

Estimating Contribution Quality in Online Deliberations Using a Large Language Model

Lodewijk Gelauff^{*1}, Mohak Goyal^{*1}, Bhargav Dindukurthi², Ashish Goel¹, and Alice Siu¹

¹Stanford University

²Cadence Design Systems

{lodewijk,mohakg}@stanford.edu, bhargavdindukurthi1@gmail.com, {ashishg,asiu}@stanford.edu

Abstract

Deliberation involves participants exchanging knowledge, arguments, and perspectives and has been shown to be effective at addressing polarization. The Stanford Online Deliberation Platform facilitates large-scale deliberations. It enables video-based online discussions on a structured agenda for small groups without requiring human moderators. This paper’s data comes from various deliberation events, including one conducted in collaboration with Meta in 32 countries, and another with 38 post-secondary institutions in the US.

Estimating the quality of contributions in a conversation is crucial for assessing feature and intervention impacts. Traditionally, this is done by human annotators, which is time-consuming and costly. We use a large language model (LLM) alongside eight human annotators to rate contributions based on justification, novelty, expansion of the conversation, and potential for further expansion, with scores ranging from 1 to 5. Annotators also provide brief justifications for their ratings. Using the average rating from other human annotators as the ground truth, we find the model outperforms individual human annotators. While pairs of human annotators outperform the model in rating justification and groups of three outperform it on all four metrics, the model remains competitive.

We illustrate the usefulness of the automated quality rating by assessing the effect of nudges on the quality of deliberation. We first observe that individual nudges after prolonged inactivity are highly effective, increasing the likelihood of the individual requesting to speak in the next 30 seconds by 65%.

Using our automated quality estimation, we show that the quality ratings for statements prompted by nudging are similar to those made without nudging, signifying that nudging leads to more ideas being generated in the conversation without losing overall quality.

1 Introduction

Deliberation is a practice that can help bring together stakeholders and ensure an exchange of opinions, facts and understanding. Deliberation is used in various contexts, including mini-publics (Goodin and Dryzek 2006), citizen juries (Smith and Wales 2000) and deliberative polls (Fishkin 1997), to better understand a population’s opinions on a complex topic or inform a decision. Some of these designs

have reported promising spillover effects such as depolarization (Van Der Does and Jacquet 2023; Fishkin et al. 2021).

Conducting deliberation online offers numerous advantages. In-person deliberation can be prohibitively expensive due to transportation and venue costs (Luskin, Iyengar, and Fishkin 2014). Additionally, organizing unbiased and reproducible moderation is challenging and requires extensive training for human moderators. Online deliberation, however, can leverage automated processes that combine group intelligence and artificial intelligence to replace human moderators (Gelauff et al. 2023). This not only reduces costs but also ensures consistent and scalable facilitation.

Facilitating online deliberation presents both opportunities for data analysis and the responsibility to ensure the process aligns with its objectives. Key factors to consider include equitable participation, conversation quality, participant experience, and outcome quality. The online setting simplifies data collection and makes evaluation more affordable. As the scale of deliberation increases, rapid evaluation becomes crucial for organizers to adjust recruitment, materials, or process design promptly. In this paper, we design a scalable method for quantifying deliberative quality and show how it can assist platform designers in evaluating and improving moderation interventions and configurations.

Different aspects of quality can be established by analyzing participant surveys, third-party observations, or the content of the deliberation. Especially the content (transcripts) of the deliberation provide a promising avenue to evaluate interventions and moderation decisions that affect the conversation locally. Traditionally, content analysis relies on human annotators to categorize or score contributions – a resource-intensive method. The arrival of LLMs makes it feasible to use automated analysis to estimate the quality of a contribution, exchange or entire deliberation. We define criteria of interest, and design a query for the LLM to report scores for statements given the preceding contributions and other relevant context. We validate scores from these queries with human annotators.

This approach not only allows analysis at a much larger scale but also enables the use of such scores in real-time evaluation, information systems for human supervisors, and even moderation interventions. We illustrate one use of this approach by analysing the effect of sending targeted nudges to inactive participants on the quality of the conversation.

^{*}These authors contributed equally.

