

PLAYPEN: An Environment for Exploring Learning From Dialogue Game Feedback

Anonymous ACL submission

Abstract

Interaction between learner and feedback-giver has come into focus recently for post-training of Large Language Models (LLMs), through the use of reward models that judge the appropriateness of a model’s response. In this paper, we investigate whether Dialogue Games—goal-directed and rule-governed activities driven predominantly by verbal actions—can also serve as a source of feedback signals for learning. We introduce PLAYPEN, an environment for off- and online learning through Dialogue Game self-play, and investigate a representative set of post-training methods: supervised fine-tuning; direct alignment (DPO); and reinforcement learning with GRPO. We experiment with post-training a small LLM (Llama-3.1-8B-Instruct), evaluating performance on unseen instances of training games as well as unseen games, and on standard benchmarks. We find that imitation learning through SFT improves performance on unseen instances, but negatively impacts other skills, while interactive learning with GRPO shows balanced improvements without loss of skills. We release the framework and the baseline training setups to foster research in this promising new direction of “learning in (synthetic) interaction”.¹

1 Introduction

Dialogue Games of the kind illustrated in Figure 1 have been used to evaluate interactive reasoning capabilities of LLMs, requiring as they do the strategic acquisition of information through multi-turn interaction (Bertolazzi et al., 2023; Chalamalasetti et al., 2023; Qiao et al., 2023; Li et al., 2023; Guertler et al., 2025; Cui et al., 2025). In this paper, we show that it is also possible to derive a learning signal from such games, which we call Dialogue Game Feedback (DGF). DGF offers a unique combination of properties: A) Unlike feedback from

A (quietly reads game card): Target word: ugly. Taboo words: displeasing, despicable, unattractive.
(1a) A (to B): Something that is not pleasing to the eye, often making a bad impression.
(1b) B: Eyesore
(2a) A: More about the quality of a person or thing, rather than a nuisance or distraction.
(2b) B: Ugly.
(3a) A: Yay! Well done us.

Figure 1: A successful episode of the two-player Dialogue Game Taboo.

the reward models used in Reinforcement Learning from Human Feedback (RLHF; Ouyang et al., 2022), DGF is *teleological* or goal-directed, in that it judges not just the appropriateness of a response, but also its effectiveness in advancing toward a desired outcome. This is a property DGF shares with feedback from process- and outcome-based reward models used to optimize reasoning models (PRM, Setlur et al. 2024; ORM, Hosseini et al. 2024; Cobbe et al. 2021; respectively). B) Unlike these aforementioned prior methods, DGF is *objective*, in that it can be computed programmatically, rather than needing a learned, “subjective” judgement model; a property it shares with the “verifiable rewards” of Lambert et al. (2025). C) Unlike all of these other methods, it can be derived from *inter-subjective*, multi-turn linguistic interaction. As DGF is defined in terms of *task success*, and task success here is conditional on *communicative success*—i.e., players are required to produce mutually intelligible language—the feedback signal implicitly carries linguistic normative pressure. D) Lastly, where other methods focus on alignment with respect to *desirability* (RLHF) or specifically on reasoning skills in domains such as maths and coding (PRM, ORM, verifiable rewards), DGF rewards general backwards- and forwards-looking language use. At the same time, insofar as they are required by the underlying game, it also targets specific skills such as spatial reasoning or referential language.

¹The framework is available here: REDACTED.

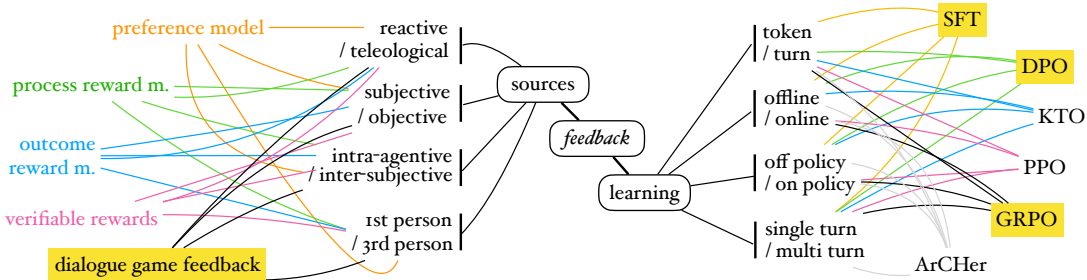


Figure 2: Placing Dialogue Game Feedback in the wider research context. Highlighted on the right the selection of learning methods that we provide baselines for.

Figure 1 illustrates this. It shows an example of two players playing the Dialogue Game Taboo where a *clue giver* needs to describe a concept to a *guesser*, while avoiding certain “taboo” words. Both players produce natural language strings, which can be judged in two ways: *teleologically*, where both 1a and 1b are appropriate but do not lead to success, while 2a and 2b succeed; and *objectively* as it can be programmatically verified that 1b and 2b comply with the rules and that 2b is the correct answer. Moreover, for the game to progress, both clue giver and guesser must produce mutually intelligible language, placed in the context of the interaction as whole. E.g. turn 2a must consider the previous guess and how it failed (backwards-looking) in order to produce an utterance that aims to elicit a better guess (forward-looking).

Contribution 1: We introduce PLAYPEN, an environment where LLMs engage in dialogue games (without human intervention) from which a feedback signal for learning is derived; either online during gameplay or offline from rollouts, and based either on single trajectories or sets of ranked alternatives. **Contribution 2:** We demonstrate how to leverage this feedback signal for post-training.² We investigate three learning approaches: imitation learning (supervised fine-tuning); an offline alignment approach, Direct Preference Optimization (DPO; Rafailov et al., 2023); and an online learning algorithm, Group Relative Policy Optimization (GRPO; Shao et al., 2024)—thereby establishing strong baselines for this learning environment. We evaluate the resulting models on a range of tests, including held-out dialogue games to assess skill generalisation, a comprehensive suite of tests assessing broader linguistic and cognitive abilities (Momentè et al., 2025), and standard NLP bench-

marks such as MMLU-Pro (Wang et al., 2024), Big-bench Hard (Suzgun et al., 2023), and IFEval (Zhou et al., 2023b). **Contribution 3:** We show that imitation learning through SFT improves performance on unseen instances but negatively impacts other skills, whereas interactive learning with GRPO achieves balanced improvements without skill degradation.

2 Related Work

Our work builds on several threads of research, which we now briefly review (visualised in Fig. 2).

Sources of Learning Feedback. Post-training methods typically rely on there being *feedback* on the appropriateness of the production R of a model relative to a context C , provided by a feedback function f_F . We have already outlined in the introduction the main approaches to designing a feedback function, via trained judge models or verifiers in formal domains (also see Kumar et al. 2025; Lambert 2024 for recent detailed overviews)—and we have positioned our proposed *Dialogue Game Feedback* within this landscape. We have also argued that DGF represents a form of *teleological* (goal-directed) and *objective* feedback, meaning that the function f_F is realised programmatically and reflects the specific goals of the game. A similar kind of feedback has been used recently by Gul and Artzi (2024), but only in the context of a single game, and with specialised learning methods. Similarly, Sadler et al. (2024) used feedback from a cognitively inspired programmatic partner to bootstrap language capabilities of a collaborative neural agents in a single visual reference task.

Methods for Learning from Feedback. Assuming that a function f_F for judging model contributions is at hand, how can this be used to improve the model (the policy) from which these are sampled?

²Leaving to future work the exploration of these types of interactive settings for language acquisition from scratch.

Various methods have been developed in recent years to do this (see the surveys cited above). For our purposes, we categorise them as follows. On the level of granularity, there are *token-based* methods, for example standard supervised fine-tuning. In such a learning setup, f_F is just used to filter out good material for imitation learning (rejection sampling); the hypothesis then would be that becoming more likely to produce token sequences as in the training material generalises to getting better at making strategic decisions. *Turn-level* methods, on the other hand, make direct use of the turn-level feedback. Within these, *offline* methods such as Direct Preference Optimization (DPO; Rafailov et al., 2023) and Kahneman-Tversky Optimization (KTO; Ethayarajh et al., 2024) work with judgements on already collected material, which does not necessarily has to have been produced by the policy that is to be improved (and in that sense are *off-policy* methods as well). Recently, due to the popularity of the next-token prediction objective used in LLMs, decision-making tasks in Embodied AI have also been formulated as a sequence modelling problem, where state, action, and feedback are interleaved in a single sequence (Decision Transformers; Chen et al., 2021; McCallum et al., 2023). *Online* methods such as Proximal Policy Optimization (PPO; Schulman et al., 2017) and Group Relative Policy Optimization (GRPO; Shao et al., 2024), on the other hand, attempt to directly improve the policy that produces the samples. Lastly, a new development that is relevant for our setup are multi-turn methods such as ArCher (Zhou et al., 2024b), ReSpect (Chen et al., 2024), and REFUEL (Gao et al., 2024), which can natively handle the multi-turn nature of conversational interactions. We leave exploring these to future work and concentrate here on a representative selection of methods: SFT, DPO, and GRPO.

LLMs and Dialogue Games. Conversational interactions framed as games have long been used to investigate language use; see discussion in (Schlangen, 2019, 2023; Suglia et al., 2024), which also proposed to use them for evaluating language use capabilities of NLP models. This idea has been implemented by various frameworks in recent years. An early precursor to this was TextWorld (Côté et al., 2019), which however operated with a single genre of game (text adventures / interactive fiction), and with the expectation of training specialist models. Only with the advent of generalist

models that can be *prompted* into being specialists (Brown et al., 2020; Wei et al., 2021) did it become possible to implement this idea at a larger scale, for single games (Bertolazzi et al., 2023) and more generally in frameworks in which various games can be implemented (Chalamalasetti et al., 2023; Qiao et al., 2023; Li et al., 2023; Gong et al., 2023; Wu et al., 2024; Zhou et al., 2024a; Duan et al., 2024; Guertler et al., 2025; Cui et al., 2025). Among these, we chose to build on clembench (Chalamalasetti et al., 2023) as the longest-running continuously maintained effort, that also comes with an extensive archive of dialogue game transcripts spanning a wide range of models.

The Cognitive Plausibility of Learning from Interaction. It is well established in the developmental literature that human language acquisition requires social interaction (Clark, 2016; Bruner and Watson, 1983; Kuhl, 2007; Hiller and Fernández, 2016; Saxton, 2000; Bloom, 2000). It has also been claimed for the machine acquisition of language that social interaction can be beneficial or might even be required (Fernández et al., 2011; Bisk et al., 2020; Bender and Koller, 2020), in particular in light of the vast differences in terms of sample efficiency (Hart and Risley, 1995; Cristia et al., 2019; Linzen, 2020). Our work connects to this literature and speaks to this hypothesis, by exploring the use of a learning signal derived from (an approximation of) social linguistic interaction. Incidentally, this is also the motivation of the BabyLM challenge (Charpentier et al., 2025), which in its latest incarnation explicitly encourages the use of synthetic interaction. A few recent pioneering works (Nikolaus and Fourtassi, 2021; Ma et al., 2025) have begun to explore this direction, demonstrating the potential benefits of interaction even in learning from scratch. While we are sympathetic to this line of work, our current approach assumes that the model is already capable of following instructions well enough to engage in gameplay. Accordingly, we focus on post-training.

3 Playpen: Dialogue Games & Feedback

3.1 Dialogue Games with LLMs

Figure 1 above provided an example of a Dialogue Game. How can we enable LLMs to play such games effectively? One of the surprising insights of the “LLM revolution” was that, at previously unseen scale, these models can be *prompted* to perform a wide range of tasks (Wei et al., 2021; Brown

et al., 2020). As the frameworks described above have shown, this extends to prompting LLMs to act as policies for playing conversational games—albeit with varying degrees of success. For example, a simple prompt such as “We are playing a word guessing game. Your task is to describe the word, but you are not allowed to use some other words. The word to describe is ‘ugly’, and the words to avoid are ‘displeasing’, ‘despicable’, and ‘unattractive’.” can induce (at least some) LLMs to act as a policy π_{taboo} capable of playing that specific role in the game reasonably well.³ To enable self-play with LLMs, each player must be separately prompted, often with distinct information states. Following Chalamalasetti et al. (2023); Smith et al. (2024), we use a programmatic *Game Master* (GM) to mediate the interaction. In the case of Figure 1, for example, the GM would insert a turn between 1a and 1b, delivering instructions to player B and relaying the clue from player A. See Appendix C for full transcripts of such interactions.

All of the games used here (see Section 3.3 below) involve some form of reasoning. Crucially, however, the reasoning involved is fundamentally different from that required in standard applications of reasoning models (Besta et al., 2025). Unlike conventional reasoning tasks such as math word problems (Hendrycks et al., 2021), which are *well-posed*—i.e., the problem is fully specified and the challenge lies in deriving a solution through a correct sequence of steps—the games studied here require *multi-turn* and *interactive* reasoning, as they are *ill-posed* at the outset. They only become tractable through iterative exchanges between players. Consider the starting prompt for a Wordle-type word guessing game: “Guess a 5 letter word”. Only through making guesses, receiving feedback, and updating beliefs accordingly does the task become solvable and the identity of the target word recoverable. The reasoning at play in such settings involves managing uncertainty and coordinating with another agent under conditions of imperfect information.

3.2 Dialogue Game Feedback

We refer to a specific type of game, such as Taboo or Wordle, as a *Dialogue Game* (DG). A particular instantiation of a DG—created by filling in a prompt template with specific parameters (for ex-

ample, the exact target and taboo words in Taboo)—is called a *Dialogue Game Realisation Instance*, or simply an *instance*. We denote instance i of game g as $x_{g,i}$. Each instance defines a *game tree*, which originates from the initial game description and branches out at every turn into all possible actions available at that point. In other words, the tree contains all possible gameplays for that particular instance. If the DG allows verbal actions of unbounded length (i.e., compositional and infinite action spaces), the corresponding game tree will have an infinite number of nodes and edges. A (complete) *trajectory* is a path from the root node to a leaf node. Each player in the game is represented by a *policy* π that determines their action at each decision point. When all players required by a game are instantiated with policies, this collectively induces a distribution over trajectories. An *episode* of gameplay for a given instance can then be viewed as a sample from that distribution. The resulting interaction can be recorded as a *transcript* t as follows (this description applies to a two-player game, with straightforward extensions to games involving more players).

Definition 3.1 (Transcript). We represent a trajectory through a game instance tree as a *transcript* t , as follows:

$$t = (S_0, C_1^A, R_1^A, S'_0, C_1^B, R_1^B, S_1, \dots, S_{F-1}, C_F^A, R_F^A, S'_{F-1}, C_F^B, R_F^B, S_F)$$

where C_n^P is the information (context) that is presented at turn n to player P ; R_n^P is the response produced by player P at turn n ; S_n is the abstract game state at turn n . C_1^P has a special status, as it contains the game description prompt. A *complete* trajectory ends with a turn t_F that either leads to a final state S_F for a completed game or to an abort state S_X . The *Dialogue Game Feedback* is computed by a game-specific scoring function s which evaluates the trajectory—typically a complete one, although for some games described below, incomplete trajectories may also be assessed. Note that a trajectory t as defined above contains all interleaved interactions between player and GM. However, everything that a player P needs to know about the actions of other players is contained in the C^P presented to them, and they never *directly* observe any $R^{P'}$ from other players. To reconstruct what one side experienced, a perspective function p_P reduces a trajectory to contain only context-response sequences $(C_i^P, R_i^P), 0 \leq i < n$ for a

³See Chalamalasetti et al. (2023) for an example of a full prompt that can be used, which needs to contain additional formatting instructions.

given player P at turn n .

3.3 The PLAYPEN Environment

The PLAYPEN environment we introduce here builds upon the Dialogue Game benchmark (clembench; Chalamalasetti et al., 2023), transforming it into an interactive playground in which LLMs can learn to be language users. As of now, PLAYPEN includes 15 clembench Games, testing language and world knowledge (e.g., in games such as Taboo, Wordle, Codenames); the ability to perform conversational grounding (e.g. Private-Shared, GuessWhat); and spatial and causal reasoning (e.g., Adventure Games orMap Navigation). We provide the full list of games with further details in Appendix A.

By recording trajectories as defined in Section 3.2, PLAYPEN supports both offline and online learning, as well as the representation of branching subtrees within the overall game tree through repeatedly sampling from player policies. This flexibility enables the learning experiments that we turn to now.

4 Experimental Setup

Our experiments focus on leveraging Dialogue Game Feedback to post-train language models, building on the premise that this feedback signal is most effectively used when the model is already capable of prompted gameplay to a certain extent.

4.1 Models

We selected Llama-3.1-8B and Llama-3.1-70B (Meta, 2024), both in the instruct variant, as they have generally shown to be receptive to further training (Taori et al., 2023; Zhou et al., 2023a), and have performed well within their size classes on the public clembench leaderboard.⁴ In some experiments, we also used a 4-bit quantised version of the model⁵ for more efficient training and inference.

4.2 Evaluation

What improvements can we expect from learning with Dialogue Game Feedback? We hypothesise that we will see improvements on gameplay on unseen instances of the games encountered during training, as well as generalisation to new game types. To assess the broader impact of this learning,

⁴<https://clembench.github.io/leaderboard.html>, retrieved 2025-05-09.

