) with referring expression objectives. For the target prediction task, VLMs significantly struggle and do not process repairs. GPT-4o is the only exception, which shows a more fine-grained understanding of the task yet falls behind on more difficult dialogues.

7 Conclusion

This paper explores VLM’s capabilities of handling repairs with a new dataset, BLOCKWORLD-REPAIRS. We collect 795 collaborative dialogues in a highly ambiguous tabletop manipulation environment that requires complex multi-modal task instructions as well as repairs. In these scenarios, referential ambiguities are common and thus being able to process TPRs becomes essential to complete the task. We validate these dialogues through an in-person human study where participants attempt to solve the task, providing a baseline for model performance.

We then show that VLMs struggle to process TPRs alongside other dialogues out of the box and propose improving their fine-tuning regime with different loss criteria. Our results indicate that VLMs benefit from partially masking the input tokens when learning to process TPRs along with the task, resulting in more generalisable models across single-turn instructions and TPRs for one of the tasks. However, this needs to be balanced as masking too many tokens hurts models when data is smaller due to a lack of training signal.

We finish with an error analysis comparing these models and human participants, finding stark differences in how they process repairs. Surprisingly, models struggle with the dialogues that human participants found easiest, highlighting a large gap in their ability to process TPRs, despite these being fundamental to effective dialogue coordination. We believe that future research should focus on two major improvements for VLMs: 1) design training objectives that facilitate *learning from interaction*, including TPRs; 2) equip models with more object-centric visual representations that facilitate visual grounding tasks (Parekh et al., 2024) which are of paramount importance for situated collaborative tasks.

8 Limitations

One of the limitations of the paper is that we collected the data for the dataset as an annotation task. This allowed us to divide the task into instances

that participants could annotate separately (instructions or corrections), which makes the annotation conceptually easier to carry out and mimics how the initial Block World instructions were collected. This allowed us to focus on the dialogue phenomena that we care about, corrections in situated contexts, at the cost of the overall collaborative nature of the task. Ideally, it would be better to design a Wizard-of-Oz or human-human experiment to collect the annotations during the interaction, at a much greater time and cost.

We also acknowledge that our environment is visually simple compared to images that can be found on the Internet which are commonly used for pre-training VLMs. However, we don’t consider this as a weakness but as a strength: our BLOCKWORLD-REPAIRS is built using the Unity game engine and allows us to test models’ ability to process TPRs ignoring other confounding variables such as processing complex visual scenes.

Another limitation of our work is that we do not attempt to model the entire spectrum of a robotic manipulation task (e.g., Octo Model Team et al., 2024). Instead, we follow previous work in Embodied AI (e.g., ALFRED (Shridhar et al., 2020), Simbot Arena (Gao et al., 2023), *inter alia*) which casts object manipulation as API call generation of an action label combined with a bounding box around the object to be manipulated. Although we don’t generate actions, the prediction tasks that we explore in BW-R involve generating the coordinates of a bounding box that can be used for picking up a specific block or putting it down onto a specific box on the board.

To the best of our knowledge, we have used state-of-the-art VLMs that are open-weight and can be easily used for fine-tuning and inference. However, due to the fast pace of research in this space, it is possible that by the time reviewers read this manuscript, a new open-weight VLM will be announced. For this reason, we decided to also report the performance of a frontier model such as GPT-4o which has superior performance to many text-only as well as visual+language tasks.

9 Potential Risks

Embodied AI and human-robot collaboration offer exciting possibilities, but also introduce a new range of potential risks. Embodied AI systems must be designed and implemented with rigorous safety protocols to minimise the risk of accidents

for human workers. This includes robust sensor systems, clear communication protocols, and fail-safe mechanisms to prevent harm in case of malfunctions. In this paper, we do not use an actual robot but we simulate the use case of a robot arm operating in a factory and collaborating with a user to complete a pick&place task. In this scenario, we argue that it is essential that the robot is able to process clarification exchanges, and TPRs are an important mechanism to ensure that models can robustly correct mistakes when they arise.

Acknowledgements

Chiyah-Garcia's PhD is funded under the EPSRC iCase with Siemens (EP/T517471/1). We thank the anonymous reviewers for their insightful comments that helped improve this paper. Finally, the authors acknowledge the use of the Heriot-Watt University high-performance computing facility (DMOG) and associated support services in the completion of this work.

References

- Angus Addlesee and Arash Eshghi. 2021. [Incremental graph-based semantics and reasoning for conversational AI](#). In *Proceedings of the Reasoning and Interaction Conference (ReInAct 2021)*, pages 1–7, Gothenburg, Sweden. Association for Computational Linguistics.
- Angus Addlesee, Oliver Lemon, and Arash Eshghi. 2024. [Clarifying completions: Evaluating how LLMs respond to incomplete questions](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 3242–3249, Torino, Italia. ELRA and ICCL.
- Shubham Agarwal, Trung Bui, Joon-Young Lee, Ioannis Konstas, and Verena Rieser. 2020. [History for visual dialog: Do we really need it?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8182–8197, Online. Association for Computational Linguistics.
- Mohammad Aliannejadi, Julia Kiseleva, Aleksandr Chuklin, Jeff Dalton, and Mikhail Burtsev. 2021. [Building and evaluating open-domain dialogue corpora with clarifying questions](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4473–4484, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- Vevake Balaraman, Arash Eshghi, Ioannis Konstas, and Ioannis Papaioannou. 2023. [No that's not what I meant: Handling third position repair in conversational question answering](#). In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 562–571, Prague, Czechia. Association for Computational Linguistics.
- Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. 2009. [Measurement Instruments for the Anthropomorphism, Animacy, Likeability, Perceived Intelligence, and Perceived Safety of Robots](#). *International Journal of Social Robotics*, 1(1):71–81.
- Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828.
- Luciana Benotti and Patrick Blackburn. 2021. [A recipe for annotating grounded clarifications](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4065–4077, Online. Association for Computational Linguistics.
- Yonatan Bisk, Deniz Yuret, and Daniel Marcu. 2016. [Natural language communication with robots](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 751–761, San Diego, California. Association for Computational Linguistics.
- Okko Buß and David Schlangen. 2011. Dium : An incremental dialogue manager that can produce self-corrections. In *Proceedings of SemDial 2011 (Los Angeles)*, Los Angeles, CA, pages 47–54.
- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. 2023. [Shikra: Unleashing multimodal llm's referential dialogue magic](#). *arXiv preprint arXiv:2306.15195*.
- Francisco Javier Chiyah Garcia, José Lopes, Xingkun Liu, and Helen Hastie. 2020. [CRWIZ: A framework for crowdsourcing real-time Wizard-of-Oz dialogues](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 288–297, Marseille, France. European Language Resources Association.
- Javier Chiyah-Garcia, Alessandro Suglia, Arash Eshghi, and Helen Hastie. 2023. ['what are you referring to?' evaluating the ability of multi-modal dialogue models to process clarificational exchanges](#). In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 175–182, Prague, Czechia. Association for Computational Linguistics.