Roadmap After discussing related work, we will introduce our data, platform, and methodology. We will explain how we use an LLM to estimate the quality of contributions, and how we evaluate and utilize these estimates. Next, we will compare the performance of our model to that of human evaluators. Finally, we will demonstrate how these quality estimations can be used to optimize the performance of the deliberation platform by comparing the quality of statements made after nudges to those made without nudges.

2 Related Work

Evaluating a deliberation or discourse has long been a question of interest, with the aspects of analysis evolving over time. Qualities often sought after include equality, rational reasoning and engagement (Dutwin 2003). Black et al. (2014) identified various approaches to evaluating a deliberation. Besides indirect measurements (looking at antecedents and outcomes), they distinguish between micro-analytic and macro-analytic approaches. Micro-analytic approaches include content analysis (annotating units of conversation, usually based on a codebook by humans) and discourse analysis (qualitative analysis of the discourse). Macro-analytic approaches use a wider range of methods to gain insights into the overall deliberation event with commonly used self-assessments from participants, moderators and/or observers or case studies with an in-depth evaluation of some of the groups. For our purposes, content analysis is particularly relevant, and we will focus on that.

The criteria used in content analysis are not universally agreed upon and depend on what the evaluator is trying to understand. For example, Gerhards (1997) looked at the abortion debate in major German newspapers between 1970 and 1994, and evaluated representativeness, respect, level of justification and rationality. Steenbergen et al. (2003) looked at a parliamentary setting with (long, formal) spoken contributions. They evaluated them on a *Discourse Quality Index* with indicators of openness of participation, justification of demands, common good/empathy, respect (for the group, the topic and arguments) and constructiveness. In line with this work, Gold et al. (2015) defined four indicators and combined the content analysis with a visualization, defining participation (distribution of contributions), respect (frequency of interruptions), argumentation (looking at sentence structure) and persuasiveness (as an indirect metric).

In online education, a lot of work discusses how to evaluate individual contributions to evaluate learning. While this is a slightly different setting, the considered criteria overlap with those of deliberation. Mason (1992) proposed a list of six questions that an educational analyst could be asked in content analysis to evaluate the quality of student learning:

- do the participants build on previous messages?
- do they draw on their own experience?
- do they refer to course material?
- do they refer to relevant material outside the course?
- do they initiate new ideas for discussion?
- does the course tutor control, direct or facilitate?

Henri (1992) proposed an analytical framework with five dimensions (participative, social, interactive, cognitive and

metacognitive), of which the cognitive is most relevant to content analysis, with detailed indicators. Newman (1995) built on both works and proposes a set of indicators on relevance, importance, novelty, bringing outside knowledge, ambiguity clarification, linking ideas, justification, critical assessment, practical utility and width of understanding.

Friess and Eilders (2015) gave an overview of the research on online deliberation platforms. Esau, Friess, and Eilders (2017) evaluated the design of (asynchronous) news discussion platforms through content analysis on rationality, reciprocity (engagement), respect and constructiveness. We adapt this literature to come up with four representative quality criteria, which we describe in §4.1.

Challenges to content analysis are myriad. The manual annotation of statements requires considerable resources, and even then, reasonable and qualified annotators can disagree on quality indicators (Black et al. 2014).

LLMs have been a transformative technology (Devlin et al. 2018; Vaswani et al. 2017; Brown et al. 2020), driving significant advancements in natural language understanding and generation. Historically, LLMs have been employed in various coding and annotation tasks, consistently outperforming humans in speed and often in quality (Goel et al. 2023; Park, Wisniewski, and Singh 2024; Refuel AI 2023). While most applications focus on factual annotations, we utilize LLMs to tackle the subjective task of assessing the quality of deliberative contributions, a process made more challenging by the lack of an established ground truth.

Several researchers have explored the use of LLMs in democratic processes. Small et al. (2023) discuss integrating LLMs into their text-based deliberation platform, Polis, for tasks such as moderation, summarization, vote prediction, and consensus identification. Other studies, such as those by (Bakker et al. 2022; Fish et al. 2023; Ding and Ito 2023), provide fine-tuning approaches to generate consensus statements from text-based opinions. Particularly relevant to our work, Fish et al. (2023) use LLMs as generative and discriminative oracles to develop consensus statements and evaluate participants' utility from it. We use LLMs as oracles for the quality of statements and use them to inform the platform's design. Li, Patel, and Du (2023) demonstrated using LLMs to evaluate subjective contributions made by other LLMs.