⁵<https://huggingface.co/unsloth/Meta-Llama-3.1-8B-Instruct-bnb-4bit>

we additionally evaluate the post-trained models for their formal and functional linguistic competence, as well as on general NLP benchmarks.

Interactive Dialogue Games. Performance on interactive dialogue games is evaluated using the clemscore metric (Chalamalasetti et al., 2023), which captures both the ability to adhere to the formal rules of a game and the quality of the gameplay. Specifically, the clemscore is obtained by multiplying the macro-average percentage of games that were validly played with the macro-average quality score (typically task success) in those valid attempts. We use seven of the 15 available Dialogue Games for training: Taboo, Private-Shared, ImageGame, Reference game, and three variants of Wordle. For evaluation, we generate new instances of these games to form the *in-domain* test set. The remaining eight games—Codenames, Adventure Game, Guess-What, MatchIt, and three variants of Map Navigation—serve as an additional *out-of-domain* test set.

Formal Linguistic Competence. We evaluate formal linguistic competence (Mahowald et al., 2024), such as the ability to recognize morphosyntactic agreement or lexical entailment, using the GLUE Diagnostic dataset (Wang et al., 2018).

Functional Linguistic Competence. We evaluate cognitive abilities required for verbal interaction such as working memory, theory of mind, common-sense reasoning, and socio-emotional skills, using a subset of the tasks curated by Momentè et al. (2025). Specifically, we use a sample from Natural Plan (Zheng et al., 2024), LogiQA 2.0 (Liu et al., 2023), CLadder (Jin et al., 2023), WinoGrande (Sakaguchi et al., 2021), EQ-Bench (Paech, 2023), LM-Pragmatics (Hu et al., 2023), SocialIQA (Sap et al., 2019), SimpleToM (Gu et al., 2025).

Knowledge and Instruction Following. Finally, we evaluate models on two widely used LLM benchmarks: MMLU-Pro (Wang et al., 2024) and Big-bench Hard (BBH; Suzgun et al., 2023). In addition, we report performance on IFEval (Zhou et al., 2023b), which assesses general instruction-following capabilities.

4.3 Training Regimes and Interaction Data

We experiment with several training regimes enabled by PLAYPEN, and compare their outcomes with the unmodified Llama-3.1 (Baseline). Additional details on training setups, hyperparameters, and data generation procedures are in Appendix E.

4.3.1 Imitation Learning / SFT

We begin by investigating the potential of pure imitation learning / supervised fine-tuning (SFT). We create a training dataset \mathcal{D} by collecting episodes of gameplay from a large variety of models listed on the public clembench leaderboard. For our experiments with SFT, we filter \mathcal{D} to retain only the 7000 successful episodes (i.e., we discard lost and aborted episodes) and convert the transcripts into separate trajectories for all player perspectives: $\mathcal{D}_{SFT} = \{p_A(t), p_B(t) \mid t \in \mathcal{D}, s(t) > \tau_g\}$, for a game-specific threshold τ_g , and using player-specific perspective functions p_A and p_B . We fine-tune the Baseline model using different data mixtures, containing only interaction data or also instruction following examples. After initial experimentation with plain SFT on 7000 transcripts over 1100 training steps, we found that more sample-efficient variants offered better generalisation. Details on this analysis are reported in Appendix E.1. In what follows, we report results from the most effective configuration, which we refer to as SFT (Cold Start, CS). This variant uses only 700 training steps and focuses exclusively on interaction data.

4.3.2 Direct Alignment

Building on the best SFT checkpoint, SFT (CS), we apply Direct Preference Optimization (DPO; Rafailov et al., 2023). DPO offers a middle ground between SFT and full online reinforcement learning by leveraging contrastive learning on offline data. We consider two variants of DPO training:

DPO Dialogue: For every positive sample in the filtered dataset \mathcal{D}_{SFT} , we find a negative counterpart in \mathcal{D} that starts from the same initial state (prompt and game instance S_0) but ends without reaching a successful final state. This includes both aborted and completed but unsuccessful dialogues. The dataset \mathcal{D}_{DPO} consists of ca. 10K pairs of positive and negative trajectories.

DPO Turn: For each turn in a successful dialogue, we find a negative counterpart that shares the same conversational history (the prompt, the game instance and the history up to a branching point) to yield $S_0, C_1^A, R_1^A, \dots, C_n^A, (R_n^A, R_n'^A)$, i.e., paired samples identical up to a penultimate branching node in the game tree. The dataset consists of ca. 86K pairs of positive and negative trajectories.

4.3.3 Online Learning

While both SFT and DPO recover useful learning signals, neither method captures the interactive nature of dialogue games. For this reason, we also performed experiments using GRPO (Shao et al., 2024). Unlike the above methods, GRPO does not rely on a fixed training dataset \mathcal{D} . Instead, for each game instance x_i (initial prompt), we use the very same model to interactively produce 8 samples of full gameplay (with temperature set to 0.75). This leads to a group G_i of alternative transcripts. Each trajectory $g \in G_i$ is evaluated using a game-specific reward function that corresponds to the quality score computation for that game in the clembench benchmark (see Appendix D.4). We test two configurations of GRPO, i.e., training either starts from the base or the SFT (CS) model.

5 Results

We now present the results of running our evaluation suite (Section 4.2) on the resulting models. A high-level take-away is that only GRPO, the one interactive and turn-based learning method that we tried, was able to learn from the Dialogue Game Feedback in a way that generalises to unseen games (where the other methods broke down, presumably due to overfitting) and without negative impact on other skills.

5.1 Interactive Gameplay

Table 1 presents results on dialogue game performance across all evaluated models. We begin with the Llama-3.1-8B baseline, which demonstrates a basic ability to follow game rules—managing to complete just over half of the games in both the in-domain and out-of-domain settings—with generally low gameplay quality. Interestingly, the baseline achieves higher quality and clemscore on out-of-domain games, indicating that the in-domain set may pose greater difficulty for this model. Training with SFT on successful episodes improves in-domain performance but at the cost of generalisation, with a decrease in all dimensions of performance on out-of-domain games. This aligns with prior observations that SFT tends to overfit to the training distribution and struggles to generalise beyond it (Zeng et al., 2023; Chu et al., 2025; Setlur et al., 2025). Turning to DPO, we observe that both its variants (turn- or dialogue-level preference pairs) improve in-domain performance over the Baseline but fail to outperform

















































Model	In Domain			Out of Domain		
	Clemscore	% Played	Quality	Clemscore	% Played	Quality
Llama-3.1-8B						
Baseline	19.39	58.50	33.15	24.58	54.53	45.08
SFT (CS)	40.11 	70.48 	56.91 	22.53 	50.55 	44.58 
SFT (CS) + DPO (Dial.)	32.33 	63.54 	50.89 	19.50 	51.67 	37.74 
SFT (CS) + DPO (Turn.)	33.29 	71.90 	46.29 	20.96 	48.26 	43.43 
GRPO	24.89 	57.55 	43.25 	33.92 	67.38 	50.34 
SFT (CS) + GRPO	24.30 	63.22 	38.44 	31.81 	67.26 	47.29 
Llama-3.1-70B						
Baseline	37.24	64.57	57.67	47.37	82.29	57.57
SFT (CS)	53.60 	81.57 	65.71 	54.40 	85.57 	63.57 
SFT (CS) + DPO (Dial.)	36.92 	52.44 	70.41 	45.46 	73.66 	61.71 
SFT (CS) + DPO (Turn.)	38.68 	67.59 	57.20 	50.65 	86.29 	58.70 

Table 1: **Main results (Games).** Clemscore, average percentage of completed games, and average quality score.

the best SFT model. Moreover, they suffer from even stronger degradation on out-of-domain games. We believe this could be a result of “likelihood displacement”—a weakness of DPO-based training strategies (Razin et al., 2024). The most robust training regime is GRPO. Applied directly on the base model, GRPO leads to consistent improvements (+5.50 in in-domain clemscore and +9.34 in out-of-domain) with the exception of a slight decrease in the number of in-domain games played. When GRPO is applied on top of the SFT model, it recovers even from this slight drop in game completion, likely due to the SFT model’s stronger ability to adhere to game instructions, at the cost of more modest gains in quality score.

What happens at the 70B scale? Larger models are known to possess stronger instruction-following capabilities, a skill that is especially relevant for our benchmark, where accurate interpretation of game prompts is critical to gameplay. We therefore conducted additional experiments with Llama-3.1-70B, applying the same training regimes used for the 8B model, but excluding GRPO due to its high computational cost and our resource constraints. It is worth noting that simply using this model as a starting point nearly doubles the overall clemscore on both in-domain and out-of-domain games—highlighting how better instruction-following capabilities indeed contribute to improved gameplay. When applying SFT to the larger base model, we observed diminishing returns for in-domain games, with an improvement of only 16.36 points over the Baseline—lower than the increase obtained when applying SFT on the base 8B model. However, on out-of-domain games, we record our best-scoring model with a clemscore of

54.4. To calibrate this result, this is still far below the top leaderboard clemscore of 70, achieved by o3-mini-2025-01-31.⁶ Finally, we find that applying DPO on top of SFT reverses some of the gains of SFT alone. Between the two DPO variants, DPO with dialogue-level preference data obtains lower scores on both in-domain and out-of-domain games—another possible case of overfitting after the preliminary SFT phase. Conversely, the turn-level variant of DPO yields modest improvements over the baseline on both in-domain and out-of-domain games, though still falling short of the best SFT model.

5.2 Non-Interactive Benchmarks

While our main focus is on performance in dialogue games, we also evaluate models across a broad set of other tasks. This helps identify whether training on dialogue games leads to regressions in general language skills (e.g., formal competence) or, conversely, contributes to improvements in language use (e.g., functional competence). Table 2 summarises the results across these evaluations.

Among the training regimes, GRPO seems to be the most balanced overall, with lower oscillations—either improvements or regressions—in non-interactive task performance. Tables 11 and 12 in the Appendix give a complete overview of the results. We highlight here that training on dialogue games seems to provide a modest improvement on the “Executive” task category for the 70B model trained with SFT (CS), suggesting that learning from dialogue games may enhance a model’s ability to integrate and reason over context-

⁶Based on <https://clembench.github.io/leaderboard.html>, accessed May 14th 2025.

























































Model	In Domain	Out of Domain	Functional, Formal, General, Instruction Following				
	Clemscore	Clemscore	Executive	Socio-Emo	GLUE D.	General QA	IFEval
Llama-3.1-8B							
Baseline	19.39	24.58	39.24	57.16	38.06	41.86	76.88
SFT (CS)	40.11 	22.53 	39.93 	59.51 	40.43 	29.95 	67.25 
SFT (CS) + DPO (Dial.)	32.33 	19.50 	39.80 	55.10 	36.20 	26.57 	68.39 
SFT (CS) + DPO (Turn)	33.29 	20.96 	38.50 	59.48 	36.07 	28.97 	70.76 
GRPO	24.89 	33.92 	39.39 	57.51 	38.68 	41.52 	76.67 
SFT (CS) + GRPO	24.30 	31.81 	33.35 	58.67 	37.31 	42.82 	75.77 
Llama-3.1-70B							
Baseline	37.24	47.37	52.42	71.37	46.16	60.56	85.16
SFT (CS)	53.60 	54.40 	55.17 	69.25 	47.72 	44.91 	79.38 
SFT (CS) + DPO (Dial.)	36.92 	45.46 	48.94 	67.89 	37.73 	38.78 	82.26 
SFT (CS) + DPO (Turn)	38.68 	50.65 	50.22 	70.21 	39.23 	44.86 	85.68 

Table 2: **Main results.** Clemscore, average percentage of completed games, and average quality score. We report the best SFT variant Cold Start (CS). Executive includes LogiQA 2.0, CLadder, and WinoGrande. Socio-emotional includes EQ-Bench, LM-Pragmatics, SocialIQA, and SimpleToM. General QA includes MMLU-Pro and BBH, while IFEval targets instruction-following specifically. Formal capabilities are evaluated in GLUE Diagnostics. Colored bars indicate whether there is a positive (*green*) or negative difference (*red*) wrt. the Baseline model.

tual information—an ability [Momentè et al. \(2025\)](#) identified as critical for these tasks.

Another relevant finding from this evaluation is the substantial drop in the ability to follow instructions, as measured by IFEval. This calls for further investigation into instruction-following training regimes that are more robust to interactive settings, allowing models not only to handle single-turn prompts—as is common in current instruction-following regimes—but also to participate effectively in complex, goal-oriented, rule-governed, and multi-turn tasks such as dialogue games.

5.3 Qualitative Discussion

An outcome of our evaluation is that current LLMs are lacking true instruction-following abilities (also see Appendix G for a detailed error analysis on gameplay abilities). Thanks to language prompts describing a game g , we should be able to derive a game policy π_g on the fly. In some cases, this might not be enough, and therefore, it is possible to use SFT to learn how to play the game by mimicking transcripts. However, because the model does not have the chance to play by itself, it might miss some nuances of the game and overfit on specific rules/formats of the game at hand. On the other hand, thanks to the online training regime of GRPO, it is possible to acquire general-purpose instruction following abilities that allow models to perform better in out-of-domain games as well as retain abilities required for more general-purpose NLP tasks—a result in line with test-time compute

analysis for RL algorithms reported in the literature ([Chu et al., 2025](#); [Setlur et al., 2025](#)).

6 Conclusion

In this paper, we explore to what extent synthetic interaction in what we call Dialogue Games—goal-directed and rule-governed activities driven by verbal actions—can provide a learning signal for LLMs. We created PLAYPEN, an environment that facilitates synthetic data generation of dialogue transcripts that can be used to train LLMs. We provide an extensive evaluation of a variety of state-of-the-art post-training methods such as SFT, DPO, and GRPO, and show how GRPO is a more stable training regime that prevents overfitting to in-domain games and facilitates generalisation to out-of-domain dialogue games. Additionally, we demonstrate that when leveraging dialogue games, it is possible to observe an improvement in performance when completing more traditional, non-interactive general instruction-following tasks such as MMLU-Pro.

The framework and the baselines presented here can form the basis for exciting future work, for example investigating novel training regimes based on reinforcement learning to truly leverage the multi-turn nature of dialogue games, or exploring the use of intermediate language feedback that can be acquired as part of the interaction (along the lines of [Sumers et al. 2021](#)), and further exploring the potential of the “learning in conversational interaction” paradigm.

Limitations

Our study makes significant strides in demonstrating the potential of dialogue games as a valuable source of feedback signals for training LLMs. The PLAYPEN environment offers a versatile platform for exploring both off- and online learning paradigms, and our comparative analysis of post-training methods, including SFT, DPO, and reinforcement learning with GRPO, provides a strong foundation for future research.

However, the current study has several limitations that warrant further investigation. Firstly, in our Direct Preference Optimization (DPO) experiments, we utilize a seed dataset of successful dialogues from which we derive positive and negative pairs. However, for the turn variant, we assume that all turns within these dialogues are successful. This assumption may not hold true in real-world scenarios, particularly when corrections or clarifications are present within the dialogue (Chiyah-Garcia et al., 2024). Secondly, our work does not explore multi-turn training methods, which could be crucial for more complex dialogue games and real-world applications where it is important to perform credit-assignment across multiple turns (e.g., Zhou et al., 2024b).

Additionally, we did not incorporate reasoning models (e.g., DeepSeek-AI et al., 2025) or chain-of-thought prompting techniques (Wei et al., 2022), which have shown promise in enhancing LLM performance in other tasks. Furthermore, our evaluation of GRPO is limited to a smaller 8B LLM. Evaluating the effectiveness of GRPO on larger models, such as the 70B parameter model, would provide valuable insights into the scalability of our findings. Unfortunately, due to limited computational resources, we leave this exploration for future work.

The current set of dialogue games in PLAYPEN provides a foundation for our research, but it is not exhaustive. Future work should aim to expand the set of games to be more representative of the diverse range of language games encountered in real-world scenarios. This is especially important considering that Momentè et al. (2025) has demonstrated that dialogue games are actually more discriminative than other benchmarks because they likely require important underlying capabilities such as working memory.

Finally, our focus is on dialogue game feedback, which is inherent to the task itself. We do not

consider additional forms of feedback, such as explicit corrective feedback, which could potentially enhance learning, as explored in prior work (e.g., Sumers et al., 2021; McCallum et al., 2023; Xi et al., 2024).

Ethical considerations

Our work broadly falls under the rubric of “self-improvement” of language models. There is a small, but non-zero chance that such self-improvement, if run unsupervised and in recursive loops, might lead to uncontrolled gains. Our advice hence would be to define clear stopping criteria for learning runs. Additionally, we created PLAYPEN as a synthetic and simulated learning environment where the model doesn’t have access to external tools or, more broadly, it doesn’t have the ability to execute actions in the real world.

