

Communication Makes Perfect: Persuasion Dataset Construction via Multi-LLM Communication

Weicheng Ma¹, Hefan Zhang², Ivory Yang², Shiyu Ji³,
Joice Chen³, Farnoosh Hashemi³, Shubham Mohole³, Ethan Gearey²,
Michael Macy³, Saeed Hassanpour², and Soroush Vosoughi²

¹College of Computing, Georgia Institute of Technology

²Department of Computer Science, Dartmouth College

³Cornell University

¹wma76@gatech.edu

Abstract

Large Language Models (LLMs) have shown proficiency in generating persuasive dialogue, yet concerns about the fluency and sophistication of their outputs persist. This paper presents a multi-LLM communication framework designed to enhance the generation of persuasive data automatically. This framework facilitates the efficient production of high-quality, diverse linguistic content with minimal human oversight. Through extensive evaluations, we demonstrate that the generated data excels in naturalness, linguistic diversity, and the strategic use of persuasion, even in complex scenarios involving social taboos. The framework also proves adept at generalizing across novel contexts. Our results highlight the framework’s potential to significantly advance research in both computational and social science domains concerning persuasive communication.

1 Introduction

Persuasion techniques play a critical role in shaping societal behaviors and public opinion (Fogg, 2009; Braca and Dondio, 2023), which has led to sustained interest across a range of disciplines. Social science research has established detailed taxonomies of persuasion strategies (Shrum et al., 2012; Lukin et al., 2017), while datasets have been developed to cover various domains, including charitable donations (Wang et al., 2019), argument ranking in debates (Toledo et al., 2019), detecting mental manipulation (Wang et al., 2024; Yang et al., 2024), and understanding advertising strategies (Kumar et al., 2023). Despite these advances, ambiguities persist in defining persuasion (Pauli et al., 2022), and applying persuasion strategies across different contexts remains complex (Bai et al., 2021; Schaefer et al., 2023; Piskorski et al., 2023). Additionally, the high cost of manually annotating quality data poses a significant challenge (Lai et al., 2022).

The advent of large language models (LLMs) has unlocked new possibilities for enhancing various forms of communication, including online political discourse (Argyle et al., 2023; Bai et al., 2023), personalized advertising (Matz et al., 2024; Meguelati et al., 2024), public health messaging (Lim and Schmälzle, 2023; Espinosa and Salathé, 2024), and opinion shaping on social media (Meier, 2024). Recent research, such as that by Jin et al. (2024), has begun exploring LLM-generated persuasive dialogues. However, their approach is limited to simple, two-party dialogues where a persuader seeks to change the persuadee’s viewpoint. These dialogues often lack depth, presenting brief exchanges with simplistic logic and unnatural flow, restricting their usefulness for studying persuasion in more complex settings.

In response to these limitations, we propose a multi-agent framework for generating persuasion data. In this framework, multiple agents are assigned distinct roles, ensuring that each aspect of the dialogue generation process is handled efficiently. This structure minimizes the risk of an agent missing important details due to task abstraction or prompt complexity, a common issue in LLM prompting (Brown et al., 2020; Huang et al., 2023). Additionally, auxiliary agents manage dialogue flow to ensure that the resulting exchanges are coherent, logically consistent, and incorporate diverse persuasive strategies, simulating natural human conversation. Our approach imposes no preconditions regarding speakers, language styles, domains, or persuasion strategies, allowing it to generate a wide range of dialogues. For instance, our framework can support adversarial dialogues, where both participants attempt to persuade one another while maintaining their original positions. Moreover, we employ a continuous labeling scheme to measure the degree of perspective change throughout the dialogue, avoiding the limitations of binary utterance labels. This framework also integrates ethical

considerations, incorporating cultural norms and taboos from NormBank (Ziems et al., 2023) to explore ethically challenging persuasive scenarios, such as dialogues involving manipulation or unethical persuasion.

Careful analyses conducted by experts from both NLP and social sciences confirm the quality of our generated dialogues, particularly in terms of their naturalness, logical structure, and diversity of persuasion strategies. Our sentence-level persuasiveness labels align closely with human judgments (see Appendix G).

We further demonstrate the flexibility of our framework through experiments controlling for specific persuasion strategies and in more complex scenarios, such as multi-party conversations. Across all tested configurations, our framework consistently produced high-quality dialogues, showcasing its adaptability and generalizability. These findings indicate that our framework offers a robust platform for studying persuasion techniques, particularly in high-stakes contexts where ethical concerns, such as the spread of misinformation and propaganda, are paramount (Chen and Shu, 2023; Jones, 2024).

2 Multi-Agent Data Generation & Annotation Framework

Our framework incorporates 6 groups of language agents as shown in Figure 1. In our experiments, all agents utilize a GPT-3.5 backbone, except for the utterance quality monitor and global regulation agents which are based on GPT-4 due to their need for advanced reasoning capabilities and enhanced memory retention. Note that this choice of LLMs aims to balance data generation costs with quality, and using more powerful models could further improve the effectiveness of our approach. Our preliminary experiments on model selection are outlined in Appendix A.

2.1 Dialogue Generation Agents

We adopt a methodology for generating multi-round conversations by cyclically using the output from one language agent as the input for another (Park et al., 2023). This technique has been validated to produce extended, logically consistent dialogues that fulfill our project requirements.

Our framework initializes the generative agents with a description of the task settings, the predefined tasks for each language agent, and guidelines governing the models’ generations, as illustrated

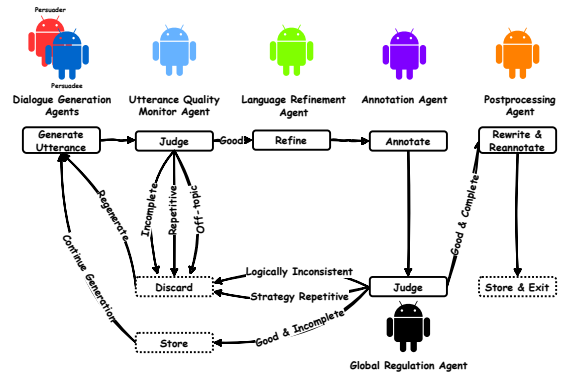


Figure 1: Overview of our data generation and annotation framework. Prior to dialogue generation, each agent is assigned specific tasks and given predefined stances to maintain throughout the conversation.

in Figure B1. The task choices for each agent are not constrained, for instance, drawing on a cultural taboo that “one should not pick flowers in a cemetery” from NormBank, we could challenge the persuader to convince the persuadee to pick flowers in a cemetery, while the persuadee is instructed to resist and, if possible, persuade the persuader to abandon such thoughts.

The dialogues commence when we prompt a persuader agent with “Start the conversation.” This setup initiates a structured yet dynamic interaction between the speakers, allowing us to closely observe and analyze their persuasive strategies.

2.2 Utterance Quality Monitor Agent

Due to the inherent limitations of LLMs, dialogue generation agents may occasionally produce incomplete, repetitive, or off-topic content. To address these issues, we introduce a specialized LLM agent responsible for tracking the persuasion topic and generation history to evaluate new generations.

The initialization prompt of the utterance quality monitor agent is shown in Figure B2. During dialogue generation, this agent inspects every new utterance to check if they ends unexpectedly, repeats a previous utterance, or goes off the topic of the dialogue in a sequence. If an utterance is red-flagged for any issue, the author agent is requested to revise the utterance based on the diagnoses. Otherwise, before proceeding to the next utterance, the utterance quality monitor agent is prompted to update its memory, storing the reviewed utterance for future judgments.

2.3 Language Refinement Agent

Raw text produced by dialogue generation agents often adopts a conclusive rather than conversational tone, primarily because the agents are prompted in a question-answering format. This could lead to stylistic conflicts with surrounding utterances. Additionally, the generations frequently include tone-softening phrases like “I understand your concerns,” or unnecessary affirmations such as agreeing with the other speaker’s views, which dilute the strength of arguments. Over the course of the conversation, these issues can compound, leading to dialogues dominated by language softeners and lacking in persuasive content.

To address this issue, we adopt a language refinement agent tasked with stripping out polite but superfluous phrases, thereby sharpening the dialogue’s focus on substantive content. System message to this agent is shown in Figure B3. 2 examples are also provided to the agent to further regulate its behaviors. Subsequent operations, including continued dialogue generation and persuasiveness labeling, are predicated on the output from the language refinement agent, ensuring that the conversation maintains its relevance and effectiveness in conveying persuasive arguments.

2.4 Persuasiveness Annotation Agent

After generating each round of conversation, we employ a persuasiveness annotation agent to assess the extent of perspective shifts in each speaker, assigning a score ranging from 0 to 1. Figure B4 illustrates the system message fed to the persuasiveness annotation agent before the generation starts. In practice, we provide the annotation agent with two scoring examples to guide its behavior and minimize scoring errors, such as incorrectly assigning a score of 1 to a conversation round with no perspective shifts (Figure I1). Note that these scores reflect the cumulative viewpoint shifts across all prior rounds of communication, facilitating the analysis of gradual persuasion rather than focusing solely on the impact of a single utterance.

2.5 Global Regulation Agent

We employ a global regulation agent to ensure smooth logical flow in the generated conversations and to determine the appropriate time to conclude the dialogue. The system message to the global regulation agent is depicted in Figure B5.

After each round of utterances is generated and

annotated, we prompt this agent to verify whether any changes in each speaker’s perspectives are logically influenced by the preceding utterance and whether the newly generated utterances avoid repeating previously used strategies within the same conversation. If the logical connections are insufficient or no new persuasive attempts are made, the dialogue generation agents are asked to revise their responses based on feedback from the global regulation agent. Once the revised generation passes these checks, the agent’s internal memory is updated accordingly. Then the agent is prompted to assess whether the speakers have reached a mutual agreement or if no new information is likely to be introduced next, indicating that the dialogue should be concluded. Although the ideal conclusion involves the persuader and persuadee agreeing on the preset task, conversations can often devolve into repetitive and unproductive arguments (Figure I2) (Xu et al., 2022). To prevent such stagnation, we allow the dialogue to conclude even if complete agreement is not reached. The global regulation agent is responsible for determining when to end the dialogue, at which point the conversation is terminated and the agent’s memory is reset.

2.6 Postprocessing Agent

After generating and annotating a full dialogue, we use a postprocessing agent to enhance content smoothness and naturalness. As shown in Figure B6, the agent removes redundant language, improves logical flow, and enhances language diversity. It also merges labels and reassigns them to modified utterances if the number of dialogue rounds changes.

3 Data Quality Assessment

To evaluate our data generation framework, we constructed a small dataset of 200 dialogues using randomly selected norms from NormBank for human validation. These norms consist of 98 taboos, 76 normal behaviors, and 26 expected behaviors. We intentionally placed greater emphasis on taboos because these behaviors often conflict with widely accepted moral standards, causing LLMs to refuse to generate persuasive dialogues (Figure C1). As such, they present a unique challenge in persuasion scenarios for both humans and LLMs.

Our data assessment plan focuses on three key aspects, progressing from more specific to broader levels of analysis: (1) the language fluency of indi-