- Herbert H. Clark. 1996. *Using Language*. Cambridge University Press.
- Herbert H Clark and Susan E Brennan. 1991. Grounding in communication. In *Perspectives on socially shared cognition.*, pages 127–149. American Psychological Association.
- M. Colman and P. G. T. Healey. 2011. The distribution of repair in dialogue. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society*, pages 1563–1568, Boston, MA.
- Amanda Cercas Curry, Ioannis Papaioannou, Alessandro Suglia, Shubham Agarwal, Igor Shalymov, Xu Xinnuo, Ondrej Dusek, Arash Eshghi, Ioannis Konstas, Verena Rieser, and Oliver Lemon. 2018. Alana v2: Entertaining and informative open-domain social dialogue using ontologies and entity linking. In *1st Proceedings of Alexa Prize (Alexa Prize 2018)*.
- Soham Dan, Michael Zhou, and Dan Roth. 2021. **Generalization in instruction following systems**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 976–981, Online. Association for Computational Linguistics.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José MF Moura, Devi Parikh, and Dhruv Batra. 2017. Visual dialog. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 326–335.
- Emanuele De Pellegrin and Ronald PA Petrick. 2024. Planning domain simulation: An interactive system for plan visualisation. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 34, pages 133–141.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. **Qlora: Efficient finetuning of quantized llms**. In *Advances in Neural Information Processing Systems*, volume 36, pages 10088–10115. Curran Associates, Inc.
- Mark Dingemanse, Seán G. Roberts, Julija Baranova, Joe Blythe, Paul Drew, Simeon Floyd, Rosa S. Gísladóttir, Kobin H. Kendrick, Stephen C. Levinson, Elizabeth Manrique, Giovanni Rossi, and N. J. Enfield. 2015. **Universal principles in the repair of communication problems**. *PLOS ONE*, 10(9):1–15.
- Simon Dobnik, John D. Kelleher, and Christine Howes. 2020. **Local Alignment of Frame of Reference Assignment in English and Swedish Dialogue**, page 251–267. Springer International Publishing.
- N. J. Enfield, Mark Dingemanse, Julija Baranova, Joe Blythe, Penelope Brown, Tyko Dirksmeyer, Paul Drew, Simeon Floyd, Sonja Gipper, Rósa Gísladóttir, Gertie Hoymann, Kobin H. Kendrick, Stephen C. Levinson, Lilla Magyari, Elizabeth Manrique, Giovanni Rossi, Lila San Roque, and Francisco Torreira. 2013. *Huh? What? –A first survey in twenty-one languages*, pages 343–380. Cambridge University Press, Cambridge.
- Arash Eshghi and Arash Ashrafzadeh. 2023. **Learning to generate and correct repair language in real-time**. In *Proceedings of SemDial’23 (MariLogue)*.
- Arash Eshghi, Christine Howes, and Eleni Gregoromichelaki. 2022. **Action coordination and learning in dialogue**. In J-P Bernardy, R Blanck, S Chatzkyriakidis, S Lappin, and A Maskharashvili, editors, *Probabilistic Approaches to Linguistic Theory*, Probabilistic Approaches to Linguistic Theory, pages 361–422. CSLI Publications.
- Qiaozhi Gao, Govind Thattai, Suhaila Shakiah, Xiaofeng Gao, Shreyas Pansare, Vasu Sharma, Gaurav Sukhatme, Hangjie Shi, Bofei Yang, Desheng Zhang, Lucy Hu, Karthika Arumugam, Shui Hu, Matthew Wen, Dinakar Guthy, Shunan Chung, Rohan Khanna, Osman Ipek, Leslie Ball, Kate Bland, Heather Rucker, Michael Johnston, Reza Ghanadan, Dilek Hakkani-Tur, and Prem Natarajan. 2023. **Alexa arena: A user-centric interactive platform for embodied ai**. In *Advances in Neural Information Processing Systems*, volume 36, pages 19170–19194. Curran Associates, Inc.
- Xiaofeng Gao, Qiaozhi Gao, Ran Gong, Kaixiang Lin, Govind Thattai, and Gaurav S. Sukhatme. 2022. **Dialfred: Dialogue-enabled agents for embodied instruction following**. *IEEE Robotics and Automation Letters*, 7(4):10049–10056.
- Felix Gervits, Gordon Briggs, Antonio Roque, Genki A. Kadamatsu, Dean Thurston, Matthias Scheutz, and Matthew Marge. 2021. **Decision-theoretic question generation for situated reference resolution: An empirical study and computational model**. In *Proceedings of the 2021 International Conference on Multimodal Interaction, ICMI ’21*, page 150–158, New York, NY, USA. Association for Computing Machinery.
- Charles Goodwin. 1981. *Conversational organization: Interaction between speakers and hearers*. Academic Press, New York.
- Patrick G. T. Healey, Gregory J. Mills, Arash Eshghi, and Christine Howes. 2018. **Running Repairs: Coordinating Meaning in Dialogue**. *Topics in Cognitive Science (topiCS)*, 10(2).
- Julian Hough. 2015. *Modelling Incremental Self-Repair Processing in Dialogue*. Ph.D. thesis, Queen Mary University of London.
- Julian Hough and Matthew Purver. 2012. **Processing self-repairs in an incremental type-theoretic dialogue system**. In *Proceedings of the 16th SemDial Workshop on the Semantics and Pragmatics of Dialogue (SeineDial)*, pages 136–144, Paris, France.
- Julian Hough and David Schlangen. 2015. **Recurrent neural networks for incremental disfluency detection**.