References

- Emily M Bender and Alexander Koller. 2020. Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 5185–5198. Issue: 2.
- Leonardo Bertolazzi, Davide Mazzaccara, Filippo Merlo, and Raffaella Bernardi. 2023. *ChatGPT’s information seeking strategy: Insights from the 20-questions game*. In *Proceedings of the 16th International Natural Language Generation Conference*, pages 153–162, Prague, Czechia. Association for Computational Linguistics.
- Maciej Besta, Julia Barth, Eric Schreiber, Ales Kubicek, Afonso Catarino, Robert Gerstenberger, Piotr Nyczyk, Patrick Iff, Yueling Li, Sam Houlston, Tomasz Sternal, Marcin Copik, Grzegorz Kwaśniewski, Jürgen Müller, Łukasz Flis, Hannes Eberhard, Hubert Niewiadomski, and Torsten Hoeffler. 2025. *Reasoning Language Models: A Blueprint*. *arXiv preprint*. ArXiv:2501.11223 [cs].
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. 2020. *Experience grounds language*. *EMNLP 2020 - 2020 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference*, pages 8718–8735. ArXiv: 2004.10151 ISBN: 9781952148606.
- Paul Bloom. 2000. *How children learn the meanings of words*. Learning, development, and conceptual change. MIT Press, Cambridge.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind

771	Neelakantan, Pranav Shyam, Girish Sastry, Amanda	Marc-Alexandre Côté, Akos Kádár, Xingdi Yuan, Ben	827
772	Askeell, Sandhini Agarwal, Ariel Herbert-Voss,	Kybartas, Tavian Barnes, Emery Fine, James Moore,	828
773	Gretchen Krueger, Tom Henighan, Rewon Child,	Matthew Hausknecht, Layla El Asri, Mahmoud	829
774	Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens	Adada, and 1 others. 2019. Textworld: A learn-	830
775	Winter, and 12 others. 2020. Language Models are	ing environment for text-based games. In <i>Computer</i>	831
776	Few-Shot Learners . In <i>Advances in Neural Informa-</i>	<i>Games: 7th Workshop, CGW 2018, Held in Con-</i>	832
777	<i>tion Processing Systems</i> , volume 33, pages 1877–	<i>junction with the 27th International Conference on</i>	833
778	1901. Curran Associates, Inc.	<i>Artificial Intelligence, IJCAI 2018, Stockholm, Swe-</i>	834
779	Jerome S. Bruner and Rita Watson. 1983. <i>Child’s talk:</i>	<i>den, July 13, 2018, Revised Selected Papers 7</i> , pages	835
780	<i>learning to use language</i> , 1st ed edition. W.W. Norton,	41–75. Springer.	836
781	London, New York. Format: E-Book.		
782	Kranti Chalamalasetti, Jana Götze, Sherzod Hakimov,	Alejandrina Cristia, Emmanuel Dupoux, Michael Gur-	837
783	Brielen Madureira, Philipp Sadler, and David	ven, and Jonathan Stieglitz. 2019. Child-directed	838
784	Schlangen. 2023. clmbench: Using game play to	speech is infrequent in a forager-farmer popula-	839
785	evaluate chat-optimized language models as conversa-	tion: A time allocation study. <i>Child development</i> ,	840
786	tional agents . In <i>Proceedings of the 2023 Confer-</i>	90(3):759–773.	841
787	<i>ence on Empirical Methods in Natural Language</i>		
788	<i>Processing</i> , pages 11174–11219, Singapore. Association	Christopher Zhang Cui, Xingdi Yuan, Zhang Xiao,	842
789	for Computational Linguistics.	Prithviraj Ammanabrolu, and Marc-Alexandre Côté.	843
790	Lucas Charpentier, Leshem Choshen, Ryan Cotterell,	2025. Tales: Text adventure learning environment	844
791	Mustafa Omer Gul, Michael Hu, Jaap Jumelet, Tal	suite. <i>arXiv preprint arXiv:2504.14128</i> .	845
792	Linzen, Jing Liu, Aaron Mueller, Candace Ross,		
793	Raj Sanjay Shah, Alex Warstadt, Ethan Wilcox, and	Michael Han Daniel Han and Unsloth team. 2023. Un-	846
794	Adina Williams. 2025. BabyLM Turns 3: Call for pa-	sloth .	847
795	pers for the 2025 BabyLM workshop . <i>arXiv preprint</i> .		
796	ArXiv:2502.10645 [cs].	DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang,	848
797	Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee,	Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,	849
798	Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind	and et al. 2025. Deepseek-r1: Incentivizing rea-	850
799	Srinivas, and Igor Mordatch. 2021. Decision trans-	soning capability in llms via reinforcement learning .	851
800	former: Reinforcement learning via sequence mod-	<i>Preprint</i> , arXiv:2501.12948.	852
801	eling. <i>Advances in neural information processing</i>		
802	<i>systems</i> , 34:15084–15097.	Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and	853
803	Zizhao Chen, Mustafa Omer Gul, Yiwei Chen, Glo-	Luke Zettlemoyer. 2023. Qlora: Efficient finetuning	854
804	ria Geng, Anne Wu, and Yoav Artzi. 2024. Retro-	of quantized llms. <i>arXiv preprint arXiv:2305.14314</i> .	855
805	spective learning from interactions. <i>arXiv preprint</i>		
806	<i>arXiv:2410.13852</i> .	Jinhao Duan, Renming Zhang, James Diffenderfer,	856
807	Javier Chiyah-Garcia, Alessandro Suglia, and Arash Es-	Bhavya Kailkhura, Lichao Sun, Elias Stengel-Eskin,	857
808	ghhi. 2024. Repairs in a block world: A new bench-	Mohit Bansal, Tianlong Chen, and Kaidi Xu. 2024.	858
809	mark for handling user corrections with multi-modal	Gtbench: Uncovering the strategic reasoning lim-	859
810	language models. In <i>Proceedings of the 2024 Confer-</i>	itations of llms via game-theoretic evaluations .	860
811	<i>ence on Empirical Methods in Natural Language</i>	<i>Preprint</i> , arXiv:2402.12348.	861
812	<i>Processing</i> , pages 11523–11542.		
813	Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Sheng-	Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff,	862
814	bang Tong, Saining Xie, Dale Schuurmans, Quoc V	Dan Jurafsky, and Douwe Kiela. 2024. Kto: Model	863
815	Le, Sergey Levine, and Yi Ma. 2025. Sft mem-	alignment as prospect theoretic optimization. <i>arXiv</i>	864
816	orizes, rl generalizes: A comparative study of	<i>preprint arXiv:2402.01306</i> .	865
817	foundation model post-training. <i>arXiv preprint</i>		
818	<i>arXiv:2501.17161</i> .	Raquel Fernández, Staffan Larsson, Robin Cooper,	866
819	Eve V. Clark. 2016. First language acquisition , 3rd ed	Jonathan Ginzburg, and David Schlangen. 2011. Re-	867
820	edition. Cambridge University Press, Cambridge.	ciprocal learning via dialogue interaction: Challenges	868
821	Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian,	and prospects . In <i>Proceedings of the IJCAI 2011</i>	869
822	Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias	<i>Workshop on Agents Learning Interactively from Hu-</i>	870
823	Plappert, Jerry Tworek, Jacob Hilton, Reiichiro	<i>man Teachers (ALIHT)</i> , Barcelona, Catalonia.	871
824	Nakano, Christopher Hesse, and John Schulman.		
825	2021. Training Verifiers to Solve Math Word Prob-	Zhaolin Gao, Wenhao Zhan, Jonathan D Chang, Gokul	872
826	lems . <i>arXiv preprint</i> . ArXiv:2110.14168 [cs].	Swamy, Kianté Brantley, Jason D Lee, and Wen Sun.	873
		2024. Regressing the relative future: Efficient pol-	874
		icy optimization for multi-turn rlhf. <i>arXiv preprint</i>	875
		<i>arXiv:2410.04612</i> .	876
		Ran Gong, Qiuyuan Huang, Xiaojian Ma, Hoi Vo, Zane	877
		Durante, Yusuke Noda, Zilong Zheng, Song-Chun	878
		Zhu, Demetri Terzopoulos, Li Fei-Fei, and Jianfeng	879
		Gao. 2023. Mindagent: Emergent gaming interaction .	880
		<i>Preprint</i> , arXiv:2309.09971.	881

882	Yuling Gu, Oyvind Tafjord, Hyunwoo Kim, Jared	Patricia K. Kuhl. 2007. Is speech learning 'gated' by	936
883	Moore, Ronan Le Bras, Peter Clark, and Yejin Choi.	the social brain? <i>Developmental Science</i> , 10(1):110–	937
884	2025. Simpletom: Exposing the gap between explicit	120.	938
885	tom inference and implicit tom application in llms.		
886	Leon Guertler, Bobby Cheng, Simon Yu, Bo Liu,	Komal Kumar, Tajamul Ashraf, Omkar Thawakar,	939
887	Leshem Choshen, and Cheston Tan. 2025. TextArena.	Rao Muhammad Anwer, Hisham Cholakkal,	940
888	<i>arXiv preprint</i> . ArXiv:2504.11442 [cs].	Mubarak Shah, Ming-Hsuan Yang, Phillip H. S.	941
889	Mustafa Omer Gul and Yoav Artzi. 2024. CoGen:	Torr, Salman Khan, and Fahad Shahbaz Khan.	942
890	Learning from Feedback with Coupled Comprehen-	2025. LLM Post-Training: A Deep Dive into Rea-	943
891	sion and Generation . In <i>Proceedings of the 2024 Con-</i>	soning Large Language Models . <i>arXiv preprint</i> .	944
892	<i>ference on Empirical Methods in Natural Language</i>	ArXiv:2502.21321 [cs].	945
893	<i>Processing</i> , pages 12966–12982, Miami, Florida,	Nathan Lambert. 2024. Reinforcement Learning from	946
894	USA. Association for Computational Linguistics.	Human Feedback . Online.	947
895	Kshitij Gupta, Benjamin Thérien, Adam Ibrahim,	Nathan Lambert, Jacob Morrison, Valentina Pyatkin,	948
896	Mats L. Richter, Quentin Anthony, Eugene	Shengyi Huang, Hamish Ivison, Faeze Brahman,	949
897	Belilovsky, Irina Rish, and Timothée Lesort. 2023.	Lester James V. Miranda, Alisa Liu, Nouha Dziri,	950
898	Continual pre-training of large language mod-	Shane Lyu, Yuling Gu, Saumya Malik, Victoria	951
899	els: How to (re)warm your model? <i>CoRR</i> ,	Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le	952
900	abs/2308.04014.	Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini,	953
901	Betty Hart and Todd R. Risley. 1995. <i>Meaningful differ-</i>	and 4 others. 2025. Tulu 3: Pushing Frontiers in	954
902	<i>ences in the everyday experience of young American</i>	Open Language Model Post-Training . <i>arXiv preprint</i> .	955
903	<i>children</i> . Paul H. Brookes Publishing Co.	ArXiv:2411.15124 [cs].	956
904	Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul	Jiatong Li, Rui Li, and Qi Liu. 2023. Beyond static	957
905	Arora, Steven Basart, Eric Tang, Dawn Song, and	datasets: A deep interaction approach to llm evalua-	958
906	Jacob Steinhardt. 2021. Measuring Mathematical	tion . <i>Preprint</i> , arXiv:2309.04369.	959
907	Problem Solving With the MATH Dataset .		
908	Sarah Hiller and Raquel Fernández. 2016. A data-driven	Tal Linzen. 2020. How can we accelerate progress	960
909	investigation of corrective feedback on subject omis-	towards human-like linguistic generalization? In	961
910	sion errors in first language acquisition . In <i>Proceed-</i>	<i>ings of the 58th Annual Meeting of the Asso-</i>	962
911	<i>ings of the 20th SIGNLL Conference on Computa-</i>	<i>ciation for Computational Linguistics</i> , pages 5210–	963
912	<i>tional Natural Language Learning</i> , pages 105–114,	5217, Online. Association for Computational Lin-	964
913	Berlin, Germany. Association for Computational Lin-	guistics.	965
914	guistics.	Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan	966
915	Arian Hosseini, Xingdi Yuan, Nikolay Malkin, Aaron	Duan, Ming Zhou, and Yue Zhang. 2023. Logiqa	967
916	Courville, Alessandro Sordoni, and Rishabh Agarwal.	2.0—an improved dataset for logical reasoning in	968
917	2024. V-STaR: Training Verifiers for Self-Taught	natural language understanding . <i>IEEE/ACM Trans-</i>	969
918	Reasoners .	<i>actions on Audio, Speech, and Language Processing</i> ,	970
919	Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan	31:2947–2962.	971
920	Allen-Zhu, Yuezhi Li, Shean Wang, Lu Wang,	Ziqiao Ma, Zekun Wang, and Joyce Chai. 2025. Babysit	972
921	Weizhu Chen, and 1 others. 2022. Lora: Low-rank	A Language Model From Scratch: Interactive Lan-	973
922	adaptation of large language models. <i>ICLR</i> , 1(2):3.	guage Learning by Trials and Demonstrations . In	974
923	Jennifer Hu, Sammy Floyd, Olessia Jouravlev, Evelina	<i>Proceedings of the 2025 Conference of the Nations</i>	975
924	Fedorenko, and Edward Gibson. 2023. A fine-	<i>of the Americas Chapter of the Association for Com-</i>	976
925	grained comparison of pragmatic language under-	<i>putational Linguistics: Human Language Technolo-</i>	977
926	standing in humans and language models . In <i>Pro-</i>	<i>gies (Volume 1: Long Papers)</i> , pages 991–1010, Al-	978
927	<i>ceedings of the 61st Annual Meeting of the Associa-</i>	buquerque, New Mexico. Association for Computa-	979
928	<i>tion for Computational Linguistics (Volume 1: Long</i>	tional Linguistics.	980
929	<i>Papers)</i> , pages 4194–4213, Toronto, Canada. Associ-	Kyle Mahowald, Anna A Ivanova, Idan A Blank, Nancy	981
930	ation for Computational Linguistics.	Kanwisher, Joshua B Tenenbaum, and Evelina Fe-	982
931	Zhijing Jin, Yuen Chen, Felix Leeb, Luigi Gresele,	dorenko. 2024. Dissociating language and thought in	983
932	Ojasv Kamal, Zhiheng Lyu, Kevin Blin, Fernando	large language models. <i>Trends in cognitive sciences</i> .	984
933	Gonzalez, Max Kleiman-Weiner, Mrinmaya Sachan,	Sabrina McCallum, Max Taylor-Davies, Stefano Al-	985
934	and Bernhard Schölkopf. 2023. CLadder: Assessing	brecht, and Alessandro Suglia. 2023. Is feedback all	986
935	causal reasoning in language models . In <i>NeurIPS</i> .	you need? leveraging natural language feedback in	987
		goal-conditioned rl. In <i>NeurIPS 2023 Workshop on</i>	988
		<i>Goal-Conditioned Reinforcement Learning</i> .	989
		Meta. 2024. The Llama 3 Herd of Models.	990