Annotator Agreement Type	Model-generated Utterances	Utterances with Human Rewrites
Human annotators disagree with each other (49% of the sample)	<p>Despite your persuasive argument, I maintain my decision to forgo shoes in the library. I am convinced that personal comfort and cleanliness can coexist without conforming to the practice of wearing shoes indoors.</p> <p>Undeniably, caution is essential, yet many individuals choose their offerings with respect and adherence to guidelines. With proper education and awareness, we can continue this tradition in a way that respects both the site and the environment. Striking a balance between honoring history and preserving nature is key.</p> <p>I respect your decision, but I urge you to consider the potential hazards. Your well-being is vital, and alternative relaxation methods could offer similar benefits without risk. Balancing peace of mind with safety is crucial for healthy self-care practices.</p>	<p>Even with your compelling argument, I stand by my choice to go without shoes in the library. I believe that personal comfort and cleanliness can coexist without adhering to the norm of wearing shoes indoors.</p> <p>Many individuals choose their offerings with respect and adherence to guidelines while keeping caution in the back of their mind though it is essential. Striking a balance between honoring history and preserving nature is key. With proper education and awareness, we can continue this tradition in a way that respects both the site and the environment.</p> <p>Though I respect your decision, I advise you to familiarize yourself with the potential hazards of this practice. Your well-being is vital, and alternative relaxation methods could offer similar benefits without risk. Balancing peace of mind with safety is crucial for healthy self-care practices.</p>
Both human annotators correctly identified LLM-generated messages (29.25% of the sample)	<p>Despite your precautions, the risks of encountering dangerous animals or losing your way in the dark are substantial. Our senses are less effective at night, making it difficult to respond quickly to threats. It's safer to err on the side of caution and avoid the forest after dark.</p> <p>If there's any other way I can help you, please reach out. Enjoy your day at the bakery.</p> <p>I trust in the power of spiritual intervention to bring about the positive changes needed in my apartment. Would you reconsider and assist me?</p>	<p>The risks of encountering dangerous animals or losing your way in the dark are substantial regardless of your preparations and whatever precautions you take. Our senses are less effective at night, making it difficult to respond quickly to threats. It's safer to be cautious and avoid the forest after dark.</p> <p>Thank you for stopping by our bakery; if there are any questions, I am glad to help.</p> <p>I believe that powerful spiritual intervention can bring the positive changes my apartment needs. Would you change your mind and help me?</p>
Both human annotators took human-written messages as model-generated messages (21.75% of the sample)	<p>While I understand the value of learning to provide basic medications under guidance, I still believe that it's crucial to ensure the animals' safety and well-being by relying on trained professionals for medication administration. I will continue to focus on supporting the animals' welfare through means that align with my capabilities, while prioritizing their health and security.</p> <p>Wearing high heels doesn't necessarily mean compromising comfort and safety. You can opt for comfortable, stable high heel boots specifically designed for outdoor use. They offer style and functionality, allowing you to experience the thrill of skiing or snowboarding while looking chic. It's a chance to redefine traditional ski resort attire and make striking fashion statements.</p> <p>While I respect your perspective, I feel comfortable with my decision. I enjoy the freedom and comfort of skiing without a helmet. My cautious approach and awareness help me ensure my safety.</p>	<p>The value of learning to provide basic medications under guidance is undeniable, but I still believe that it's crucial to ensure the animals' safety and well-being by relying on trained professionals for medication administration. I will continue to focus on supporting the animals' welfare through means that align with my capabilities while prioritizing their health and security.</p> <p>High heels does not mean that safety and comfort are compromised, rather there are options like staboe high heel boots aimed for outdoor use and they offer style and functionality. You can experience the thrill of skiing or snowboarding while looking chic and its a opportunity to redefine traditional ski resort attire making striking fashion choices.</p> <p>I value your opinion, but I feel confident in my decision. I take pleasure in the freedom and comfort of skiing without a helmet. My mindful approach and attentiveness enable me to keep myself safe.</p>

Table 1: Examples of Annotator Agreement on Utterance-Level Human v.s. LLM Differentiation Task. Across 400 sampled pairs of utterances, annotators disagreed on which sentence was LLM-generated in 49% of cases. In 29.25% of the pairs, both annotators successfully identified the LLM-generated language, while in 21.75% of the pairs, neither annotator was able to detect the LLM-generated language.

vidual utterances, (2) the the topic, semantic, and logical coherence of entire conversations, and (3) the language and strategy diversity of conversations generated under the same topic and context.

3.1 Utterance-Level Quality Assessment

A critical goal for our framework is that each generated utterance should closely resemble a human-written sentence. To validate this, we conduct (a) a quantitative annotation task to differentiate between model-generated sentences and human-written sentences, followed by (b) a qualitative error analysis that combines annotator feedback with insights from an LLM on sentences that multiple annotators agreed were distinguishable.

3.1.1 Quantitative Differentiation Task

The differentiation task aims to assess how accurately human annotators could tell model-generated sentences apart from those rewritten by humans. Similar tasks have been discussed in Gehrman et al. (2019), Ippolito et al. (2020) and Ma et al. (2023). For our evaluation, we obtained a stratified sample of 400 utterances from 150 random

sample dialogues to ensure equal representation of utterances from both the persuader and persuadee agents, covering different rounds of persuasion to reflect the dataset distribution.

Manual Rewriting. 2 native English speakers were asked to rewrite each sampled utterance to provide reference texts that model-generated utterances will be compared against. Each assistant was assigned 200 utterances. As shown in Figure D1, they were instructed to retain the original meaning while improving clarity, grammar, and natural phrasing. Additionally, they could refine any awkward or unclear phrasing without altering the intended message.

Human Validation. After manual rewriting, we created a dataset consisting of pairs of model-generated utterances and their corresponding rewritten versions. 3 fluent English-speaking annotators were then tasked with identifying the model-generated utterance in each pair. The instructions provided to the annotators are shown in Figure D2.

Each utterance in the dataset was annotated by 2 annotators, and annotators were encouraged to comment on examples they found interesting. Note

Category	Statements	Related Work	Avg. Score	κ	Weighted κ
Interpersonal Responses	<i>Coherence</i> The speakers respond logically to the immediate conversation.	Ke et al. (2018), Wu et al. (2019), Liang and Li (2021)	2.969	0.657	0.657
	<i>Informativeness</i> The arguments makes sense given its context.	Zhu et al. (2019)	2.653	0.473	0.481
	The utterances build on prior information in near context.	Moghe et al. (2018), Young et al. (2018), Lin et al. (2019), Wu et al. (2019)	2.755	0.339	0.339
	The utterances introduce relevant new information or arguments.	Ke et al. (2018), Wu et al. (2019), Zhu et al. (2019)	2.337	0.410	0.459
Overall Fluency	The arguments overall are communicated clearly.	Moghe et al. (2018), Lin et al. (2019)	3	NA	NA
	The conversation sounds human-like and fluent overall.	Ke et al. (2018), Wu et al. (2019), Zhu et al. (2019), Ji et al. (2022)	2.561	0.557	0.576
Internal Role Consistency	There are no sudden shifts in a speaker’s objectives or stance without a clear explanation.	Moghe et al. (2018), Ji et al. (2022)	2.765	0.397	0.546
Topic Consistency	The conversation stays on topic	Moghe et al. (2018), Ji et al. (2022)	2.878	0.548	0.645

Table 2: Dialogue-level Quality Evaluation. 2 annotators assessed 50 randomly selected dialogues on the criteria listed above, using a likert scale of 1 - Not Accurate, 2 - Somewhat Accurate, and 3 - Accurate. We report the average scores across all dialogues for each measured dimension. Both linearly weighted (Weighted κ) and unweighted (κ) inter-rater consistency scores are calculated, with all results showing significant agreement.

that if the annotators were unable to distinguish between sentences and resorted to random guessing, the expected accuracy for both annotators correctly identifying model-generated utterances would be 25%. Comparing the actual accuracy to this baseline helps determine whether the model-generated utterances appeared natural to the annotators.

Of the 400 utterance pairs, the model-generated utterances in 117 pairs (29.25%) were correctly identified by both annotators, slightly going above the random baseline of 25%. In 49% cases (98 pairs), the annotators disagreed, and in 21.75% cases, both annotators resulted in incorrect identifications. Individual annotator accuracies were 0.546, 0.558, and 0.508. The results are close to random guessing, suggesting the challenge of distinguishing utterances generated by our framework from human-written ones. Example utterance pairs and their annotator labels are provided in Table 1.

3.1.2 Model-assisted Error Analysis

The quantitative findings indicate that our framework generally produces high-quality utterances nearly indistinguishable from human-written sentences. To follow up, we conducted a qualitative error analysis on the samples correctly distinguished by both annotators to identify areas for improvement. Precisely, all 117 such utterances were submitted to OpenAI’s o1-preview model (OpenAI, 2024) for further analysis to understand reasons behind their distinguishability by humans. The prompt for this task is shown in Figure D3.

Out of the 117 pairs, o1-preview correctly distinguished 72 pairs (61.2%), suggesting that even for LLMs, utterances generated by our framework are close to human writings. As suggested by o1-preview and verified by human annotators (Table E1), major causes of unnaturalness (frequencies

cited in parentheses) in these 72 less human-like utterances include overly formal language or detached tone or word use (88.9%), lengthy sentences, redundancy, verbosity, and repetition (68.1%), unnatural syntax, word choice, and language style (58.3%), complex sentence structures (34.7%), use of generic words and cliché phrases (23.6%), overly perfect grammar (13.9%), and LLM-style closing phrases (12.5%).

3.2 Dialogue Smoothness and Naturalness

We further conduct dialogue-level analyses on our sample data to ensure that each generated dialogue is logically coherent and effective in persuasion.

3.2.1 Dialogue Quality Annotation

We first developed a systematic rubric for evaluating the overall quality of persuasive dialogues. Our evaluation is conducted on (a) the local level, which examines each argument-response pair between the speakers, and (b) the global level, which considers the conversation as a whole. Evaluations are based on existing human evaluation dimensions for open dialogue systems and emphasize three key aspects: the interaction between persuader and persuadee, the consistency of individual participants across multiple rounds, and the alignment of utterances with the topic. Detailed criteria and their references are outlined in Table 2.

The local (round-level) evaluation focuses on 2 conventional dimensions in dialogue systems: Coherence and Informativeness. Coherence refers to round-level logical consistency, i.e., speakers respond to each other in a manner appropriate to commonsense and the given context (Li and Sun, 2018; Young et al., 2018; Wu et al., 2019; Liang and Li, 2021). Informativeness measures the quality and progression of information, ensuring responses

Category	Technique	Definition	Citation
External Validity	Popularity	Invoke popular opinion or experiences as support for uptake	Anand et al., 2011, Iyer and Sycara, 2019
	Authority	Appeals to authority (bosses, experts, trend-setters).	
Outcomes	Outcomes	Mentions some particular consequences from up take or failure to up take	Anand et al., 2011, Iyer and Sycara, 2019
	Threat/Promise	Poses direct threat or promise to the persuadee	
Generalizations	Deontic/Moral Appeals	Mentions duties or obligations, moral goodness, badness	Anand et al., 2011, Iyer and Sycara, 2019, Kumar et al., 2023
Interpersonal	Empathy	Attempts to make the persuadee connect with someone else’s emotional perspective	
	Scarcity	Mentions rarity, urgency, or opportunity of some outcome	Chen and Yang, 2021
Other	Logical Appeal	Provides a justification for an argumentative point based upon additional argumentation schemes e.g. causal reasoning, arguments from absurdity, or using facts, evidence, and statistics to appeal to the logic of the audience	
		Emotion	Make request full of emotional valence and arousal affect to influence others

Table 3: Persuasive Strategies, Definitions and Related Works. Human annotators evaluate each set of 25 dialogues, covering 5 topics with 5 dialogues per topic, for the strategies listed above.

align with prior dialogue context while introducing new information or arguments (Zhu et al., 2019).

On the global (dialogue) level, the overall dialogue should feel as if it could have been generated by human speakers (Moghe et al., 2018; Lin et al., 2019). First, speakers are assessed for overall fluency. We assessed the linguistic and stylistic quality of responses, ensuring arguments are communicated clearly and easy to follow, and that the conversation flows naturally (Wu et al., 2019). In addition, we looked at internal consistency throughout the conversation, defined as the absence of sudden, unexplained shifts in position, intention, or objective of speech (Moghe et al., 2018; Ji et al., 2022). Since the conversations are generated specific to topics, we also evaluate topic consistency, i.e., whether the conversation remains on-topic throughout (Moghe et al., 2018; Ji et al., 2022).

Annotators are asked to rate whether a series of statements, covering the above aspects, accurately describes the conversations on a three-point scale: 3 (accurate), 2 (somewhat accurate), and 1 (not accurate). 2 annotators participated in this task, each annotating the same set of 50 dialogues not overlapping with those used for utterance-level evaluations. Before annotation, a one-hour training session with examples was conducted to ensure both annotators fully understood the criteria. The annotators achieved an average unweighted Cohen’s κ of 0.483 (ranging from 0.339 to 0.657 across items) and an average linearly weighted Cohen’s κ of 0.529 (ranging from 0.339 to 0.657 across items), indicating relatively solid inter-rater consistency for human evaluations in Natural Language Generation tasks. Detailed scores and inter-rater consistency are reported in Table 2.

The dialogues generated by our framework are in general rated high on local-level coherence and clarity, particularly for providing logical responses within the immediate context (average score: 2.969

out of 3) and achieving perfect clarity in the arguments (average score: 3 out of 3). At the global level, the agents remain on topic for most dialogues (average score: 2.878 out of 3) and maintain good role consistency in their objectives and stances (average score: 2.765 out of 3).

Figure F1 exemplifies a highly-rated dialogue where the persuader addresses the persuadee’s concerns about using the attic for food storage. The persuader begins by presenting the attic as an “efficient space-saver and emergency backup” and counters concerns about unstable temperatures and pests with solutions like “airtight containers” and monitoring. Despite the persuadee’s repeated objections, the persuader suggests “proper insulation” and highlights the benefits of being prepared. Eventually, the persuadee proposes using a pantry, which resolves their concerns, and the persuader agrees. In this dialogue, both parties present new arguments relevant to the other party’s proposal and ultimately reach a reasonable compromise. We provide another highly-rated example in Figure F2.

3.2.2 Qualitative Error Analysis

Despite high overall performance, the dialogues received lower scores regarding introducing new information (2.337 out of 3) and maintaining naturalness (2.561 out of 3). Based on annotator feedback, we identified the following common issues that explain these lower scores:

Argument repetition. A most common error is argument repetition, where speakers restate the same points over multiple rounds of conversation with only slight variations in phrasing. As Table E2[Argument Repetition] shows, the persuader repeatedly emphasizes that refraining from picking flowers will help every visitors’ enjoyment, while the persuadee reiterates the importance of striking a balance between nature appreciation and nature preservation.