- In *INTER_SPEECH 2015, 16th Annual Conference of the International Speech Communication Association, Dresden, Germany, September 6-10, 2015*, pages 849–853.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Nikolai Ilinykh and Simon Dobnik. 2021. [What does a language-and-vision transformer see: The impact of semantic information on visual representations](#). *Frontiers in Artificial Intelligence*, 4.
- Nikolai Ilinykh and Simon Dobnik. 2022. [Attention as grounding: Exploring textual and cross-modal attention on entities and relations in language-and-vision transformer](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 4062–4073, Dublin, Ireland. Association for Computational Linguistics.
- Nikolai Ilinykh, Yasmeen Emampoor, and Simon Dobnik. 2022. [Look and answer the question: On the role of vision in embodied question answering](#). In *Proceedings of the 15th International Conference on Natural Language Generation*, pages 236–245, Waterville, Maine, USA and virtual meeting. Association for Computational Linguistics.
- Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. 2023. [VIMA: Robot manipulation with multimodal prompts](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 14975–15022. PMLR.
- Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. 2017. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*.
- Satwik Kottur, Seungwhan Moon, Alborz Geramifard, and Babak Damavandi. 2021. [SIMMC 2.0: A task-oriented dialog dataset for immersive multimodal conversations](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4903–4912, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Staffan Larsson. 2018. [Grounding as a Side-Effect of Grounding](#). *Topics in Cognitive Science (topiCS)*, 10(2).
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024. [What matters when building vision-language models?](#) *Preprint*, arXiv:2405.02246.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023. Improved baselines with visual instruction tuning.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Brielen Madureira and David Schlangen. 2023. [Instruction clarification requests in multimodal collaborative dialogue games: Tasks, and an analysis of the Co-Draw dataset](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2303–2319, Dubrovnik, Croatia. Association for Computational Linguistics.
- Brielen Madureira and David Schlangen. 2024. [Taking action towards graceful interaction: The effects of performing actions on modelling policies for instruction clarification requests](#). In *Proceedings of the Third Workshop on Understanding Implicit and Underspecified Language*, pages 1–21, Malta. Association for Computational Linguistics.
- Nikhil Mehta and Dan Goldwasser. 2019. [Improving natural language interaction with robots using advice](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1962–1967, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gregory J. Mills. 2007. *Semantic co-ordination in dialogue: the role of direct interaction*. Ph.D. thesis, Queen Mary University of London.
- Gregory J. Mills. 2014. [Dialogue in joint activity: Complementarity, convergence and conventionalization](#). *New Ideas in Psychology*, 32:158–173.
- Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Lawrence Yunliang Chen, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. 2024. Octo: An open-source generalist robot policy. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands.
- OpenAI. 2024. GPT-4o. <https://openai.com/index/hello-gpt-4o/>. Accessed: 2024-06-12.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al.