991	Filippo Momentè, Alessandro Suglia, Mario Giulianelli,	Matthew Saxton. 2000. Negative evidence and negative	1048
992	Ambra Ferrari, Alexander Koller, Oliver Lemon,	feedback: Immediate effects on the grammaticality	1049
993	David Schlangen, Raquel Fernández, and Raffaella	of child speech. <i>First Language</i> , 20(60):221–252.	1050
994	Bernardi. 2025. Triangulating LLM progress through		
995	benchmarks, games, and cognitive tests. <i>arXiv</i>	David Schlangen. 2019. Language tasks and language	1051
996	<i>preprint arXiv:2502.14359</i> .	games: On methodology in current natural language	1052
		processing research . <i>CoRR</i> , abs/1908.10747.	1053
997	Mitja Nikolaus and Abdellah Fourtassi. 2021. Mod-		
998	eling the interaction between perception-based and	David Schlangen. 2023. Dialogue games for bench-	1054
999	production-based learning in children’s early acqui-	marking language understanding: Motivation, taxon-	1055
1000	sition of semantic knowledge . In <i>Proceedings of</i>	omy, strategy . <i>CoRR</i> , abs/2304.07007.	1056
1001	<i>the 25th Conference on Computational Natural Lan-</i>		
1002	<i>guage Learning</i> , pages 391–407, Online. Association	John Schulman, Filip Wolski, Prafulla Dhariwal,	1057
1003	for Computational Linguistics.	Alec Radford, and Oleg Klimov. 2017. Proxi-	1058
		mal policy optimization algorithms. <i>arXiv preprint</i>	1059
1004	Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida,	<i>arXiv:1707.06347</i> .	1060
1005	Carroll Wainwright, Pamela Mishkin, Chong Zhang,		
1006	Sandhini Agarwal, Katarina Slama, Alex Ray, and 1	Amrith Setlur, Chirag Nagpal, Adam Fisch, Xinyang	1061
1007	others. 2022. Training language models to follow in-	Geng, Jacob Eisenstein, Rishabh Agarwal, Alekh	1062
1008	structions with human feedback. <i>Advances in neural</i>	Agarwal, Jonathan Berant, and Aviral Kumar. 2024.	1063
1009	<i>information processing systems</i> , 35:27730–27744.	Rewarding Progress: Scaling Automated Process Ver-	1064
		ifiers for LLM Reasoning .	1065
1010	Samuel J. Paech. 2023. Eq-bench: An emotional		
1011	intelligence benchmark for large language models .	Amrith Setlur, Nived Rajaraman, Sergey Levine, and	1066
1012	<i>Preprint</i> , arXiv:2312.06281.	Aviral Kumar. 2025. Scaling Test-Time Compute	1067
		Without Verification or RL is Suboptimal . <i>arXiv</i>	1068
1013	Dan Qiao, Chenfei Wu, Yaobo Liang, Juntao Li, and	<i>preprint</i> . ArXiv:2502.12118 [cs].	1069
1014	Nan Duan. 2023. Gameeval: Evaluating llms on		
1015	conversational games . <i>Preprint</i> , arXiv:2308.10032.	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu,	1070
		Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan	1071
1016	Rafael Rafailov, Archit Sharma, Eric Mitchell, Christo-	Zhang, YK Li, Y Wu, and 1 others. 2024. Deepseek-	1072
1017	pher D Manning, Stefano Ermon, and Chelsea Finn.	math: Pushing the limits of mathematical reason-	1073
1018	2023. Direct preference optimization: Your lan-	ing in open language models. <i>arXiv preprint</i>	1074
1019	guage model is secretly a reward model. <i>Advances in</i>	<i>arXiv:2402.03300</i> .	1075
1020	<i>Neural Information Processing Systems</i> , 36:53728–		
1021	53741.	Chandler Smith, Rakshit Trivedi, Jesse Clifton, Lewis	1076
		Hammond, Akbir Khan, Sasha Vezhnevets, John P.	1077
1022	Noam Razin, Sadhika Malladi, Adithya Bhaskar, Danqi	Agapiou, Edgar A. Duéñez-Guzmán, Jayd Matyas,	1078
1023	Chen, Sanjeev Arora, and Boris Hanin. 2024. Un-	Danny Karmon, Marwa Abdulhai, Dylan Hadfield-	1079
1024	intentional unalignment: Likelihood displacement	Menell, Natasha Jaques, Joel Z. Leibo, Oliver Slum-	1080
1025	in direct preference optimization. <i>arXiv preprint</i>	bers, Tim Baarslag, and Minsuk Chang. 2024. The	1081
1026	<i>arXiv:2410.08847</i> .	Concordia Contest: Advancing the Cooperative Intel-	1082
		ligence of Language Agents .	1083
1027	Philipp Sadler, Sherzod Hakimov, and David Schlangen.		
1028	2024. Sharing the cost of success: A game for evalu-	Alessandro Suglia, Ioannis Konstas, and Oliver Lemon.	1084
1029	ating and learning collaborative multi-agent instruc-	2024. Visually Grounded Language Learning: a re-	1085
1030	tion giving and following policies . In <i>Proceedings of</i>	view of language games, datasets, tasks, and models .	1086
1031	<i>the 2024 Joint International Conference on Comput-</i>	<i>Journal of Artificial Intelligence Research</i> , 79:173–	1087
1032	<i>ational Linguistics, Language Resources and Evalu-</i>	239. ArXiv: 2312.02431.	1088
1033	<i>ation (LREC-COLING 2024)</i> , pages 14770–14783,		
1034	Torino, Italia. ELRA and ICCL.	Theodore R Sumers, Mark K Ho, Robert D Hawkins,	1089
		Karthik Narasimhan, and Thomas L Griffiths. 2021.	1090
1035	Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavat-	Learning rewards from linguistic feedback. In <i>Pro-</i>	1091
1036	ula, and Yejin Choi. 2021. Winogrande: an adver-	<i>ceedings of the AAAI Conference on Artificial Intelli-</i>	1092
1037	sarial winograd schema challenge at scale . <i>Commun.</i>	<i>gence</i> , volume 35, pages 6002–6010.	1093
1038	<i>ACM</i> , 64(9):99–106.		
		Mirac Suzgun, Nathan Scales, Nathanael Schärli, Se-	1094
1039	Maarten Sap, Hannah Rashkin, Derek Chen, Ronan	bastian Gehrmann, Yi Tay, Hyung Won Chung,	1095
1040	Le Bras, and Yejin Choi. 2019. Social IQa: Com-	Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny	1096
1041	monsense reasoning about social interactions . In	Zhou, and Jason Wei. 2023. Challenging BIG-bench	1097
1042	<i>Proceedings of the 2019 Conference on Empirical</i>	tasks and whether chain-of-thought can solve them .	1098
1043	<i>Methods in Natural Language Processing and the</i>	In <i>Findings of the Association for Computational Lin-</i>	1099
1044	<i>9th International Joint Conference on Natural Lan-</i>	<i>guistics: ACL 2023</i> , pages 13003–13051, Toronto,	1100
1045	<i>guage Processing (EMNLP-IJCNLP)</i> , pages 4463–	Canada. Association for Computational Linguistics.	1101
1046	4473, Hong Kong, China. Association for Computa-		
1047	tional Linguistics.		

1102	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann	Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang,	1159
1103	Dubois, Xuechen Li, Carlos Guestrin, Percy Liang,	Haofei Yu, Zhengyang Qi, Louis-Philippe Morency,	1160
1104	and Tatsunori B Hashimoto. 2023. Alpaca: A	Yonatan Bisk, Daniel Fried, Graham Neubig, and	1161
1105	strong, replicable instruction-following model. <i>Stan-</i>	Maarten Sap. 2024a. SOTOPIA: Interactive Evalua-	1162
1106	<i>ford Center for Research on Foundation Models.</i>	tion for Social Intelligence in Language Agents. In	1163
1107	https://crfm.stanford.edu/2023/03/13/alpaca.html ,	<i>Proceedings of ICLR 2024</i> , pages 1–45.	1164
1108	3(6):7.		
1109	Alex Wang, Amanpreet Singh, Julian Michael, Felix	Yifei Zhou, Andrea Zanette, Jiayi Pan, Sergey Levine,	1165
1110	Hill, Omer Levy, and Samuel R Bowman. 2018.	and Aviral Kumar. 2024b. ArCher: Training Lan-	1166
1111	Glue: A multi-task benchmark and analysis platform	guage Model Agents via Hierarchical Multi-Turn RL.	1167
1112	for natural language understanding. <i>arXiv preprint</i>	<i>arXiv preprint</i> . ArXiv:2402.19446 [cs].	1168
1113	<i>arXiv:1804.07461</i> .		
1114	Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni,		
1115	Abhranil Chandra, Shiguang Guo, Weiming Ren,	Appendix	1169
1116	Aaran Arulraj, Xuan He, Ziyang Jiang, and 1 oth-	A Dialogue Games	1170
1117	ers. 2024. Mmlu-pro: A more robust and challenging		
1118	multi-task language understanding benchmark. <i>arXiv</i>	• Taboo: two-player game where one player	1171
1119	<i>preprint arXiv:2406.01574</i> .	gives a clue and the other player needs to	1172
		guess a target word based on the clue.	1173
1120	Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin		
1121	Guu, Adams Wei Yu, Brian Lester, Nan Du, An-	• Privateshared: two-player scorekeeping	1174
1122	drew M. Dai, and Quoc V. Le. 2021. Finetuned	game where an answerer agent goes through	1175
1123	Language Models Are Zero-Shot Learners. pages	a form with a questioner. The Gamemaster	1176
1124	1–46. ArXiv: 2109.01652.	keeps track of which information has been al-	1177
1125	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	ready shared. If a piece of information has	1178
1126	Bosma, Brian Ichter, Fei Xia, Ed H Chi, Quoc V Le,	not yet been shared, the answerer should not	1179
1127	and Denny Zhou. 2022. Chain-of-thought prompt-	expect the questioner to know it.	1180
1128	ing elicits reasoning in large language models. In		
1129	<i>Proceedings of the 36th International Conference</i>	• Imagegame: two-player instruction giving	1181
1130	<i>on Neural Information Processing Systems</i> , pages	and following game where one player de-	1182
1131	24824–24837.	scribes an image (a matrix in ASCII format)	1183
1132	Yue Wu, Xuan Tang, Tom M. Mitchell, and Yuanzhi Li.	and the other player draws (generates some	1184
1133	2024. Smartplay: A benchmark for llms as intelligent	matrix again in ASCII format) the image	1185
1134	agents. <i>Preprint</i> , arXiv:2310.01557.	based on the description. Then the constructed	1186
1135	Jiajun Xi, Yinong He, Jianing Yang, Yinpei Dai, and	image is compared with the drawn one. I also	1187
1136	Joyce Chai. 2024. Teaching embodied reinforcement	think we should consider if we wan	1188
1137	learning agents: Informativeness and diversity of lan-		
1138	guage use. In <i>Proceedings of the 2024 Conference on</i>	• Referencegame: two-player game where one	1189
1139	<i>Empirical Methods in Natural Language Processing</i> ,	player is given three different images (one	1190
1140	pages 4097–4114.	of them is selected as the target) and asked	1191
1141	Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao	to generated a referring expression that de-	1192
1142	Liu, Yuxiao Dong, and Jie Tang. 2023. Agenttuning:	scribes the target image by differentiating it	1193
1143	Enabling generalized agent abilities for llms. <i>arXiv</i>	from other two (distractors). Another player	1194
1144	<i>preprint arXiv:2310.12823</i> .	is then given the same images (orders are shuf-	1195
1145	Huaxiu Steven Zheng, Swaroop Mishra, Hugh Zhang,	flled) and asked to guess which image is the	1196
1146	Xinyun Chen, Minmin Chen, Azade Nova, Le Hou,	target based on the given referring expression.	1197
1147	Heng-Tze Cheng, Quoc V. Le, Ed H. Chi, and Denny	It is the only single-turn game in the bench-	1198
1148	Zhou. 2024. Natural plan: Benchmarking llms on	mark.	1199
1149	natural language planning.		
1150	Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer,	• Wordle: popular single-player game where	1200
1151	Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping	the task is to guess a 5-letter word. In each	1201
1152	Yu, Lili Yu, and 1 others. 2023a. Lima: Less is	turn, feedback is provided based on the place-	1202
1153	more for alignment. <i>Advances in Neural Information</i>	ments of characters in guessing attempts.	1203
1154	<i>Processing Systems</i> , 36:55006–55021.		
1155	Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha	• Wordle with clue: slightly changed version	1204
1156	Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and	of base wordle with the addition of a clue for	1205
1157	Le Hou. 2023b. Instruction-following evaluation for	the target word.	1206
1158	large language models. <i>Preprint</i> , arXiv:2311.07911.		

1207	• Wordle with critic: a two-player version	this way, the player explores the map. The	1256
1208	of the base wordle where the second player	player is required to stop exploring the map	1257
1209	(critic) provides feedback on the guesses of	on its own.	1258
1210	the first player.		
1211	• Codenames: a popular cooperative game with		1259
1212	two teams that try to uncover their agent’s	• Mapnavigation - specific room: changed ver-	1260
1213	code names (words) where one player de-	sion of the game above. The player is given	1261
1214	scribes clues that could strategically corre-	a target room (kitchen, garage, living room,	1262
1215	spond to more than one word. The team is	etc.) and is required to stop exploring once it	1263
1216	composed of Spymaster that provides clues	explores the target.	
1217	and Field operative that guesses.		
1218	• Adventuregame: single-player text adventure		
1219	game where a player is placed in a random	• Mapnavigation with graph reasoning:	1264
1220	location of an environment and given a task	change version of the base game where in	1265
1221	(pick up the flower and place it on the table).	each step the player is required to draw	1266
1222	The player explores the environment by giving	(generate JSON-like structure) the graph	1267
1223	commands (“go to the kitchen”, “open the	it has explored up to that point (besides	1268
1224	cupboard”, etc.) and the environment provides	providing the direction).	1269
1225	feedback where the command can be realised		
1226	and its outcome or not. The player explores		
1227	the world (with multiple rooms and objects in		
1228	them) and needs to decide on its own when to		
1229	stop the game.		
1230	• Guesswhat: two-player information seeking	B License for Artifacts	1270
1231	game where one player needs to guess the		
1232	target word (out of eight options) by asking	PLAYPEN’s implementation reuses important com-	1271
1233	questions about the target. The other player	ponents from clemcore, and all the games used	1272
1234	knows the target word and answers the ques-	come from clembench (Chalamalasetti et al., 2023).	1273
1235	tions with “yes” or “no”.	Therefore, we inherit its MIT License for PLAYPEN.	1274
1236	• Matchit: two-player game where each player	We acknowledge that the transcripts used for SFT	1275
1237	is given an image (ASCII representation, e.g.	and DPO training are derived from the clembench-	1276
1238	grid), which is not revealed to the other player,	runs ⁷ repo which has CC-By 4.0 License.	1277
1239	and asked to discuss what each player sees.		
1240	Players are allowed to ask questions to each		
1241	other about the image and provide responses		
1242	to the questions. At the end of the game, play-		
1243	ers need to decide whether they are given the	C Dialogue Games Transcripts	1278
1244	same or different images.		
1245	• Mapnavigation: single-player game where a	We include below transcripts for GuessWhat,	1279
1246	player needs to navigate the map by choosing	Taboo and Codename games where the baseline	1280
1247	the directions to go. Initially, the player is	model played the episodes of GuessWhat and	1281
1248	placed in a random room and given informa-	Taboo games and reaches the goal state (correct	1282
1249	tion where adjacent rooms are located (east,	guesses). However, it fails to follow instructions	1283
1250	west, north, south) and what is the name of	for Codenames episode and the game gets aborted.	1284
1251	the current room (kitchen, bathroom, garage,		
1252	etc.). The player gets the same feedback (adja-		
1253	cent rooms and the name of the current room)		
1254	whenever it makes a step. The player makes		
1255	a step by choosing the direction to go. Doing		

⁷<https://github.com/clembench/clembench-runs/blob/main/LICENSE>

C.1 GuessWhat game Transcript

#	A	Game Master
1		<p>[A]GM Let's play a guessing game. Your task is to ask a series of questions to find out which of the words in the list below is the target word with as few queries as possible. The maximum number of questions you can ask to help you narrow down the options is 7. Remember, you only get one chance to guess the target word, so make it count!</p> <p>The rules of the game are:</p> <p>1. Only ask one question per turn. DO NOT ASK MORE THAN ONE QUESTION.</p> <p>2. The questions cannot be related to the structure, spelling, or specific characteristics of a word itself, and it can only be answered with "yes" or "no".</p> <p>3. Use the following format for asking questions and DO NOT ADD ANY TEXT OTHER THAN THIS:</p> <p>QUESTION: <your question></p> <p>4. You have only one opportunity to guess the target word. To make your guess, use the following format and DO NOT ADD ANY TEXT OTHER THAN THIS:</p> <p>GUESS: <target word></p> <p>Remember you cannot ask for a hint. Make your question brief. Let's start!</p> <p>List of candidate words: ['Mackerel', 'Grouper', 'Carp', 'Catfish', 'Halibut', 'Flounder', 'Minnow', 'Goldfish']</p>
2		<p>[A]GM QUESTION: Is the target word a freshwater fish?</p>
3		<p>[GM]B Let's play a guessing game! Your task is to answer the other player's questions. Based on your knowledge of the word: Catfish, respond to the following questions or guesses. Limit your response to only 'yes' or 'no' with no explanation or other words. Never reveal the answer in your response.</p> <p>You must reply using the format below and DO NOT ADD ANY TEXT OTHER THAN THIS:</p> <p>ANSWER: <some text></p> <p>Target Word: Catfish</p> <p>QUESTION: Is the target word a freshwater fish?</p>

#	B
4	[GM]B ANSWER: Yes
5	[A]GM QUESTION: Is the target word commonly kept as a pet?
6	[GM]B ANSWER: Yes
7	[A]GM QUESTION: Is the target word relatively small in size?
8	[GM]B ANSWER: No
9	[A]GM QUESTION: Is the target word native to North America?
10	[GM]B ANSWER: Yes
11	[A]GM QUESTION: Is the target word a scaleless fish?
12	[GM]B ANSWER: Yes
13	[A]GM QUESTION: Does the target word have whisker-like barbels?
14	[GM]B ANSWER: Yes
15	[A]GM GUESS: Catfish
16	[GM]GM CORRECT GUESS
17	[GM]GM End game.

C.2 Taboo game Transcript

#	A	Game Master	B
1		<p>[A]GM You are playing a collaborative word guessing game in which you have to describe a target word for another player to guess.</p> <p>Rules:</p> <p>(a) You have to reply in the form: CLUE: <some text>. Guesses from the other player will start with GUESS.</p> <p>(b) You cannot use the target word itself, parts or morphological variants of it in your description.</p> <p>(c) In addition, the same rules apply for related words which are provided below.</p> <p>End conditions:</p> <p>(i) If you use the target word or a related word in your description, then you lose.</p> <p>(ii) If the other player can guess the target word in 3 tries, you both win.</p> <p>Let us start. This is the target word that you need to describe and that the other player needs to guess:</p> <p>ordinary</p> <p>Related words are:</p> <ul style="list-style-type: none">- common- normal- plain <p>Important: You are under time pressure, give short descriptions that are to the point!</p>	

1291

1292

2 [A>GM] CLUE: Something that is usual or expected.

[GM>B] You are playing a collaborative word guessing game in which you have to guess a target word that another player describes to you.

You can make one guess at each trial. You win when you guess the target word. You lose when you cannot guess it in 3 tries.

3 After each trial you will get a new hint from the other player which starts with CLUE.

Make your guesses by just saying the word using the following form: GUESS: <a word>

Let us start.

CLUE: Something that is usual or expected.

4 [GM>B] GUESS: Norm

5 [A>GM] CLUE: Not fancy or special.

6 [GM>B] GUESS: Ordinary

7 [GM>GM] CORRECT GUESS

8 [GM>GM] End game.