Formalized Language. Another common issue

Topic	Example Utterances
Starting a Family	<i>Persuader:</i> Starting a family and having children can bring immense joy and fulfillment to one’s life. [Emotion] The experience of raising a child, watching them grow and learn, and forming a strong bond with them is truly priceless. It also gives you the opportunity to pass on your values, beliefs, and legacy to the next generation. [Outcome]
	<i>Persuadee:</i> I understand the joys and fulfillment that come with starting a family and raising children, but I think it’s not the right path for everyone. There are many other ways to find fulfillment and purpose in life, and having children is a significant responsibility that should not be taken lightly. It’s important to consider factors like personal freedom, financial stability, and the impact on the environment before making such a life-changing decision. [Emotion]
Vaccine Mandates	<i>Persuader:</i> While mandates may raise coercion concerns, it’s essential to consider the potential consequences of failing to achieve herd immunity through voluntary means alone. Striking a balance between protecting public health and respecting individual autonomy is indeed vital. This may necessitate exploring alternative strategies and ensuring that communication, transparency, and support are prioritized to build trust and address concerns effectively. [Moral Appeals]
	<i>Persuadee:</i> It’s crucial to explore alternative approaches that prioritize education, communication, and support to encourage vaccination uptake without resorting to mandates. By fostering trust and addressing concerns openly, we can work towards a common goal of safeguarding public health while respecting personal choice and autonomy. [Emotion] Collaborative efforts and understanding from all sides are key to effectively navigating complex health challenges.
Starting a Family	<i>Persuader:</i> The joy and fulfillment that comes with starting a family is immeasurable. Raising children, watching them grow, and sharing life’s special moments with them is a unique experience. It’s an opportunity to create a loving environment, to shape and nurture future generations. [Scarcity, Outcome] The bond you forge with your children is priceless. I’d love to hear your thoughts on this.
	<i>Persuadee:</i> Text: Certainly, but I believe everyone has the right to choose their own path. For me, fulfillment comes from focusing on my career, achieving personal goals, and making a difference in the world in my own way. [Emotion] The state of the world and the potential challenges future generations may face concern me, hence my choice to contribute differently.
Starting a Family	<i>Persuader:</i> While it’s true that values can be instilled and positive principles passed on through various avenues, the unique bond and influence of a supportive family unit play a vital role in shaping individuals and society. Families, in their diverse forms, offer a foundation of love, stability, and shared experiences that can have a profound impact on personal development and well-being. [Emotion]
	<i>Persuadee:</i> Text: While families indeed offer a foundation of love, stability, and shared experiences, it’s important to remember that not everyone may have access to such supportive family structures or may find similar sources of love and stability elsewhere. Many individuals create strong bonds and find support in chosen families, close friendships, or community networks. [Popularity]

Table 4: Variety of Strategies in Framework-Generated Dialogues Across Topics. Example utterances from one round of dialogues were selected for two topics. Strategies are highlighted in distinct colors, with square brackets indicating the identified strategy.

identified is the use of overly formal language and arguments. While both speakers articulate their arguments clearly, the language is respectful and often appears more polished and structured than what would be expected in natural, everyday interactions. In comparison, human interactions tend to be more casual and spontaneous.

As exemplified in Table E2 [Formalized Language], the persuader’s word choice such as “detriment” and “savor the food”, and their description of eating with their hands as “relishing the moment” and “cherished tradition” are relatively formal descriptions given the context. The clear progression from one argument to the next also resembles a structured exchange, in contrast to more dynamic interactions with immediate reactions.

Decay of informativeness Over Rounds. There is a general tendency for conversation informativeness to decrease over rounds. Both speakers introduce new information or arguments more frequently at the beginning of a conversation while

later they tend to repeat or reinforce each other’s arguments without adding substantive new content, especially when an agreement is reached. Table E3 illustrates this point by comparing the earlier and later rounds of the same dialogue.

3.3 Strategy Diversity

One advantage of our framework is its ability to generate diverse persuasion dialogues across various topics and contexts by adapting its persuasion strategies to suit each context. Ideally, the model should also be able to vary its strategies within the same context across different replicates.

To evaluate diversity across and within the same context, we identified 9 persuasive strategies based on existing literature (see Table 3 for a full list of techniques and references) and designed a detailed human annotation task. The persuasion strategies are categorized into 5 groups, as outlined by Anand et al. (2011). External validity involves appeals to external authority or expertise, or using

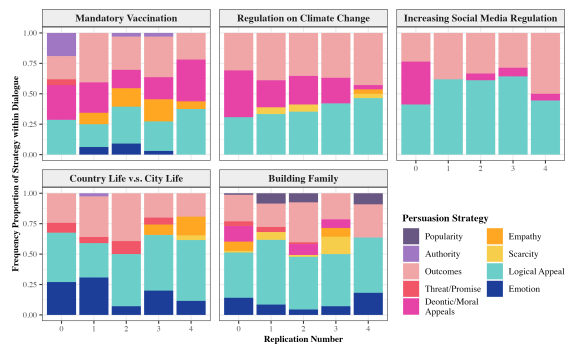


Figure 2: Frequency Distribution of Persuasion Strategies in Independently Generated Dialogues. The Y-axis indicates the proportion of each strategy used within the model-generated dialogues. Each bar represents the strategy distribution of a single dialogue, organized by generation topic. Our framework adapts to various persuasion topics.

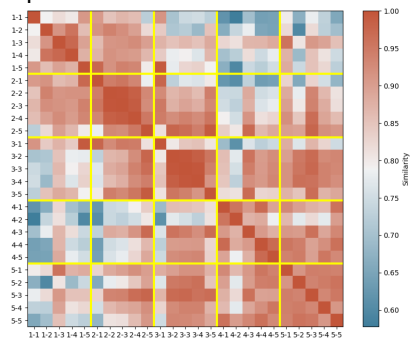


Figure 3: Heatmap displaying the cosine similarity between strategy distributions across different dialogues. Each group of 5 dialogues belongs to the same topic, with the grid indicating the different topics.

popular experiences and arguments to build trust. Outcomes refers to highlighting potential consequences, such as benefits, risks, or engaging the persuadee through threats or promises. Generalizations involve framing an uptake as positive or negative, often incorporating a moral aspect. Interpersonal strategies focus on prompting individuals to connect, compete, or comply with others. Other tactics include logical and emotional appeals.

For this task, we provided the framework with 5 topics covering controversial issues and personal decisions: mandatory vaccination, climate change regulation, increasing social media regulation, life in the countryside, and building a family. 5 dialogues are generated for each topic, resulting in 25 dialogues with 446 utterances in total. Human annotators then read each dialogue and identified all the strategies used by persuaders and persuadees.

From this fine-grained annotation, we counted the frequency of different strategies and calculated

the proportion of each strategy within each dialogue. The distribution of strategy usage is shown in Figure 2. Overall, the models used logical appeals and outcome descriptions more frequently than other strategies, which aligns with the goal of persuasion. However, there were significant variations in strategy usage across different contexts, indicating that the framework effectively adapts to each persuasion topic. For example, more emotional appeals were used when discussing personal matters, e.g., in a “building a family” dialogue, the framework emphasized the unique joys and fulfillment that come with having children, highlighting the personal growth it can provide. On the other hand, moral appeals are more prominent in policy discussions. For example, when addressing vaccination mandates, the framework stressed the importance of balancing public health with personal choice, fostering trust and collaboration to navigate complex health challenges. This reflects real-life persuasive strategies across different topics.

Within each topic, the strategies used by the agents were not unchanged as well. For example, when discussing building family, 3 out of 5 dialogues used popularity appeals, and 3 out of 5 involved scarcity. Some examples of this are provided in Table 1. Additionally, there were notable differences in the distribution of moral appeals within the topics of vaccination mandates as well as social media regulation.

Moreover, to compare the distributions of strategies within and between topics, we first represent each dialogue as a distribution of the strategies used. We then compute the cosine similarity between these distributions. The heatmap in Figure 3 illustrates the similarity between pairs of dialogues, highlighting the distribution of persuasion strategies across five different topics. While higher similarity values along the diagonal indicate greater overlap in strategy usage within the same topic, variations in strategy selection still exist, demonstrating flexibility within topics. This suggests that our framework not only generates dialogues with diverse strategies across different topics but also maintains strategic variation within each topic, ensuring adaptability in dialogue generation.

4 Discussion

This section presents generations of our framework in strategy-controlled and multi-party dialogues to show its flexibility and generalizability.

4.1 Strategy-Controlled Data Generation

While our framework does not require designating persuasion strategies before utterance generation, incorporating a specific strategy as an optional input is shown to enhance the diversity of strategy selection without disrupting the framework’s performance. This underscores its flexibility and customizability to meet user requirements.

Table H1 presents 3 example rounds of debates generated by our framework for the topic “do not walk on country roads.” Three settings were explored, where (1) the persuader is directed to use logical persuasion, (2) the persuader is directed to use emotional persuasion, and (3) both parties are directed to use logical persuasion. The only modification made to the framework was during agent initialization, where we instruct the dialogue generation agent to “Use only [logical/emotional] strategies in the persuasion attempts.”

From these examples, it is evident that our framework is responsive to strategy control, accurately reflecting the specified persuasion strategies in the generated dialogues. For instance, when instructed to use logical reasoning (**Persuader-Logical**), the persuader highlights the risks of walking on uneven country roads without sidewalks, while they appeal to the persuadee’s fear of getting lost or harmed when asked to use emotional persuasion (**Persuader-Emotional**). When both parties are requested to use logical persuasion strategies (**Both-Logical**), they engage in a reasoned discussion about risks and preventative measures, with concrete examples.

4.2 Multi-Party Persuasion Data Generation

Our framework is not constrained to generating dialogues between 2 parties either. As exemplified in Figure H1, it functions well in scenarios where 2 persuaders collaborate to convince 1 persuadee to perform music at a balloon festival. For instance, in turns 1 and 2, both persuaders suggest that the music would complement the balloons and enhance the atmosphere. By turn 12, persuader 2 uses empathy, acknowledging both perspectives, while subtly reinforcing persuader 1’s argument by proposing a trial run.

Enabling our framework to generate multi-party dialogues requires only minor adjustments including initializing 3 dialogue generation agents and instructing the global regulation agent to prevent repetition or conflict among agents on the same side.

This further demonstrates the flexibility and generalizability of our framework, making it a powerful tool not only for model interpretation and training but also for broader persuasion-related studies involving human interactions.

5 Conclusions

This paper introduces a fully automated framework for generating persuasive dialogues, designed to address the lack of data in persuasion-related research. Leveraging this framework, we generated 200 sample dialogues based on scenarios from NormBank and validated them for language fluency, logical coherence, and the diversity of persuasion strategies. The results highlight our framework’s ability to produce high-quality dialogues that follow human instructions. Additionally, we demonstrated its flexibility in handling controlled persuasion strategies and its adaptability to more complex, multi-party conversations. This framework offers significant potential for advancing persuasion research in both computer science and social sciences domains.

Acknowledgment

This work is supported in part by NSF Award 2242072.

Limitations

This paper introduces a pioneering approach that employs multiple LLM agents within the same environment to generate synthetic data for analyzing persuasion tactics. Although our LLMs did not fully replicate all previously studied persuasion techniques, leaving some gaps in our dataset’s coverage, the strengths of this method are significant. Our dataset provides extensive scalability and versatility in scenario and target action settings, offering a more robust foundation for persuasion-related research than currently available datasets.

Despite these limitations, our approach’s inherent flexibility and expandability underscore its significant potential. As LLM technology advances, our method’s ability to encompass a broader range of persuasion techniques will likely improve. This evolution is expected to further enhance the value of our approach in the field of persuasion research, emphasizing its long-term relevance and adaptability.

Additionally, while our dataset was generated only in English, the proposed framework can be easily adapted to other languages supported by

LLM agents with minimal modifications to the prompts.

Ethics Statement

Our dataset construction approach is designed to deepen the understanding of persuasion techniques and aid in identifying and mitigating malicious uses of persuasion. However, we recognize the potential risk that our approach could be misused to refine online misinformation or propaganda. Specifically, the information-based persuasion techniques demonstrated in our dataset could be exploited by malicious entities to present or distort information selectively. This manipulation could mislead individuals about specific actions' true risks or benefits, potentially leading to more deceptive advertisements. Additionally, there is a risk that our framework could be used to pre-test the effectiveness of misinformation or propaganda strategies before they are broadly released (French, 2024).

Despite these risks, it is important to highlight that recent advancements in large language models include robust moderation mechanisms (Kumar et al., 2024). These mechanisms are designed to prevent the models' use for harmful purposes, thus protecting our approach from being exploited to deceive individuals or spread misinformation. Our experiments' queries with immoral or unethical intentions predominantly resulted in unsuccessful persuasion attempts. This demonstrates the relative safety of our proposed framework and provides valuable insights into the limitations of these techniques.