2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Aishwarya Padmakumar, Jesse Thomason, Ayush Shrivastava, Patrick Lange, Anjali Narayan-Chen, Spananda Gella, Robinson Piramuthu, Gokhan Tur, and Dilek Hakkani-Tur. 2022. **Teach: Task-driven embodied agents that chat**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(2):2017–2025.
- Georgios Pantazopoulos, Alessandro Suglia, Oliver Lemon, and Arash Eshghi. 2024. **Lost in space: Probing fine-grained spatial understanding in vision and language resamplers**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 540–549, Mexico City, Mexico. Association for Computational Linguistics.
- Amit Parekh, Nikolas Vitsakis, Alessandro Suglia, and Ioannis Konstas. 2024. **Investigating the role of instruction variety and task difficulty in robotic manipulation tasks**. *Preprint*, arXiv:2407.03967.
- Matthew Purver. 2004. *The Theory and Use of Clarification Requests in Dialogue*. Ph.D. thesis, University of London.
- Matthew Purver and Jonathan Ginzburg. 2004. Clarifying noun phrase semantics. *Journal of Semantics*, 21(3):283–339.
- Matthew Purver, Julian Hough, and Christine Howes. 2018. Computational models of miscommunication phenomena. In Patrick G. T. Healey, Jan de Ruiter, and Gregory J. Mills, editors, *Topics in Cognitive Science (topiCS)*, volume 10.
- Verena Rieser and Oliver Lemon. 2006. **Using machine learning to explore human multimodal clarification strategies**. In *Proceedings of the COLING/ACL 2006 Main Conference Poster Sessions*, pages 659–666, Sydney, Australia. Association for Computational Linguistics.
- Verena Rieser and Johanna Moore. 2005. Implications for generating clarification requests in task-oriented dialogues. In *Proceedings of the 43rd Annual Meeting of the ACL*, pages 239–246, Ann Arbor. Association for Computational Linguistics.
- Kepa Rodríguez and David Schlangen. 2004. Form, intonation and function of clarification requests in German task-oriented spoken dialogues. In *Proceedings of the 8th Workshop on the Semantics and Pragmatics of Dialogue (SEMDIAL)*, Barcelona, Spain.
- Emmanuelle Salin, Badreddine Farah, Stéphane Ayache, and Benoit Favre. 2022. **Are vision-language transformers learning multimodal representations? a probing perspective**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):11248–11257.
- Ruben San-Segundo, Juan M. Montero, J. Ferreiros, R. Córdoba, and José M. Pardo. 2001. Designing confirmation mechanisms and error recover techniques in a railway information system for Spanish. In *Proceedings of the 2nd SIGdial Workshop on Discourse and Dialogue*, pages 136–139, Aalborg, Denmark. Association for Computational Linguistics.
- Antonios Saravanos, Stavros Zervoudakis, Dongnanzi Zheng, Neil Stott, Bohdan Hawryluk, and Donatella Delfino. 2021. *The Hidden Cost of Using Amazon Mechanical Turk for Research*, page 147–164. Springer International Publishing.
- E.A. Schegloff. 1992. Repair after next turn: The last structurally provided defense of intersubjectivity in conversation. *American Journal of Sociology*, pages 1295–1345.
- E.A. Schegloff, Gail Jefferson, and Harvey Sacks. 1977. The preference for self-correction in the organization of repair in conversation. *Language*, 53(2):361–382.
- David Schlangen. 2023. **On general language understanding**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8818–8825, Singapore. Association for Computational Linguistics.
- Maximilian Seitzer, Max Horn, Andrii Zadaianchuk, Dominik Zietlow, Tianjun Xiao, Carl-Johann Simon-Gabriel, Tong He, Zheng Zhang, Bernhard Schölkopf, Thomas Brox, and Francesco Locatello. 2023. **Bridging the gap to real-world object-centric learning**. In *Proceedings of the Eleventh International Conference on Learning Representations*.
- Igor Shalymov, Arash Eshghi, and Oliver Lemon. 2017. Challenging neural dialogue models with natural data: Memory networks fail on incremental phenomena. In *Proceedings of the 21st Workshop on the Semantics and Pragmatics of Dialogue (SemDial 2017 - SaarDial)*, Barcelona.
- Zhengxiang Shi, Yue Feng, and Aldo Lipani. 2022. **Learning to execute actions or ask clarification questions**. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 2060–2070, Seattle, United States. Association for Computational Linguistics.
- Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. 2020. **Alfred: A benchmark for interpreting grounded instructions for everyday tasks**. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10737–10746.
- Alessandro Suglia, Ioannis Konstas, and Oliver Lemon. 2024. Visually grounded language learning: a review of language games, datasets, tasks, and models. *Journal of Artificial Intelligence Research*, 79:173–239.
- Hao Tan and Mohit Bansal. 2018. **Source-target inference models for spatial instruction understanding**.

Proceedings of the AAAI Conference on Artificial Intelligence, 32(1).

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. [Opa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#). In *International Conference on Machine Learning*, pages 23318–23340. PMLR.

Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. 2016. [Modeling context in referring expressions](#). In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pages 69–85. Springer.

Yi Zhu, Yue Weng, Fengda Zhu, Xiaodan Liang, Qixiang Ye, Yutong Lu, and Jianbin Jiao. 2021. [Self-motivated communication agent for real-world vision-dialog navigation](#). In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1574–1583.

A Related Datasets

We compare BLOCKWORLD-REPAIRS with related works in Table 5. Our work builds on a robotic manipulation task involving common pick&place actions and rich referring expressions in a situated dialogue. Similar to our benchmark, both TEACH (Padmakumar et al., 2022) and DialFRED (Gao et al., 2022) extend the popular ALFRED (Shridhar et al., 2020) to include dialogues with clarificational exchanges. From a high-level ALFRED directive (e.g., “prepare coffee”), these works crowd-source possible questions for low-level actions that may be underspecified in the original task (e.g., “where is the coffee mug?”). Thus, they often operate on the last level of Clark’s (1996) joint action ladder, *action and consideration*, where the goal is already clear and lack of knowledge (underspecified location of objects or actions available) represents the primary ambiguity (e.g., VIMA-Bench; Jiang et al., 2023)). These works rarely feature bi-directional information flow (*human* \leftrightarrow *agent*) in their dialogues, unlike BW-R which ensures dialogue dependency with partial robot movements and candidate responses. DialFRED, Alexa Arena (Gao et al., 2023) and CoDraw iCRs (Madureira and Schlangen, 2023) contain questions for underspecified actions/objects (e.g., “Which spoon should I pick up?”) but, in these situations, the ambiguity is reduced to a minimum so that candidates are easily distinguishable with strong visual representations (e.g., “yellow or blue spoon?” or “left or right?”).

By contrast, SIMMC 2.0 (Kottur et al., 2021) includes a large number of visually ambiguous objects (e.g., 5 identical red t-shirts in view) within long dialogues that reference multiple items (averaging 4.5 ± 2.4 unique objects per dialogue). This makes it rich in cross-modal coordination phenomena, with terse referring expressions that mix spatial, historical and visual cues (e.g., “Pick the red shirt. Which one? The one on the right wardrobe, above the blue jumper”). Similarly, Cups (Dobnik et al., 2020) aligns with our work, as it also explores coordination in highly ambiguous environments where many objects share the same shape/colour but are in different positions, requiring complex referring expressions. However, Cups focuses on Frame of Reference coordination rather than manipulation or resolving coreferences. BW-R is aimed at exploring this coordination as part of a collaborative human-robot tabletop manipulation task.

B BLOCKWORLD-REPAIRS

This section provides additional details about the dataset.