C.3 Codenames game Transcript

1293

#	A	Game Master	B
		[A>GM] Provide one single-word clue that relates to some of your team's words from the list below. You can choose to how many and to which words your clue relates to. The goal is to have your partner guess all of your team's words based on the clues you give them, before the other team has uncovered all of theirs. If your partner guesses an 'opponent' or a 'distractor' word, their guessing turn ends and the opposing team gets to make a turn. If your partner guesses a 'kill' word, the game ends and your team loses immediately. The more words your clue relates to, the quicker you can win the game, but the harder the associations may be to guess for your partner, so choose your clue word wisely. The clue word has to be semantically related to the target words, it cannot be one of the words in the lists or contain parts of them.	
1		Always give your single-word clue and your comma-separated list of related target words in the following format and make your answers as short as possible, never include any other text than is required in this form:	
		CLUE: <WORD> TARGETS: <TARGETS>	
		Your team words are: hail, neck, mountain, poland, sister, couch, stool, captain, officer. Your opponent's team words are: hair, teeth, england, sergeant, table, fog, desert, rye. Distractor words are: daughter, recliner, private, son, cousin, thigh, ireland. The kill words are: liquor	
2		[A>GM] CLUE: storm TARGETS: hail, mountain	
		[GM>B] Provide a comma-separated list of up to 2 words from the following list that best relate or are most closely associated with the word 'storm'. Always start your list of guess(es) with 'GUESS: ' and do not include any other text in your answer.	
3		sister, hair, hail, couch, table, desert, recliner, poland, daughter, thigh, teeth, cousin, stool, sergeant, private, england, son, fog, ireland, liquor, neck, officer, rye, mountain, captain	

1294

4 [GM(B)] GUESS: hail, storm
 5 [GM|GM] Guessed word **storm** is the
 6 [GM|GM] Game aborted.

D Training Data Generation

D.1 Supervised Fine-tuning

Data Collection In a first step, the data had to be extracted from the clembench-runs repository and transformed into a format that allowed further filtering. The clembench-runs repository is a collection of all generated interaction data produced by the models evaluated on the benchmark⁸. The extraction included the removal of all game-master specific logging and parsing prompts that were not part of the conversation pieces that were directly prompted to the model. Besides the conversational data, also the information about the model, played game, benchmark-version, experiment, and episode number were extracted as well as the information of whether the episode was a successfully played, lost or aborted episode.

Data Filtering For the supervised fine-tuning, only successful episodes were considered for training data. Lost and aborted episodes were discarded.

Data Transformation After filtering, some game-specific data transformations were added to mitigate changes between the different benchmark versions and to improve training performance. Most of the transformations were necessary due to changes in the prompts between the benchmark versions and changes in the parsing rules for model answers. All of the transformations are listed below with a reasoning for the necessity of change:

Imagegame While clembench versions 0.9 and 1.0 allowed the player to add "*what is your next instruction*" to its answers, the same behavior led to parsing errors in clembench version 1.6, which resulted in abortion of all image game episodes. To address the problem caused by the new parsing rule, all model answers from player-1 were truncated to only contain the correct format required by clembench version 1.6. This change was necessary, because when training with the original data, the model learns to always include additional information, such as "*what is your next instruction*". And even though the model correctly answers the game instructions, all episodes would be aborted.

Before: Instruction: Put a B in the first column of all rows
 what is your next instruction
 After: Instruction: Put a B in the first column of all rows

Wordle games For the three wordle game variants, there existed a few successfully played episodes that contained an "*INVALID_FORMAT*" token inside the prompts. These instances were removed since the model should not reproduce outputs with invalid format. Further, all wordle games have a mechanic to reprompt the model when e.g. the target word has five letters but the model produces a six-letter word as a guess. In these cases, the model is prompted with the instruction to guess a word with exactly five letters.

To mitigate the models learning to produce output that leads to re-prompting, all re-prompted actions were filtered to not be part of the training data. The episodes were not excluded completely but only the model answer that led to the re-prompting as well as the prompted text were excluded.

Referencegame For referencegame, the initial prompt between the older clembench versions (0.9 and 1.0) completely changed in version 1.6. While the older versions contained multiple examples (few-shot prompting), in version 1.6 there are no examples present. While this alone would not be a problem, the examples in the old clembench versions directly implied a strategy on how to play the game and describe the grids (see following snippet from the old version of the prompt).

Here is an example with grids.
 The first grid is the target grid and the following two grids are distractors.

Target grid:

X X X X X
 O O X O O
 O O X O O
 O O X O O
 O O X O O

...

The referring expression for the given target grid is like so:

Expression: Filled as T.

Here, the model is directly instructed to describe the whole grid as a letter or shape. The problem found in this strategy is that not all grids follow

⁸<https://github.com/clembench/clembench-runs>

this pattern and even for humans it is not trivial to describe all of the grids. Combined with the fact that reference game is a two-player game and the second player has to guess the described grid out of three presented grids, meaning that despite a bad description of player-1, player-2 has a 33% chance of guessing correctly.

This resulted in a situation where about 53% of the successful episodes, player-1 described the target grid as "Filled as T" while except for the prompt example, there is no grid in the data that is actually shaped as a T.

This meant that reference game data from the old benchmark versions could not be used for the training process due to the low quality. To mitigate this problem, data from clembench version 1.6 was used in the training data while 20-30% of the episodes of each experiment were held out for testing.

Privateshared For privateshared, after the first experiments it appeared that in most cases, the trained model answered with "*ASIDE: No*" to all probe-questions. In privateshared, the model should act like its e.g. a customer at a travel agency that wants to travel. The agent asks question about destination, time and other related properties of the inquiry. Bit by bit the model has to tell the agent all the information the agent needs. After every question there is a block of probing questions where the model is asked whether or not specific information has been shared already and the model has to answer with "*ASIDE: yes*" or "*ASIDE: no*".

Considering the structure of privateshared, the model does have to answer with "*ASIDE: no*" to all probe-questions in the beginning of the game dialogue which changes to more and more "*ASIDE: yes*" during the course of the game play depending on what information has already been shared.

This observation led to the following transformation: Instead of including all probe-questions answer pairs as individual samples in the dataset, all probe questions were used together in one sample. This reduced the amount of probing-question blocks within the privateshared data to prevent over-fitting on answering all probe-questions with "*ASIDE: no*".

Instead of adding six to ten samples for each probing block all containing answer "*ASIDE: no*", now there is only one sample containing all probing questions of a probing-block.

It has to be mentioned that these specific changes privateshared, imagegame and referencegame were

partially detected through experiments that will be explained further below. For all other games, no data transformation was needed.

Iterative Data-Processing While some of the previously described data transformations were motivated by observations during the data preparation and collection phase, different experiments were conducted to iteratively improve the data to optimize the model fine-tuning performance.

The findings of the experiments were combined into a final data set that contains all improvements discovered during the experiment phase.

In total, more than 30 different experiments were conducted with different dataset configurations. The experiments were structured in nine main experiments. Each main experiment has one or multiple sub-experiments. In the following, the main experiments will be explained, reasoned and most important findings will be stated.

D1 Contains, as an initial experiment, all successfully played episodes of all models. The dialogues are not processed in any way and just parsed into the model-specific chat-templates.

D2 Contains only successfully played episodes from the top k models. The tier list was derived from having the most successful episodes. This directly reflects the models with the best clemcores.

It is to be expected that the quality of the played episodes from better models is higher than the models that only succeeded in a small number of episodes. The idea behind this experiment was to determine whether the difference in quality is reflected by the fine-tuned model.

It appears that training only on the successfully played episodes of the top 10 models has a positive impact on the quality score compared to using all available data.

D3 In the previous experiments, a training sample consisted of a complete episode. This means, the whole conversation over multiple turns was served as one sample. This implies that intermediate turns were not available as individual training samples in the data.

This experiment was designed to determine the impact of using individual conversation pieces as training samples rather than the whole conversation at once.

Therefore, every episode was split into individual continuously growing training samples that started with the first question answer pair and was then extended with each question answer pair until

the end of the conversation was reached.

This means the first conversation piece contains only the initial instruction and first answer while the last piece contains the whole conversation held in one episode.

The configurations of experiments D1 and D2 were repeated, but the data was split into conversation pieces. It is important to note that the data was shuffled before splitting to ensure that the conversation parts of one episode remain in the correct order and are trained on together.

It was shown that most of the experiments from D3 outperformed their respective experiments from D1 and D2.

D4 Was conducted to test different balancing strategies. In the previous experiments, the data was not balanced between games. Data can be balanced before or after splitting the conversation parts (as described in D3). The downsampling can be done by random selection or by considering the leader board positions of the models. Further, there can be oversampling for games with only few available episodes. While balancing overall showed a positive impact, the best performance was achieved when the data was balanced before splitting and the sampling was based on the leader board without oversampling. This was also shown in D2, where using the data from the best performing models showed a positive impact on the fine-tuned models' performance.

D5 & D6 These were two complementary experiments where for D5 the model is only trained on one game while D6 consists of the opposite experiment and can be described as leave-one-game-out. While this experiment did not yield meaningful insights in the dataset-configuration, it led to some improvements described previously. As an example, it provided insights into overfitting of the probe questions in privateshared which led to a different treatment of the probing blocks.

D7 & D8 Consisted of experiments that meant to verify or reject the improvement of game-specific data-transformations. This includes the transformations on privateshared and referencegame.

D9 While D1-D8 were completely focused on the data, D9 comprises a hyperparameter tuning of the QLoRA parameters.

Final Dataset Overview: As can be seen in table 3, depending on the game, the number of samples after splitting varies heavily. Eventually the total number of samples available for the training is

Game	Samples Train	Samples Test
Before Splitting		
taboo	434	18
referencegame	324	36
wordle_withcritic	302	12
wordle_withclue	295	5
imagegame	278	12
wordle	230	5
privateshared	214	5
After Splitting		
taboo	560	22
referencegame	324	36
wordle_withcritic	1192	105
wordle_withclue	717	12
imagegame	1579	52
wordle	1038	19
privateshared	1669	45
total	7079	291

Table 3: Final Dataset Constellation Before and After Splitting.

about 7000 while the number of samples for evaluation during the training is about 300.

D.2 Synthetic data set for warm start and rehearsal training regimes

In an attempt to enforce instruction following capabilities when fine-tuning, we designed a synthetic data set to use for training regimes such as warm-up training and rehearsal training. In contrast to the data derived from clembench runs, this data consists of mostly single turn user-assistant interactions and was completely programatically created with the help of human-made templates. Thus, both the user input as well as the assistant output were created without a language model in the loop. These templates consist of short representations of instruction-answer-interactions, what we call 'minigames'. Such a minigame could for example be:

```
[{"role": "user", "content": "Sum these numbers: 14, 26, and give the answer after the tag SUM:"}, {"role": "assistant", "content": "SUM: 40"}]
```

We included a number of minigames based on letters/words, numbers, transforming inputs into json format, and making choices between different options. What mattered more than the content of the artificial interactions was the form. Since we wanted to focus on enhancing instruction following capabilities and learning of certain input/output formats (e.g. "start your response with ...", "Do not produce anything else"). We achieved this by keeping model answers short, using a variety of

formatting instructions, most importantly a variation of “Start your answer with \$KEYWORD”, where the keywords were varied in kind and format and using a variety of different inputs and outputs. We also included one multiturn game where the expected answer was a choice from items from a list that were not chosen in previous turns. The goal of this was to enhance attention to the overall context. With different combinations of fillers for the open slots for the templates of 26 different minigames, we had about 20000 datapoints. The two training subsets (warmup and rehearsal) with 400 and 900 datapoints are balanced by type of minigame with 16 and 36 random samples per game, respectively⁹.

D.3 DPO

DPO requires paired preference data, i.e., samples sharing the same context before positive vs negative continuations. For DPO Dialogue and DPO Turn with clembench runs, positive continuations are obtained from successful games’ interactions, and negative continuations from unsuccessful and aborted games’ interactions. Since the SFT models obtained top performances in % Played for all games except Wordle and its variants, we only integrated aborted interactions for this game (and variants). Unsuccessful and aborted interactions have been collected and transformed from the same sources and with the same procedure as SFT data. For multi-player games like Taboo, data for both player 1 (i.e., giving clues) and player 2 (i.e., making guesses) have been integrated into training. The DPO Dialogue dataset consists of around 10K samples as in Fig. 3; the DPO Turn dataset consists of around 86K samples as in Fig. 4.

For DPO Dialogue we experiment with two variables: the number of negative samples per positive sample and the model source for negative samples. Each positive interaction is paired with n unsuccessful – and n aborted interactions for Wordle and variants – where n is manipulated to find the optimal number of negative trajectories to learn from. The source for negative samples falls into three conditions: negative samples from all the models, from only the 10 best-performing models, and only models from the same family as the model to train. The best performances, as tested on clembench version 1.6, have been obtained by coupling 1 negative per positive from the 10 best-performing models.

⁹The full data set as well as the two subsets can be found here: It will be released after the review process.

Game	DPO Dialogue	DPO Turn
taboo	4.689	20.074
referencegame	1.712	2.056
wordle_withcritic	1.469	18.234
wordle_withclue	677	21.380
imagegame	1.291	12.094
wordle	285	13.788

Table 4: DPO datasets’ samples per game.

D.4 GRPO

Reward Functions. We employed various reward functions, depending on the training game in question (see Table 5): these reward functions were intended to reflect the quality score computation for each game in the Clembench benchmark. Wordle (including the withclue and withcritic variants) and Referencegame utilized a simple 0/1 loss function: $r_g = 1$ if the agent reached the correct answer within the turn limit (six and one turns, respectively), and $r_g = 0$ otherwise.

Game	Reward Function
Taboo	$(1/\text{len}(\text{turns})) \cdot \mathbb{I}_{\text{corr}}$
Wordle (+withclue, +withcritic)	\mathbb{I}_{corr}
Referencegame	\mathbb{I}_{corr}
Imagegame	$F_1(G_{\text{pred}}, G_{\text{target}})$
Privateshared	$\text{Acc}(S_{\text{pred}}, S_{\text{target}})$

Table 5: GRPO reward functions by training game. $\mathbb{I}_{\text{corr}} = 1$ if the agent reached the correct answer within the maximum number of turns, and $\mathbb{I}_{\text{corr}} = 0$ otherwise.

For Imagegame, r_g is given by the F_1 score between the agent’s predicted grid and the target grid, and r_g is given by the agent’s slot-filling accuracy for Privateshared. The reward function for Taboo incorporates a length penalty: $r_g = 1/n$ if the agent guesses the correct word within $n \leq 3$ turns, and $r_g = 0$ otherwise.

For all seven games, the $r_g = 0$ if the game was aborted due to agent error, such as incorrect formatting or violation of the game’s rules (e.g. predicting a non-five-letter word in Wordle).

Teacher Model. In the offline learning experiments conducted in this paper (Section 4.3.1 and Section 4.3.2), the models were trained to play both roles in all two-player games: this was not practical for the online RL experiment. If the model is

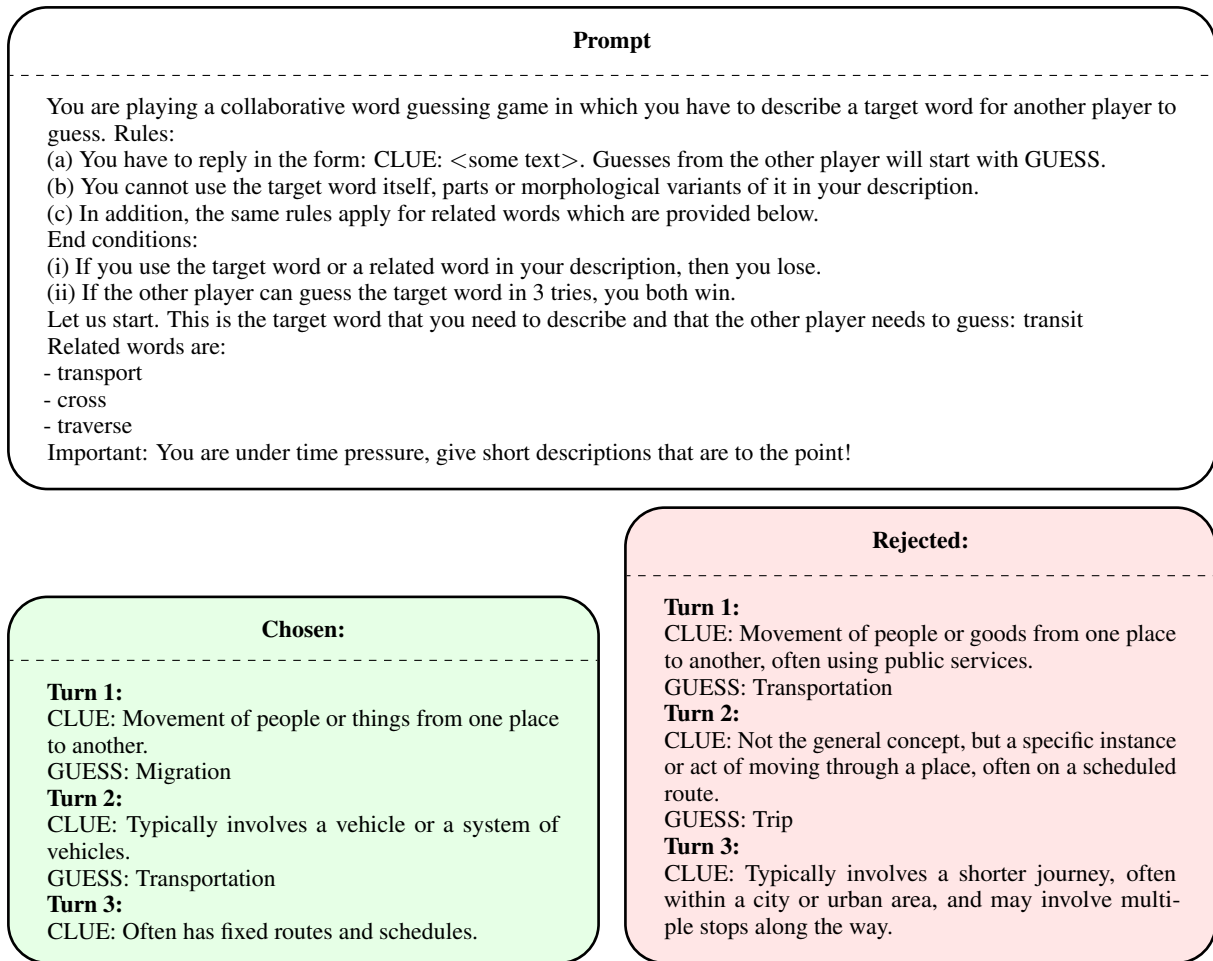


Figure 3: **DPO dialogue** dataset: the initial state (prompt and game instance) is shared, the chosen and rejected continuations are the remaining turns from the successful and unsuccessful episodes.