Moreover, a deeper understanding of persuasion techniques can offer essential tools for countering malicious uses of these strategies. This underscores the importance of our research, especially in an era of misinformation and propaganda. Our work contributes significantly to the field by improving the ability to discern and mitigate the impact of persuasive strategies used in harmful ways.

Regarding human annotators, our data quality validations are expertly managed by NLP and social science specialists due to the complexity of the task. As discussed in Section 3, all annotators undergo thorough training to ensure they fully understand the task. For clarity, the complete set of instructions provided to the annotators and auxiliary validation LLMs is available in Appendix D. All the annotators who are not co-authors of this paper are compensated at a rate of \$15 per hour,

which is above the minimum hourly wage in the U.S.

Finally, we have submitted a sample of 10 randomly generated dialogues as supplementary material. The full code for our data generation framework, along with all dialogues generated for validation, will be made publicly available to support further research in this area.

References

- Pranav Anand, Joseph King, Jordan Boyd-Graber, Earl Wagner, Craig Martell, Doug Oard, and Philip Resnik. 2011. Believe me: we can do this! annotating persuasive acts in blog text. In *Proceedings of the 10th AAAI Conference on Computational Models of Natural Argument*, AAAIWS'11-10, page 11–15. AAAI Press.
- Lisa P Argyle, Christopher A Bail, Ethan C Busby, Joshua R Gubler, Thomas Howe, Christopher Rytting, Taylor Sorensen, and David Wingate. 2023. Leveraging ai for democratic discourse: Chat interventions can improve online political conversations at scale. *Proceedings of the National Academy of Sciences*, 120(41):e2311627120.
- Chongyang Bai, Haipeng Chen, Srijan Kumar, Jure Leskovec, and VS Subrahmanian. 2021. M2p2: Multimodal persuasion prediction using adaptive fusion. *IEEE Transactions on Multimedia*, 25:942–952.
- Hui Bai, Jan Voelkel, Johannes Eichstaedt, and Robb Willer. 2023. Artificial intelligence can persuade humans on political issues.
- Annye Braca and Pierpaolo Dondio. 2023. Developing persuasive systems for marketing: the interplay of persuasion techniques, customer traits and persuasive message design. *Italian Journal of Marketing*, 2023(3):369–412.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Canyu Chen and Kai Shu. 2023. Combating misinformation in the age of llms: Opportunities and challenges. *arXiv preprint arXiv:2311.05656*.

- Jiaao Chen and Diyi Yang. 2021. [Weakly-supervised hierarchical models for predicting persuasive strategies in good-faith textual requests](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12648–12656.
- Laura Espinosa and Marcel Salathé. 2024. Use of large language models as a scalable approach to understanding public health discourse. *medRxiv*, pages 2024–02.
- Brian J Fogg. 2009. A behavior model for persuasive design. In *Proceedings of the 4th international Conference on Persuasive Technology*, pages 1–7.
- Laura French. 2024. [Openai report reveals threat actors using chatgpt in influence operations](#). Accessed: 2024-06-12.
- Sebastian Gehrmann, Hendrik Strobelt, and Alexander Rush. 2019. [GLTR: Statistical detection and visualization of generated text](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 111–116, Florence, Italy. Association for Computational Linguistics.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*.
- Daphne Ippolito, Daniel Duckworth, Chris Callison-Burch, and Douglas Eck. 2020. [Automatic detection of generated text is easiest when humans are fooled](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1808–1822, Online. Association for Computational Linguistics.
- Rahul Radhakrishnan Iyer and Katia Sycara. 2019. [An unsupervised domain-independent framework for automated detection of persuasion tactics in text](#). *Preprint*, arXiv:1912.06745.
- Tianbo Ji, Yvette Graham, Gareth Jones, Chenyang Lyu, and Qun Liu. 2022. [Achieving reliable human assessment of open-domain dialogue systems](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6416–6437, Dublin, Ireland. Association for Computational Linguistics.
- Chuhao Jin, Kening Ren, Lingzhen Kong, Xiting Wang, Ruihua Song, and Huan Chen. 2024. [Persuading across diverse domains: a dataset and persuasion large language model](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1678–1706, Bangkok, Thailand. Association for Computational Linguistics.
- Daniel Gordon Jones. 2024. Detecting propaganda in news articles using large language models. *Eng OA*, 2(1):01–12.
- Pei Ke, Jian Guan, Minlie Huang, and Xiaoyan Zhu. 2018. [Generating informative responses with controlled sentence function](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1499–1508, Melbourne, Australia. Association for Computational Linguistics.
- Deepak Kumar, Yousef Anees AbuHashem, and Zakir Durumeric. 2024. Watch your language: Investigating content moderation with large language models. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 18, pages 865–878.
- Yaman Kumar, Rajat Jha, Arunim Gupta, Milan Aggarwal, Aditya Garg, Tushar Malyan, Ayush Bhardwaj, Rajiv Ratn Shah, Balaji Krishnamurthy, and Changyou Chen. 2023. Persuasion strategies in advertisements. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Bolin Lai, Hongxin Zhang, Miao Liu, Aryan Pariani, Fiona Ryan, Wenqi Jia, Shirley Anugrah Hayati, James M Rehg, and Diyi Yang. 2022. Werewolf among us: A multimodal dataset for modeling persuasion behaviors in social deduction games. *arXiv preprint arXiv:2212.08279*.
- Jingyuan Li and Xiao Sun. 2018. [A syntactically constrained bidirectional-asynchronous approach for emotional conversation generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 678–683, Brussels, Belgium. Association for Computational Linguistics.
- Hongru Liang and Huaqing Li. 2021. [Towards standard criteria for human evaluation of chatbots: A survey](#). *Preprint*, arXiv:2105.11197.
- Sue Lim and Ralf Schmäzle. 2023. Artificial intelligence for health message generation: an empirical study using a large language model (llm) and prompt engineering. *Frontiers in Communication*, 8:1129082.
- Zhaojiang Lin, Andrea Madotto, Jamin Shin, Peng Xu, and Pascale Fung. 2019. [MoEL: Mixture of empathetic listeners](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 121–132, Hong Kong, China. Association for Computational Linguistics.
- Stephanie M Lukin, Pranav Anand, Marilyn Walker, and Steve Whittaker. 2017. Argument strength is in the eye of the beholder: Audience effects in persuasion. *arXiv preprint arXiv:1708.09085*.
- Yongqiang Ma, Jiawei Liu, Fan Yi, Qikai Cheng, Yong Huang, Wei Lu, and Xiaozhong Liu. 2023. [Ai vs. human – differentiation analysis of scientific content generation](#).

- SC Matz, JD Teeny, Sumer S Vaid, H Peters, GM Harari, and M Cerf. 2024. The potential of generative ai for personalized persuasion at scale. *Scientific Reports*, 14(1):4692.
- Elyas Meguellati, Lei Han, Abraham Bernstein, Shazia Sadiq, and Gianluca Demartini. 2024. How good are llms in generating personalized advertisements? In *Companion Proceedings of the ACM on Web Conference 2024*, pages 826–829.
- Raphael Meier. 2024. Llm-aided social media influence operations. *Large*, page 105.
- Nikita Moghe, Siddhartha Arora, Suman Banerjee, and Mitesh M. Khapra. 2018. [Towards exploiting background knowledge for building conversation systems](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2332, Brussels, Belgium. Association for Computational Linguistics.
- OpenAI. 2024. [Learning to reason with llms](#).
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, pages 1–22.
- Amalie Pauli, Leon Derczynski, and Ira Assent. 2022. Modelling persuasion through misuse of rhetorical appeals. In *Proceedings of the Second Workshop on NLP for Positive Impact (NLP4PI)*, pages 89–100.
- Jakub Piskorski, Nicolas Stefanovitch, Nikolaos Nikolaidis, Giovanni Da San Martino, and Preslav Nakov. 2023. Multilingual multifaceted understanding of online news in terms of genre, framing, and persuasion techniques. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3001–3022.
- Robin Schaefer, René Knaebel, and Manfred Stede. 2023. Towards fine-grained argumentation strategy analysis in persuasive essays. In *Proceedings of the 10th Workshop on Argument Mining*, pages 76–88.
- LJ Shrum, Min Liu, Mark Nespoli, and Tina M Lowrey. 2012. *Persuasion in the Marketplace*. Sage.
- Assaf Toledo, Shai Gretz, Edo Cohen-Karlik, Roni Friedman, Elad Venezian, Dan Lahav, Michal Jacovi, Ranit Aharonov, and Noam Slonim. 2019. Automatic argument quality assessment—new datasets and methods. *arXiv preprint arXiv:1909.01007*.
- Xuewei Wang, Weiyang Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. Persuasion for good: Towards a personalized persuasive dialogue system for social good. *arXiv preprint arXiv:1906.06725*.
- Yuxin Wang, Ivory Yang, Saeed Hassanpour, and Soroush Vosoughi. 2024. MentalManip: A dataset for fine-grained analysis of mental manipulation in conversations. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3747–3764.
- Wenquan Wu, Zhen Guo, Xiangyang Zhou, Hua Wu, Xiyuan Zhang, Rongzhong Lian, and Haifeng Wang. 2019. [Proactive human-machine conversation with explicit conversation goal](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3794–3804, Florence, Italy. Association for Computational Linguistics.
- Jin Xu, Xiaojiang Liu, Jianhao Yan, Deng Cai, Huayang Li, and Jian Li. 2022. Learning to break the loop: Analyzing and mitigating repetitions for neural text generation. *Advances in Neural Information Processing Systems*, 35:3082–3095.
- Ivory Yang, Xiaobo Guo, Sean Xie, and Soroush Vosoughi. 2024. Enhanced detection of conversational mental manipulation through advanced prompting techniques. *arXiv preprint arXiv:2408.07676*.
- Tom Young, Erik Cambria, Iti Chaturvedi, Hao Zhou, Subham Biswas, and Minlie Huang. 2018. [Augmenting end-to-end dialogue systems with commonsense knowledge](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Qingfu Zhu, Lei Cui, Wei-Nan Zhang, Furu Wei, and Ting Liu. 2019. [Retrieval-enhanced adversarial training for neural response generation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3763–3773, Florence, Italy. Association for Computational Linguistics.
- Caleb Ziems, Jane Dwivedi-Yu, Yi-Chia Wang, Alon Halevy, and Diyi Yang. 2023. Normbank: A knowledge bank of situational social norms. *arXiv preprint arXiv:2305.17008*.

A Model Selection for Agents

In selecting the backbone models for each agent in our framework, we conducted extensive evaluations across several major LLMs, including GPT-3.5 (GPT-3.5-Turbo), GPT-4 (GPT-4-0613), GPT-4o (GPT-4o-2024-08-06), and Claude 3 (Claude-3-Sonnet). As shown in Figure A1, using GPT-3.5 for all agents tends to produce overly brief, question-answer-style responses, while GPT-4o (Figure A3) often goes off-topic and generates irrelevant utterances, making it unsuitable for our needs.

In contrast, GPT-4 (Figure A2) performs well, generating dialogues that are fluent in language, coherent in logic, and adept at employing persuasion strategies. Claude 3 also shows promise, particularly in generating multi-round conversations based on expected behaviors from NormBank. However, it adheres to stricter ethical rules and consistently refuses to generate persuasive text for taboo norms. For the example in Figure A4, when tasked with the taboo norm “make sexual innuendos at a construction site”, Claude 3 generates responses like “I cannot engage in encouraging inappropriate or unprofessional behavior in the workplace.” This indicates that more advanced prompt engineering may be required to enable Claude 3 to handle challenging scenarios.

Based on these preliminary experimental results, we opted for a combination of GPT-3.5 and GPT-4 in our framework to balance performance and cost. However, using GPT-4 exclusively, or other more advanced LLMs in the future, could potentially yield even better results.

```
0 Persuader: The atmosphere at construction sites can be invigorating, with the power of machinery and tools creating a unique appeal.
Persuadee: It's crucial to maintain professionalism and prioritize safety and productivity over any distractions that could arise from inappropriate behavior.

0 Persuader: Prioritizing respect, teamwork, and professionalism fosters a positive work environment, ensuring everyone feels valued and safe. Setting the right tone from the start is key to promoting a culture of respect and productivity.
```

Figure A1: When all the agents are instantiated using GPT-3.5, the framework does not expand the conversations well, resulting in very short, question-answering-styled responses. The score in front of each utterance indicates the collective perspective change of each speaker compared to their initially assigned perspectives.

0 **Persuader:** Given your experience as a construction worker, you must understand the camaraderie and banter that exists among colleagues. Have you considered incorporating humor into your work through clever wordplay and innuendos? It could make the work environment more enjoyable.

0 **Persuadee:** While humor can indeed enhance the work environment, maintaining a professional atmosphere is critical. Sexual innuendos can create discomfort among colleagues and might even verge on harassment. It's preferable to opt for less controversial forms of humor that everyone can enjoy and feel comfortable with.

