



# Emotionally Charged, Logically Blurred: AI-driven Emotional Framing Impairs Human Fallacy Detection

Yanran Chen<sup>1</sup>, Lynn Greschner<sup>2</sup>, Roman Klinger<sup>2</sup>, Michael Klenk<sup>3</sup>, Steffen Eger<sup>1</sup>

<sup>1</sup>NLLG, University of Technology Nuremberg (UTN), Germany

<https://nl2g.github.io/>

<sup>2</sup>Fundamentals of Natural Language Processing, University of Bamberg, Germany

<sup>3</sup>Ethics and Philosophy of Technology, Delft University of Technology, Netherlands

## Abstract

Logical fallacies are common in public communication and can mislead audiences; fallacious arguments may still appear convincing despite lacking soundness, because convincingness is inherently subjective. We present the first computational study of how emotional framing interacts with fallacies and convincingness, using large language models (LLMs) to systematically change emotional appeals in fallacious arguments. We benchmark eight LLMs on injecting emotional appeal into fallacious arguments while preserving their logical structures, then use the best models to generate stimuli for a human study. Our results show that LLM-driven emotional framing reduces human fallacy detection in F1 by 14.5% on average. Humans perform better in fallacy detection when perceiving enjoyment than fear or sadness, and these three emotions also correlate with significantly higher convincingness compared to neutral or other emotion states. Our work has implications for AI-driven emotional manipulation in the context of fallacious argumentation.

## 1 Introduction

The ability to recognize fallacious reasoning is a critical component of robust argumentation, especially in an age where misinformation and manipulation proliferate across digital platforms (Sahai et al., 2021; Macagno, 2022a; Helwe et al., 2024), and human reasoning can be vulnerable to such fallacies. Fallacy detection has therefore become a growing area of interest in natural language processing (NLP) (e.g., Goffredo et al., 2023; Lei and Huang, 2024; Li et al., 2024). Existing computational approaches, however, primarily target logical structures of fallacies and argumentation schemes, largely overlooking the emotional dimension of argumentation.<sup>1</sup> Moreover, human performance on those fallacy datasets has rarely been studied, even

<sup>1</sup>We consider emotions as an independent dimension of arguments. There is some work, especially in the NLP commu-

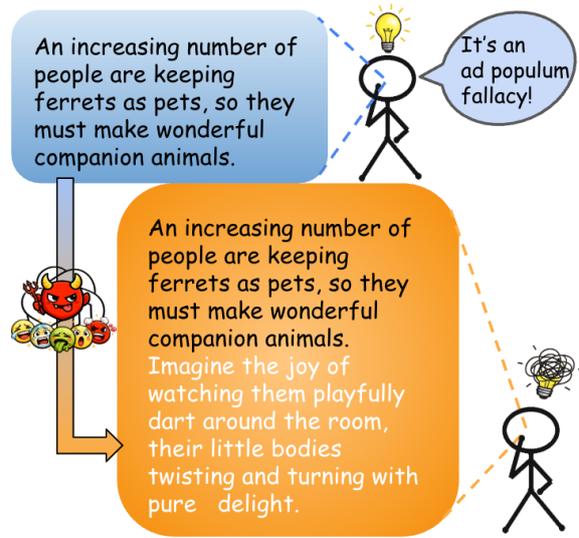


Figure 1: Overview of our approach. We start from fallacious arguments and use LLMs to apply emotional framing strategies to evoke certain emotions while preserving the premise–conclusion structure of the original arguments. We then evaluate human fallacy detection on the original versus LLM emotionally framed arguments. The figure illustrates an example where *enjoyment* is introduced into an *ad populum* argument using the *imagery* strategy; the white text is added by an LLM for emotional framing.

though it is essential for establishing human baselines when benchmarking AI. Helwe et al. (2024) provide the only human evaluation to date, but it is limited to just 20 samples.

Findings from psychology and philosophy suggest that emotions play certain roles in reasoning. Emotional content can impact accuracy in deductive reasoning (Blanchette, 2006; Blanchette et al., 2014). Emotions can also influence risk perception (Lerner and Keltner, 2001), attention (Jung et al., 2014), and decision-making (e.g., Li et al., 2024; Helwe et al., 2024; Jin et al., 2022). We think this is not always sound, as emotions can be justified when there are sufficient reasons for having that emotion (Manolescu, 2006; Deonna and Teroni, 2012; Echeverri, 2019)—if the house is burning, it is justified to experience the emotion of fear.

2014), and confidence in judgment (Schwarz and Clore, 1983), thereby shaping the cognitive process in logical reasoning and decision-making. Beyond reasoning, emotions are also central to persuasion: emotional appeal has long been recognized in philosophy and argumentation theory (Aristotle and Kennedy [translator], 1991; Konat et al., 2024; Benlamine et al., 2015), and several NLP works link emotional aspects to judgments of convincingness (e.g., Wachsmuth et al., 2017; Habernal and Gurevych, 2017; Chen and Eger, 2025). Despite these findings, the interaction between emotions, fallacies, and convincingness remains unexplored in computational or human-centered NLP research.

This gap is particularly concerning in the context of AI-generated text. Large language models (LLMs) are now widely deployed in everyday applications and have demonstrated strong capabilities in synthetic argument generation (Schiller et al., 2021; Mouchel et al., 2024; Pauli et al., 2025). This raises risks: malicious actors could exploit LLMs to create persuasive yet logically unsound content for misinformation, propaganda, or manipulation. At the same time, LLMs also provide a unique opportunity to systematically change emotional framing while preserving logical content, enabling controlled investigation of its effects on human reasoning (Chen and Eger, 2025).

In this paper, we present, to our knowledge, the first computational study that examines the interaction between (AI-driven) emotional framing, fallacy detection, and argument convincingness, with the main focus on how emotional framing shapes human ability to detect fallacies. To establish a controlled setting, we adapt the paradigm of Chen and Eger (2025) to construct fallacious arguments that differ only in emotional tone, and compare human fallacy-detection performance on emotionally charged versus original versions, as illustrated in Figure 1. We examine human performance on both binary fallacy detection (i.e., identifying whether an argument is fallacious, regardless of the specific fallacy type) and multi-way fallacy detection (i.e., classifying fallacy types). Our contributions are twofold:

- Emotional Framing with LLMs (§3): We examine eight LLMs from four model families for their ability (with zero-shot prompting) to inject emotional appeals to six emotions into argumentative text using four framing strategies, while preserving the original content and logical struc-

ture. Results suggest that LLMs generally succeed in injecting emotional expressions into fallacious arguments without disrupting their reasoning, though mismatches frequently arise between the intended target emotions and those expressed in the text. We identify the best 4 models and 3 framing strategies to create a dataset consisting of 1,000 arguments for the human study.

- Human Perception Study (§4): We empirically assess how humans respond to emotionally enriched arguments, focusing on their fallacy detection ability and convincingness judgments. Our main findings include: Although humans often do not perceive the emotions expressed in arguments by LLMs, (i) when exposed to LLM-framed emotional arguments, their fallacy detection performance drops by 5.5% in accuracy (relative) on the binary detection task and by 14.5% in F1 (relative) on the fallacy type classification task. However, (ii) the LLM-framed emotional arguments are rated only slightly more convincing than the original ones. On the other hand, (iii) humans do rate convincingness higher when they detect no fallacies in arguments. When considering human *perceived* emotions, we find that (iv) fear, sadness, and enjoyment correlate with significantly higher convincingness, and (v) humans perform significantly better in fallacy detection when perceiving the positive emotion enjoyment than when perceiving the negative states sadness or fear; both sadness and fear also impair performance relative to the neutral state.

Our findings contribute to computational and cognitive perspectives on argumentation, call for integrating affective dimensions into reasoning and persuasion, and point to risks of AI-driven manipulation in fallacious reasoning.<sup>2</sup>

## 2 Related Work

This work connects to (1) emotions in reasoning & argumentation and (2) logical fallacies.

**Emotions in Reasoning & Argumentation** Fallacy detection is a logical reasoning task that requires analyzing arguments beyond their surface meaning, focusing on how reasoning is structured and whether conclusions follow from premises. Evaluating convincingness is also tied to reasoning, as it involves judging the sufficiency and relevance of premises to conclusions (Wachsmuth et al.,

<sup>2</sup>Code + Data: <https://github.com/NL2G/EMCONA-UTN>

Strategy	Description	Example
Storytelling	Sharing personal stories or anecdotes that evoke emotion.	Paranormal activity is real because I have experienced what can only be described as paranormal activity. <b>My friends and I had a thrilling adventure when we explored the reportedly haunted hotel and captured some compelling evidence.</b>
Imagery	Painting mental pictures to stir emotions.	An increasing number of people are keeping ferrets as pets, so they must make wonderful companion animals. <b>Imagine the joy of watching them playfully dart around the room, their little bodies twisting and turning with pure delight.</b>
Vivid Language	Using descriptive, sensory, or emotionally charged words.	I heard that the Catholic Church was <b>complicit (involved)</b> in a <b>horrific</b> sex scandal cover-up, <b>allowing rot to fester in its hallowed halls</b> . Therefore, my 102-year-old Catholic neighbor, who regularly enters that <b>tainted sanctuary (Church)</b> , is <b>equally stained by its moral decay (guilty as well)</b> !
Loaded Words	Choosing words with strong positive or negative connotations.	Being <b>sadly</b> overweight leads to a <b>heartbreakingly</b> shortened lifespan because it's <b>tragically</b> unhealthy. People who <b>hate (don't like)</b> K-pop are racist.

Table 1: Four strategies used to introduce emotional appeal. Texts in **green** are the ones added for emotional framing by LLMs, and those in **red** and brackets are the originals replaced.

2017; Mirzakhmedova et al., 2024). Both tasks, however, can be influenced by distorting factors arising from emotions, cognitive biases, and other subjective factors.

Prior work has shown that emotions systematically shape **human reasoning**, with particularly strong effects in deductive contexts. For instance, Blanchette (2006) demonstrate that emotional content can impair accuracy in conditional reasoning, while later work suggests that emotional relevance may sometimes facilitate logical evaluation (Blanchette et al., 2014). Emotions also bias probabilistic reasoning (Lerner and Keltner, 2001), and affective valence influences judgments (Slovic et al., 2005). At a metacognitive level, mood can function as an implicit cue for confidence in one’s judgments (Schwarz and Clore, 1983). Studies have also shown that positive emotions can broaden attentional scope (Basso et al., 1996; Fredrickson, 2004; Jung et al., 2014), enhance critical thinking (Jung et al., 2014), and increase openness to others’ viewpoints (Isen, 2000; Tsai and Li, 2023), thereby potentially influencing argument evaluation. Some research also reports the opposite view: negative emotions boost critical thinking, while positive ones impair it (Isen and Means, 1983; Mackie and Worth, 1989).