3 Data

3.1 Deliberation Platform

The data we used for this study was collected from deliberations on the Stanford Online Platform for Deliberation. This platform is designed to support large-scale online deliberations, with multiple small-group rooms deliberating in parallel on the same agenda, all without the need for a human moderator. The platform mimics the setup of an in-person deliberative poll. Participants are assigned to rooms, each of which is part of a nested structure:

- *Room*: a small group of typically 5–12 participants who discuss with each other.
- *Roomgroup*: a set of rooms deliberating in parallel, using an identical configuration of the deliberation platform, with an identical agenda.

- *Session*: Roomgroups that discuss the same agenda, typically with nearly identical configurations, but not necessarily at the same time. This could, for example, facilitate different availability, privacy preferences or languages.
- *Deliberation event*: a series of sessions where the same group of participants discusses a set of agendas. A deliberation event typically consists of 1–5 sessions.

While the platform supports various configurations and exceptions, we will describe a typical deliberation setup. The organizer of the deliberative event recruits a representative sample of their population of interest through polling firms and shares background materials. The agenda has 3-6 agenda items, each with a balanced set of arguments in favor and against the proposals. Participants often receive these materials before the event. Participants are then invited with details on the date and time of each session, along with a link to the landing page for their assigned room group. An admin team is available to assist participants with technical issues, supervise discussions, and intervene in case of technical or behavioral problems.

The participants are in advance, at the starting time or on arrival (randomly) distributed across rooms, where the software considers a minimum size and target size for each room. When the first small-group session begins, participants in each room are greeted with a video message and a reintroduction to the conversation topic. They watch an instructional video demonstrating the platform’s buttons and functions. Participants then introduce themselves, appearing in small video screens with the speaker’s video enlarged. Speaker turns are managed by an automated moderator using a visible queue that participants can join by requesting to speak. Each participant can speak for up to 45 seconds, after which the next person in the queue takes their turn.

Each agenda item is introduced with an audio or video prompt, accompanied by a description and key arguments for and against the proposals, visible on the screen. Participants can propose to move to the next agenda item if they feel the topic has been sufficiently discussed. If a majority agrees, or if time runs out, the automated moderator introduces the next topic and presents the new arguments.

Once all agenda items have been discussed, the room enters the question development phase. Participants propose questions about the discussed topics and rank them by importance. The group then discusses the top-ranked questions, and the original proposer can edit their question. After discussing all questions or when time expires, the group votes on two questions to send to the panel of experts in the following plenary session. Participants are then redirected to the plenary session, usually after a short break, where a panel of experts answers their questions. After this, participants may be asked to join a second session, where they might be randomly reassigned to different groups.

3.2 Deliberation Events

The sets E1, E2 and E3 refer to rooms that were part of three deliberation events. E1 (“Shaping Our Future”) was an event in May 2021, conducted in collaboration with 38 post-secondary institutions with a representative sample of participants from the United States, discussing climate change

in one session and voting systems in another (Beyene et al. 2021). E2 and E3 were two sets of roomgroups, part of a larger event in 2022, conducted in collaboration with Meta (formerly Facebook), with a sample of participants from North America (E2) and the United Kingdom (E3), both discussing the same agenda on regulating bullying and harassment in virtual reality (the “Metaverse”) (Fishkin et al. 2023). The entire event was held in 32 countries in their respective languages. These events were selected because of the availability of matched demographic information. When we were limited in how much data we could analyze, especially involving human annotators, we chose E1 for most of our analysis, as it used the most representative agenda items.

In the set E4, we consider roomgroups on the platform between July 19th, 2022, and March 26th, 2024 configured with language settings English (United States) and at least five rooms with a transcript. This set includes rooms in the Meta deliberation in E2 and several other events and agendas. We use this set of rooms to study the nudge effectiveness, and only include roomgroups that were active after the introduction of a change in the software where some nudges would be skipped (as explained in §4.3).

We filtered contributions by agenda item (removing any contributions during the introductions phase and the question development phase) and removed statements of less than 100 characters in the transcript (short statements are often less reliable and may include more clarifications and technical queries). We refer to these as filtered contributions. Table 1 describes some statistics of these roomsets.