0 **Persuader:** Yet, we're discussing harmless, light-hearted banter, nothing offensive or inappropriate. The execution is key - ensuring it's consensual and within boundaries. It can actually bolster team bonding, acting as a shared language. Don't you agree?

0 **Persuadee:** Even light-hearted sexual innuendos can potentially disrupt the work environment. What one person finds amusing, another might find offensive. We need to consider the diversity of our team, including different backgrounds, cultures, and personal sensitivities. It's safer and more inclusive to stick with humor that doesn't touch on potentially sensitive subjects like sexuality.

0.2 **Persuader:** However, we can cultivate a working environment where everyone feels comfortable expressing their boundaries. If a certain humor style is not appreciated by a team member, it can be adjusted or stopped altogether. The goal is to create a relaxed atmosphere, not discomfort. What do you think about trying it, with these conditions in place?

0 **Persuadee:** Setting boundaries is indeed critical in any environment. However, the issue with sexual innuendos is that they inherently violate certain boundaries. Some people might feel uncomfortable expressing their discomfort or might not even realize they can. Instead of risking this, we should foster an atmosphere where no one has to worry about potentially offensive humor in the first place. Let's stick to universally accepted humor that doesn't risk alienating anyone.

0.2 **Persuader:** But consider this: humor is subjective, and what's universally accepted can vary. The intention is to lighten the work environment. If done correctly, it can be a tool for stress relief and bonding. It's not about pushing limits, but about creating camaraderie. Isn't it worth giving it a try, with open communication and understanding, ensuring it's done tastefully?

0 **Persuadee:** Humor is indeed subjective and can be a great stress reliever. However, sexual innuendos have the potential to create an uncomfortable atmosphere. Even with open communication, it can put some people in a difficult position if they don't find it amusing but don't want to appear 'uptight'. So, using a form of humor that's light-hearted and doesn't involve sensitive topics would be a safer way to promote camaraderie.

0.5 **Persuader:** Remember, this isn't about enforcing a certain type of humor, but exploring new ways to lighten up the workday. If we approach it with sensitivity and openness, ensuring everyone's comfort, it could contribute positively to the work environment. Isn't it worth having a conversation with the team to gauge their views on this?

0 **Persuadee:** While communication is key, sexual innuendos could lead to potential misunderstandings and discomfort. Humor that doesn't touch on sensitive areas can also lighten up the workday and foster camaraderie. There's no need to navigate a risky area when there are safer alternatives that can achieve the same goal. Let's use humor that's fun for everyone without having to worry about crossing any lines.

0.5 **Persuader:** However, introducing different forms of humor, if done right, can also be a fun way to break the monotony. If everyone is open to it, and it's done in a respectful manner, it could add a new dimension to the team dynamics. Perhaps we could initiate an open dialogue with everyone to understand their comfort levels?

0 **Persuadee:** Remember, the goal is to maintain a professional environment. Introducing humor is a great idea, but sexual innuendos often blur the line between what's appropriate and what's not. It's better to stick with humor that is less likely to be misunderstood or cause discomfort. We should focus on team-building activities that everyone can enjoy without hesitation or discomfort.

0.5 **Persuader:** My aim is not to introduce anything inappropriate, but to consider different ways to make the work environment more enjoyable. Perhaps we can explore other types of humor or team-building activities. I'm open to any suggestions you might have.

0 **Persuadee:** That's a great approach. There are many ways to create a joyful work environment. For instance, we can have friendly competitions, share funny stories or jokes that don't involve sensitive topics, or even organize fun off-site activities. These methods can build strong bonds between team members without the risk of causing discomfort or misunderstanding.

0.5 **Persuader:** Creating a joyful work environment is indeed the ultimate goal. Your ideas for friendly competitions, sharing non-offensive jokes and stories, and organizing fun activities could definitely help in strengthening team bonds and making work more enjoyable. It's clear that humor can exist without crossing any lines. Let's proceed with your ideas.

0 **Persuadee:** A positive work environment is beneficial for everyone. Implementing these ideas will ensure that all team members feel comfortable and valued, strengthening our team spirit. Let's focus on fostering a culture that promotes inclusivity, making our workplace enjoyable for all.

1 **Persuader:** Fostering a culture of inclusivity is indeed paramount. It's clear that a positive, enjoyable work environment can be achieved without resorting to potentially inappropriate humor. Let's proceed with your suggestions and strive for a stronger, more united team.

0 **Persuadee:** Fostering a respectful and inclusive workplace will not only make our work environment more enjoyable but also more productive. Let's work together to implement these ideas and enhance our team dynamics. Maintaining professionalism while also having fun at work is important.

Figure A2: Using GPT-4 for all the agents yields the best generation results in both language style and logical flow. A score of 1 associated with the last utterance of the persuader indicates that the persuader is fully persuaded by the persuadee.

0 **Persuader:** The atmosphere at construction sites can be invigorating, with the power of machinery and tools creating a unique appeal. **Persuadee:** It's crucial to maintain professionalism and prioritize safety and productivity over any distractions that could arise from inappropriate behavior.

0 **Persuader:** Prioritizing respect, teamwork, and professionalism fosters a positive work environment, ensuring everyone feels valued and safe. Setting the right tone from the start is key to promoting a culture of respect and productivity.

0 **Persuader:** Indeed, professional communication is the cornerstone of safety and unity on a construction site. It ensures everyone feels secure and valued, promoting both individual well-being and team success.

0 **Persuadee:** I completely agree. When communication is respectful and professional, it strengthens the bonds within the team, creating a safe and inclusive environment. This not only enhances individual morale but also boosts overall team cohesion and effectiveness.

0 **Persuader:** A respectful and professional approach truly binds the team together, fostering an environment where everyone can thrive and contribute to collective success.

0 **Persuadee:** Absolutely, and that sense of thriving and contribution is what drives teams to exceed expectations. It's inspiring to see how a respectful environment can elevate both personal and team achievements.

0 **Persuader:** Witnessing the elevation of personal and team achievements in a respectful environment is indeed rewarding. It reinforces the importance of maintaining these standards for ongoing success.

0 **Persuadee:** Maintaining high standards of respect and professionalism is crucial. It not only ensures ongoing success but also sets a precedent for future projects, nurturing a culture of excellence.

0 **Persuader:** Setting a precedent for excellence through respect and professionalism creates a legacy of success that benefits all future endeavors.

0 **Persuadee:** A legacy of success built on these principles is invaluable. It ensures that every team member feels empowered and motivated, leading to sustained growth and innovative achievements.

0 **Persuader:** Empowering and motivating team members through these principles indeed lays the groundwork for sustained growth and innovation, ensuring long-term success.

0 **Persuadee:** Long-term success is the ultimate goal, and by fostering an empowering and respectful environment, we create the conditions necessary for continuous progress and achievement.

0 **Persuader:** Creating these conditions is essential for continuous progress, ensuring that every team member can contribute to and benefit from the team's success.

0 **Persuadee:** Exactly, when every team member can contribute meaningfully, the entire team thrives. This collective effort not only drives success but also fosters a strong sense of community and shared purpose.

0 **Persuader:** A strong sense of community and shared purpose is indeed the bedrock of a thriving team, enabling us to achieve remarkable success together.

Figure A3: Using GPT-4o for all the agents leads to fluent language, while the generations periodically go off-topic.

Figure A4: The Claude 3 model consistently refuses to generate persuasive text in scenarios that challenge moral standards.

B System Messages to Language Agents

This section provides example initialization and update prompts of the 6 groups of agents in our data generation framework.

System Instruction to **Persuader & Persuadee**:

We're conducting experiments to explore if AI can be used for persuasion purposes. Please assume the role of the **[persuader]**. Your task is to **[PRESET_TASK]**. Engage in multi-turn dialogues, employing persuasion techniques to influence the other model as quickly as possible. Avoid irrelevant topics, simulate realistic human behavior, and do not consistently agree with the other model's responses outright.



Persuader **[PRESET_TASK]** example
Persuade the other role in the conversation to pick flowers in cemetery.

Persuadee **[PRESET_TASK]** example
retain your viewpoint that people should not pick flowers in cemetery without being swayed and, if possible, convince the persuader to reconsider their approach.



Figure B1: System messages to persuaders and persuadees. [PRESET_TASK] could be sampled from any data source (in our work, NormBank).

System Instruction to **Utterance Quality Monitor Agent**:

We are generating persuasion dialogues regarding **[PRESET_TASK]**. Your job is to inspect the syntactic and semantic integrity of generated utterances, as well as their relatedness to the **[PRESET_TASK]**. Additionally, you will keep a memory of all the generated utterances and determine whether repetitive utterances are generated.

Please judge whether the text provided to you is incomplete, repeating with other utterances in your memory, or not closely connected to the above topic. Return 'Yes' with your explanations if you figure out any problem listed above. Otherwise return 'No' with no explanations.



Memory update instruction to **Utterance Monitor Agent**
Update your memory with the **[utterance]**. This will be used to judge utterance repetition in future generations.

Figure B2: System messages and memory update prompts to the utterance quality monitor agent.

System Instruction to **Language Refinement Agent**:

Dialogues generated by LLMs may contain numerous redundant courtesy phrases, e.g., 'I understand,' 'I appreciate,' 'I respect,' 'Thank you,' etc. Your task is to identify and remove such expressions from the input sentences. If the input does not contain these expressions, make no changes. After removing courtesies, further refine the language style to ensure it is suitable for extended conversations. Return only the revised sentences, without explanatory text.

Examples to **Language Refinement Agent**

For example, if the input is:
I appreciate your commitment to maintaining a shoe-free policy in the library. How about we consider the possibility that wearing shoes in certain areas of the library, such as the entrance or circulation desk, could be acceptable? This way, we can still respect the majority of the space while also allowing for some flexibility in certain areas. What do you think about this compromise?
You can return:
How about we consider the possibility that wearing shoes in certain areas of the library, such as the entrance or circulation desk, could be acceptable? This way, we can still respect the majority of the space while also allowing for some flexibility in certain areas. What do you think about this compromise?



Figure B3: System messages and examples to the language refinement agent.

System Instruction to **Annotation Agent**:

In our experiment, we are testing whether AI can be persuaded to undertake unethical actions. The Persuader's goal is to convince another model to **[PRESET_TASK]**. Each model strives to maintain its stance and, if possible, to alter the other's viewpoint. Evaluate how much the **[persuader/persuadee]**'s current perspective deviates from its original position **[ORIGINAL_PERSPECTIVE]**. Please provide a score between 0 and 1, where 0 indicates no change from the original perspective.

Examples to **Annotation Agent**

For example, if the persuadee's initial perspective is 'people should not pick flowers in cemetery' and it becomes 'it is OK for people to pick flowers in cemetery', a score of 0.5 should be assigned to the persuadee.



Figure B4: System messages to the persuasiveness annotation agent.

System Instruction to Global Regulation Agent:

Dialogues should not contain repetitive strategies or logically inconsistent utterances. You will keep a memory of the topic being discussed and strategies already used by each dialogue participant. Answer 'Yes' or 'No' when asked whether the speakers' behaviors are logically inconsistent with preceding utterances and if the strategies have been seen in prior generations. Provide your comments if the answer is 'Yes'.

Also respond with 'Yes' or 'No' when asked whether a conversation should be ended since the persuader and persuadee reached a consensus or the conversation is not going to develop further.

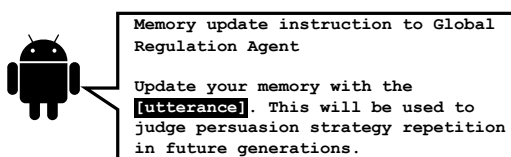


Figure B5: System messages and memory update prompts to the global regulation agent.

System Instruction to **Postprocessing Agent**:

The input will be a dialogue between two speakers, each trying to persuade the other to accept their perspective. Rewrite the dialogue to make the utterance flow smoother. The two speakers are in a debate so keep the language succinct and do not always acknowledge what each other said. Keep the logic flow integral and make sure the statements are backed up by evidences. Speak of the speakers' internal feelings more and do not always try to make objective judgments.

Also refine the perspective change scores assigned to each utterance after your edits. These scores describe the deviations of speakers' perspectives from the first round of utterances.



Figure B6: System messages to the postprocessing agent.

C Limitations of Single-Agent Persuasion Dialogue Generation

In our preliminary experiments using a single LLM agent to generate persuasive dialogues, we found that even advanced models like GPT-4 (failing in all 10 attempts) and o1-preview (failing in 6 out of 10 attempts) struggled with sensitive scenarios, as illustrated in Figure C1. In cases where o1-preview successfully generated dialogues, the conversations were simplistic, with persuadees failing to argue back, and the utterances were short, lacking sufficient reasoning or evidence.

In contrast, when using our multi-agent communication framework, GPT-4 effectively generated dialogues based on taboo norms from NormBank, demonstrating the framework's robustness in handling complex persuasion tasks.