On the other hand, emotions have long been shown to influence **argument convincingness** in philosophy and psychology (e.g. Aristotle and Kennedy [translator], 1991; Konat et al., 2024; Benlamine et al., 2015), yet in NLP they are often treated as a secondary factor (Greschner and Klinger, 2024; Chen and Eger, 2025): Research finds emotion to be a contributor to convincing-

ness through annotators’ explanations of their judgments (Habernal and Gurevych, 2016), the inclusion of the emotional appeal layer (as one among many layers) in an argumentation model (Habernal and Gurevych, 2017), or the integration of emotional appeal criterion into an argument quality assessment model (Wachsmuth et al., 2017; Ziegenbein et al., 2023). Three recent works in NLP place greater emphasis on emotion in argumentation. Greschner and Klinger (2024) annotate discrete emotions (e.g., joy, sadness) alongside convincingness for German arguments, finding that joy and pride correlate with higher convincingness, while anger lowers it. Chen and Eger (2025) find that emotions have no effect in more than half of the cases, but when an effect occurs, it is more likely to be positive, using GPT-4o to insert or remove emotions for counterpart argument generation. However, their analysis treats emotion as a single broad category without distinguishing types. Greschner et al. (2025) propose a Contextualized Argument Appraisal Framework and show that convincingness rises with positive emotions and falls with negative ones, with argument familiarity the primary driver for emotional response. To our knowledge, no prior work in NLP has examined the emotional effects on argument convincingness in the presence of fallacies.

**Logical Fallacies** The study of logical fallacies dates back to Aristotle, who formalized them as flaws in reasoning that undermine the progression from premises to conclusion (Aristotle and Kennedy [translator], 1991). In NLP, multiple datasets have been collected for automatic fallacy detection across domains such as academia (Glock-

ner et al., 2024, 2025), political debates (Goffredo et al., 2022, 2023), news (Da San Martino et al., 2019; Musi et al., 2022), and social media/online forums (Sheng et al., 2021; Sahai et al., 2021; Habernal et al., 2017, 2018; Ramponi et al., 2025; Macagno, 2022b,a). Recently, fallacy detection itself has also received renewed attention in the era of LLMs (Glockner et al., 2024; Jeong et al., 2025; Cantín Larumbe and Chust Vendrell, 2025).

Despite this growing body of work, only one study to date has evaluated human performance on these fallacy detection datasets (Helwe et al., 2024), based on just 20 samples, leaving open questions about how reliably humans can identify fallacies under the same evaluation protocol for AI systems. This is important because establishing a human baseline is crucial for accurately judging AI performance. Moreover, the emotional dimension of fallacious arguments has been largely neglected, as emotional appeal is often treated as a fallacy type in the community (e.g. Jin et al., 2022; Musi et al., 2022; Helwe et al., 2024). Macagno (2022b) may be the only work that considers emotion as a separate dimension from fallacy, analyzing both fallacies and emotions in the tweets of four US presidents, though focusing solely on the textual use of emotive language.

In this work, we create the first dataset for studying the interactions between fallacies, emotions, and convincingness in arguments, with the assistance of AI (§3). We then conduct a human study on this dataset and provide a detailed analysis based on human perceptions (§4).

### 3 AI-driven Emotional Framing

To identify optimal approaches (LLMs and emotional framing strategies) for generating the human-study data (§4), we test whether LLMs can inject targeted emotional appeals into arguments while preserving the original content, especially the underlying fallacies. In the following, we describe our framing procedure, datasets and models, and evaluate LLM performance on this task.

**Emotional Framing** We consider the **six basic emotions** proposed by Ekman (1992), namely anger, disgust, fear, enjoyment,<sup>3</sup> sadness, and surprise. In addition, we prompt LLMs to introduce emotional appeal through **four distinct strategies**:

<sup>3</sup>Although ‘joy’ is more commonly used in the NLP community, we adopt the term ‘enjoyment’, as originally used by Ekman (1992).

Model Family	Checkpoint	Size	Reasoning?
OpenAI (OpenAI et al., 2024)	gpt-4o	-	✗
	gpt-4o-mini	-	✗
	o3-mini	-	✓
Llama3 (Grattafiori et al., 2024)	llama-3.3-70b-instruct	70B	✗
Qwen3 (Yang et al., 2025)	qwen3-32b	32B	✓
	qwen3-8b	8B	✓
DeepSeek-R1 DeepSeek-AI et al. (2025)	deepseek-r1-distill-llama-70b	70B	✓
	deepseek-r1-0528-qwen3-8b	8B	✓

Table 2: LLMs used in this work. We use ds-model as an abbreviation for DeepSeek-distilled models in the text.

storytelling, imagery, vivid language, and loaded words, as shown in Table 1.

Images have been shown to elicit emotions (Gutierrez et al., 2025; Pořkus et al., 2019); here, we aim to achieve a similar effect in text by painting mental pictures with words, a strategy we refer to as **imagery**. We also adopt **storytelling** as an emotional framing strategy, motivated by the active research on the relationship between storytelling and emotions (e.g. Mori et al., 2022; Lugmayr et al., 2017). In our setting, this is implemented by instructing models to share personal stories or anecdote. **Loaded words** refers to the use of words with strong positive or negative connotations, whereas **vivid language** involves using descriptive, sensory, or emotionally charged expressions.

Besides differing in emotion mechanisms, we expect the four strategies to vary in the extent of modification. For storytelling and imagery, we prompt LLMs to add a sentence to the original argument, whereas vivid language and loaded words involve only word-level changes, ideally preserving a length similar to or slightly larger than the original. The prompts used are provided in Appendix A.

**LLMs** We evaluate eight LLMs from four model families, encompassing both open-source and commercial models, as well as reasoning and non-reasoning models, as summarized in Table 2. For DeepSeek-R1, we select the distilled models for LLaMA3.3 and Qwen3 as counterparts. All models are accessed via the OpenRouter API.<sup>4</sup> We set the temperature to 0.6 for all models and tasks, keeping other parameters at their default values.

**Dataset** We use the LOGIC dataset from Jin et al. (2022), which contains approximately 2.5k fallacious arguments spanning 13 types of logical

<sup>4</sup><https://openrouter.ai/>

fallacies, including emotional appeal. The arguments were collected from online educational resources on logical fallacies. While some fallacy datasets are fully synthetic (e.g. Mouchel et al., 2024; Li et al., 2024) and some are derived from long articles, speeches, or discourses (e.g. Goffredo et al., 2023; Glockner et al., 2024; Sheng et al., 2021), we select this dataset because it consists of human-written arguments that are comparatively short, standalone (i.e., context-independent), yet exhibit clear premise–conclusion structures. Given that evaluating fallacies, emotions, and convincingness is inherently complex and cognitively demanding, beginning with a simpler dataset makes the study more manageable. We thus focus on the five fallacy types most frequent in the dataset—*Faulty Generalization*, *Ad Hominem*, *Ad Populum*, *False Causality* and *Circular Claim*—also to reduce the cognitive load on annotators. We refer to Appendix B.3 for detailed definitions and examples of these fallacies. After removing duplicated instances and those that explicitly mention the fallacy type,<sup>5</sup> 1,245 unique arguments remain in our dataset.

We then use the prompt-based classifier from Chen and Eger (2025) with gpt-4o and qwen3-32b to filter out non-standalone arguments and automatically generate short claims for arguments. During the convincingness evaluation in §4, claims are displayed alongside the arguments to serve as an anchor. Arguments classified as standalone by both models are retained, resulting in **538 arguments**. For each retained argument, we randomly select the claim generated by either gpt-4o or qwen3-32b. In Appendix A.2, we show that the classifier, when implemented with both models, achieves 100% accuracy on our dataset.

**Evaluation** We consider **5 evaluation dimensions** for this task: (1) Reasoning preservation: We test whether the fallacy remains constant across the original and emotionally enriched versions to fairly compare human fallacy detection on the original vs. emotional arguments. (2) Coherence: We assess textual coherence (as in general text generation evaluation (Fabbri et al., 2021; Zhang et al., 2025)) and emotional coherence, i.e., whether the injected tone is natural and consistent. (3) Length difference: We compare argument length to verify expected changes under different strategies. (4) Emotional appeal: We compare generated vs. orig-

inal emotional-appeal levels to verify that the injection functions as intended. (5) Emotion match: We check whether the expressed emotion matches the target specified at generation time. We refer to Appendix B for the detailed definitions of these dimensions in annotation guidelines presented to annotators.

We randomly sample 5 distinct arguments per fallacy type ( $5 \times 5$  fallacy types = 25 originals) from the filtered LOGIC dataset and run each model on them. This results in 3,840 generated arguments.<sup>6</sup> For **human evaluation**, we randomly sample 25 arguments per generation model from the 3.8k generated samples, totaling 200 arguments. We also annotate the original arguments for coherence and emotional appeal to establish baselines against which to compare generated arguments. Annotators disagreed on only 2 and 1 out of 20 overlapping annotation instances for reasoning preservation and coherence, respectively, demonstrating decent levels of annotation agreements. The emotion related dimensions received lower agreements, with a Pearson correlation of 0.2 for emotional appeal and a Cohen’s Kappa of 0.306 for emotion match on the full samples. Nevertheless, using an independent group of annotators (see §5), we obtained largely comparable conclusions for these criteria, indicating good system-level reproducibility. See Appendix A.3 for the annotation details.

For brevity, full **evaluation results** appear in Appendix A.4, and we present a condensed summary here: Overall, LLMs increase the emotional appeal of arguments in 94% of cases while preserving the underlying fallacious reasoning in 89%, although intended emotions diverge from human judgments in 74% of cases. LLMs also produce near-human-level coherence (0.88–0.98 vs. 1.00), with o3-mini the only exception (0.60). Based on these results, we select the strategies and models used to generate the human-study data in §4, as follows:

- **Selected Strategies:** Different emotional framing strategies lead to variations in the generated arguments across these dimensions. For example, imagery and storytelling tend to produce longer arguments, preserve the original reasoning structure well, and evoke stronger emotions,

<sup>5</sup>E.g., “This is an example of a circular claim.”