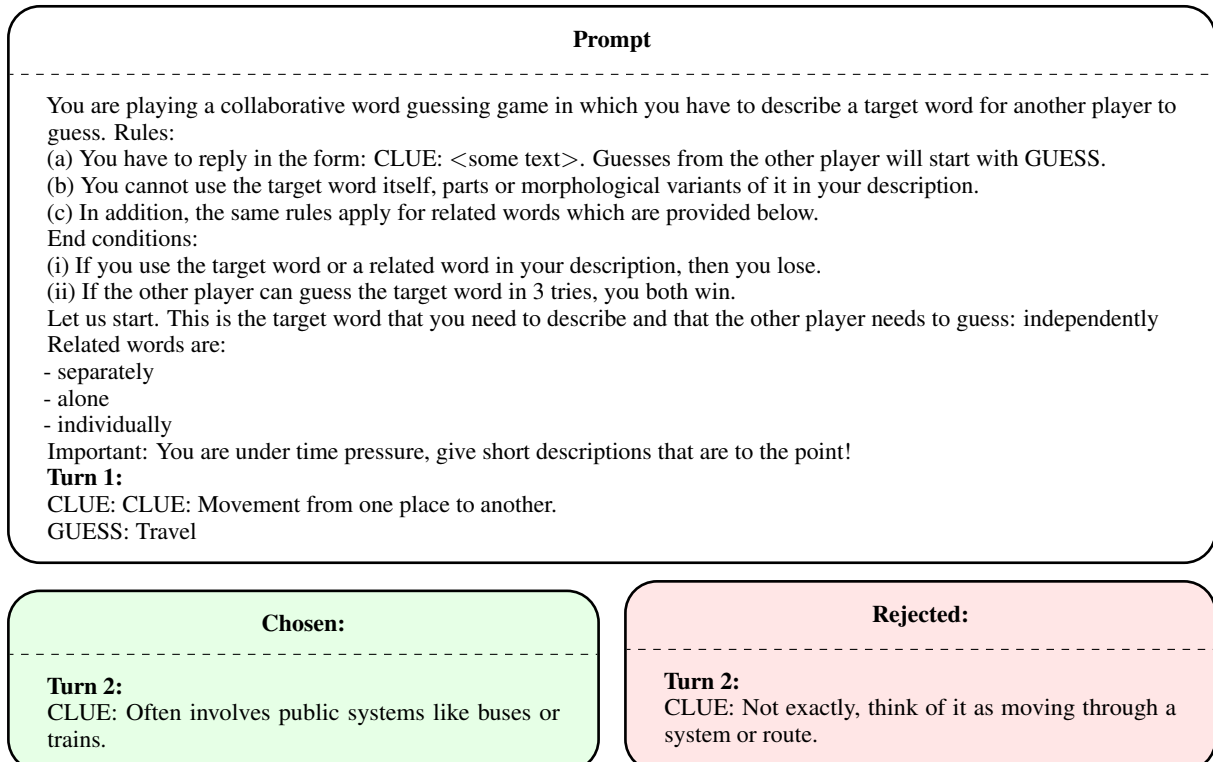


Figure 4: **DPO turn** dataset: the initial state and conversational history are shared, the chosen and rejected continuations are single turns from the successful and unsuccessful episodes.

trained in both roles, approximation of the optimal policy is intractable: updates to the current policy are dependent on the reward, which is dependent on the environment, which in turn is dependent on the current policy (via the generations of the current policy playing in the other role).

On the other hand, if we employ a *different*, frozen model as the second player (the *teacher* model), approximation of the optimal policy becomes tractable. However, evaluating the agent model against itself (i.e. playing both roles) introduces a mismatch between the train and test splits: the agent will have approximated the optimal policy for the environment in which the teacher is the second player.

Despite this train-test mismatch, we trained the GRPO agent with GPT-4o-mini¹⁰ as the teacher model: in line with our goal of *learning* from interaction, we employed a more advanced model than our agent (Llama-3.1-8B), to enable the agent to learn from its teacher/caregiver. A list of the two-player games—and the roles played by the teacher and agent in each—is given in Table 9.

Challenges and Adaptations of the Playpen Environment. We adapted the Playpen environment to online RL applications by re-configuring Playpen to allow individual game instances to be played separately: this allows for the tuning of batch size as a hyperparameter, and the random permutation of game instances across batches. We additionally implemented non-agent token masking, so that the agent’s loss is only computed with respect to its own generated tokens.

Teacher-Aborted Episodes In the Playpen environment, an episode can be aborted if there is a rule violation from either the agent or teacher model: for example, if the teacher model includes the target word in its clue during a Taboo game instance.

In the case of teacher error, the agent model should not be negatively rewarded due to the aborted episode. To account for this, we set a *retry limit* ρ , such that a teacher-aborted episode will be replayed up to ρ times in the case of teacher error¹¹.

If a single instance $g \in G_i$ has been aborted ρ times due to teacher error, we replace g with an-

other randomly-selected $g' \neq g \in G_i$ from the same group for loss computation and backpropagation. If *every* instance in the group G_i is aborted ρ times due to teacher error, we replace G_i with another group $G_{k \neq i}$ in the same batch.

Privateshared The Privateshared game was particularly problematic for online RL, as the quality score for this game is primarily computed from probes that are conducted adjacent to the actual game, and the transcripts from these probes are removed from the agent’s observations after they are completed.

Including the probes in the instance trajectory during training results in a mismatch between the train and test splits, as the agent only sees the *current* probe at test time. Conversely, removing the probes from the trajectory leads to unpredictable rewards from the environment: if the agent’s reward is negatively affected by its performance in a probe, the reason for the negative reward will not be reflected in the trajectory.

For these reasons, we made no reference to the agent’s probing-task performance in the computation of the Privateshared reward function. This has a severe negative effect on test-set performance for this game: online RL substantially degrades the model’s quality score for Privateshared (see Appendix F.3), even when beginning online RL from the SFT Llama model.

E Training details

E.1 Supervised Fine-tuning

The SFT models are fine-tuned using QLoRA (Dettmers et al., 2023) adapters ($r = 64$, $\alpha = 32$, $dropout = 0.05$) on all linear layers. The models were trained with the following arguments ($optim = adamw_8bit$, $lr = 2e - 4$, $lr_scheduler = linear$, $decay = 0.01$, $batch_size = 4$, $steps = 600 - 700$ and fixed $seed = 7331$). The models were quantized in 4-bit using the *unsloth* (Daniel Han and team, 2023) library and the following bits-and-bytes configuration ($use_4bit = True$, $bnb_4bit_compute_dtype = float16$, $bnb_4bit_quant_type = nf4$, $use_nested_quant = False$). As a stopping criterion, the first checkpoint before the minimal eval-loss that has a distance of less or equal than 0.015 from the minimal eval-loss was chosen. Hence a full epoch must be trained to

¹⁰<https://platform.openai.com/docs/models/gpt-4o-mini>

¹¹In practice, we set $\rho = 1$ for all experiments due to computational resource limitations.

determine the optimal checkpoint. The most important libraries and their versions are (*torch* = 2.4.0, *unsloth* = 2024.8, *transformers* = 4.47.1, *bitsandbytes* = 0.43.3, *trl* = 0.9.6, *accelerate* = 0.34.2).

Training Setup All previously described experiments were conducted on a quantized version of Llama-3.1-8B (instruct). All models were fine-tuned using Unsloth (Daniel Han and team, 2023) with 4-bit quantization and QLoRA (Detmers et al., 2023) for a more efficient and resource-optimized fine-tuning.

Hardware The training was conducted on one NVIDIA A100 GPU with 80 GB of VRAM and one NVIDIA H100 GPU with 95 GB of VRAM. It must be noted that technically a multi-GPU setup was possible, but every experiment was only conducted on a single-GPU setup.

Training Procedure In the first step, the models were trained on all available training data. Based on the training statistics (train and evaluation loss), a second model was trained using the number of steps with the lowest evaluation loss.

To address the issue of over-fitting, a third model was trained using significantly fewer steps. The number of steps was chosen based on the evaluation loss, with a threshold set to 0.015. The third model was trained until the evaluation loss reached a value within or equal to this threshold relative to the best evaluation loss.

As an example: The first model is trained for 1700 steps (all available data), but the minimal evaluation loss is reached at around 1100 steps. With a minimal eval-loss of 0.2315, the second model is trained for 1100 steps while the third model is trained until the eval-loss reaches the threshold of $0.2315 + 0.015 = 0.2430$. The final model required 700 training steps. This approach helps prevent over-fitting, as continuing training beyond the threshold (where evaluation loss increases by 0.015) provides diminishing returns while potentially reducing generalization capabilities. This is a strategy we call SFT (Cold Start, CS) in the main paper.

We also experimented with other variants that we report below:

- SFT (Warm Start, WS): Before training on the interaction data, the model was trained on 100 steps (400 samples) of synthetic instruction following tasks, using the findings from

(Gupta et al., 2023) and focusing on instruction following abilities.

- SFT (Rehearsal, R): during training, we interleave the gameplay training dataset with basic instruction following data following a similar approach to Lambert et al. (2025).

E.2 DPO

For both DPO-dialogue and DPO-turn, an SFT QLoRA adapter has been mounted on top of the base model Llama3.1-8B. To merge the base model and the SFT adapter, three merging strategies have been tested before DPO: merging the full-precision Llama3.1-8B model with the adapter, merging the unsloth 4-bit quantized Llama3.1-8B version with the adapter in 16bit, and merging the unsloth 4-bit quantized Llama3.1-8B with the adapter in 4bit. As reported in Tab. 6, the first strategy outperforms the others, showing comparable results to the unmerged adapter.

Model	ClemScore	pp	qs
unmerged	46.82	75.24	62.23
full-precision	47.79	74.88	63.82
16bit	33.52	70.19	47.76
4bit	30.14	60.00	50.23

Table 6: Comparison of merging strategies in terms of Clemscore, average % played (pp) and quality score (qs).

DPO training is performed on top of the 4-bit quantized SFT model, with the same bits-and-bytes configuration as the SFT models (*use_4bit* = *True*, *bnb_4bit_compute_dtype* = *float16*, *bnb_4bit_quant_type* = *nf4*, *use_nested_quant* = *False*). QLoRA adapters are employed on the same modules as for SFT (with *r* = 64, *α* = 64, and *dropout* = 0). The models have been trained with the *adamw_8bit* optimizer, a learning rate of $5e - 6$, with *linear* lr scheduler and the *beta* = 0.1 (*decay* = 0, *batch size* = 2, *gradient accumulation steps* = 3 and fixed *seed* = 42). During training, we evaluate the model every 20% on held-out training samples. At the end of training, only the best-performing checkpoints on the dev sets were saved. The libraries versions are: *torch* = 2.5.1, *unsloth* = 2024.12.4, *transformers* = 4.46.3, *bitsandbytes* = 0.45.0, *trl* = 0.12.2, *accelerate* = 1.2.0.

In terms of hardware, DPO development has been performed on 2xA5000s. Large differences

have been observed when comparing results obtained on the A5000 and A100. The final training for clembench v2.0 has been performed on an RTX3090 with 24GB RAM.

E.3 GRPO

We conducted two online RL experiments: one pure RL experiment, in which we initialized the agent from the baseline Llama-3.1-8b model (GRPO); and a second experiment in which the RL agent was initialized from the best-performing SFT model (SFT(CS)+GRPO).

The training set for both experiments consisted of game instances from Clembench V0.9 and V1.0 for Taboo (90 instances), Wordle (60), Wordle-withclue (60), Wordle-withcritic (60), Referencegame (256), Imagegame (80), and Private-shared (80), for a total of 686 instances. The validation split consisted of 420 Clembench V1.6 game instances (total) for the training games.

Both GRPO models were trained on four NVIDIA H100 GPUs with 80 GB of VRAM (each): for speedup, trajectory generation was parallelized across the four GPUs.

Both GRPO models were tuned using LoRA (Hu et al., 2022) adapters ($r = 64$, $\alpha = 128$, $dropout = 0$) on their Q , K , V , and O attention projection matrices. We trained the models for five epochs on 686 game instances with a temperature of 0.75, a batch size of 16, a group size of 8, KL regularization $\beta = 0.04$, and a learn rate of 10^{-6} using the Adam optimizer (for GRPO $seed = 250329152534053703$, for SFT(CS)+GRPO $seed = 250327114458100881$).

E.4 Evaluation Details

The evaluation of the models on non-interactive datasets have been conducted on Ampere-architecture GPUs (A100, A40). The experiments have been conducted by extending the *lm-eval* framework with the tasks which were not present in the original version (i.e. CLadder, LM-Pragmatics, NATURAL PLAN, GLUE Diagnostics, SimpleToM). Out of these, CLadder, NATURAL PLAN and SimpleToM have been taken without any modification from those implemented by (Momentè et al., 2025). NATURAL PLAN has been also taken from there, but the upper and lower bound to the amount of token that the model was allowed to generate have been removed. GLUE Diagnostics have been implemented from scratch.

To ensure comparability of the results, all evaluations on clembench v2.0 were carried out exclusively on a H100 GPU. It appears that when using different GPUs, the results can differ by up to 5 percentage points in some models. The H100 was chosen due to its higher inference speeds to save time on evaluation.

F Results

F.1 Supervised Fine-tuning

Table 7 depicts the clembench version 2 results for the baseline Llama-3.1-8B as well as the three variants cold-start (cs), warm-start (ws), and rehearsal (r). The games are separated by their domain. The upper half of the table represents the in-domain games while the lower half contains the out-of-domain games.

It becomes visible that in-domain all three versions appear to have quite substantial performance gains. Out-of-domain on the other hand, for most games, all three models show substantial performance decreases.

Comparing the three fine-tuning versions, it appears that overall, the Llama-3.1-8B with a cold-start outperforms the other two models. While ws and r come along with slight decreases for private-shared and referencegame (in-domain games) that all models were fine-tuned on, the cs version is the only one that shows continuous improvements for all in-domain games.

For the out-of-domain games, the performance in- and decreases shifted between the models but some patterns (e.g. codenames and tmw improvements) still remain. A larger discrepancy can be seen for textmapworld-graphreasoning (tmw_graph). Here, ws and r seem to negatively impact the model performance. tmw_grap is the only game that requires the model to produce a valid JSON-object. This game in particular it is crucial to output the exact required output since a malformed JSON leads to an aborted game.

Interestingly, Llama-3.1-8B cs shows out-of-domain a more substantial decrease in the % played score than the quality scores. For the other two models it appears that the performance loss is more balanced across the two scores. For Llama-3.1-8B cs this indicates, that the fine-tuning negatively impacts the models ability to properly play the game. The % played is an indicator of what % of episodes were actually played and how many were aborted. This is tightly bound to game specific output for-

Game/Model	L3-8B (Baseline)	PP L3-8B CS	PP L3-8B WS	PP L3-8B R
<i>In Domain</i>	pp/qs	pp/qs	pp/qs	pp/qs
imagegame	67.8/54.62	32.20/39.87	32.20/37.65	32.20/39.19
prvtshared	100/23.48	0.00/73.65	0.00/73.05	-4.00/69.16
referencegame	100/38.89	0.00/7.78	0.00/-3.33	0.00/-4.45
taboo	98.33/31.92	1.67/5.58	1.67/3.91	1.67/9.19
wordle	36.67/0	20.00/1.18	30.00/5.00	16.66/8.12
wordle_wcl	0/-	23.33/71.43	10.00/16.67	6.67/0.00
wordle_wcr	6.67/50	6.66/0.00	-3.34/50.00	-3.34/-16.67
<i>Out-of-Domain</i>	pp/qs	pp/qs	pp/qs	pp/qs
adv-game	35.94/33.85	-17.97/-18.23	-15.63/-26.93	-15.36/-17.70
codenames	43.08/16.07	-17.70/5.14	-26.93/17.26	-17.7/5.14
tmw	36/55.46	32.00/-8.75	24.00/-2.31	32.00/-0.11
tmw_graph	20/44.33	-3.33/-7.54	-16.67/-15.76	-13.33/-13.38
tmw_room	56.67/94.12	-16.67/-2.45	-6.67/-7.45	16.66/-16.85
matchit_ascii	100/60	-2.50/9.23	-10.00/-26.67	0.00/7.50
guesswhat	90/11.73	-1.67/19.09	-18.33/20.05	-13.33/-3.67

Table 7: Gains and losses w.r.t baselines of average % played and quality score of individual games; L3: llama-3.1, PP: playpen, CS: Cold Start, WS: Warm Start, R: Rehearsal

ments especially for the in-domain games. For the out-of-domain games, the played score sometimes will be also negatively impacted if the model reaches a turn-limit. Even though the model knows how to structure the output the episode will be counted as aborted. This makes it difficult to pinpoint the exact reason for the decrease in % played to one particular cause.