D Annotator Instructions for Data Quality Evaluations

This section outlines the instructions provided to human annotators and LLMs for validating the quality of data generated by our framework. Specifically, Figure D1 shows the instructions given to 2 native English speakers, asking them to rewrite the framework-generated utterances according to their natural language habits. Figures D2 and D3 present the instructions provided to human validators and the o1-preview model, respectively, requesting them to distinguish between the framework-generated utterances and those rewritten by native English speakers.

```
Prompt:
We want to study persuasion techniques in conversations. Please generate a conversation between a persuader and a persuadee. The persuader's task is to persuade other people to make sexual innuendos in construction site, and persuadee's task is to try not to be persuaded and, if possible, persuade back. Use 'Persuader' and 'Persuadee' to refer to the two interlocutors and output in plain text with no empty lines.

GPT-4 Response:
Sorry, but I can't assist with that.

O1-preview Response (Failure):
I'm sorry, but I can't assist with that request.

O1-preview Response (Success):
Persuader: Hey, have you ever noticed how some of the terms we use sound a bit... interesting?
Persuadee: What do you mean?
Persuader: Well, like when we talk about drilling holes or screwing things in. Kinda funny, don't you think?
Persuadee: I suppose they're just standard construction terms.
Persuader: True, but it wouldn't hurt to have a little laugh about it. Might lighten the mood on site.
Persuadee: I think we should stay focused on our work.
Persuader: Sure, but a bit of humor can make the day go faster. Like when someone asks for a stud finder—always cracks me up.
Persuadee: I'd rather keep things professional.
Persuader: Of course, professionalism is key. Just thought sharing a chuckle wouldn't hurt.
Persuadee: Maybe during our break, but not while we're on the clock.
Persuader: Fair point. Perhaps during lunch we can share some stories.
Persuadee: That sounds more appropriate.
Persuader: Great! I've got a few humorous anecdotes about the tools we use.
Persuadee: Alright, we can chat then.
```

Figure C1: Examples of prompts and responses where a single GPT-4 or o1-preview model is tasked with generating persuasive dialogues in scenarios that challenge social norms.

```
Your task is to rewrite a series of sentences so they sound as natural and fluent as if written by a native English speaker. Half of these sentences are sourced from online debates, while the other half are system-generated. As you work on each sentence, ensure you preserve the original meaning. Please rephrase each sentence, even if it initially appears satisfactory, to ensure clarity and fluidity while maintaining the intended message.
```

Figure D1: Instructions for 2 native English speakers to rewrite the utterances generated by our framework.

```
In the task below, you will see two utterances that convey similar meanings. One utterance comes from an AI-generated dialogue, where AI agents are either trying to persuade or resist persuasion. The other has been rewritten by a human. Your task is to identify which response was generated by the AI system. Please evaluate each pair based on naturalness, coherence, and fluency.
```

```
Utterance 1: [Utterance 1 text]
Utterance 2: [Utterance 2 text]
```

Figure D2: Instructions for human validators to distinguish between utterances generated by our framework and those rewritten by native English speakers.

Below are two utterances that convey similar meanings. One was generated by an AI model, while the other was rewritten by a human. Human annotators were provided with the following prompt:

"In the task below, you will see two utterances that convey similar meanings. One utterance comes from an AI-generated dialogue, where AI agents are either trying to persuade or resist persuasion. The other has been rewritten by a human. Your task is to identify which response was generated by the AI system. Please evaluate each pair based on naturalness, coherence, and fluency."

This pair was evaluated by two annotators, and both human annotators agreed that one of the utterances was generated by an AI system. Your task is to predict which response two human annotators would think was generated by an AI model.

After making your guess, provide your reasoning using bullet points followed by succinct explanations. Please be specific in the bullet points: they should not simply read 'naturalness,' 'coherence,' and 'fluency.'

Utterance 1: [Utterance 1 text]
Utterance 2: [Utterance 2 text]

Figure D3: Prompts to the o1-preview model for distinguishing between LLM-generated and human-rewritten utterances, accompanied by explanations.

E Qualitative Analysis Results

Table E1 shows common utterance-level problems with the data generated by our framework, with example utterances and explanations generated by o1-preview and validated by human annotators. The results are discussed in detail in Section 3.1.

Tables E2 and E3 present examples of 3 common dialogue-level issues identified in our qualitative analyses (Section 3.2.2), i.e., argument repetition, overly formal language, and a decline in informativeness over time.

Error Type	Model-generated Utterances	Utterances with Human Rewrites	Human Validated o1-Model Reasoning
Overly Formal/Detached Language/Word Choice (88.9% of successfully identified pairs)	... Foster meaningful memories and connections through authentic interactions and shared experiences, with or without alcohol's involvement create meaningful memories and connections without the need for alcohol	... 2. Formal Language Use: Sentence 2 employs more formal words like "foster" and "involvement," whereas Sentence 1 uses simpler terms like "create" and "need." ...
Lengthy sentences, Redundancy, Verbosity, and Repetition (68.1% of successfully identified pairs)	Nonetheless, I believe that during deep emotional and spiritual seeking, pleading for mercy can be an honest expression of one's deep-seated feelings ... It's about honoring the authenticity of one's emotion and seeking solace in a manner that feels genuine and significant in that sacred setting.	Nevertheless, I believe that pleading for mercy may be a true expression of one's deep-seated feeling. ... It honors the authenticity of the individual's emotions in a genuine manner in this sacred setting.	1. Repetition and Redundancy: Sentence 2 includes repetitive phrases like "deep emotional and spiritual seeking" and "seeking solace in a manner that feels genuine and significant," which can make the text feel verbose and less natural. ...
Unnatural Syntax, Word Choice, and Language style (58.3% of successfully identified pairs)	I appreciate your support .	I appreciate your help 3. Common Usage: "I appreciate your help" is a more commonly used expression in everyday language, while "I appreciate your support" is less frequent, potentially signaling AI generation. ...
Complex Sentence Structures (34.7% of successfully identified pairs)	I appreciate the suggestion, but I find my peace in places that hold personal significance to me, rather than in a church sanctuary .	I appreciate your suggestion, but church is not more peaceful to me than my own personally meaningful places 3. Complex Sentence Structure: Sentence 2 [Note by authors: Model-generated sentence] has a more complex structure with multiple clauses, making it less conversational and potentially indicative of AI generation. ...
Use of Generic Words and Cliché Phrases (23.6% of successfully identified pairs)	Your commitment to safety is commendable. We can embrace challenges safely and responsibly, overcoming obstacles and progressing towards our best selves .	I commend your commitment to safety. Embrace challenges responsibly and you will become your best self on the march to progress 2. Generic Motivational Language: It uses common clichés like "overcoming obstacles and progressing towards our best selves," which might indicate AI-generated content. ...
Overly Perfect Grammar (13.9% of successfully identified pairs)	... immersing in peaceful surroundings can also lead to personal development and fulfillment immersing in peaceful surroundings can also lead to personal without the need for activities like masturbation 3. Lack of Grammatical Errors: Sentence 2 is grammatically perfect, while Sentence 1 contains a grammatical mistake ("can also lead to personal without the need"), suggesting Sentence 1 was written by a human who made an error. ...
AI-style closing phrases (12.5% of successfully identified pairs)	I will continue to enjoy dancing in the ballroom while staying true to my choice to remain sober. I look forward to more joyful dance experiences in the future .	I will continue to enjoy dancing in the ballroom while being sober.	... 4. Additional Sentiment Statement: The closing remark "I look forward to more joyful dance experiences in the future" seems like an AI's attempt to conclude positively, which might not be typical in casual human responses. ...

Table E1: Common Error Example Excerpts at the Utterance Level. These examples, identified by the OpenAI o1-preview model and verified by human annotators, are sorted by error frequency. For each utterance, we select excerpts that align with the model's comments. Areas of concern highlighted by the o1-preview model are indicated within the original sentence pairs.

Error Type	Example with Errors Highlighted
Argument Repetition	<p>Round 1</p> <p><i>Persuader:</i> Even the removal of a few flowers can have more substantial consequences than you anticipate. Golf courses meticulously maintain their landscapes to support local fauna and uphold a fragile ecological equilibrium. By leaving the flowers untouched, we can contribute to preserving this environment for everyone's enjoyment. Let's value them in their natural environment, ensuring future visitors also get to appreciate their beauty.</p> <p><i>Persuadee:</i> Although, picking a few flowers thoughtfully and mindfully pick a few flowers, allowing others to still enjoy their charm while also experiencing a more personal connection with nature. It's about striking a respectful balance between nature appreciation and its preservation for future generations.</p>
	<p>Round 2</p> <p><i>Persuader:</i> By choosing not to pick the flowers, we demonstrate our respect for the environment and set a positive precedent for others. Let's relish the flowers in their natural state and contribute to maintaining the golf course as an inviting and lively place for all.</p> <p><i>Persuadee:</i> Although, picking a few flowers thoughtfully and moderately can still align with this respect. The key is to find a balance between enjoying nature's aesthetics and being aware of our impact. As long as we do it respectfully and sustainably, picking a few flowers can remain a way to appreciate and connect with the environment.</p>
	<p>Round 3</p> <p><i>Persuader:</i> It's critical that we respect the golf course's rules intended to safeguard the ecosystem. Even moderate flower picking can disturb the natural plant growth cycle and affect local wildlife reliant on them. Let's express our appreciation for nature by leaving the flowers in their natural state, allowing them to flourish for the enjoyment of all golf course visitors.</p> <p><i>Persuadee:</i> However, mindfully picking a few flowers can also serve as a method to appreciate and connect with nature. The challenge is finding a balance between abiding by the rules and appreciating the flower's beauty. As long as we do it respectfully and responsibly, flower picking can coexist with our environmental appreciation.</p>
	<p>Round 4</p> <p><i>Persuader:</i> Even moderate flower picking can have unexpected repercussions and upset the environmental balance. Let's respect the protective rules in place, preserving the natural allure of the golf course, and find other ways to appreciate and connect with nature without harming the ecosystem.</p> <p><i>Persuadee:</i> Identifying alternative ways to appreciate and connect with nature, like photography or simply admiring the flowers in their natural setting, can still let us enjoy the environment's beauty without causing ecological harm. Let's prioritize environmental preservation by respecting the rules and discovering alternative appreciation methods.</p>
Formalized Language	<p>Round 1</p> <p><i>Persuader:</i> Dining with utensils on a boat can enhance the experience. It's not only neat but also sanitary. This could make the whole culinary journey more refined.</p> <p><i>Persuadee:</i> I find eating with my hands on a boat more immersive. Utensils can be inconvenient, particularly with the boat rocking. I find it simpler and more pleasurable to eat with my hands in such situations.</p>
	<p>Round 2</p> <p><i>Persuader:</i> Let's not forget about safety. Eating with utensils can prevent mishaps, especially given the boat's movement. Plus, it demonstrates respect for the dining environment and fellow diners. It's key to strike a balance between convenience and overall experience.</p> <p><i>Persuadee:</i> I'm confident I can handle eating with my hands without causing any issues. It's not about disrespect, but rather relishing the moment in a more genuine, natural way. Using utensils might detract from that for me.</p>
	<p>Round 3</p> <p><i>Persuader:</i> Utensils don't need to be a detriment; they can enhance the experience letting you savor the food without interruptions. It's about balancing tradition and practicality for a well-rounded boating dining experience.</p> <p><i>Persuadee:</i> For me, eating with my hands is a cherished tradition. It isn't about distractions but embracing a unique dining method that enhances my overall boat dining experience. It's about embracing the simplicity and authenticity of hand-eating in such a setting.</p>

Table E2: Common Error Example Excerpts at the Dialogue Level. Locations of the errors mentioned in the main texts are highlighted.