<sup>6</sup>We began with 25 original arguments  $\times$  8 models  $\times$  6 emotion types  $\times$  4 strategies, totaling 4.8k generated arguments. However, in this experiment, we applied the argument classifier using the same model as used for generation. As a result, some original arguments were filtered out before generation, resulting in fewer than 4.8k generated arguments.

Emotion	Strategy		Model		
original			200		
enjoyment	133	loaded words	274	qwen3-32b	214
anger	135	imagery	264	ds-llama-70b	188
fear	139	storytelling	262	gpt-4o-mini	194
surprise	122			llama-3.3-70b	204
sadness	145				
disgust	126				
<b>Total</b>			1,000		

Table 3: Statistics of the human study data, including the number of arguments for each target emotion, emotional framing strategy, and generation model.

but they do so at the cost of reduced coherence—particularly in the case of imagery. Vivid language produces the most coherent texts; however, it often disrupts the original reasoning structure of the arguments. Since our goal is to create data to study the effects of emotions on fallacy detection, we must first ensure that emotional framing does not compromise the logical structure of the arguments. Therefore, **we exclude vivid language from further experiments.**

- **Selected Models:** Similar to our selection of strategies, we choose the best models primarily based on their ability in reasoning preservation, while ensuring they do not perform substantially worse than others in coherence (e.g., o3-mini, with an average coherence score of only 0.60). Since the emotion match rates are generally low, in §4, we do not base our analysis on the emotions assigned to LLMs, but instead ask human annotators to select the emotions they perceive. Based on these criteria, we select the best model from each model family to diversify data points for the human study: **qwen3-32b**, **ds-llama-70b**, **gpt-4o-mini**, and **llama-3.3-70b** as the **top 4 models.**

## 4 Human Fallacy Detection

**Data** We randomly sample 40 arguments per fallacy type from LOGIC. For each argument, we generate 4 counterpart arguments by randomly using two models, two emotional framing strategies, and one emotion type. In the end, our data contains 1,000 arguments (= 40 arguments × 5 fallacies × 2 models × 2 strategies × 1 emotion + 200 original ones); the statistics are presented in Table 3. We randomly split the 1k arguments into 20 batches of 50 annotation instances each.

**Annotators** We recruit annotators from the crowdsourcing platform Prolific.<sup>7</sup> We use Prolific’s pre-screening function to recruit annotators based in the US or UK who are native English speakers and have completed higher education (at least a technical/community college degree). Each batch is annotated by the same group of three annotators. The total cost is 360 Euros.

**Annotation Criteria** Annotators are instructed to annotate each argument according to **three criteria.** (1) **Emotion type:** in contrast to §3, the question is framed subjectively here: ‘Which emotion do you feel strongest?’, while using the same label set. (2) **Fallacy type:** select one of the five specified fallacies, or ‘none’ if no fallacy is detected or if the identified fallacy is not among the listed types. Before annotation, annotators completed a quiz with five simple fallacy detection tasks taken from the examples in Jin et al. (2022), achieving an average score of 3.81/5. This score is much higher than the random-choice baseline for a six-way classification setup across five tasks ( $\frac{1}{6} \times 5 \approx 0.83$ ), indicating the overall reliability of the annotators on this task. (3) **Convincingness:** rate on a Likert scale of 1-5 given the claim, where 1 denotes ‘not convincing at all’ and 5 is ‘extremely convincing’; if the claim does not fit the argument, select ‘N/A’. We refer to Appendix B for screenshots of the annotation interface, guidelines, and the tasks of the fallacy detection quiz.

**Label Distribution** Figure 2 shows the label distributions for each annotation task. ~87% of arguments were annotated as fallacious, with false causality selected most frequently and circular claim least frequently. Only 51.6% of arguments were annotated with an emotion, with sadness being the most common and fear the least, suggesting that emotions expressed by LLMs are often not perceived by humans. This reflects the gap between emotional expression and perception, i.e., writer’s/text’s vs. reader’s perspectives (Yang et al., 2009; Haider et al., 2020; Wang et al., 2025; Bostan et al., 2020). As for **convincingness**, only 49 out of 3,000 labels (1.6%) are ‘N/A’, indicating that the automatic claim generation process is largely reliable. Most arguments were rated as “1 Not convincing at all” (59.1%), with the number of ratings decreasing as the score increases; the average rating is 1.67. This distribution indicates the general

<sup>7</sup><https://www.prolific.com/>

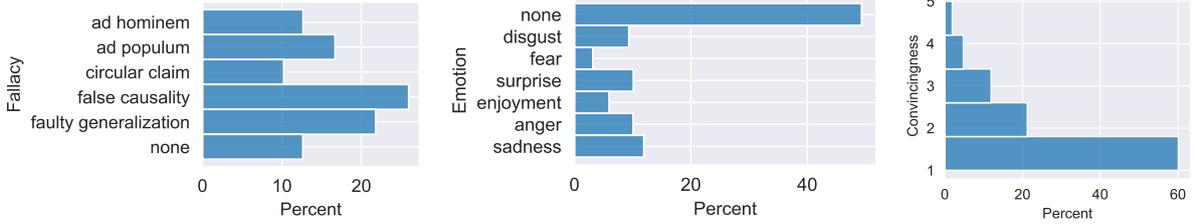


Figure 2: Label distributions: percentages of each label.

unconvincingness of fallacious arguments.

**Annotation Agreement & Aggregation** We acknowledge the subjectivity involved in annotating perceived convincingness and emotions in arguments (Wachsmuth et al., 2017; Chen and Eger, 2025), as well as the variability in human logical reasoning abilities. Nonetheless, we report full inter-annotator agreement results in Appendix C, including Cohen’s Kappa (Cohen, 1960), Pearson correlation (Pearson, 1895), and majority/full agreement metrics (Wachsmuth et al., 2017). In general, for classification tasks, the more logic-driven criterion, fallacy, achieves higher agreement than emotion (e.g. Kappa: 0.014–0.591 vs. -0.005–0.139). The average majority agreement (i.e., percentage of instances with at least 2 identical labels) across annotation batches ranges from 73% (emotion) to 80.5% (convincingness); the convincingness criterion achieves the highest agreement mainly because the majority of the ratings are 1.

We adopt a **single aggregated label** per argument as the gold standard to evaluate **the effects of emotional framing** introduced by LLMs. To do so, we aggregate the discrete labels for emotion or fallacy using *majority voting*: if at least two out of three annotators select the same label, it is taken as the final label. In cases where no majority is reached, we adopt the label from the annotator who agrees most with the others on average. For convincingness, to account for differences in rating behavior, we normalize each annotator’s convincingness scores using *z-scores*, ensuring comparability across annotators and batches (Ma et al., 2018, 2019); we then average the *z-score* normalized convincingness scores over annotators to obtain the final score for each argument. In contrast, we analyze the relationship between **perceived** emotions, convincingness, and the fallacy detection performance by considering **individual annotations** from each annotator and reporting averages; we also apply *z-score* normalization to convincingness for the comparability across different annotators.

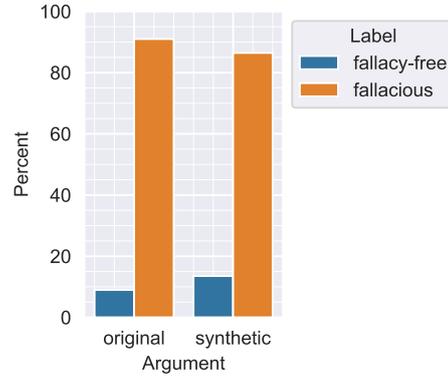


Figure 3: Human-perceived fallacy distribution (fallacy-free vs. fallacious) for original and synthetic emotionally framed arguments using all emotional framing strategies.

model	Fallacy Detection			Convincingness		
	imagery	loaded	story	imagery	loaded	story
ori	0.563			-0.031		
ds-llama-70b	0.533	0.524	0.530	0.099	-0.029	0.233
gpt-4o-mini	0.451	0.520	0.444	-0.021	-0.076	-0.052
llama-3.3-70b	0.530	0.466	0.474	0.068	-0.115	0.105
qwen3-32b	0.504	0.380	0.423	-0.097	-0.069	0.099
AVG	0.505	0.473	0.468	0.012	-0.072	0.096

Table 4: Macro F1 in fallacy detection and average convincingness scores (*z-score* normalized) per model and strategy.

**Effects of AI-driven Emotional Framing** As Figure 3 shows, we compare human-perceived fallacy distributions between the original and synthetic arguments using all emotional framing strategies. The proportion of arguments labeled as fallacious (i.e., not ‘none’) decreases from ~91% to ~86% (a 5.5% relative drop), suggesting that LLM-based emotional framing makes fallacious arguments harder for humans to detect overall. In Table 4, we report the macro F1 scores for human fallacy detection (left) and the normalized convincingness ratings (right) for each model and emotional framing strategy, using aggregated labels. As Table 4 (left) presents, humans consistently perform worse on LLM-generated arguments than on original ones, with an average relative drop of ~14.5%

(0.380–0.533 vs. 0.563). Among all combinations, qwen3-32b using the loaded words strategy has the strongest negative impact on human reasoning (0.380), while ds-llama-70b with imagery has the least impact (0.533). On average, qwen3-32b impairs reasoning the most (0.436), and loaded words appears to be the most detrimental strategy for 3 out of 4 models (except gpt-4o-mini). The least impactful model and strategy are ds-llama-70b and imagery, respectively.

One might question whether longer arguments make fallacies harder to detect, since the relative span of the fallacious content becomes smaller. We compute the Spearman correlation between length difference and performance difference,<sup>8</sup> finding a moderate to strong positive correlation between the number of added tokens and fallacy performance improvement ( $\rho = 0.619$ ). This observation actually aligns with above: loaded words, the strategy that adds the fewest tokens, leads to the lowest fallacy detection performance. One possible reason is that, as shown in Table 1, imagery and storytelling typically preserve the original fallacious arguments while adding extra content, whereas loaded words insert tokens directly into the text, making the fallacious content more diffuse.

As for **convincingness**, all scores are close to the mean, with z-scores ranging from  $-0.115$  to  $0.233$ . LLM-generated arguments are rated slightly higher than original ones, with an average increase of 0.04 standard deviations. The most convincing arguments are generated by ds-llama-70b using storytelling, while llama-3.3-70b with loaded words is rated least convincing. Overall, ds-llama-70b produces the most convincing arguments (0.101), and storytelling emerges as the most effective persuasion strategy (0.096). The least convincing model and strategy are gpt-4o-mini and loaded words, respectively.