4 Methodology

We design a prompting scheme for using LLMs to evaluate the quality of statements made in online video-conferencing-based deliberations. We then conduct experiments with human annotators to validate the usefulness of the LLM evaluations. Using these quality evaluations, we study the effect of nudges on our deliberation platform on the quality of the statements made. We use the current state-of-the-art GPT-4 model from OpenAI (Achiam et al. 2023), which we will refer to as “the model”.

4.1 Estimating the Quality of a Statement

There are many possible definitions for the “quality” of a contribution in the context of a deliberation, as discussed through the various examples in §2. In our study, we build on the goal of a deliberative poll, which is to provide participants with optimal conditions to form an informed opinion on the topic at hand (Fishkin 2009).

We based three of our criteria on the questions and indicators proposed by Mason and Newman (Mason 1992; Newman 1995): whether the contribution *provides justification*, *is novel*, *expands the conversation*. We added a fourth criterion whether the contribution *allows further expansion* because we’re interested in how contributions in a deliberative conversation fit together. Specifically, we ask the annotators to score the statement on following criteria prompts:

- Q1 This statement includes examples or anecdotes to support the speaker’s point.

ID	Topics	Rooms	Unique Participants	Filtered Contributions	Sessions	Median Room Size	Mean Room Size	Average Length (Characters)
E1	Climate change; voting systems	116	583	6232	2	9	8.66	513.64
E2	Regulating bullying and harassment in virtual reality (Meta)	221	698	7015	4	8	8.34	386.56
E3	Regulating bullying and harassment in virtual reality (Meta)	47	149	1824	4	9	8.57	376.24
E4	(various, including E2)	1368	4654	54990	-	8	7.68	425.31

Table 1: Sets of rooms used for analysis with general topic and statistics. A participant typically participates in one room per session within the same deliberation event.

- Q2 This statement introduces novel ideas, perspectives, or solutions.
- Q3 This statement builds on top of the previous statements and the proposal.
- Q4 This statement raises points which will likely improve the quality of the following discussion.

In the rest of the paper, we will refer to these prompts as Q1, Q2, Q3, and Q4 for brevity. Each time, we provided the annotator with the contribution, the topic of the current agenda item, and all preceding statements within the same agenda item. We then asked the annotator to rate their agreement with Q1–4 on a 5-point Likert scale (Strongly Disagree – Strongly Agree) and provide a justification for each criterion. The exact query given to the model is as follows:

Query for Estimating the Quality of a Statement

System Instructions: Your expertise is required to help a research study to evaluate the quality of statements made in a deliberation on a topic of social interest.

Task: You are observing a live deliberation on the following proposals *<topic>*.

The input is from a noisy speech-to-text system. Your task is to evaluate a statement in the context of the ongoing discussion. Specifically, you have to rate it on the standard Likert scale on 1 to 5 on whether: *<critierium statement>*. This is what the rating on the standard Likert scale means. 1: Strongly Disagree. 2: Disagree. 3: Undecided. 4: Agree. 5: Strongly Agree.

Please give a succinct justification in one short sentence. Format your answer as follows:

Rating: *x/5*. Justification: [One short sentence].

Here is the transcript of the deliberation from before the statement which is to be evaluated. This serves as the context for the evaluation. *<previous contributions>*.

Here is the statement which you have to evaluate. *<contribution>*.

The system instructions aim to improve the quality of responses by *assigning a role* (Wu et al. 2023). The query was refined by experimenting with the model’s evaluations of ten randomly chosen contributions. This process helped us develop a query that enables the model to consistently distinguish between good and bad statements and provide meaningful justifications for its ratings.

4.2 Evaluating Automated Quality Estimations

In order to evaluate the quality estimates created by our quality model, we need to compare these estimates with a baseline. If we were presented with the challenge to estimate the quality of a contribution without access to language models, we would ask one or some research assistants (often familiar with the deliberation procedures to score them: interns, undergraduate and graduate students) to score the contributions, and if necessary take the average of those scores. For an evaluation, we want to compare the score of a human estimate with the model estimate.

For this study, we recruited eight research assistants associated with our lab, but not directly involved with this project. We identified a deliberation in English with participants based in the United States where the gender information of participants was readily available (E1 in Table 1). These rooms had 6232 filtered contributions, discussing a specific agenda item with a transcript of at least 100 characters. We randomly selected 15 of these statements from male and female participants each.