B.1 AMT Collection

Each task or Human Intelligence Task (HIT) in AMT consisted of 4 dialogues, with an initial easier, training dialogue that workers were free to experiment with. AMT workers could repeat the task for as long as there were block configurations they had not seen before. The instructions are in Figure 4, with the user interface of the task itself in Figure 5. We introduced many attention checks, particularly with the repairs, where workers were asked to select whether the robot had performed the correct action (Yes/No) before entering free-text repairs. If their answer did not match the expected answer too many times, they would get a warning and be limited from submitting any subsequent HITs. We used the CRWIZ framework (Chiyah Garcia et al., 2020) to control many of these details, including the robot’s non-verbal actions (i.e., movements, pointing), effectively running a simulation or sequence of events different to each worker in AMT.

Human-Robot Collaboration Task

In this task, you will be interacting with a robot arm that will help you arrange blocks. You will have to give commands to the robot, which will pick up and move blocks. However, this robot arm is still in training, so it may not understand well which block you want to move or where to. The robot arm will try to understand you but it is important to try to give clear and precise commands.

HIT Information

- Should last around 12 minutes (\$0.15/minute or \$9/hr).
- Requires a minimum level of English, but you won’t be rejected based on English skills alone.
- There will be a short training exercise first and successfully completing this will unlock two other similar HITs (\$1.8 each).
- Submissions with **incoherent or repeated text will be prevented from further HITs.**
- If you run into any issues, you can get in touch with us at **Hidden until HIT is accepted** quoting the following number: **USER03GLVC-ASSIGNMENTJJ61NCPUZF**

Instructions

You need to arrange blocks to achieve the target configuration, given on the right. For that, you need to find the block that needs to be moved and tell the robot where to move it, referring to the other blocks on the table. For example, you could say something like “Move the rightmost block between the blocks at the leftmost side of the table”.

Below you will see an example conversation with the robot before starting. Please aim for clear robot instructions every time. Refer to your right and your left.

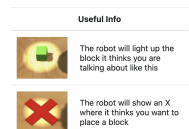


Figure 4: Instructions that AMT workers saw in the landing page of the HIT.

At the end of the task, workers gave free-form comments and feedback in scales about how easy the task was and the Godspeed Questionnaire IV for Perceived Intelligence (Bartneck et al., 2009) to ensure that the task was not too difficult and that the agent was not too unintelligent (due to the continuous “misunderstandings”).

We limited HITs to 30 minutes and the mean elapsed time for all the HITs was 12 minutes (SD=4.5 minutes), or 4 minutes per dialogue. The payment was \$1.8 per HIT, so workers were paid, on average, a rate of \$9.0/hr. This pay is above the Federal minimum wage in the US (\$7.25/hr or \$0.12/min) at the time of the data collection from

Related Work	Manipulation	Ref Exp	Dialogue	Repairs	Miscommunication
VIMA-Bench (Jiang et al., 2023)	✓	R	✗	✗	-
ALFRED (Shridhar et al., 2020)	✓	R	✗	✗	-
TEACH (Padmakumar et al., 2022)	✓	V R	✓	✓	(4) Action
DialFRED (Gao et al., 2022)	✓	V R	✓	✓	(4) Action
Alexa Arena (Gao et al., 2023)	✓	V R	✓	✗	(4) Action
SIMMC 2.0 (Kottur et al., 2021)	✗	V R D	✓	✓	(3) Understanding
CoDraw iCR (Madureira and Schlagen, 2023)	✗	V R	✓	✓	(4) Action
Cups (Dobnik et al., 2020)	✗	V R D	✓	✓	(3) Understanding
Block World (Bisk et al., 2016)	✓	V R	✗	✗	-
BLOCKWORLD-REPAIRS (Ours)	✓	V R	✓	✓	(3) Understanding

Table 5: Comparison of related works to our dataset BW-R. Columns describe whether the datasets have: (1 **Manipulation**) robotic manipulation as a task (i.e., pick&place); (2 **Ref Exp**) referring expression type following (Chiyah-Garcia et al., 2023), **R** for relational expressions (e.g., “the left one”), **V** for references to visual properties (e.g., “the green one”) or **D** when referring to the dialogue history (e.g., “the one I mentioned”); (3 **Dialogue**) dialogue with multiple turns for the same action; (4 **Repairs**) explicit repairs in the data (i.e., clarifications or corrections); and (5 **Miscommunication**) the miscommunication level based on Clark’s (1996) joint action ladder.

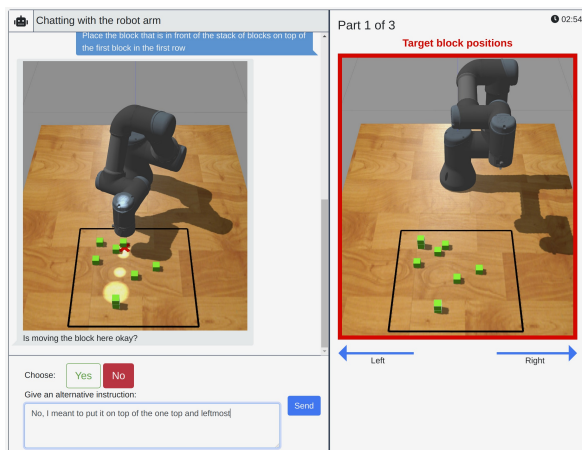


Figure 5: User interface for the data collection. The images of the robot and blocks would change depending on the task or the instructions given. The left has the chat and a text box to send messages whereas the right shows how the final placement should look in the end.

May to September 2021. 178 unique AMT workers submitted HITs for our task (58 female, 106 male, and 12 rest) and we limited our task to be only available in the United States. We manually checked data subsets for offensive, toxic language, or personal information. Participants signed a virtual consent with details about the data collection, how the data would be used and how they could withdraw if they wished to do so. We did not collect any personal or sensitive information. Ethics approval for this collection was provided by our institution’s ethics committee.