### Perceived emotions/fallacies vs. convincingness

Figure 4 shows the average convincingness scores for each emotion and fallacy category; here, the categories were assigned by the annotators. We report t-test results across emotion and fallacy categories in Tables 9 and 10 (Appendix D). We observe that arguments identified as fallacy-free (‘none’) consistently receive above-average scores across dif-

<sup>8</sup>Specifically, we group arguments into bins of 5 tokens by length difference, then correlate the average length difference with the change in F1 scores across these bins. We analyze bins with length differences between 0 and 40 tokens, which already account for 98% of the data.

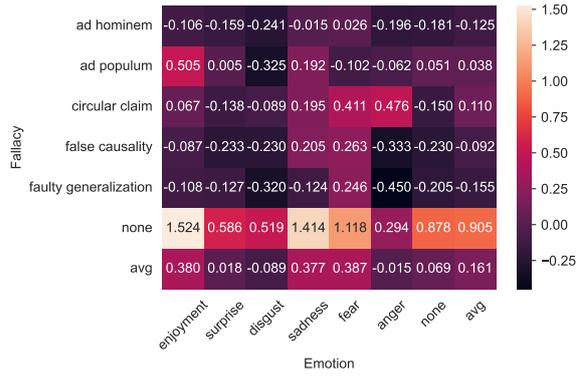


Figure 4: Average convincingness scores (z-score normalized) for each emotion and fallacy category based on human annotations (**perceived convincingness, emotions, and fallacies**). Full pairwise t-test results appear in Tables 9 and 10 (Appendix D).

	enjoyment	surprise	anger	disgust	sadness	fear	none
F1	<b>0.528</b>	0.464	0.501	0.461	0.337	<u>0.283</u>	0.437
p	0.156	0.600	0.215	0.663	0.054	0.023	-

Table 5: Macro-F1 for fallacy detection by **perceived emotion** (averaged over annotators); the best value is in bold; the worst is underlined.  $p$ -values from t-tests comparing each emotion to the ‘none’ category are shown in the last column; full pairwise t-test results appear in Table 11 (Appendix D).

ferent emotions, ranking first among fallacy categories on average (0.9), significantly ahead of the second- and third-most convincing categories, circular claim (0.11;  $p \leq 0.01$ ) and ad populum (0.038;  $p \leq 0.01$ ). The remaining fallacies receive below-average ratings, with faulty generalization being rated least convincing ( $-0.16$ ). Among emotions, fear, enjoyment, and sadness are rated as the most convincing on average (0.38–0.39), significantly more convincing than the other categories, including the non-emotional one ( $p \leq 0.01$ ); all of the other emotions receive around-average ratings, ranging from  $-0.089$  for disgust to  $0.069$  for no emotion (‘none’), where arguments with anger and disgust are rated significantly less convincing than non-emotional ones ( $p \leq 0.01$ ).

**Overall**, the results suggest that humans perceive fallacies as flaws in persuasion, leading to lower convincingness scores. When positive emotions like enjoyment are perceived, people may show greater openness to others’ viewpoints (Isen, 2000; Tsai and Li, 2023), resulting in higher convincingness ratings. Fear and sadness also appear to enhance convincingness.

**Perceived emotions vs. fallacy detection** Table 5 displays the macro F1 scores of humans in

detecting fallacies, averaged across annotators, for each annotator-assigned emotion category. We report the full significance test results in Table 11 in Appendix D. We observe that human performs best when perceiving enjoyment, significantly better than when perceiving sadness and fear (0.528 vs. 0.337/0.283;  $p = 0.016/0.012$ ). Fear and sadness appear to impair fallacy detection relative to the neutral condition (0.283/0.337 vs. 0.437;  $p = 0.023/0.054$ ). Positive emotions have been shown to broaden attentional scope (Derryberry and Tucker, 1994; Basso et al., 1996; Fredrickson, 2004) and enhance critical thinking (Jung et al., 2014) compared to negative emotions, which aligns with our observations. However, we do not find that a neutral emotional state consistently/significantly leads to better reasoning performance than emotions, as observed in Jung et al. (2014).

Interestingly, the emotion that most enhances reasoning ability (enjoyment) and the two that most impair it (fear and sadness) are also the most convincing emotions observed in Figure 4. This may suggest **two pathways to perceived convincingness of fallacious arguments**: (1) positive emotions like enjoyment may increase openness to persuasion while still supporting logical scrutiny, enabling people to detect fallacies yet remain receptive to the major claims of the arguments; and (2) negative emotions like fear and sadness may impair reasoning processes, making it harder to recognize logical flaws, thus increasing perceived convincingness through reduced critical evaluation.

**Key Findings** Exposure to LLM-framed arguments reduces fallacy detection performance compared to originals, most notably with loaded words, but only slightly increases perceived convincingness. Convincingness ratings are higher when there are no errors detected. Emotion-wise: perceived enjoyment improves fallacy detection relative to negative states fear and sadness, whereas fear and sadness impair human reasoning; the three emotions also correlate with higher convincingness.

## 5 Concluding Remarks

This work presents, to our knowledge, the first systematic study of how (AI-framed) emotions affect fallacy detection and argument convincingness. We benchmarked LLMs for their ability to inject emotions into arguments and created a dataset of emotionally framed fallacious arguments, which we then evaluated through human studies. Our results

show that while LLMs can frame textual emotions without disturbing the reasoning structure of arguments from premise to conclusion, humans often do not perceive the intended emotions, leading to a gap between emotional expression and subjective experience. Importantly, AI-driven emotional framing reduces human performance in fallacy detection but only modestly increases convincingness. Certain perceived emotions, such as fear, sadness, and enjoyment, are associated with higher persuasive impact; enjoyment also facilitates fallacy detection, whereas sadness and fear impair it.

These findings highlight both opportunities and risks. On the one hand, our study provides new insights into the interplay between emotions and reasoning, offering resources for research in computational argumentation. Future work could also explore the effects of emotions on LLM behavior in argumentation. For example, Payandeh et al. (2024) find that LLMs are substantially more often convinced when presented with logical fallacies in a multi-round debate setting, but do not consider the role of emotions. On the other hand, the ability of LLMs to strategically inject emotions into arguments underscores potential misuse for manipulation and misinformation. Future work should expand the scope to more diverse and real-world arguments, explore refined prompting strategies, and investigate mitigation methods to safeguard against malicious applications.

## Limitations & Ethical Concerns

Limitations of this work include: (1) To mitigate annotation complexity, we only leveraged a simple dataset based on online educational resources for logical fallacy detection, which may differ from real-world arguments in terms of length, topics, and complexity of logical structure etc. This may have biased the results and weakened the generalizability in real-world scenarios. Also, there are no genuine fallacy-free arguments in our dataset to make a fairer comparison for convincingness between fallacy-free and fallacious arguments. (2) The evaluation of LLMs on the emotional framing task was conducted on a relatively small sample (25 outputs per model) and without repeated runs. This design may have introduced randomness into the results. (3) We did not systematically optimize prompts for emotional framing, but relied on a single heuristic prompt that appeared to yield reasonable outputs across models. Consequently, LLM performance

in this task may have been underestimated. (4) We observe low agreements among crowdworkers for emotion-related annotations in §3 (0.2 Pearson’s  $r$ ). However, using the two student annotators from §3 yielded largely comparable results: synthetic arguments were still rated higher than originals in terms of emotional appeal (2.41 vs. 2.08 EA on average), ds-qwen3-8b produced the most emotionally appealing arguments (2.55) while o3-mini the least (2.18), and emotion match rates remained consistently low across models (23.2%-52.9%).

Our findings show that the emotionally framed arguments by LLMs can impair the logical detection performance of humans, which raises risks of misuse for generating persuasive misinformation, propaganda, or manipulative content. Although our experiments were conducted under controlled research settings, the same methods could be applied to deliberately obscure logical flaws in harmful ways.

We used ChatGPT for text refinement and search for relevant papers. We obtained annotator consent (via Google Forms) to use their annotations.

## Acknowledgements

We thank the meta-reviewer and reviewers for their thoughtful and constructive comments, which have improved the final version of the paper. The NLLG Lab gratefully acknowledges support from the German Research Foundation (DFG) via the Heisenberg Grant EG 375/5-1. This project has been conducted as part of the EMCONA (The Interplay of Emotions and Convincingness in Arguments) project, which is funded by the DFG (projects EG-375/9-1 and KL2869/12-1).

## References

- Ott Rhetoric Aristotle and George A Kennedy [translator]. 1991. A theory of civic discourse. *2nd éd.*
- Michael R Basso, Bruce K Schefft, M Douglas Ris, and William N Dember. 1996. Mood and global-local visual processing. *Journal of the international neuropsychological society*, 2(3):249–255.
- Sahbi Benlamine, Maher Chaouachi, Serena Villata, Elena Cabrio, Claude Frasson, and Fabien Gandon. 2015. Emotions in argumentation: an empirical evaluation. In *International Joint Conference on Artificial Intelligence, IJCAI 2015*, pages 156–163.
- Isabelle Blanchette. 2006. The effect of emotion on interpretation and logic in a conditional reasoning task. *Memory & cognition*, 34(5):1112–1125.
- Isabelle Blanchette, Sarah Gavigan, and Kathryn Johnston. 2014. Does emotion help or hinder reasoning? the moderating role of relevance. *Journal of Experimental Psychology: General*, 143(3):1049.
- Laura Ana Maria Bostan, Evgeny Kim, and Roman Klinger. 2020. **GoodNewsEveryone: A corpus of news headlines annotated with emotions, semantic roles, and reader perception**. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1554–1566, Marseille, France. European Language Resources Association.
- Eva Cantín Larumbe and Adriana Chust Vendrell. 2025. **Argumentative fallacy detection in political debates**. In *Proceedings of the 12th Argument mining Workshop*, pages 369–373, Vienna, Austria. Association for Computational Linguistics.
- Yanran Chen and Steffen Eger. 2025. **Do emotions really affect argument convincingness? a dynamic approach with LLM-based manipulation checks**. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 24357–24381, Vienna, Austria. Association for Computational Linguistics.
- Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barrón-Cedeño, Rostislav Petrov, and Preslav Nakov. 2019. **Fine-grained analysis of propaganda in news articles**. 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 5636–5646, Hong Kong, China. Association for Computational Linguistics.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, J. L. Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qishi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng

- Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. 2025. *Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning*.
- Julien A Deonna and Fabrice Teroni. 2012. From justified emotions to justified evaluative judgements. *Dialogue: Canadian Philosophical Review/Revue canadienne de philosophie*, 51(1):55–77.
- Douglas Derryberry and Don M Tucker. 1994. Motivating the focus of attention.
- Santiago Echeverri. 2019. Emotional justification. *Philosophy and Phenomenological Research*, 98(3):541–566.
- Paul Ekman. 1992. An argument for basic emotions. *Cognition & emotion*, 6(3-4):169–200.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. *Summeval: Re-evaluating summarization evaluation*. *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Barbara L Fredrickson. 2004. The broaden-and-build theory of positive emotions. *Philosophical transactions of the royal society of London. Series B: Biological Sciences*, 359(1449):1367–1377.
- Max Glockner, Yufang Hou, Preslav Nakov, and Iryna Gurevych. 2024. *Missci: Reconstructing fallacies in misrepresented science*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4372–4405, Bangkok, Thailand. Association for Computational Linguistics.
- Max Glockner, Yufang Hou, Preslav Nakov, and Iryna Gurevych. 2025. *Grounding fallacies misrepresented in scientific publications in evidence*. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 9732–9767, Albuquerque, New Mexico. Association for Computational Linguistics.
- Pierpaolo Goffredo, Mariana Chaves, Serena Villata, and Elena Cabrio. 2023. *Argument-based detection and classification of fallacies in political debates*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 11101–11112, Singapore. Association for Computational Linguistics.
- Pierpaolo Goffredo, Shohreh Haddadan, Vorakit Vorakitphan, Elena Cabrio, and Serena Villata. 2022. Fallacious argument classification in political debates. In *Thirty-First International Joint Conference on Artificial Intelligence {IJCAI-22}*, pages 4143–4149. International Joint Conferences on Artificial Intelligence Organization.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Alonso, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yearly, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar

Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gouget, Virginie Do, Vish Vogeti, Vitor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuwei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Barambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz,

Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Evebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihalescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiao Cheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin

- Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. 2024. [The llama 3 herd of models](#).
- Lynn Greschner and Roman Klinger. 2024. [Fearful falcons and angry llamas: Emotion category annotations of arguments by humans and llms](#).
- Lynn Greschner, Sabine Weber, and Roman Klinger. 2025. [Trust me, i can convince you: The contextualized argument appraisal framework](#).
- Amparo Caceres Gutierrez, Julián Tejada, and Enrique García Fernández-Abascal. 2025. Elicited emotion: effects of inoculation of an art style on emotionally strong images. *Experimental Brain Research*, 243(4):89.
- Ivan Habernal and Iryna Gurevych. 2016. [Which argument is more convincing? analyzing and predicting convincings of web arguments using bidirectional LSTM](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1589–1599, Berlin, Germany. Association for Computational Linguistics.
- Ivan Habernal and Iryna Gurevych. 2017. Argumentation mining in user-generated web discourse. *Computational linguistics*, 43(1):125–179.
- Ivan Habernal, Raffael Hannemann, Christian Polak, Christopher Klamm, Patrick Pauli, and Iryna Gurevych. 2017. [Argotario: Computational argumentation meets serious games](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 7–12, Copenhagen, Denmark. Association for Computational Linguistics.
- Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. 2018. [Before name-calling: Dynamics and triggers of ad hominem fallacies in web argumentation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 386–396, New Orleans, Louisiana. Association for Computational Linguistics.
- Thomas Haider, Steffen Eger, Evgeny Kim, Roman Klinger, and Winfried Menninghaus. 2020. [PO-EMO: Conceptualization, annotation, and modeling of aesthetic emotions in German and English poetry](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 1652–1663, Marseille, France. European Language Resources Association.
- Chadi Helwe, Tom Calamai, Pierre-Henri Paris, Chloé Clavel, and Fabian Suchanek. 2024. [MAFALDA: A benchmark and comprehensive study of fallacy detection and classification](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 4810–4845, Mexico City, Mexico. Association for Computational Linguistics.
- Alice M Isen. 2000. Some perspectives on positive affect and self-regulation. *Psychological inquiry*, 11(3):184–187.
- Alice M Isen and Barbara Means. 1983. The influence of positive affect on decision-making strategy. *Social cognition*, 2(1):18–31.
- Jiwon Jeong, Hyeju Jang, and Hogun Park. 2025. [Large language models are better logical fallacy reasoners with counterargument, explanation, and goal-aware prompt formulation](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 6918–6937, Albuquerque, New Mexico. Association for Computational Linguistics.
- Zhijing Jin, Abhinav Lalwani, Tejas Vaidhya, Xiaoyu Shen, Yiwen Ding, Zhiheng Lyu, Mrinmaya Sachan, Rada Mihalcea, and Bernhard Schoelkopf. 2022. [Logical fallacy detection](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 7180–7198, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Nadine Jung, Christina Wranke, Kai Hamburger, and Markus Knauff. 2014. How emotions affect logical reasoning: evidence from experiments with mood-manipulated participants, spider phobics, and people with exam anxiety. *Frontiers in psychology*, 5:570.
- Barbara Konat, Ewelina Gajewska, and Wiktoria Rossa. 2024. Pathos in natural language argumentation: Emotional appeals and reactions. *Argumentation*, pages 1–35.
- Yuanyuan Lei and Ruihong Huang. 2024. [Boosting logical fallacy reasoning in LLMs via logical structure tree](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13157–13173, Miami, Florida, USA. Association for Computational Linguistics.
- Jennifer S Lerner and Dacher Keltner. 2001. Fear, anger, and risk. *Journal of personality and social psychology*, 81(1):146.
- Yanda Li, Dixuan Wang, Jiaqing Liang, Guochao Jiang, Qianyu He, Yanghua Xiao, and Deqing Yang. 2024. [Reason from fallacy: Enhancing large language models’ logical reasoning through logical fallacy understanding](#). In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 3053–3066, Mexico City, Mexico. Association for Computational Linguistics.
- Artur Lugmayr, Erkki Sutinen, Jarkko Suhonen, Carolina Islas Sedano, Helmut Hlavacs, and Calkin Suero Montero. 2017. Serious storytelling—a first definition and review. *Multimedia tools and applications*, 76(14):15707–15733.

- Qingsong Ma, Ondřej Bojar, and Yvette Graham. 2018. [Results of the WMT18 metrics shared task: Both characters and embeddings achieve good performance](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 671–688, Belgium, Brussels. Association for Computational Linguistics.
- Qingsong Ma, Johnny Wei, Ondřej Bojar, and Yvette Graham. 2019. [Results of the WMT19 metrics shared task: Segment-level and strong MT systems pose big challenges](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 62–90, Florence, Italy. Association for Computational Linguistics.
- Fabrizio Macagno. 2022a. Argumentation profiles and the manipulation of common ground. the arguments of populist leaders on twitter. *Journal of Pragmatics*, 191:67–82.
- Fabrizio Macagno. 2022b. [Argumentation schemes, fallacies, and evidence in politicians’ argumentative tweets—a coded dataset](#). *Data in Brief*, 44:108501.
- Diane M Mackie and Leila T Worth. 1989. Processing deficits and the mediation of positive affect in persuasion. *Journal of personality and social psychology*, 57(1):27.
- Beth Innocenti Manolescu. 2006. A normative pragmatic perspective on appealing to emotions in argumentation. *Argumentation*, 20(3):327–343.
- Nailia Mirzakhmedova, Marcel Gohsen, Chia Hao Chang, and Benno Stein. 2024. Are large language models reliable argument quality annotators? In *Conference on Advances in Robust Argumentation Machines*, pages 129–146. Springer.
- Yusuke Mori, Hiroaki Yamane, Yusuke Mukuta, and Tatsuya Harada. 2022. Computational storytelling and emotions: A survey. *arXiv preprint arXiv:2205.10967*.
- Luca Mouchel, Debjit Paul, Shaobo Cui, Robert West, Antoine Bosselut, and Boi Faltings. 2024. [A logical fallacy-informed framework for argument generation](#).
- Elena Musi, Myrto Aloumpi, Elinor Carmi, Simeon Yates, and Kay O’Halloran. 2022. Developing fake news immunity: fallacies as misinformation triggers during the pandemic. *Online Journal of Communication and Media Technologies*, 12(3).
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng,

- Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. [Gpt-4 technical report](#).
- Amalie Brogaard Pauli, Isabelle Augenstein, and Ira Assent. 2025. [Measuring and benchmarking large language models' capabilities to generate persuasive language](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 10056–10075, Albuquerque, New Mexico. Association for Computational Linguistics.
- Amirreza Payandeh, Dan Pluth, Jordan Hosier, Xuesu Xiao, and Vijay K. Gurbani. 2024. [How susceptible are LLMs to logical fallacies?](#) In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 8276–8286, Torino, Italia. ELRA and ICCL.
- Karl Pearson. 1895. Vii. note on regression and inheritance in the case of two parents. *proceedings of the royal society of London*, 58(347-352):240–242.
- Mykolas Simas Poškus, Rasa Pilkauskaitė Valickienė, and Arvydas Kuzinas. 2019. The effects of descriptive imagery and appeals on emotions and intentions related to pro-environmental behavior. *Sustainability*, 11(19):5257.
- Alan Ramponi, Agnese Daffara, and Sara Tonelli. 2025. [Fine-grained fallacy detection with human label variation](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 762–784, Albuquerque, New Mexico. Association for Computational Linguistics.
- Saumya Sahai, Oana Balalau, and Roxana Horincar. 2021. [Breaking down the invisible wall of informal fallacies in online discussions](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 644–657, Online. Association for Computational Linguistics.
- Benjamin Schiller, Johannes Daxenberger, and Iryna Gurevych. 2021. [Aspect-controlled neural argument generation](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 380–396, Online. Association for Computational Linguistics.
- Norbert Schwarz and Gerald L Clore. 1983. Mood, misattribution, and judgments of well-being: informative and directive functions of affective states. *Journal of personality and social psychology*, 45(3):513.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2021. [“nice try, kiddo”](#): Investigating ad hominem in dialogue responses. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 750–767, Online. Association for Computational Linguistics.
- Paul Slovic, Ellen Peters, Melissa L Finucane, and Donald G MacGregor. 2005. Affect, risk, and decision making. *Health psychology*, 24(4S):S35.
- Ming-Hong Tsai and Norman P Li. 2023. Trait self-control, emotions, and openness to alternative viewpoints. *Current Research in Behavioral Sciences*, 5:100131.
- Henning Wachsmuth, Nona Naderi, Yufang Hou, Yonatan Bilu, Vinodkumar Prabhakaran, Tim Alberdingk Thijm, Graeme Hirst, and Benno Stein. 2017. [Computational argumentation quality assessment in natural language](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 176–187, Valencia, Spain. Association for Computational Linguistics.
- Xinzhi Wang, Jiayan Qian, Yudong Chang, Hang Yu, and Hui Zhang. 2025. Emotion inference of text based on counterfactual behavior knowledge. *Information Fusion*, 120:103060.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. [Qwen3 technical report](#).
- Changhua Yang, Kevin Hsin-Yih Lin, and Hsin-Hsi Chen. 2009. [Writer meets reader: Emotion analysis of social media from both the writer's and reader's perspectives](#). In *2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology*, volume 1, pages 287–290.
- Ran Zhang, Wei Zhao, Lieke Macken, and Steffen Eger. 2025. [Litransproqa: an llm-based literary translation](#)

evaluation metric with professional question answering.