It could be an over-fitting of the prompt structure of the in-domain games which leads to more aborted games due to malformed outputs, but it can also be that other abilities such as the contextual awareness of the model are worsened by the fine-tuning process. The latter would lead to many more games being aborted due to reaching the turn-limit.

These questions can only be answered by an in-depth analysis of the model errors while playing the games.

F.2 DPO

Compared to the base L3-8B SFT(CS), both DPO Dialogue and DPO turn appear to result in a degradation of performance on in-domain games, with the most pronounced declines observed for Wordle. On the other hand, in out-of-domain games, improvements in many games are observed for DPO Dialogue, with peaks in textmapworld_graph and textmapworld_room. DPO Turn, instead, seems to perform worst than the baseline model for most of the games.

F.3 GRPO

The performance increases and decreases for the GRPO models relative to their respective baselines are given in Table 10. Pure reinforcement learning leads to near-across-the-board improvements over the baseline Llama 3.1 8B model on all in- and out-of-domain games, although we observe slight decreases in percentage played on Wordle, Wordle-withcritic, Textmapworld, Textmapworld-specificroom, and Guesswhat, along with slight decreases in quality score for the latter three out-of-domain games.

On the other hand, GRPO struggles to improve the SFT Llama model (SFT(CS)+GRPO), and only results in slight increases in quality score for Taboo, Wordle, and Wordle-withcritic. We also observe substantial decreases in in-domain performance, in particular on Imagegame and Privateshared: the decrease in Privateshared is to be expected, as the reward function for this game is only loosely connected to the clemscore (as discussed in Appendix D.4). However, GRPO greatly improves the out-of-domain clemscores of the SFT model—with the notable exceptions of Matchit-ascii and Guesswhat.

F.4 Evaluation on General Instruction Following Benchmarks

We report in Table 11 and Table 12 a detailed breakdown of the results obtained in the evaluation on general instruction following tasks considered in

Game/Model	Baseline:		
	L3-8B SFT(CS)	SFT(CS)+DPO Dialogue	SFT(CS)+DPO Turn
<i>In Domain</i>	pp/qs	pp/qs	pp/qs
imagegame	100/94.49	-15.25/-15.07	0.0/-1.93
prvtshared	100/97.13	0.0/-4.40	0.0/-1.28
referencegame	100/46.67	0.0/-4.45	0.0/-5.56
taboo	100/37.5	0.0/+6.94	0.0/+7.22
wordle	56.67/1.18	-30.0/+2.57	-13.34/+1.13
wordle_wcl	23.33/71.43	0.0/-21.43	+10.0/-48.93
wordle_wcr	26.66/27.78	-3.33/+2.78	+13.34/-25.0
<i>Out-of-Domain</i>	pp/qs	pp/qs	pp/qs
adv-game	17.97/15.62	+2.34/-3.9	-3.91/0.0
codenames	25.38/21.21	-15.38/-5.83	-8.46/-7.57
tmw	68/46.71	+12.0/+9.0	-8.0/+5.28
tmw_graph	16.67/36.79	+43.33/+9.18	-3.34/-0.91
tmw_room	40/91.67	+23.33/-7.46	+8.33/-13.34
matchit_ascii	97.5/69.23	-70.0/-51.05	0.0/-10.26
guesswhat	88.33/30.82	+6.67/-3.92	+10.0/-4.27

Table 8: Comparison of % played (pp) and quality score (qs) on individual games for the L3-8B SFT(CS) and the further trained DPO Dialogue and DPO Turn.

Game	Agent Role	Teacher Role
Imagegame	Instruction Follower	Instruction Giver
Reference- game	Instruction Follower	Instruction Giver
Taboo	Guesser	Describer
Wordle (with critic)	Guesser	Critic

Table 9: Two-player games from the train split, and the roles played by the agent and teacher models in each for the online RL experiment.

this study.

G Qualitative Discussion

While we know that the absolute number of aborted episodes goes down from the baseline to SFT to GRPO, the distribution of reasons for those aborted episodes might change. For out-of-domain games, we investigated these reasons. Five overarching error categories were manually grouped together; the relative distributions of errors between the different models are depicted in Figure 5. The main problem for the baseline Llama-8B is exceeding the turn limit (e.g., in the textmapworld game, this could look like a player not continuing to look for new rooms but going in loops between rooms), whereas for Llama70B are output-formatting issues. This verbose behaviour by Llama-8B is reduced by all the types of training. The best performing Llama-8B version is the GRPO, achieving the lowest absolute number of aborted episodes,

with fewer turn limit errors and more game mechanic understanding ones in proportion. A similar proportion of error reasons is obtained by the best Llama-70 version out-of-domain, i.e., Llama-70 SFT, as mechanical understanding errors often require complex reasoning abilities than just following instructions (e.g., counting abilities). Both for Llama-8B and Llama-70B, the DPO-trained model shows the highest proportion of hallucination and context-related errors, a possible signal of overfitting to the training data.

In the Results Section 5, Llama-8B trained with GRPO data has been shown to generalise to out-of-domain games. One of the main reasons for this is the reduced number of aborted games due to exceeding the game’s turn limit (Fig. 5). Fig. 6 reports the absolute number of aborted episodes per possible aborted reasons in the Adventuregame for Llama-8B. The GRPO trained version drastically reduce the number of overall errors in the game, with around 1/4 of the original aborted episodes due to reaching the turn limit, and not reproducing the rambling errors of the SFT version (“next_action_missing”). Adventuregame, where the GRPO’s higher percentage of played games (+19.62) is coupled with a higher quality score (+11.28), is a good example of the stability of this out-of-domain gains by GRPO. For Codenames,

Game/Model	Baseline:		Baseline:	
	L3-8B	GRPO	L3-8B SFT(CS)	SFT(CS)+GRPO
<i>In Domain</i>	pp/qs	pp/qs	pp/qs	pp/qs
imagegame	67.8/54.62	+1.69/+2.65	100/94.49	-24.14/-43.38
prvtshared	100/23.48	0.0/+0.69	100/97.13	0.0/-76.17
referencegame	100/38.89	0.0/+4.44	100/46.67	0.0/-10.0
taboo	98.33/31.92	+1.67/+2.8	100/37.5	0.0/+4.39
wordle	36.67/0	-6.37/0.0	56.67/1.18	-3.34/+1.94
wordle_wcl	0/-	0.0/-	23.33/71.43	-20.0/-21.43
wordle_wcr	6.67/50	-3.34/+50.0	26.66/27.78	-16.66/+50.0
<i>Out-of-Domain</i>	pp/qs	pp/qs	pp/qs	pp/qs
adv-game	35.94/33.85	+19.62/+11.28	17.97/15.62	+28.7/+20.82
codenames	43.08/16.07	+3.84/+15.08	25.38/21.21	+13.85/-3.56
tmw	36/55.46	-6.0/-1.35	68/46.71	+6.0/+13.37
tmw_graph	20/44.33	+57.78/+1.36	16.67/36.79	+64.28/+13.83
tmw_room	56.67/94.12	-3.34/-0.37	40/91.67	+10.0/+1.66
matchit_ascii	100/60	0.0/+2.5	97.5/69.23	+2.5/-6.73
guesswhat	90/11.73	-10.0/-1.31	88.33/30.82	-8.33/-20.4

Table 10: Gains and losses w.r.t baseline of average % played (pp) and quality score (qs) of individual games.

Fig. 7 reports the absolute number of aborted episodes per possible aborted reasons for Llama-8B. We observe for GRPO a reduced number of hallucinations (“Target is hallucination” and “Guess word is hallucination”) compared to the base and SFT, while not decreasing in most cases compared to the baseline. A notable exception is the “Wrong number of guesses”, where the GRPO model seems not to respect the number of guesses per turn required by the game. Finally, we report a codenames episode played by all the base Llama-8B, SFT, and GRPO. As shown in Fig. 13, the base model fail due to the common error of guessing the clue word; the SFT does not encounter errors but reveals the killer word, losing the game. The GRPO, instead, is able to play the game and succeed.

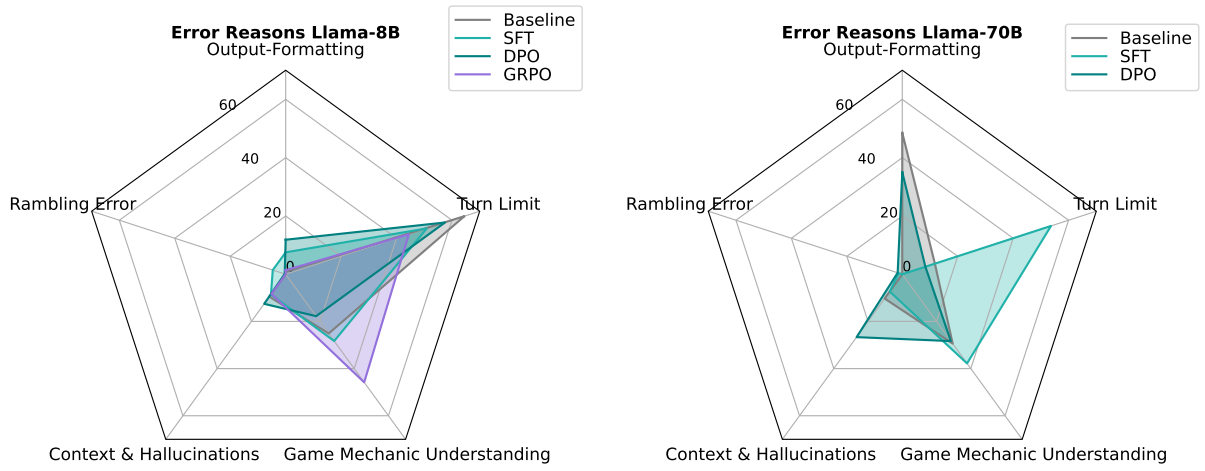


Figure 5: Relative distribution of error categories.

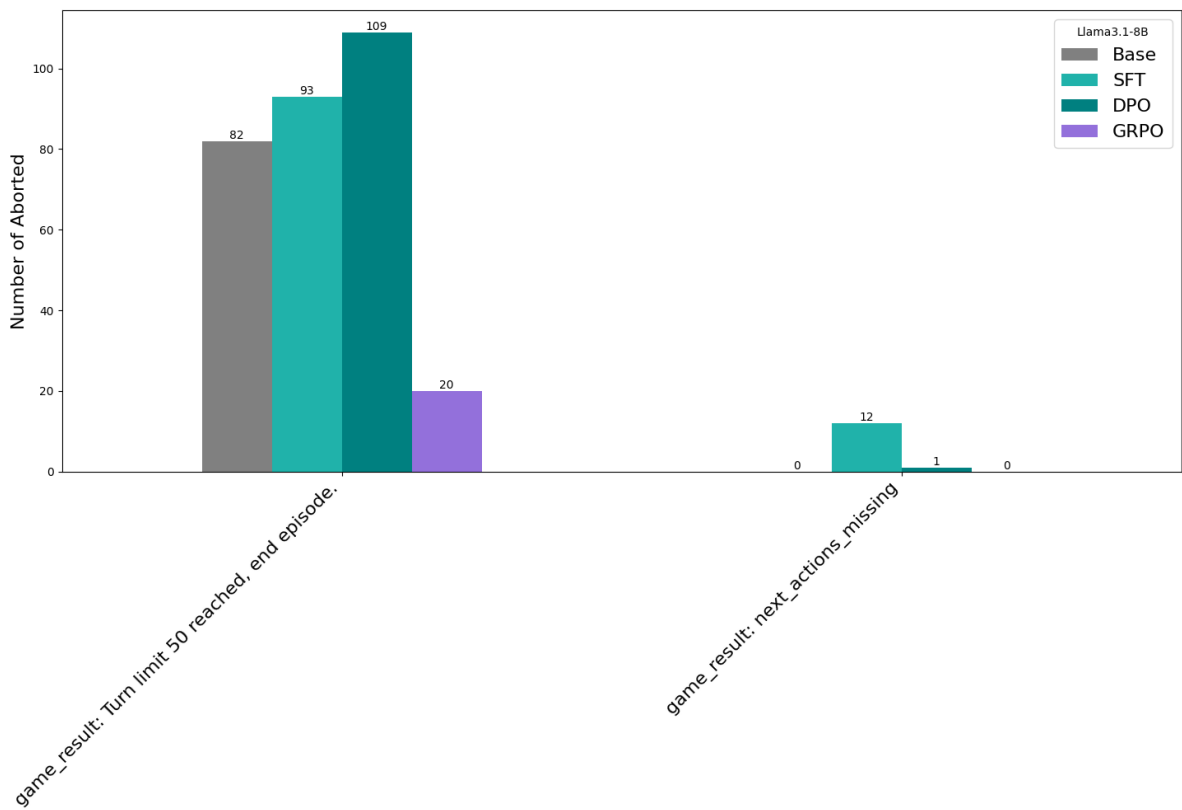


Figure 6: Number of aborted interactions per abortion reason in **Adventuregame**.

	Executive				Socio-Emotional				
	NATURAL PLAN	LogiQA2	CLadder	WinoGrande	EQBench	LM-Pragmatics	SocialIQA	SimpleToM (AJ)	SimpleToM (ToM)
Llama-3.1-8B-Instruct									
Base	06.40	32.31	50.57	67.71	67.79	65.12	48.36	46.68	57.71
SFT (CS)	09.83	31.11	54.13	64.64	61.45	62.08	47.24	38.45	88.31
SFT (WS)	06.17	32.63	51.95	69.69	51.72	45.73	47.85	34.26	71.83
SFT (R)	12.80	32.18	53.37	67.88	49.84	55.97	49.84	24.80	82.39
SFT (CS) + DPO (Dial.)	12.17	27.16	53.36	61.09	61.48	50.85	44.88	38.19	80.12
SFT (CS) + DPO (Turn)	12.06	32.06	51.91	63.14	59.54	61.71	48.52	40.89	86.75
GRPO	07.31	32.12	50.96	67.17	67.69	65.49	48.56	46.60	59.20
SFT (CS) + GRPO	05.42	32.06	29.24	66.69	68.94	65.61	48.56	50.04	60.24
Llama-3.1-70B-Instruct									
Base	29.03	51.52	56.34	72.77	82.03	80.97	55.02	44.33	94.5
SFT(CS)	32.03	53.24	56.95	78.45	76.40	80.61	54.96	37.84	96.43
SFT(WS)	30.75	48.72	52.39	76.60	75.50	76.58	54.86	36.01	94.59
SFT(R)	30.00	52.80	56.02	77.42	77.15	76.83	56.40	41.06	88.40
SFT(CS)+DPO(Dial.)	28.81	45.61	57.22	64.09	80.24	81.46	48.06	45.47	84.22
SFT(CS)+DPO(Turn)	28.75	48.09	56.15	67.88	81.93	83.05	52.87	48.26	84.92

Table 11: **Performance on Executive and Socio-Emotional Tasks.** SimpleToM (AJ) and (ToM) are grouped based on the taxonomy in [Momentè et al. \(2025\)](#).

	Formal	General		Instruction-following
	GLUE Diagnostics	MMLU-Pro	BBH	IFEval
Llama-3.1-8B-Instruct				
Base	38.06	43.35	40.37	76.88
SFT (CS)	40.23	13.16	46.75	67.25
SFT (WS)	30.74	01.70	45.52	61.40
SFT (R)	38.62	31.99	45.66	68.76
SFT (CS) + DPO (Dial.)	36.20	09.28	43.86	68.39
SFT (CS) + DPO (Turn.)	36.07	11.13	46.80	70.76
GRPO	38.68	43.73	39.31	76.97
SFT (CS) + GRPO	37.31	41.55	44.09	75.77
Llama-3.1-70B-Instruct				
Base	46.16	60.37	60.74	85.16
SFT(CS)	47.72	25.90	63.91	79.38
SFT (WS)	45.86	25.03	63.58	75.10
SFT (R)	46.51	18.30	65.63	79.68
SFT(CS)+DPO(Dial.)	37.73	38.34	39.21	82.26
SFT(CS)+DPO(Turn)	39.23	36.02	53.69	85.68

Table 12: **Model performance on formal, general and instruction-following capabilities**, as measured by GLUE Diagnostics ([Wang et al., 2018](#)), MMLU-Pro and BBH ([Wang et al., 2024](#); [Suzgun et al., 2023](#)), IFEval ([Zhou et al., 2023b](#)).