B.2 Human Baselines

Participants interacted with the interface shown in Figure 6, where the chat was already filled by a randomly selected dialogue. Participants had to read the instructions and repairs and provide their best interpretation of which block the instructions referred to and where to place it. Once satisfied with their choice, they would confirm and proceed to the next dialogue until the session ended. This interface uses the Unity engine (v2022.3)⁷ and (De Pellegrin and Petrick, 2024) to generate environments based on the original Block World configurations.

The study ran from January to April 2024 and it lasted around 40 minutes. We recruited a total of 22 participants (5 female, 17 male, most common age range was 23-29 and above undergraduate education level) through internal institution mailing lists. 12 were native English speakers and most participants self-reported a near bilingual level of English proficiency. On average, participants annotated 50 dialogues each and received £10 for taking part in the study (a rate of £15/hour). This pay is above the minimum wage in the UK at the time of the study (£11.44/hour). Participants signed consent forms with details about how the data would be used and published as well as the procedure to withdraw if they wished so. We did not collect any personal or sensitive information. Ethics approval for this study was provided by our institution’s ethics committee.

⁷<https://unity.com/>

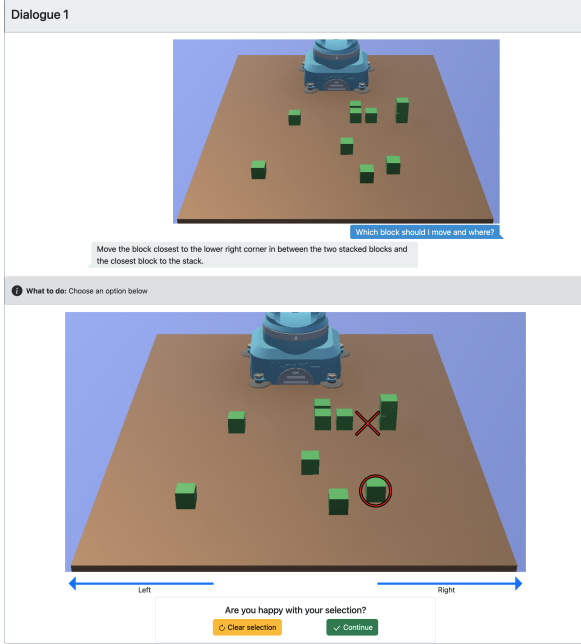


Figure 6: User interface for the human validation. Participants had to click on the large image at the bottom where they thought that the instructions referred to, both for the source block and the target position.

We collected 991 after discarding 61 low-quality annotations. The annotations are split as follows: 343 for instructions (34.6%), 327 for source block TPRs (33.0%) and 321 for target position TPRs (32.4%).

B.3 Dataset Statistics

Table 6 provides additional statistics for the BLOCKWORLD-REPAIRS dataset.

Attribute	Value
Instructions	795 (38.6%)
Source TPRs	629 (30.5%)
Target TPRs	635 (30.8%)
Total Entries	2059

Table 6: Summary of collected entries in BW-R.

For our experiments, we join the instructions from the original Block World with the collected entries to obtain a larger dataset, see Table 7. The initial instructions in our dialogues are the same as in (Bisk et al., 2016), as they contain both which block to move and where to move it. We also use the same block arrangements as in the original dataset, respecting train/test splits. Block configurations are unique across train-test splits, so models are tested on unseen world arrangements.

C Experimental Setup

C.1 Prompts

Each task has its own slightly different prompt, see Table 8. We also provide example full prompts used during zero-shot and fine-tuning in Table 9.

When parsing model generations, we try to parse a bounding box with 4 decimal numbers. However, models sometimes generate invalid boxes or additional tokens, particularly in zero-shot settings. In these cases, less than 3%, we count this prediction as ‘failed’ and use the output of the random baseline to calculate other metrics. Idefics2 and GPT-4o are more prone to generate additional tokens.

C.2 Model Fine-Tuning

For our experiments, we use the default hyperparameters for each model (Idefics2 and LLaVA). We train for 1 epoch as we did not observe benefits from training for longer. Images had the size of 1024x576 pixels, and we used the model’s default image processing pipelines. Training batches had both instruction and repair examples, sampled at random. We did not use custom sampling to balance the type of entries in batches and leave this for future work.

We fine-tune the VLMs with the recommended parameter-efficient methods: LoRA (Hu et al., 2021) for LLaVA and QLoRA (Dettmers et al., 2023) for Idefics2. We used 2x NVIDIA A40 (40GB) at most and set the maximum sequence length to 2048. Each experiment (training and testing) takes around 1 hour to fully complete. The final table of experiments takes approximately 30 GPU hours to run.

D Additional Error Analysis

As mentioned in Section 6, it seems that easier referring expressions or TPRs have a lower number of words on average than harder ones (see Table 10).

We also provide in Table 11 the results of selecting GPT-4o performance as the proxy for example difficulty. Since we only have 1 prediction value for each BW-R entry and we cannot do the average as with the human annotations, we instead order the data by distance (source and target separately) and divide it into three equal sizes at the 33rd and 66th percentiles: 1) **Easy** category contains the top 33% of entries (shortest distance to source/target); 2) **Medium** contains entries that are between the 33% and 66% distances; and 3) **Hard** is reserved for the worst performing 33% entries.

Split	Train	Dev	Test	Total
Block World (Bisk et al., 2016)	2403 (66.5%)	360 (100.0%)	810 (48.8%)	3573
BLOCKWORLD-REPAIRS	1210 (33.5%)	0 (0.0%)	849 (51.2%)	2059
Experiment Data	3613	360	1659	5632

Table 7: Data used in this paper. We use the train/test sets during our experiments of Sections 4, 5 and 6.