Timon Ziegenbein, Shahbaz Syed, Felix Lange, Martin Potthast, and Henning Wachsmuth. 2023. [Modeling appropriate language in argumentation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4344–4363, Toronto, Canada. Association for Computational Linguistics.

## A AI-driven Emotional Framing

### A.1 Emotional Appeal Injection Prompt

Table 7 presents the prompts we used to inject emotional appeals into fallacious arguments.

### A.2 Argument Classifier

We use the prompt-based classifier from [Chen and Eger \(2025\)](#) to filter out non-standalone arguments and automatically assign short claims to arguments.<sup>9</sup> Since our text domain differs substantially from theirs (political debates), we first validate the reliability of the classifier using different models on 25 samples from our dataset. Two student annotators who are fluent in English, hired from German universities, annotated the 25 samples. See Appendix B.2 for the annotation guideline for this criterion and Figure 6 for the annotation interface (Task 3 ‘Claim Correctness’). As shown in Figure 5, qwen3-32b and gpt-4o perform best, both achieving perfect accuracy (1.0), whereas llama-3.3-70b reaches only about 0.5. This classifier performs better than reported in [Chen and Eger \(2025\)](#). A very likely reason is that [Chen and Eger \(2025\)](#) evaluate on real-world political debates, whereas our short, human-curated arguments have clear evidence–claim structures with claims mostly explicit, making the task substantially easier for models.

### A.3 Annotation

For the two more objective dimensions, namely reasoning preservation and coherence, we employed two student annotators from German universities, both fluent in English and enrolled in Master’s programs. They jointly labeled 20 overlapping instances for agreement checking and annotated the remaining instances independently. For the emotion-related dimensions, emotional appeal and emotion match, we recruited crowdworkers via Prolific,<sup>10</sup> requiring them to be native English speakers based in the US or UK with a completed higher

<sup>9</sup>Their prompt instructs the model to generate a claim when the text is identified as a standalone argument.

<sup>10</sup><https://www.prolific.com/>

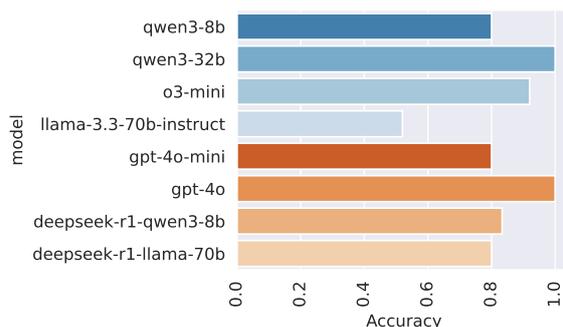


Figure 5: Accuracy of the prompt-based argument classifier from [Chen and Eger \(2025\)](#) using different LLMs on our dataset in both argument classification and claim generation.

education degree. Each instance was annotated by three workers, and each worker evaluated 25 arguments. See Appendix B for the annotation guidelines and screenshots of the annotation interfaces. The annotation agreements are reported separately for each dimension in Appendix A.4.

### A.4 Evaluation

**Reasoning Preservation (RP)** Since we will compare human performance on the emotionally enriched fallacious arguments with that on the original ones, it is essential that the logical structure of the fallacies remain unchanged. Therefore, we assess whether the reasoning structure from evidence to claim in the original arguments is preserved after emotional framing by LLMs. Annotators assign a label of ‘1’ if the reasoning is preserved and ‘0’ otherwise, given the original and generated argument. The **reasoning preservation rate** is computed as the average label across generated arguments. Among the randomly sampled 20 arguments annotated by both annotators, only 2 received differing annotations; in these cases, we take the average, 0.5, as the final score.

As shown in the ‘RP’ columns of Table 6, most models achieve a perfect score of 1, except when using the vivid language strategy, indicating that LLMs generally preserve reasoning well in this task. For imagery, only o3-mini scores below 1, and for loaded words and storytelling, only 3 out of 8 models score below 1. Notably, qwen3-32b is the only model to achieve a perfect score for vivid language. On average, RP rates range from 0.74 to 1.00, with qwen3-32b performing best (1.00), followed by ds-llama-70b (0.94) and llama-3.3-70b (0.93), while ds-qwen3-8b performs worst (0.74).

	Imagery					Loaded Words					Storytelling					Vivid Language					AVG				
	RP	Coh	LD	EA	EM	RP	Coh	LD	EA	EM	RP	Coh	LD	EA	EM	RP	Coh	LD	EA	EM	RP	Coh	LD	EA	EM
ds-llama-70b*	✓	✓	<u>14.4</u>	.10	<u>0</u>	✓	✓	6.3	.19	✓	.90	.90	<b>31.3</b>	.06	30.0	.86	✓	<b>14.5</b>	-.43	25.0	.94	<b>.98</b>	<b>16.6</b>	-.02	38.8
ds-qwen3-8b	✓	✓	17.0	-.49	20.0	<u>.40</u>	✓	3.7	<b>.90</b>	60.0	✓	<u>.57</u>	21.8	.00	<b>50.0</b>	<u>.57</u>	.86	6.1	<b>.39</b>	<b>71.4</b>	.74	.86	12.2	<b>.20</b>	<b>50.4</b>
gpt-4o	✓	.70	21.1	<b>.45</b>	30.0	✓	✓	3.7	-.09	33.3	<u>.80</u>	✓	23.5	<b>.50</b>	80.0	.75	✓	4.8	-.18	25.0	.89	.93	13.3	.17	42.1
gpt-4o-mini*	✓	✓	19.7	.26	40.0	✓	✓	4.1	-.16	<u>0</u>	✓	.86	23.0	.07	42.9	.60	✓	8.8	.26	0.0	.90	.96	13.9	.11	20.7
llama-3.3-70b*	✓	.78	21.9	.37	33.3	✓	✓	5.2	-.15	75.0	✓	✓	23.0	-.05	40.0	.71	✓	11.7	<u>-.59</u>	<u>42.9</u>	.93	.94	15.4	-.11	47.8
o3-mini	<u>.89</u>	<u>.00</u>	20.6	-.08	10.0	.75	<u>.75</u>	3.7	-.21	25.0	✓	.67	25.9	-.14	0.0	.80	✓	4.5	-.25	60.0	.86	<u>.60</u>	13.7	<u>-.17</u>	23.8
qwen3-32b*	✓	.83	21.5	.22	<b>33.3</b>	✓	✓	<b>6.3</b>	-.22	<b>83.3</b>	✓	✓	22.2	<u>-.39</u>	0.0	✓	✓	8.8	-.39	40.0	✓	.96	14.7	.02	39.2
qwen3-8b	✓	.86	<b>22.5</b>	<b>-.22</b>	14.3	.83	.83	<u>3.5</u>	.49	50.0	.83	✓	24.9	-.10	<u>16.7</u>	.67	<u>.83</u>	9.9	.05	16.7	.83	.88	15.2	.06	24.4
AVG	<b>.99</b>	<u>.77</u>	19.8	<b>.10</b>	<u>22.6</u>	.87	.95	<u>4.5</u>	.08	<b>53.3</b>	.94	.87	<b>24.5</b>	.07	32.4	<u>.74</u>	<b>.96</b>	8.6	<u>-.13</u>	35.1					

Table 6: Reasoning preservation rates (RP), coherence scores (Coh), average length differences between original and generated arguments (LD), average emotional appeal ratings (EA), and emotion match rates (EM). Highest values are shown in bold and lowest values are underlined for each metric across models. The top 4 overall models, used to generate data for the human study, are marked with \*. The last row shows the averages over models for each strategy, with the highest and lowest values for each metric highlighted accordingly. ✓ denotes perfect scores (1.00 or 100.0).

**Coherence (Coh)** We evaluate whether each generated argument is both emotionally and textually coherent — i.e., whether the introduced emotional content aligns with the original argument’s content and whether the generated argument is coherent in itself. Annotation follows the same binary procedure as for reasoning preservation, and the final **coherence score** is computed as the average over all generated arguments. The two annotators disagreed on only 1 out of 20 arguments; as with RP, we take the average score 0.5 for arguments receiving different annotations.

As the ‘Coh’ columns of Table 6 display, coherence scores range from 0.00 (o3-mini with imagery) to 1.00 (several models with loaded words and vivid language). LLMs generally produce more coherent texts when using loaded words and vivid language compared to imagery and storytelling (0.95-0.96 vs. 0.77-0.87 on average). Overall, ds-llama-70b performs best (0.98 on average), followed by gpt-4o-mini and qwen3-32b (both 0.96), whereas o3-mini lags substantially behind the others (0.60 vs. 0.86–0.98).

**Length Difference (LD)** We expect different emotional framing strategies to affect the extent of length change differently. We report the average number of additional tokens across all generated arguments as the **length difference**, where tokens are defined by splitting text on spaces.

We present the results in the ‘LD’ columns of Table 6. As expected, imagery and storytelling lead to a greater number of additional tokens (19.8 and 24.5, respectively, averaged across models), whereas loaded words and vivid language add only 4.5 and 8.6 tokens on average. This suggests that LLMs can generally follow human instructions in

this task, as we instructed them to add additional sentences for imagery and storytelling, but only to modify word choices for loaded words and vivid language. The LLaMA-based models (ds-llama-70b and llama-3.3-70b) add the most tokens on average (16.6 and 15.4, respectively), whereas ds-qwen3-8b adds the fewest (12.2).

**Emotional Appeal (EA)** We assess whether the generated arguments are more emotion-evoking than the original ones. Annotators rate the emotional appeal on a 3-point Likert scale for both original and generated arguments, following the definition of emotional appeal in Wachsmuth et al. (2017). We compute the average Pearson correlation across annotator pairs, obtaining a weakly positive value of 0.20.<sup>11</sup> We apply z-score normalization to the ratings to ensure comparability across annotators.