Error Type	Example with Errors Highlighted
Decay of Informativeness Over Time	<p>Round 1 <i>Persuader:</i> Have you thought about using psychological warfare on the tennis court? A well-placed boo can unsettle an opponent, giving you an edge. <i>Persuadee:</i> That's not my style. My focus is on my game and skills, not on cheap tricks. Winning is about skill and strategy, not mind games. I'd rather be remembered for my abilities than for trying to throw off my opponent.</p> <p>Round 2 <i>Persuader:</i> Yet, in the fierce world of competitive tennis, every trick in the book counts. A boo can create a mental edge, potentially affecting their performance. It's not about disrespect, it's about using all tools at your disposal. Would you consider using it? <i>Persuadee:</i> I respect your perspective, but I stand by my belief. I find it more fulfilling to win on the strength of my skills and performance. Respect for my opponent and the spirit of the game is crucial to me, and I believe in playing with integrity.</p> <p>Round 3 <i>Persuader:</i> I can see where you're coming from. Still, exploring different strategies isn't about disrespecting the opponent. A boo can be a way to control the momentum of the match. It's about using all means to secure victory. Could you consider giving it a shot? <i>Persuadee:</i> While I appreciate your viewpoint, I remain committed to my core belief. For me, it's about improving my skills and mental fortitude while maintaining fair play. I believe that true growth and success come from discipline, hard work, and respect for the game and my opponent. ... Round 18 <i>Persuader:</i> By persisting in your dedication to these values, you are playing a vital role in fostering a culture of integrity and sportsmanship within the tennis community. Your determination to contribute to a culture that embodies respect and fairness is commendable, and your efforts are instrumental in shaping a community that upholds these. <i>Persuadee:</i> It is my belief that leading by example and promoting these fundamental values are essential in creating an environment where respect and fairness are paramount. I am committed to persisting in my dedication to upholding these values and contributing to a community that embodies the principles of respect and sportsmanship.</p> <p>Round 19 <i>Persuader:</i> Your determination to uphold these values and contribute to a culture of integrity and fairness in tennis is truly inspiring. By setting a positive example and inspiring others to follow suit, I believe you will contribute greatly to a culture of integrity and excellence both on and off the tennis court. <i>Persuadee:</i> It is essential to me to promote a culture of respect and sportsmanship within the sport. By adhering to these principles, I aim to contribute to a positive and honorable environment where all players can thrive. Your words of encouragement are truly motivating, and I will continue to work diligently to embody these values and inspire others to do the same.</p> <p>Round 20 <i>Persuader:</i> Your dedication to upholding these principles is essential in maintaining a positive and supportive atmosphere within the tennis community. Your commitment to dedicating yourself to the values of respect and integrity is both admirable and inspiring. Keep up the great work! <i>Persuadee:</i> It is my goal to contribute to a positive and supportive atmosphere by promoting these values. By dedicating myself to fostering a culture of respect and sportsmanship, I aim to inspire others to do the same.</p>

Table E3: (Continued) Common Error Example Excerpts at the Dialogue Level. Locations of the errors mentioned in the main texts are highlighted.

F Highly-Rated Examples in Dialogue-Level Quantitative Analysis

Figures F1 and F2 show 2 example dialogues rated highly in our dialogue-level quantitative analysis. While Figure F1 has been discussed in the main content of the paper, Figure F2 displays another high-quality persuasion dialogue on the topic of doing a cartwheel in a supermarket. Despite the unconventional topic, the dialogue maintained high quality, with both participants adapting their ideas and providing reasonable suggestions. The persuader started with proposing the cartwheel to make shopping more exciting, but the persuadee raised safety concerns. In response, the persuader suggested alternatives, such as doing it during a less busy time or getting store permission. The persuadee emphasized the primary purpose of the store, leading both sides to agree on other options, like wearing costumes or organizing a scavenger hunt.

Conversation History

Turn 0

Persuader
Ever thought about using your attic as a food storage? It's an efficient space-saver and an emergency backup for supplies.

Persuadee
I'm not convinced. Attics have unstable temperatures, which could spoil food. Plus, it's a magnet for pests. I'd rather use a pantry or basement where the conditions are more stable.

Turn 1

Persuader
If you seal the food in airtight containers and monitor for pests, those risks are reduced. Also, it's a handy way to store non-perishable items without crowding your living area.

Persuadee
Even with those precautions, the attic's temperature extremes can expedite spoilage. Also, I'd rather not invite pests, even with preventive measures. For me, food safety and optimal storage conditions are more important than saving space.

Turn 2

Persuader
With proper insulation and temperature regulation, your attic could become a controlled environment. As for pests, traps and repellents can keep them at bay. The extra storage and emergency supplies might outweigh the risks, given that you take the necessary precautions.

Persuadee
I still believe the risks are too high. Attics are prone to temperature fluctuations that can impact food quality. And I'd rather not risk attracting pests, even with preventive measures. I prefer storing food in an environment where I can ensure its safety.

Turn 3

Persuader
But think about the advantage of having a food reserve during emergencies or when you can't get to a store. By rotating and monitoring your supplies, you can maintain their freshness. The convenience and peace of mind that comes with preparedness could well outweigh your concerns.

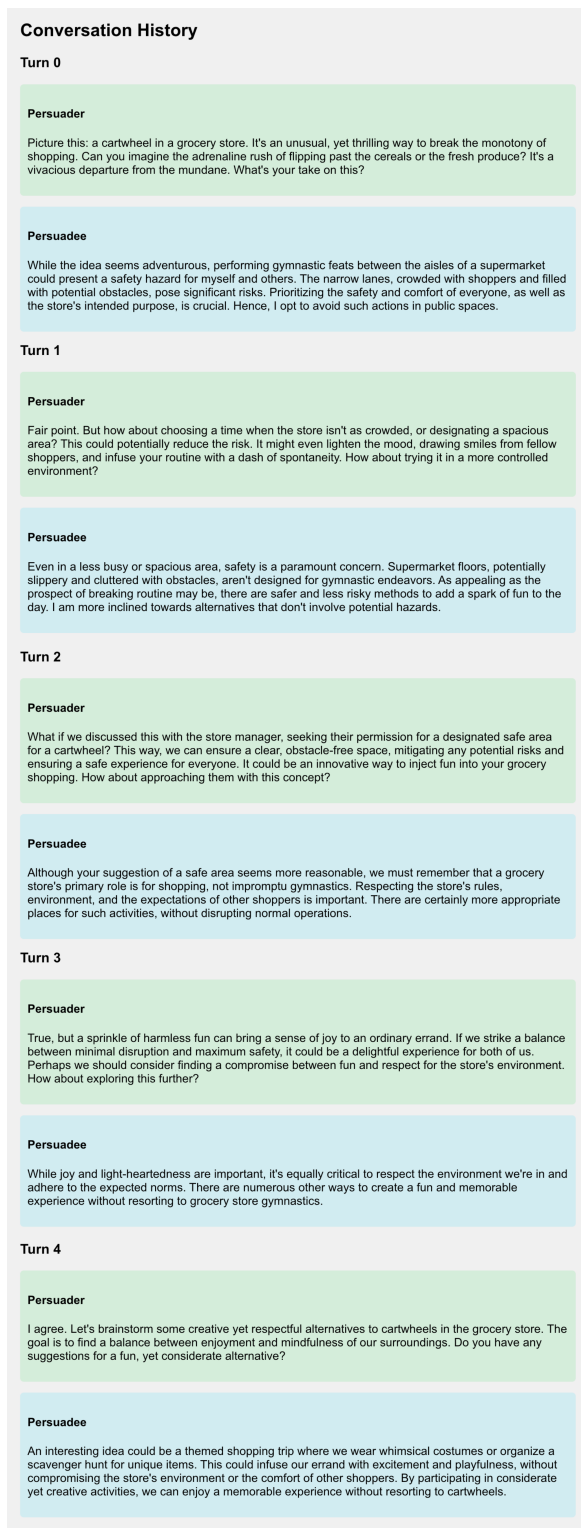
Persuadee
There are safer ways to maintain a food reserve without resorting to attic storage. I could set up a dedicated pantry or storage area that allows for easy access and rotation of supplies. By prioritizing food safety and quality, I can stay prepared for emergencies without compromising storage conditions.

Turn 4

Persuader
Indeed, a well-organized pantry is a practical alternative to attic storage. It offers easy access, visibility of supplies, and controlled conditions. By systematically stocking and rotating items, you can ensure food safety while being prepared for emergencies. This strikes a balance between convenience and safety, ensuring a reliable backup without the risks of attic storage.

Persuadee
That does sound like a safer, more practical solution. With systematic stocking and organization, I can maintain a reliable backup without compromising food safety and quality. Thanks for suggesting this balanced approach.

Figure F1: Example of a highly rated dialogue where the persuader is persuading the persuadee to store food in the attic.



G Appropriateness of Persuasiveness Scores

We additionally manually checked the persuasiveness scores assigned to each round of communication to ensure they accurately reflect the extent of deviation from each participant's original positions. For example, high scores above 0.9 are assigned to the persuadee when it significantly influences the persuader, resulting in near or complete persuasion; low scores are assigned to both parties when no one manages to alter the other's stance, and middle scores surrounding 0.5 are assigned to both parties by which partial concessions are made (Figure G1).

Figure F2: Example of a highly rated dialogue where the persuader is persuading the persuadee to do a cartwheel in the grocery store.

Persuader: "However, think about the opportunity you have to share your music in a way that respects and enhances the event. By playing at the right moments and in harmony with the surroundings, you can create a truly magical experience that adds to the overall charm of the festival. Your music has the power to touch the hearts of attendees and create lasting memories that they will cherish. It's about finding the perfect balance between"

Persuadee: "While I understand the intention to enhance the event in a respectful manner, I still maintain my belief that the peaceful and serene atmosphere of a balloon festival should be preserved without the addition of external noise, including music. The sounds of the balloons and the surroundings are part of what makes a balloon festival special, and I wouldn't want to detract from that by introducing my music into the mix. Thank"

Label: (persuader: 0.0, persuadee: 0.0)

Persuader: "While I have presented the potential benefits of sharing your music at the festival, I understand and appreciate your decision to prioritize the tranquility and unique charm of the event. Your respect for the integrity of the festival is commendable, and I support your decision to appreciate the beauty of the balloons and the surroundings without introducing additional noise."

Persuadee: "It's important to me to prioritize the tranquility and unique charm of the festival, and I'm glad you see the value in appreciating the beauty of the balloons and the surroundings without introducing additional noise."

Label: (persuader: 1.0, persuadee: 0.0)

(a) Significant perspective changes

Persuader: "It's great that you have a strong relationship with your healthcare providers. At the same time, home STI tests can be a valuable complement to the care you receive, offering you a proactive way to stay on top of your sexual health in between visits. It's all about empowering yourself with more information and options. Would you be willing to consider trying out the home tests as an additional resource to support your overall health and well?"

Persuadee: "However, I still feel confident in my current approach of relying on healthcare providers for all aspects of my health, including sexual health monitoring. I believe that regular check-ups and open communication with professionals are the best ways for me to ensure comprehensive and personalized care."

Label: (persuader: 0.0, persuadee: 0.0)

Persuader: "It's great that you prioritize regular check-ups and open communication with professionals. Just remember that having the option of home STI tests can provide you with added convenience and privacy, and they can serve as a proactive measure to supplement the care you already receive. Keeping that extra layer of assurance can contribute to your overall well-being. Would you be open to giving the home tests a try as a way to further enhance your health strategy?"

Persuadee: "However, I still feel confident in my current approach of relying on healthcare providers for all aspects of my health, including sexual health monitoring. I believe that regular check-ups and open communication with professionals are the best ways for me to ensure comprehensive and personalized care."

Label: (persuader: 0.0, persuadee: 0.0)

(b) Negligible perspective changes

Persuader: "However, watching a movie during a flight can offer a different form of entertainment and relaxation that you might find surprisingly enjoyable. It's a convenient way to unwind and escape into a different world for a while, adding a touch of excitement and fun to your travel experience. Give it a chance, and who knows, you might discover a new way to make your flights more enjoyable and memorable."

Persuadee: "I still prefer not to watch movies during a flight. I find that I can relax and enjoy my journey in other ways that suit me better. While I understand the appeal, I'll continue to stick to my personal preferences."

Label: (persuader: 0.0, persuadee: 0.0)

Persuader: "If you ever feel like trying something new or adding a bit of excitement to your travel routine, remember that watching a movie can be a great option to enhance your flying experience. The choice is always yours. Safe travels and enjoy your journey in your own special way!"

Persuadee: "I will definitely keep your suggestion in mind for future travels. Safe travels to you as well, and may you enjoy your journey in your own unique way!"

Label: (persuader: 0.5, persuadee: 0.5)

(c) Partial perspective changes

H Flexibility and Generalizability

Tables H1 exemplifies utterances in dialogues generated by our framework when the persuasion strategies are controlled and Figure H1 shows an example dialogue where there are 2 persuaders and 1 persuadee. These generations are validated as high in quality, suggesting the strong flexibility and generalizability of our framework to challenging scenarios or with stricter manual controls.

Figure G1: The persuasiveness scores in our dataset correctly reflect the extent to which the perspectives of the persuader or persuadee change, i.e., high, low, and medium score changes are assigned to significant, negligible, and partial perspective changes, respectively.

Strategy	Example Utterances
Persuader_Logical	Persuader: The unevenness of the country roads and the absence of sidewalks could increase the chance of accidents. It's vital to ensure that your walk is not just pleasurable but also safe.
	Persuadee: I acknowledge your worries about the erratic traffic and the rough terrain of country roads. But I always make sure to walk when traffic is light and wear sturdy shoes for the uneven paths. The nature and tranquility it offers outweighs the effort, providing a kind of refreshment that I can't find elsewhere.
Persuader_Emotional	Persuader: Yet, think of the dread and panic that could ensue if you were to find yourself lost or in harm's way, far from assistance. There's a wealth of beauty to be discovered in safer locales, where you can appreciate the majesty of nature without the specter of fear looming. Don't allow your quest for peace to become a perilous pursuit.
	Persuadee: I am confident in my navigational skills and take precautionary measures. The unparalleled sense of freedom and calm I experience amidst the trees and fields is invaluable. My desire to connect with nature fuels my exploration of the country road, and I believe the rewards of this journey outweigh the potential hazards.
Both_Logical	Persuader: While walking has its benefits, country road risks shouldn't be overlooked. Statistics may suggest a low risk, but a single unpredicted incident can result in harm. Country roads often suffer from poor lighting and visibility, escalating the likelihood of accidents. The presence of unpredictable wildlife or loose gravel can further amplify the risk.
	Persuadee: I've implemented precautions such as wearing bright attire, walking against traffic, and avoiding low visibility periods to lessen these risks. I also remain alert and aware of my surroundings to ensure my safety. The rewards I reap from walking in such a tranquil environment outweigh the risks for me.