Task	Task Instruction
Source Block Selection	Answer only with the bounding box of the block mentioned. The bounding box consists of 4 values between 0 and 1. Here is an example: [0.000, 0.123, 0.075, 0.204]
Target Position Prediction	Answer only with the final bounding box location mentioned. The bounding box consists of 4 values between 0 and 1. Here is an example: [0.000, 0.123, 0.075, 0.204]

Table 8: Task instructions used to prompt models.

We find that Idefics2 and LLaVA have low performance on all levels, whilst GPT-4o only surpasses humans on easy dialogues (not surprising), but quickly deteriorates to distances worse than random on hard dialogues. Interestingly, humans perform well across all difficulties except the hard target predictions. We also observe that the mean word trends get stronger (Table 10), although we leave the analysis to prove a correlation for future work. These results further show that models do not process TPRs as consistently as humans.

E Additional Experimental Results

We completed a comprehensive set of ablations testing for different ways of masking (as defined in the main paper), different VLMs, and different data regimes as well. Table 12 presents the full set of results.

F Other Notes

Object detection and predictions Splitting the task into object detection and then prediction out of a list of candidates did not show any advantages in our initial tests, possibly because these models are already pretrained with object detection objectives (Liu et al., 2023; Laurençon et al., 2024). We al-

ready see better performance with source block prediction than with target location prediction. Thus, we unified both tasks into a single common task format (bounding box prediction) for easier comparison. Additionally, we would like to point out that target prediction would be a harder task because it requires the model to derive a bounding box for an empty board location, which models particularly struggle as shown in Section 6.

Repairs facilitate the task Repairs are fundamental capabilities in human conversations, which our human study further shows with the increasing performance with TPRs (see Table 3). TPRs have the dual role of narrowing down the candidate pool (i.e., “not this block”) and providing new information to uniquely identify the correct referent (i.e., “just left of there”). Our paper shows that current out-of-the-box models do not have the capabilities to exploit repairs to facilitate the task (unlike humans), but we should try to distil these skills with alternative training regimes that facilitate learning from interaction data. Handling TPRs is essential for real-world human-robot collaboration in novel scenarios and this paper explores this gap between models and humans.

Model	Sample
Idefics2 8B	<p>User: <image> Answer only with the bounding box of the block mentioned. The bounding box consists of 4 values between 0 and 1. Here is an example: [0.000, 0.123, 0.075, 0.204]. take the block in the very middle and move it slightl left and down.<end_of_utterance></p> <p>Assistant: location bounding box [0.329, 0.724, 0.339, 0.741] I recall you wanting the block here, is this correct?<end_of_utterance></p> <p>User: move it to the right a little<end_of_utterance></p> <p>Assistant:</p> <p>Model generation: location bounding box [0.394, 0.585, 0.436, 0.682]</p>
LLaVA v1.5 7B	<p><s> A chat between a curious user and an artificial intelligence assistant. The assistant gives helpful, detailed, and polite answers to the user’s questions. Answer only with the bounding box of the block mentioned. The bounding box consists of 4 values between 0 and 1. Here is an example: [0.000, 0.123, 0.075, 0.204].</p> <p>USER: <image> There is a block in the center of the board. Take that block and move it to the left in front of the other block.</p> <p>ASSISTANT: block bounding box block bounding box [0.391, 0.635, 0.434, 0.735] Can you confirm it is this block?</s></p> <p>USER: The block that is a row closer to you and to the right.</p> <p>ASSISTANT:</p> <p>Model generation: block bounding box [0.391, 0.635, 0.434, 0.735]</p>
GPT-4o	<p>System: You are a helpful assistant that responds with only with the bounding box of the location mentioned. The bounding box consists of 4 values between 0 and 1. Here is an example: [0.000, 0.123, 0.075, 0.204].</p> <p>User: <image> Take the block that is furthest to the right in the fourth row and place it in front of all the blocks on the table.</p> <p>Assistant: block bounding box [0.749, 0.714, 0.800, 0.818] Is it this block?</p> <p>User: Take the block that is directly above this block and to the left of it.</p> <p>Model generation: block bounding box [0.624, 0.571, 0.675, 0.675]</p>

Table 9: Full prompts given to models, including special tokens.

Difficulty Level	Source Block	Target Position
<i>Human Performance</i>		
Easy	28.69 (±12.20)	32.19 (±14.66)
Medium	28.80 (±12.97)	31.27 (±12.86)
Hard	27.62 (±12.19)	27.46 (±12.35)
<i>GPT-4o Performance</i>		
Easy	29.95 (±12.09)	34.29 (±13.81)
Medium	29.62 (±12.02)	31.46 (±12.45)
Hard	26.52 (±12.31)	24.54 (±11.73)

Table 10: Mean and Standard Deviation (SD) for test dialogues in BW-R by difficulty level and task. We differentiate between using human performance (Section 6) and GPT-4o performance (Appendix D) as the difficulty proxies.

Difficulty	Model	Source		Target
		Acc ↑	Mean (SD) ↓	Mean (SD) ↓
Easy	Idefics2	0.43	2.60 (±2.6)	6.11 (±2.9)
	LLaVA	0.48	2.02 (±2.0)	4.20 (±2.9)
	Humans	0.78	1.07 (±1.1)	2.05 (±2.1)
	GPT-4o	0.93	.00 (±.0)	1.27 (±.5)
Medium	Idefics2	0.47	2.29 (±2.3)	4.97 (±2.6)
	LLaVA	0.47	2.01 (±2.0)	4.25 (±2.7)
	Humans	0.77	1.05 (±1.1)	2.23 (±2.3)
	GPT-4o	0.58	1.26 (±1.3)	2.63 (±.4)
Hard	Idefics2	0.04	4.91 (±4.9)	5.92 (±3.1)
	LLaVA	0.22	4.61 (±4.6)	5.69 (±3.4)
	Humans	0.64	2.52 (±2.5)	4.84 (±3.7)
	GPT-4o	0.00	8.33 (±8.3)	6.31 (±3.1)

Table 11: Model performances across subsets with increasing levels of complexity (*easy, medium, hard*) using GPT-4o performance as the difficulty proxy.