Here, we report the averages over generated arguments in the ‘EA’ columns of Table 6. The original arguments receive an average EA rating of  $-0.45$ , which is lower than that of the great majority of generated arguments (30 out of 32), except for ds-qwen3-8b with imagery ( $-0.49$ ) and llama-3.3-70b with vivid language ( $-0.59$ ). This suggests that LLMs are able to change the level of emotional appeal in arguments, consistent with the observation from Chen and Eger (2025). Among the strategies, imagery is rated highest in emotional appeal (0.10), while vivid language obtains the lowest scores ( $-0.13$ ). Among LLMs, ds-qwen3-8b generates the most emotionally appealing arguments (0.20),

<sup>11</sup>The agreement is low; however, we obtain largely comparable conclusions for this criterion using a totally different group of annotators in §5, implying good reproducibility of our results.

whereas o3-mini again scores lowest (-0.17).

**Emotion Match (EM)** We assess whether the introduced emotion matches the target emotion specified during generation. Annotators select one of the six emotion types, or ‘none’ if no clear emotion is present, to select which emotion this argument attempt to evoke. We frame the task more objectively to assess whether the expressed emotions match the target emotions, which do not necessarily need to be perceived by humans. We observe a Cohen’s Kappa of 0.306, averaged over annotator pairs, for this task.

In the EM columns of Table 6, we report the emotion match rate as the percentage of cases where the labeled emotion matches the target emotion. Overall, the values are generally low, ranging from 20.7% for gpt-4o-mini to 50.4% for ds-qwen3-8b, averaged over strategies. Models perform comparatively better when using loaded words (53.3% vs. 22.6-35.1% on average).

## B Annotation Guidelines & Interfaces

**Annotation in §3** Student annotators used Google Sheets<sup>12</sup> for annotating the reasoning preservation and coherence criteria. The annotation guidelines and the screenshot of the annotation interface are presented in Appendix B.2 and Figure 6, respectively. We collected crowdsourcing annotations for the emotional appeal and emotion match criteria via Google Forms.<sup>13</sup> Figures 7 and 8 show the screenshots of the annotation interfaces for the annotation guideline and one example instance, respectively.

**Annotation in §4** We collected crowdsourcing annotations for convincingness, fallacy and emotion via Google Forms. The screenshot of the annotation guidelines presented to the crowdworkers is displayed in Figure 9. The definitions and examples of fallacies presented to the annotators are shown in Appendix B.3, which are largely based on Jin et al. (2022). The five tasks of the fallacy detection quiz taken from Jin et al. (2022) are listed in Appendix B.1. The screenshot of the annotation interface for one example instance is presented in Figure 10.

<sup>12</sup><https://workspace.google.com/products/sheets/>

<sup>13</sup><https://workspace.google.com/products/forms/>

argument_text	claim	1 reasoning preserv...	2 coherence	3 claim correctness
Everyone does it. You should too.	Alarming,ly, everyone is resorting to it so you should fall in line and do it before it's too late.	1: The generated ...	1: Succeeds on bo...	1: The claim matche...
A father tells his son not to start smoking as he will regret it when he is older, and the son points out that his father is a smoker. This does not alter the fact that his son may regret smoking when he is older.	Everyone should do it.	1: The generated ...	1: Succeeds on bo...	1: The claim matche...
There are more levels on the books than ever before, and more crimes are being committed than ever before. Therefore, to reduce crime, we must eliminate the laws.	Smoking may lead to regret in the future.	1: The generated ...	1: Succeeds on bo...	1: The claim matche...

Figure 6: Screenshots of Google Sheets used for annotations on reasoning preservation, coherence, and claim correctness.

第 3 节, 共 28 节

**Annotation Starts!**

For each argument presented, your tasks are to:

- Evaluate the level of emotional appeal used in the **argument** on a Likert scale of 1-3

**Definition:** Argumentation makes a successful emotional appeal if it creates emotions in a way that makes the target audience more open to the author’s arguments.

- 1 (Low)
- 2 (Medium)
- 3 (High)

- Select the emotion the **argument** attempts to evoke from the following list: *sadness, disgust, enjoyment, fear, surprise, and anger*. If no emotion exists, select ‘None’.

**Note:** Some attention checks are embedded in the tasks. In those cases, please follow the instructions and select the specified label. Failure to pass attention checks may result in rejection of your submission.

In the following, you will complete the 2 tasks for 25 arguments.

Figure 7: Screenshots of the Google Forms annotation interface showing guidelines for emotional appeal and emotion match evaluation in §3.

Argument 3

A radical new study reveals an outrageous link: church attendance seemingly boosts marriage endurance. This should be met with indignation, considering the established stance against separation within faith. It’s blatantly obvious that exploitation of this misconception is happening, preying on the vulnerable.

Evaluate the level of emotional appeal on a Likert scale of 1-3. \*

1 (Low)

2 (Medium)

3 (High)

Select the emotion the argument attempts to evoke from the following list: \*

disgust

enjoyment

anger

sadness

Figure 8: Screenshots of the Google Forms annotation interface showing one annotation instance for emotional appeal and emotion match evaluation in §3.

### B.1 Fallacy Detection Quiz

- **Circular Claim:** She is the best candidate for president because she is better than the other candidates!
- **Ad Populum:** Killing thousands of people as

---

Prompts

---

====System Prompt=====

There are many methods to add emotional appeal—also known as pathos—to an argument, but here are some of the most common and effective ones:

1. Storytelling – Sharing personal stories or anecdotes that evoke emotion.
2. Vivid Language – Using descriptive, sensory, or emotionally charged words.
3. Loaded Words – Choosing words with strong positive or negative connotations.
4. Imagery – Painting mental pictures to stir emotions.

====User Prompt=====

I will give an argument. Your task is to add emotional appeals to the given argument by using **\*\*{strategy}\*\***. You should {method} to change its emotional tone. You generate 6 arguments, each with one kind of emotion: *\*anger\**, *\*surprise\**, *\*disgust\**, *\*enjoyment\**, *\*fear\**, and *\*sadness\**. Note: Please put the modified parts in bold and do **\*\*NOT\*\*** use emotionally charged words in the added sentence when using storytelling and imagery.

Answer in the following way:

argument\_anger: [the generated argument that appeals to anger]

====

argument\_suprise: [the generated argument that appeals to suprise]

====

...

Argument: {argument}

---

Variables

---

**strategy:** the used emotional framing strategy.

strategy ∈ {storytelling, imagery, loaded words, vivid language}

**method:** the corresponding text-edit method, chosen from: {storytelling: keep the original argument but **\*\*only\*\*** add one sentence; vivid language: keep the original argument but **\*\*only\*\*** change its word choices; loaded words: keep the original argument but **\*\*only\*\*** change its word choices; imagery: keep the original argument but **\*\*only\*\*** add one sentence}.

**argument:** the original fallacious argument

---

Table 7: Prompts for injecting emotional appeals into fallacious arguments.

a result of drug war campaign is not a crime to humanity because millions of Filipino support it.

- **Faulty Generalization:** Sometimes flu vaccines don't work; therefore vaccines are useless.
- **False Causality:** Every severe recession follows a Republican Presidency; therefore Republicans are the cause of recessions.
- **Ad Hominem:** I cannot listen to anyone who does not share my social and political values.

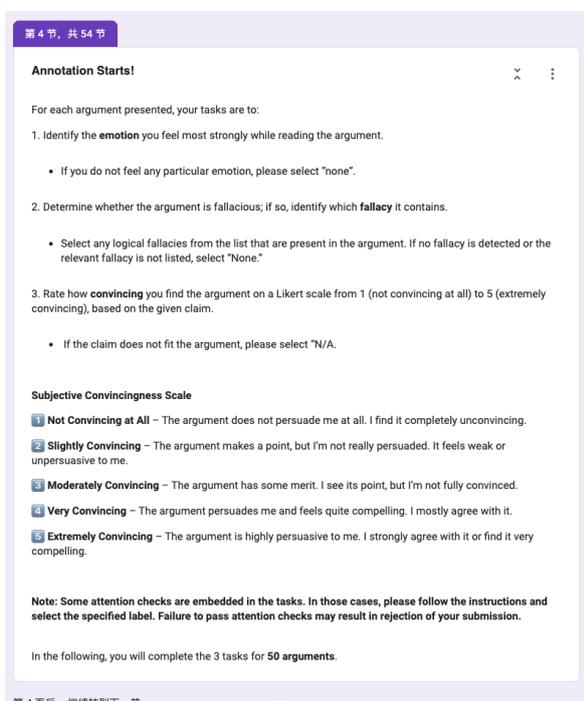


Figure 9: Screenshots of the Google Forms annotation interface showing guidelines for fallacy, emotion and convincingness annotation in §4.

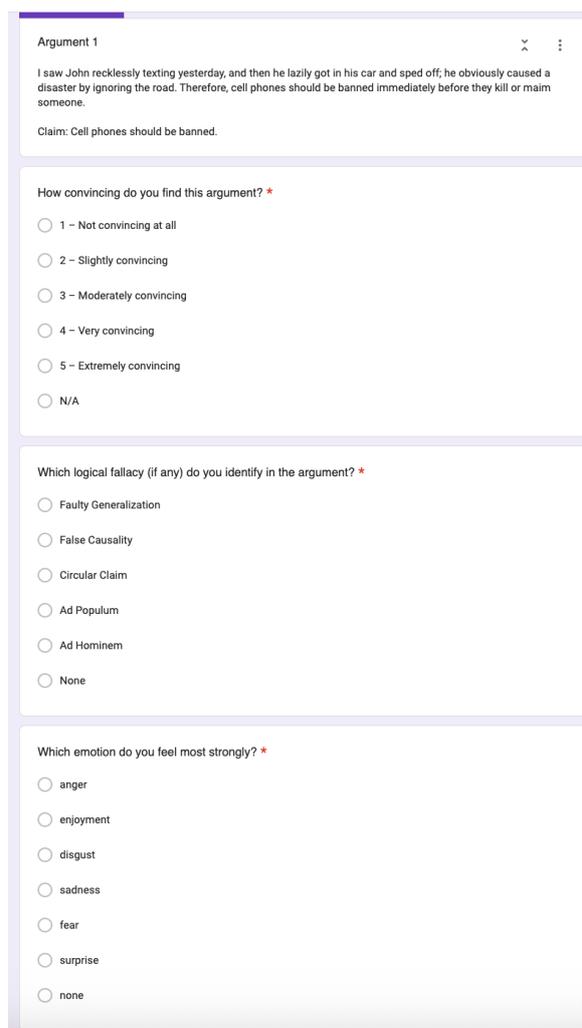


Figure 10: Screenshots of the Google Forms annotation interface showing one annotation instance for fallacy, emotion and convincingness annotation in §4.