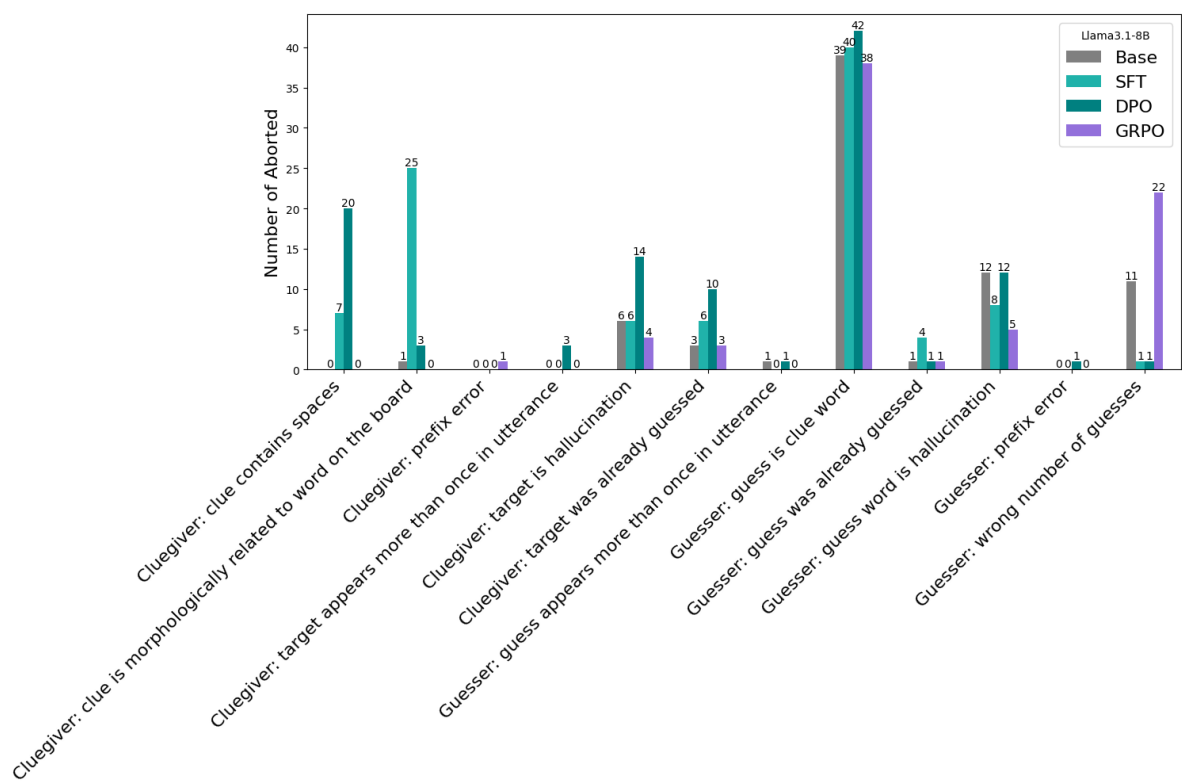


Figure 7: Number of aborted interactions per abortion reason in **Codenames**.

Table 13: Codenames sample by the Base model, the SFT and GRPO trained ones.

1	<p>[A]GM]</p> <p>Provide one single-word clue that relates to some of your team's words from the list below. You can choose to how many and to which words your clue relates to.</p> <p>The goal is to have your partner guess all of your team's words based on the clues you give them, before the other team has uncovered all of theirs. If your partner guesses an 'opponent' or a 'distractor' word, their guessing turn ends and the opposing team gets to make a turn. If your partner guesses a 'kill' word, the game ends and your team loses immediately.</p> <p>The more words your clue relates to, the quicker you can win the game, but the harder the associations may be to guess for your partner, so choose your clue word wisely. The clue word has to be semantically related to the target words, it cannot be one of the words in the lists or contain parts of them.</p> <p>Always give your single-word clue and your comma-separated list of related target words in the following format and make your answers as short as possible, never include any other text than is required in this form:</p> <p>CLUE: <WORD> TARGETS: <TARGETS></p> <p>Your team words are: gospel, shears, thunder, fog, bluegrass, attic, entrance, basement, hose.</p> <p>Your opponent's team words are: cotton, head, buttocks, running, back, wool, swimming, neck.</p> <p>Distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv.</p> <p>The kill words are: silver.</p>		
	Base	SFT	GRPO
2	<p>[A]GM]</p> <p>CLUE: farm</p> <p>TARGETS: bluegrass, shears, gospel</p>	<p>[A]GM]</p> <p>CLUE: roof</p> <p>TARGETS: attic, entrance, basement</p>	<p>[A]GM]</p> <p>CLUE: farm</p> <p>TARGETS: bluegrass, shears, gospel</p>
3	<p>[GM]B]</p> <p>Provide a comma-separated list of up to 2 words from the following list that best relate or are most closely associated with the word 'farm'. Always start your list of guess(es) with 'GUESS: ' and do not include any other text in your answer.</p> <p>attic, running, neck, bluegrass, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, head, oregano, fog, swimming, wool, gospel, cotton, basement, hose, shears, back, thunder, silver</p>	<p>[GM]B]</p> <p>Provide a comma-separated list of up to 3 words from the following list that best relate or are most closely associated with the word 'roof'. Always start your list of guess(es) with 'GUESS: ' and do not include any other text in your answer.</p> <p>attic, running, neck, bluegrass, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, head, oregano, fog, swimming, wool, gospel, cotton, basement, hose, shears, back, thunder, silver</p>	<p>[GM]B]</p> <p>Provide a comma-separated list of up to 3 words from the following list that best relate or are most closely associated with the word 'farm'. Always start your list of guess(es) with 'GUESS: ' and do not include any other text in your answer.</p> <p>attic, running, neck, bluegrass, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, head, oregano, fog, swimming, wool, gospel, cotton, basement, hose, shears, back, thunder, silver</p>
4	<p>[GM]B]</p> <p>GUESS: wool, shears, cotton</p>	<p>[GM]B]</p> <p>GUESS: attic, head, back</p>	<p>[GM]B]</p> <p>GUESS: bluegrass, cotton, shears</p>

Continued on next page

Table 13

#	Base	SFT	GRPO
5	<p>[A]GM]</p> <p>The word wool was guessed but is an opponent word. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, thunder, fog, bluegrass, attic, entrance, basement, hose.</p> <p>Remaining words for your opponent are: cotton, head, buttocks, running, back, swimming, neck.</p> <p>Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv.</p> <p>Remaining kill words are: silver.</p>	<p>[A]GM]</p> <p>The word attic was guessed correctly. The word head was guessed but is an opponent word. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, thunder, fog, bluegrass, entrance, basement, hose.</p> <p>Remaining words for your opponent are: cotton, buttocks, running, back, wool, swimming, neck.</p> <p>Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv.</p> <p>Remaining kill words are: silver.</p>	<p>[A]GM]</p> <p>The word bluegrass was guessed correctly. The word cotton was guessed but is an opponent word. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, thunder, fog, attic, entrance, basement, hose.</p> <p>Remaining words for your opponent are: head, buttocks, running, back, wool, swimming, neck.</p> <p>Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv.</p> <p>Remaining kill words are: silver.</p>
6	<p>[A]GM]</p> <p>CLUE: storm</p> <p>TARGETS: thunder, fog, hose</p>	<p>[A]GM]</p> <p>CLUE: storm</p> <p>TARGETS: thunder, fog, hose</p>	<p>[A]GM]</p> <p>CLUE: storm</p> <p>TARGETS: thunder, fog</p>
7	<p>[GM]B]</p> <p>The word wool was guessed but is an opponent word. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 3 words from the following list of words that best relate or are most closely associated with the word 'storm'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>attic, running, neck, bluegrass, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, head, oregano, fog, gospel, cotton, basement, hose, shears, back, thunder, silver</p>	<p>[GM]B]</p> <p>The word attic was guessed correctly. The word head was guessed but is an opponent word. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 3 words from the following list of words that best relate or are most closely associated with the word 'storm'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>neck, bluegrass, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, oregano, fog, swimming, wool, gospel, cotton, basement, hose, shears, back, thunder, silver</p>	<p>[GM]B]</p> <p>The word bluegrass was guessed correctly. The word cotton was guessed but is an opponent word. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 2 words from the following list of words that best relate or are most closely associated with the word 'storm'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>attic, running, neck, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, head, oregano, fog, swimming, wool, gospel, basement, hose, shears, thunder, silver</p>
8	<p>[GM]B]</p> <p>GUESS: fog, thunder, basement</p>	<p>[GM]B]</p> <p>GUESS: fog, thunder, back</p>	<p>[GM]B]</p> <p>GUESS: fog, thunder</p>

Continued on next page

Table 13

#	Base	SFT	GRPO
9	<p>[A(GM)] The words fog, thunder, basement were guessed correctly. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, bluegrass, attic, entrance, hose. Remaining words for your opponent are: cotton, head, buttocks, running, back, neck. Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv. Remaining kill words are: silver.</p>	<p>[A(GM)] The words fog, thunder were guessed correctly. The word back was guessed but is an opponent word. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, bluegrass, entrance, basement, hose. Remaining words for your opponent are: cotton, buttocks, wool, swimming, neck. Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv. Remaining kill words are: silver.</p>	<p>[A(GM)] The words fog, thunder were guessed correctly. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, attic, entrance, basement, hose. Remaining words for your opponent are: head, buttocks, running, wool, swimming, neck. Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv. Remaining kill words are: silver.</p>
10	<p>[A(GM)] CLUE: music TARGETS: gospel, bluegrass</p>	<p>[A(GM)] CLUE: home TARGETS: entrance, basement, hose</p>	<p>[A(GM)] CLUE: home TARGETS: attic, entrance, basement</p>
11	<p>[GM(B)] The words fog, thunder, basement were guessed correctly. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 2 words from the following list of words that best relate or are most closely associated with the word 'music'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>attic, running, neck, bluegrass, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, head, oregano, gospel, hose, shears, back, silver</p>	<p>[GM(B)] The words fog, thunder were guessed correctly. The word back was guessed but is an opponent word. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 3 words from the following list of words that best relate or are most closely associated with the word 'home'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>neck, bluegrass, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, oregano, swimming, gospel, cotton, basement, hose, shears, silver</p>	<p>[GM(B)] The words fog, thunder were guessed correctly. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 3 words from the following list of words that best relate or are most closely associated with the word 'home'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>attic, running, neck, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, oregano, swimming, wool, gospel, basement, hose, shears, silver</p>
12	<p>[GM(B)] GUESS: bluegrass, gospel</p>	<p>[GM(B)] GUESS: entrance, china, basement</p>	<p>[GM(B)] GUESS: attic, entrance, basement</p>

Continued on next page

Table 13

#	Base	SFT	GRPO
13	<p>[A(GM)] The words bluegrass, gospel were guessed correctly. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: shears, attic, entrance, hose. Remaining words for your opponent are: head, buttocks, running, back, neck. Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv. Remaining kill words are: silver.</p>	<p>[A(GM)] The word entrance was guessed correctly. The word china was guessed but is an innocent word. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, bluegrass, basement, hose. Remaining words for your opponent are: cotton, buttocks, swimming, neck. Remaining distractor words are: seasoning, brazil, oregano, portugal, truck, atv. Remaining kill words are: silver.</p>	<p>[A(GM)] The words attic, entrance, basement were guessed correctly. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, hose. Remaining words for your opponent are: buttocks, running, wool, swimming, neck. Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck, atv. Remaining kill words are: silver.</p>
14	<p>[A(GM)] CLUE: yard TARGETS: shears, hose</p>	<p>[A(GM)] CLUE: tool TARGETS: shears, hose</p>	<p>[A(GM)] CLUE: door TARGETS: entrance, attic</p>
15	<p>[GM(B)] The words bluegrass, gospel were guessed correctly. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 2 words from the following list of words that best relate or are most closely associated with the word 'door'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>attic, running, neck, atv, brazil, portugal, seasoning, entrance, truck, china, buttocks, oregano, hose, shears, back, silver</p>	<p>[GM(B)] The word entrance was guessed correctly. The word china was guessed but is an innocent word. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 2 words from the following list of words that best relate or are most closely associated with the word 'tool'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>neck, bluegrass, atv, brazil, portugal, seasoning, truck, oregano, swimming, gospel, cotton, basement, hose, shears, silver</p>	<p>[GM(B)] The words attic, entrance, basement were guessed correctly. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 2 words from the following list of words that best relate or are most closely associated with the word 'yard'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>running, atv, brazil, portugal, seasoning, truck, china, buttocks, oregano, swimming, wool, gospel, hose, shears, silver</p>
16	<p>[GM(B)] GUESS: entrance, door</p>	<p>[GM(B)] GUESS: atv, shears</p>	<p>[GM(B)] GUESS: atv, hose</p>

Continued on next page

Table 13

#	Base	SFT	GRPO
17	<p>[GM GM] Guessed word 'door' is the same word as the provided clue word, you should only select words from the provided list.</p> <p>[GM GM] Game Aborted</p>	<p>[A GM]</p> <p>The word atv was guessed but is an innocent word. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, bluegrass, basement, hose. Remaining words for your opponent are: cotton, swimming, neck. Remaining distractor words are: seasoning, brazil, oregano, portugal, truck. Remaining kill words are: silver.</p>	<p>[A GM]</p> <p>The word atv was guessed but is an innocent word. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears, hose. Remaining words for your opponent are: buttocks, running, wool, swimming. Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck. Remaining kill words are: silver.</p>
18		<p>[A GM]</p> <p>CLUE: music</p> <p>TARGETS: gospel, bluegrass</p>	<p>[A GM]</p> <p>CLUE: water</p> <p>TARGETS: hose</p>
19		<p>[GM B]</p> <p>The word atv was guessed but is an innocent word. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 2 words from the following list of words that best relate or are most closely associated with the word 'music'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>neck, bluegrass, brazil, portugal, seasoning, truck, oregano, swimming, gospel, basement, hose, shears, silver</p>	<p>[GM B]</p> <p>The word atv was guessed but is an innocent word. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 1 words from the following list of words that best relate or are most closely associated with the word 'water'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>running, brazil, portugal, seasoning, truck, china, buttocks, oregano, swimming, gospel, hose, shears, silver</p>
20		<p>[GM B]</p> <p>GUESS: bluegrass, gospel</p>	<p>[GM B]</p> <p>GUESS: hose</p>

Continued on next page

Table 13

#	Base	SFT	GRPO
21		<p>[A]GM]</p> <p>The words bluegrass, gospel were guessed correctly. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: shears, basement, hose. Remaining words for your opponent are: swimming, neck. Remaining distractor words are: seasoning, brazil, oregano, portugal, truck. Remaining kill words are: silver.</p>	<p>[A]GM]</p> <p>The word hose was guessed correctly. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: gospel, shears. Remaining words for your opponent are: buttocks, running, swimming. Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck. Remaining kill words are: silver.</p>
22		<p>[A]GM]</p> <p>CLUE: storage TARGETS: basement, hose</p>	<p>[A]GM]</p> <p>CLUE: music TARGETS: gospel</p>
23		<p>[GM]B]</p> <p>The words bluegrass, gospel were guessed correctly. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 2 words from the following list of words that best relate or are most closely associated with the word 'storage'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>brazil, portugal, seasoning, truck, oregano, swimming, basement, hose, shears, silver</p>	<p>[GM]B]</p> <p>The word hose was guessed correctly. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 1 words from the following list of words that best relate or are most closely associated with the word 'music'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>running, brazil, portugal, seasoning, truck, china, buttocks, oregano, gospel, shears, silver</p>
24		<p>[GM]B]</p> <p>GUESS: basement, silver</p>	<p>[GM]B]</p> <p>GUESS: gospel</p>

Continued on next page

Table 13

#	Base	SFT	GRPO
25		<p>[GM GM] Team revealed.</p> <p>[GM GM] Lose</p>	<p>[A GM] The word gospel was guessed correctly. Your teammate's turn ended there. Now provide another clue relating to some of your remaining team words and a list of the related target words. Remember to start your clue with 'CLUE: ', put a new line, and start your comma-separated list of target words with 'TARGETS: '. Notice: some words have been removed from the lists compared to previous requests.</p> <p>Your remaining team words are: shears. Remaining words for your opponent are: buttocks, running. Remaining distractor words are: seasoning, brazil, oregano, china, portugal, truck. Remaining kill words are: silver.</p>
26			<p>[A GM] CLUE: garden TARGETS: shears</p>
27			<p>[GM B] The word gospel was guessed correctly. Your turn ended there. Now provide another comma-separated list of at least 1 and up to 1 words from the following list of words that best relate or are most closely associated with the word 'garden'. Remember to start your answer with 'GUESS: '. Notice: some words have been removed from the list compared to previous requests.</p> <p>brazil, portugal, seasoning, truck, china, buttocks, oregano, shears, silver</p>
28			<p>[GM B] GUESS: shears</p>
29			<p>[GM GM] Game Success</p>