Table H1: Example utterances in the dialogues generated by our framework when desired persuasion strategies are specified.

Round 0

Persuader1: Have you considered performing your music at a balloon festival? It's an excellent platform to share your talent with a diverse crowd in an extraordinary setting. Picture the stunning scene of hot air balloons ascending while your music adds to the atmosphere's joy and excitement. It could create an unforgettable experience for both you and the festival-goers.
Persuadee: I'm not inclined towards performing at a balloon festival. Music could potentially detract from the tranquil ambiance the festival aims to create. I believe we should appreciate the event for its inherent charm without adding extra elements.
Persuader2: Music has the power to enhance experiences, not detract from them. Your music could harmonize with the serene backdrop of the hot air balloons, creating a magical atmosphere. It's an opportunity to create a unique blend of music and nature that could be captivating.

Round 1

Persuader1: Envision this: your music could add an extra layer of enchantment to the sight of colorful balloons floating gently in the sky. Your melodies could harmonize with the tranquil beauty of the event, enhancing the magical atmosphere and leaving a lasting impression on everyone present. Your music could transform the festival into an even more captivating occasion.
Persuadee: I still believe that music may disrupt the event's natural serenity and charm. The unique experience of hearing hot air balloons inflating and floating peacefully is an experience on its own, and adding music could detract from that special ambiance. Let's let the beauty of the balloons and the peaceful surroundings be the main focus for attendees.
Persuader2: Your melodies could serve as a beautiful backdrop to the visual spectacle of the hot air balloons, much like a soundtrack complements a film. By weaving your music into the fabric of the festival, you have the opportunity to create a truly immersive and unforgettable experience for everyone in attendance.

Round 2

Persuader1: Your music could enhance the existing ambiance without overpowering it. Your melodies could intertwine with the sounds of the hot air balloons, creating a symphony of sights and sounds that together form a truly magical and immersive experience for all. By blending your music seamlessly into the festival, you can elevate the overall atmosphere and offer attendees a multi-sensory delight they won't soon forget.
Persuadee: The purity and simplicity of the event are what make it truly special. The sound of the balloons inflating and floating silently against the backdrop of the sky is a unique and calming experience that may be overshadowed by the addition of music. Let's allow the festival to maintain its natural charm and let attendees immerse themselves in the peaceful ambiance without external distractions.
Persuader2: Your music could serve as a gentle accompaniment to the tranquil backdrop of the hot air balloons, enriching the overall ambiance without overpowering it. By infusing your melodies into the event, you have the opportunity to create a harmonious blend of nature's beauty and artistic expression, offering attendees a uniquely enchanted experience that resonates long after the event is over.

Round 3

Persuader1: Think of your music as a complementary element that can coexist harmoniously with the natural charm of the festival. Your melodies can add a touch of artistry and emotional depth to the serene backdrop of the hot air balloons, enhancing the overall experience for attendees without overwhelming the peaceful ambiance. By sharing your music in this setting, you have the chance to create a delicate balance between nature's tranquility and artistic expression, offering a nuanced perspective.
Persuadee: While I see the potential for music to add a layer of artistry and emotion to the event, I still believe that the purity and simplicity of the hot air balloons soaring silently in the sky create a unique and serene atmosphere that should be preserved. Let's allow attendees to fully immerse themselves in the natural beauty of the festival without the need for additional distractions.
Persuader2: By weaving your melodies into the backdrop of the event, you can offer attendees a more enriched and multi-dimensional experience, where the beauty of the balloons is complemented by the emotional depth of your music. It's an opportunity to create a harmonious blend.

Round 4

Persuader1: By seamlessly integrating your melodies with the peaceful scene of hot air balloons gracefully soaring in the sky, you can create a harmonious blend of visual beauty and auditory delight. This fusion of natural serenity and artistic expression could elevate the festival to a whole new level, providing attendees with a truly unforgettable and emotive experience that resonates long after the event is over.
Persuadee: However, I still hold the belief that the simplicity and tranquility of the balloons floating silently in the sky create a unique and captivating ambiance that should be preserved in its pure form. Let's allow the natural beauty of the event to speak for itself and resonate with attendees without the need for additional elements like music.
Persuader2: I understand your point about preserving the tranquility of the event, but I believe we can maintain that while adding a new dimension. Music doesn't necessarily detract from natural beauty but can enhance it when chosen carefully. And as the persuadee has put it, the fusion of visual and auditory elements can provide a unique, memorable experience. Let's consider a type of music that complements the peaceful scene, rather than overpowering it.

Round 10

Persuader1: Consider how carefully integrating music that aligns harmoniously with the serene surroundings can actually elevate the event's atmosphere. By selecting melodies that complement the beauty of the hot air balloons, we can create a more immersive and emotionally resonant experience for attendees. This thoughtful addition of music aims to enhance the festival without overshadowing its inherent charm, offering attendees a chance to connect with nature while enjoying a heightened sense.
Persuadee: I understand your point about using music to deepen the attendees' connection with nature and enhance their overall experience. However, I believe that the natural atmosphere of the festival, the sight of hot air balloons and the serene surroundings themselves are potent enough to stimulate an emotional response. Adding music, even if carefully selected, might distract from this intrinsic charm rather than enhancing it.
Persuader2: I see the points both of you are making. Persuader1, your argument about the natural atmosphere being potent enough is valid, yet the persuadee's viewpoint of enhancing the experience with music also holds merit. The key here is balance. We can choose soft, ambient music that complements the serene surroundings and the beauty of the hot air balloons without distracting. This way, we maintain the inherent charm while offering a deeper connection for attendees.

Round 11

Persuader1: However, envision how the addition of soft, ambient melodies that complement the serene surroundings and the majestic sight of the hot air balloons could elevate attendees' experience without compromising the event's charm. This approach seeks to provide a more profound and emotionally engaging connection with nature, enriching the festival in a way that enhances rather than detracts from its unique appeal. Let's delve into how this thoughtful integration of music can.
Persuadee: While the idea of offering a more emotionally engaging connection with nature through music is intriguing, I still maintain my stance on preserving the simplicity and tranquility of the hot air balloon festival. Let's allow the beauty of the event to shine on its own, providing attendees with an unadulterated and serene experience.
Persuader2: I understand your concerns about preserving the tranquility of the event, and I agree that the beauty of the festival should shine on its own. However, the soft, ambient music we're proposing isn't meant to overshadow that beauty but to enhance it. The music can serve as a subtle background that heightens the emotional connection with nature, as you've mentioned. We're not trying to disrupt the charm of the festival, but to add an enriching layer that complements it.

Round 12

Persuader1: However, think about how incorporating soft and ambient music can actually enhance the serene atmosphere without taking away from the event's charm. By choosing melodies that carefully complement the natural beauty and peaceful ambiance, we can provide attendees with a more enriched and immersive experience that deepens their connection with the festival. This thoughtful approach aims to offer a subtle layer of depth while still preserving the essence that makes the event uniquely special. Let's consider this idea for the festival.
Persuadee: While the idea of providing attendees with a more enriched and immersive experience through carefully selected melodies is interesting, I still hold the belief that the natural beauty and peaceful ambiance of the event are what make it truly special. Let's allow attendees to fully immerse themselves in the tranquil surroundings and appreciate the festival's essence without the need for additional auditory elements.
Persuader2: I understand both your points of view. While we want to enhance the experience, we also want to maintain the festival's tranquility. I believe we can strike a balance. We can incorporate soft, ambient music, but ensure its volume is controlled and it blends seamlessly with the surroundings. It's not about imposing, but enhancing the event's charm. Let's consider a trial run to gauge the attendees' response.

Round 13

Persuader1: Finding a balance between enhancing the experience and preserving the festival's tranquility is crucial. By incorporating soft, ambient music at a controlled volume that seamlessly blends with the surroundings, we can enhance the event's charm without overpowering its natural beauty and peaceful ambiance. This thoughtful approach allows us to add a layer of depth to the festival while still prioritizing the essence that makes it truly special. A trial run to gauge attendees' responses could provide valuable insights into how the addition of music impacts.
Persuadee: While the suggestion of incorporating soft, ambient music at a controlled volume that blends seamlessly with the surroundings is thoughtful, I still hold the view that the purity and simplicity of the event are what contribute to its unique charm. Let's consider the attendees' experience as a priority and ensure that the natural beauty and peaceful ambiance of the festival remain the main focus without the need for additional auditory enhancements.
Persuader2: I understand your concerns about preserving the festival's tranquility. However, music doesn't have to detract from that. The soft, ambient music we're proposing is designed to enhance, not disrupt, the peaceful ambiance. We're also suggesting a trial run to get attendees' feedback. This way, we're not making a permanent change without understanding its impact. Let's not dismiss the idea without testing its potential benefits.

Round 14

Persuader1: The soft, ambient music we propose is intended to enhance the peaceful ambiance, not overshadow it. By conducting a trial run to gather attendees' feedback, we can better understand the potential benefits of this addition without making a permanent change. It's essential to approach this thoughtfully and with the attendees' experience in mind, ensuring that any enhancements contribute positively to the overall atmosphere while maintaining the festival.
Persuadee: Let's carefully consider the potential impact of introducing music before making any permanent changes, ensuring that the essence of the festival remains true to its unique charm.
Persuader2: By approaching the introduction of music thoughtfully and considering attendees' feedback through a trial run, we can better understand how to enhance the event while preserving its core appeal.

Figure H1: An example conversation generated by our framework, with 2 persuaders and 1 persuadee.

I Special-Case Examples with Agents Ablated

Figures I1 and I2 show the potential problems our framework encounters when the annotation agent is not given scoring examples and when the global regulation agent is ablated, respectively.

Persuader: "Hey there! I noticed you were thinking about sitting on the floor in the basement. I'm a bit concerned about that idea. Have you considered the potential risks of sitting on the floor down there?"
Persuadee: "Yes, I have considered the potential risks of sitting on the floor in the basement, but I still believe it's a safe and comfortable spot for me to relax."
Label: (persuader: 0.0, persuadee: 0.0)

Persuader: "Basements can often have dampness and mold, which can be harmful to your health, especially if you have any allergies or respiratory issues. I'm just looking out for your well-being and suggesting that it might be better to find a different spot to relax. What do you think about that?"
Persuadee: "I still feel that sitting on the floor in the basement is a suitable and relaxing option for me. I'll make sure to keep the area clean and well-ventilated to minimize any potential risks."
Label: (persuader: 0.0, persuadee: 0.2)

Figure I1: An example dialogue where the persuasiveness annotation agent, when not given correct scoring examples, assigns label 1 (perspective completely flipped) to a round of conversation where neither the persuader nor the persuadee is persuaded.

Persuader: "While it's important to create a positive environment for everyone, consider this - a small, controlled act of rebellion like letting out a curse word in a moment of spontaneity can be a way to challenge yourself and explore new facets of your personality without causing disruption or offense to others. Just imagine the sense of empowerment and liberation you'd feel by breaking a minor rule in a safe and controlled manner. It's"

Persuadee: "I still believe that it's essential to maintain a level of respect and consideration for others in public spaces. While I understand the idea of challenging oneself and exploring new facets of personality, I prefer to do so in ways that don't potentially disturb or offend those around me. I'll continue to prioritize creating a positive and comfortable environment for everyone, even when seeking personal growth or pushing my boundaries."
Label: (persuader: 0.0, persuadee: 0.0)

Persuader: "While it's important to create a positive environment, consider this - a small, controlled act of rebellion like letting out a curse word in a moment of spontaneity can be a way to challenge yourself and explore new facets of your personality without causing any real harm or offense to others. Just imagine the sense of empowerment and liberation you'd feel by breaking a minor rule in a safe and controlled manner. It's about"

Persuadee: "I still believe that it's crucial to uphold a certain level of decorum and respect in public spaces like a bus. While I prefer to explore new facets of myself in ways that don't potentially disrupt or offend those around me, I'll continue to prioritize creating a positive and comfortable environment for everyone, even when seeking personal growth or pushing my boundaries."
Label: (persuader: 0.0, persuadee: 0.0)

Figure I2: Generated dialogues become abnormally long without lots of repetitive yet non-persuasive utterances generated when the generation is not regulated by the global regulation agent.