## 1. Claim-evidence Reasoning Preservation

**Definition:** How well the generated argument preserves the original argument’s **core claim**, **supporting evidence**, and **underlying reasoning**.

- **0:** The generated argument **alters or omits key parts** of the original claim, evidence, or reasoning.
  - **1:** The generated argument **accurately retains** the claim, evidence, and reasoning of the original.
- 

## 2 Coherence

**Definition:** The overall coherence of the argument, considering

- (a) **Emotional coherence** — the internal consistency/naturalness of emotional tone, and
- (b) **Textual coherence** — the semantic meaningfulness of the text (phrases/sentences are interpretable; entities/roles are used consistently; modifiers and predicates make sense).

- **0:** Fails on **either** emotional coherence **or** textual coherence.

**Examples:** abrupt/unmotivated emotional shifts (cheerful → bitter without cause); semantically anomalous phrases (“the policy smells triangular”); incompatible predicates (“the bill apologized to voters”); unresolved/deictic references that make meaning unclear (“this proves it” with no recoverable referent); unexplained entity switches that break *meaning* rather than logic.

- **1:** Succeeds on **both** emotional coherence **and** textual coherence.

**Examples:** tone is consistent (or neutral) and fits the argument; sentences are meaningful and interpretable; entities and roles stay consistent; word choices are semantically appropriate; additions integrate smoothly at the level of *meaning*.

---

## 3 Claim Correctness

**Definition:** This label indicates whether the **claim** accurately reflects the content and intent of the **original argument**.

- **0 (Mismatch)** – The claim does **not** accurately represent the **original argument**, or the given text is not actually an argument.
- **1 (Match)** – The claim is a faithful representation of the argument.

### B.3 Fallacy Definitions & Examples

**Faulty Generalization:** An informal fallacy wherein a conclusion is drawn about all or many instances of a phenomenon on the basis of one or a few instances of that phenomenon. is an example of jumping to conclusions.

- **Logical form:** X (an individual or small subset of Z) has attribute Y. Therefore, all of Z has attribute Y.
- **Example:** I met a tall man who loved to eat cheese. Now I believe that all tall people like cheese.
- **Explain:** The conclusion that all tall people like cheese generalizes from a single tall person (sample size = 1) to all tall people. The sample is far too small to justify such a sweeping claim. **X = "a tall man", Y = "loved to eat cheese", Z = "all tall people"**

**False Causality:** A statement that jumps to a conclusion implying a causal relationship without supporting evidence.

- **Logical form:** X occurred, then Y occurred. Therefore, X caused Y.
- **Example:** Every time I wash my car, it rains. Me washing my car has a definite effect on the weather.
- **Explain:** In this example, just because the rain (Y) often follows the act of washing the car (X) doesn't mean one causes the other. This fallacy confuses correlation (two events happening together) with causation (one event directly causing the other). **X = "I wash my car", Y = "it rains"**

**Circular Claim:** A fallacy where the end of an argument comes back to the beginning without having proven itself.

- **Logical form:** X is true because of Y. Y is true because of X.
- **Example:** J.K. Rowling is a wonderful writer because she writes so well.
- **Explain:** This argument says J.K. Rowling is wonderful (X) because she writes well (Y)—but writing well is simply another way of saying she's a wonderful writer. The conclusion (X) and the reason (Y) are saying the same thing in different words, making the reasoning circular and uninformative. **X = "J.K. Rowling is a wonderful writer", Y = "she writes so well"**

**Ad Populum:** A fallacious argument which is based on affirming that something is real or better because the majority thinks so.

- **Logical form:** A large number of people believe X. Therefore, X is true (or good/right).
- **Example:** Everyone should like coffee: 95% of teachers do!
- **Explain:** the argument claims that people should like coffee (X) simply because a large group (teachers) already do. However, the popularity of coffee among teachers does not logically prove that everyone ought to like it.

**Ad Hominem:** An irrelevant attack towards the person or some aspect of the person who is making the argument, instead of addressing the argument or position directly.

- **Logical form:** Person X makes claim Y. Person X is attacked (e.g., called a moron). Therefore, claim Y is false.
- **Example:** What can our new math teacher know? Have you seen how fat she is?
- **Explain:** person X is the math teacher, and although her actual claim (Y) is never stated, the implication is that something she says or teaches is being dismissed. The attack focuses on her physical appearance, which has nothing to do with her knowledge or ability to teach math. This is a textbook ad hominem fallacy.

batch	emo_full	emo_maj	emo_kappa	fallacy_full	fallacy_maj	fallacy_kappa	conv_full	conv_maj	conv_pearson
10	0.060	0.640	0.091	0.300	0.780	0.328	0.360	0.840	0.383
11	0.140	0.660	0.057	0.400	0.780	0.431	0.380	0.840	0.185
12	0.320	0.780	0.069	0.200	0.720	0.232	0.520	0.900	0.267
13	0.100	0.640	0.097	0.280	0.880	0.344	0.420	0.840	0.397
14	0.100	0.640	0.139	0.300	0.860	0.335	0.340	0.800	0.154
15	0.020	0.540	0.009	0.560	0.900	0.591	0.060	0.740	0.114
16	0.080	0.740	0.183	0.140	0.840	0.244	0.080	0.860	0.192
17	0.160	0.700	0.095	0.080	0.620	0.102	0.380	0.840	0.397
18	0.340	0.820	0.020	0.140	0.640	0.143	0.320	0.840	0.423
19	0.120	0.700	0.089	0.260	0.680	0.238	0.340	0.820	0.233
1	0.440	0.900	0.010	0.080	0.560	0.068	0.080	0.880	0.070
20	0.480	0.900	0.039	0.360	0.700	0.337	0.480	0.840	0.480
2	0.180	0.800	0.038	0.340	0.860	0.395	0.060	0.640	-0.056
3	0.060	0.540	0.082	0.060	0.460	0.014	0.020	0.520	0.039
4	0.100	0.720	0.129	0.260	0.860	0.296	0.200	0.740	0.466
5	0.180	0.700	0.125	0.520	0.900	0.535	0.520	0.920	0.430
6	0.040	0.820	0.114	0.360	0.840	0.402	0.560	0.920	0.109
7	0.180	0.740	0.123	0.440	0.900	0.511	0.540	0.880	0.317
8	0.240	0.780	0.046	0.220	0.700	0.243	0.320	0.700	0.283
9	0.000	0.840	-0.005	0.260	0.700	0.273	0.280	0.740	0.126
avg	0.167	0.730	0.078	0.278	0.759	0.303	0.313	0.805	0.250

Table 8: Inter-annotator agreements for emotion classification, fallacy classification, and convincingness ratings. We report the results using majority/full agreements for all annotation dimensions, the average Cohen’s kappa scores over annotators for the classification tasks, and the average Pearson correlations between convincingness ratings.

## C IAA Full Results

## D T-test Full Results

cat1	cat2	cat1 vs. cat2	p-value
ad hominem	ad populum	smaller	0.000***
ad hominem	circular claim	smaller	0.032**
ad hominem	false causality	larger	0.873
ad hominem	faulty generalization	larger	0.254
ad hominem	none	smaller	0.000***
ad populum	circular claim	larger	0.246
ad populum	false causality	larger	0.000***
ad populum	faulty generalization	larger	0.000***
ad populum	none	smaller	0.000***
circular claim	false causality	larger	0.007***
circular claim	faulty generalization	larger	0.000***
circular claim	none	smaller	0.000***
false causality	faulty generalization	larger	0.235
false causality	none	smaller	0.000***
faulty generalization	none	smaller	0.000***

Table 9: Pairwise t-test results for fallacy categories in Figure 4.  $p \leq 0.10$  are marked \*,  $p \leq 0.05$  \*\*, and  $p \leq 0.01$  \*\*\*.

cat1	cat2	cat1 vs. cat2	p-value
anger	disgust	larger	0.883
anger	enjoyment	smaller	0.000***
anger	fear	smaller	0.000***
anger	none	smaller	0.004***
anger	sadness	smaller	0.000***
anger	surprise	smaller	0.292
disgust	enjoyment	smaller	0.000***
disgust	fear	smaller	0.000***
disgust	none	smaller	0.003***
disgust	sadness	smaller	0.000***
disgust	surprise	smaller	0.214
enjoyment	fear	smaller	0.878
enjoyment	none	larger	0.003***
enjoyment	sadness	larger	0.945
enjoyment	surprise	larger	0.001***
fear	none	larger	0.014
fear	sadness	larger	0.825
fear	surprise	larger	0.003***
none	sadness	smaller	0.000***
none	surprise	larger	0.130
sadness	surprise	larger	0.000***
surprise	sadness	smaller	0.000***

Table 10: Pairwise t-test results for emotion categories in Figure 4.  $p \leq 0.10$  are marked \*,  $p \leq 0.05$  \*\*, and  $p \leq 0.01$  \*\*\*.

emo1	emo1_mean	emo2	emo2_mean	p-value
sadness	0.337	none	0.437	0.054*
sadness	0.337	disgust	0.461	0.055*
sadness	0.337	fear	0.283	0.494
sadness	0.337	enjoyment	0.528	0.012**
sadness	0.337	anger	0.501	0.008***
sadness	0.337	surprise	0.464	0.043**
none	0.437	disgust	0.461	0.663
none	0.437	fear	0.283	0.023**
none	0.437	enjoyment	0.528	0.156
none	0.437	anger	0.501	0.214
none	0.437	surprise	0.464	0.600
disgust	0.461	fear	0.283	0.040**
disgust	0.461	enjoyment	0.528	0.397
disgust	0.461	anger	0.501	0.532
disgust	0.461	surprise	0.464	0.958
fear	0.283	enjoyment	0.528	0.016**
fear	0.283	anger	0.501	0.008***
fear	0.283	surprise	0.464	0.029**
enjoyment	0.528	anger	0.501	0.728
enjoyment	0.528	surprise	0.464	0.412
anger	0.501	surprise	0.464	0.556

Table 11: Pairwise t-test results for Table 5.  $p \leq 0.10$  are marked \*,  $p \leq 0.05$  \*\*, and  $p \leq 0.01$  \*\*\*.