Test Data	Train Data	Model	Loss	Source		Target
				Accuracy \uparrow	Mean Distance (SD) \downarrow	Mean Distance (SD) \downarrow
Instructions		Idefics2 8B <i>zeroshot</i>	Default	0.18	5.06 (± 3.50)	9.55 (± 2.50)
	Instructions	Idefics2 8B	Default	0.33	3.33 (± 3.02)	4.93 (± 2.64)
	Instructions	Idefics2 8B	User-turn	0.33	3.33 (± 3.02)	4.93 (± 2.64)
	Instructions	Idefics2 8B	Completions	0.32	3.44 (± 3.04)	4.75 (± 2.62)
	Repairs	Idefics2 8B	Default	0.21	4.36 (± 3.13)	5.38 (± 2.44)
	Repairs	Idefics2 8B	User-turn	0.22	4.59 (± 3.25)	8.03 (± 3.21)
	Repairs	Idefics2 8B	Completions	0.14	4.88 (± 3.07)	6.19 (± 2.97)
	Full	Idefics2 8B	Default	0.30	3.39 (± 2.92)	4.75 (± 2.60)
	Full	Idefics2 8B	User-turn	0.36	3.12 (± 3.00)	4.59 (± 2.59)
	Full	Idefics2 8B	Completions	0.33	3.33 (± 2.99)	4.52 (± 2.58)
		LLaVA v1.5 7B <i>zeroshot</i>	Completions	0.24	4.40 (± 3.34)	5.74 (± 3.08)
	Instructions	LLaVA v1.5 7B	Default	0.25	3.77 (± 2.99)	5.06 (± 2.92)
	Instructions	LLaVA v1.5 7B	User-turn	0.27	3.79 (± 3.03)	5.02 (± 2.86)
	Instructions	LLaVA v1.5 7B	Completions	0.27	3.76 (± 3.00)	4.40 (± 2.69)
	Repairs	LLaVA v1.5 7B	Default	0.14	4.97 (± 3.24)	6.24 (± 2.91)
	Repairs	LLaVA v1.5 7B	User-turn	0.18	4.19 (± 2.82)	6.37 (± 2.91)
	Repairs	LLaVA v1.5 7B	Completions	0.09	5.12 (± 3.00)	6.27 (± 2.86)
	Full	LLaVA v1.5 7B	Default	0.22	3.82 (± 2.91)	4.46 (± 2.49)
	Full	LLaVA v1.5 7B	User-turn	0.24	3.65 (± 2.86)	4.60 (± 2.54)
	Full	LLaVA v1.5 7B	Completions	0.19	3.93 (± 2.73)	4.51 (± 2.82)
	GPT-4o <i>zeroshot</i>	Completions	0.30	3.83 (± 3.38)	4.22 (± 2.65)	
Repairs		Idefics2 8B <i>zeroshot</i>	Completions	0.00	3.74 (± 1.95)	4.13 (± 2.05)
	Instructions	Idefics2 8B	Default	0.21	4.19 (± 2.75)	5.11 (± 2.60)
	Instructions	Idefics2 8B	User-turn	0.26	3.74 (± 2.84)	4.68 (± 2.12)
	Instructions	Idefics2 8B	Completions	0.28	3.66 (± 2.77)	3.98 (± 2.58)
	Repairs	Idefics2 8B	Default	0.58	2.19 (± 3.06)	7.07 (± 2.77)
	Repairs	Idefics2 8B	User-turn	0.58	2.16 (± 2.92)	6.45 (± 2.22)
	Repairs	Idefics2 8B	Completions	0.57	1.55 (± 1.95)	7.35 (± 2.02)
	Full	Idefics2 8B	Default	0.26	3.43 (± 2.40)	5.81 (± 2.35)
	Full	Idefics2 8B	User-turn	0.32	2.94 (± 2.56)	6.00 (± 1.95)
	Full	Idefics2 8B	Completions	0.47	2.29 (± 2.36)	5.90 (± 2.71)
		LLaVA v1.5 7B <i>zeroshot</i>	Completions	0.13	3.29 (± 2.09)	3.45 (± 1.20)
	Instructions	LLaVA v1.5 7B	Default	0.18	3.43 (± 2.49)	3.49 (± 1.34)
	Instructions	LLaVA v1.5 7B	User-turn	0.16	3.47 (± 2.42)	3.46 (± 1.39)
	Instructions	LLaVA v1.5 7B	Completions	0.25	2.79 (± 2.47)	3.76 (± 1.88)
	Repairs	LLaVA v1.5 7B	Default	0.07	3.20 (± 1.79)	3.67 (± 1.31)
	Repairs	LLaVA v1.5 7B	User-turn	0.13	3.28 (± 2.05)	3.89 (± 1.35)
	Repairs	LLaVA v1.5 7B	Completions	0.50	2.33 (± 2.64)	5.57 (± 1.67)
	Full	LLaVA v1.5 7B	Default	0.44	2.66 (± 2.76)	3.69 (± 2.18)
	Full	LLaVA v1.5 7B	User-turn	0.37	2.69 (± 2.48)	4.04 (± 2.41)
	Full	LLaVA v1.5 7B	Completions	0.54	2.01 (± 2.47)	4.56 (± 2.69)
	GPT-4o <i>zeroshot</i>	Default	0.41	2.75 (± 3.06)	2.95 (± 2.17)	

Table 12: Model performance in source and target sub-tasks across all losses.