We asked the human annotators to perform two tasks: first, to estimate the quality of a contribution based on each of the four criteria and justify their score in a sentence; second, to evaluate a quality estimation, which includes a rating and its justification. We divided the human evaluators randomly into two equal cohorts, and split the statements by gender. In the first round, the first cohort estimated the quality of the first set of 15 statements and scored a rating-justification pair for each statement generated by the model but unknown to them. Meanwhile, the second cohort performed the same tasks on the second set of statements. In the second round, the first cohort estimated the quality of the second set of statements. After each contribution, they scored five different rating-justification pairs (one from the model and four from the other cohort of evaluators, in random order). The second cohort did the same for the first set of statements. For each evaluation of a rating-justification pair, the human evaluators did not know its source.

4.3 Nudging to Encourage Participation

The platform implements nudges as pop-up overlays on the user’s screen. The pop-up includes a phrase encouraging the user to contribute to the conversation or asking whether the group is ready to move on to the next agenda item. In this paper, we will focus on the nudges that encourage users to

participate (a “Speak” nudge) when they have not participated for a period of three minutes. The textual content of the nudge is randomly drawn from a list of possible nudges. This list varies between events and can be broadly categorized in a *general nudge* (e.g. “What would someone who disagrees with you say about this discussion”), a *personalized nudge* that addresses the user by their screen name (e.g. “XYZ, you haven’t spoken in a while, is there something you would like to share with the group?”) or a *pro-con nudge* that explicitly refers to the listed argument in the agenda item that is estimated to have been discussed least. Since 2022, the platform has also included dummy nudges for research purposes: the platform would randomly “skip” a speak nudge altogether (typically 15%). The platform also emits Speak-Room nudges to all participants when no one has spoken for over 30 seconds. We do not study the effect of Speak-Room nudges in this work.

5 Evaluating Automated Quality Estimation

We will first establish the level of consensus between the human annotators (§ 4.2), and then use this as a benchmark for the model’s performance.

5.1 Inter-Rater Agreement of Human Annotators

Inter-rater reliability (IRR) or inter-rater agreement measures the consensus among different annotators of a dataset. A common statistic for studying IRR is within-group reliability (r_{wg}) (James, Demaree, and Wolf 1993). While there are multiple variants of r_{wg} (O’Neill 2017), we report r_{wg}^* as it is best suited for our setting where multiple annotators report Likert scale scores for many statements.

$$r_{wg}^* = 1 - \frac{S_x^2}{\sigma_{eu}^2}. \quad (1)$$

Higher r_{wg}^* indicates better IRR. Here S_x^2 is the mean of the observed variance in annotators’ ratings on each item, and σ_{eu}^2 is the variance of the rectangular, uniform null distribution, $(m^2 - 1)/12$, where m is the number of discrete Likert-type response options. Here $m = 5$.

The r_{wg}^* values for the human annotators are as follows: Q1: 0.519; Q2: 0.449; Q3: 0.605, and Q4: 0.496. Although interpreting these values can vary by field, we follow the standards recommended by LeBreton and Senter (2008): 0–0.30 (lack of agreement), 0.31–0.50 (weak agreement), 0.51–0.70 (moderate agreement), 0.71–0.90 (strong agreement), and 0.91–1.0 (very strong agreement). The human annotators show weak agreement on novelty (Q2) and their potential for further expansion (Q4). They have moderate agreement on whether contributions justify their claims (Q1) and expand the conversation (Q3).

5.2 Ratings from Humans and the Model

We report the humans’ and model’s average scores across the 30 statements in Table 2. We obtained a 95% confidence interval for the difference of the means by taking 10,000 bootstrap resamples of the original data and calculating the mean difference for each resample. We observe that the model consistently gives lower scores than humans on

Q	Model Mean	Humans Mean	Mean Difference	Mean Difference 95% CI
Q1	2.50	3.14	-0.64	[-0.98, -0.34]
Q2	2.93	3.35	-0.42	[-0.67, -0.15]
Q3	4.03	3.79	0.24	[0.05, 0.42]
Q4	4.10	3.79	0.31	[0.04, 0.54]

Table 2: Average scores given by the model and humans with 95% confidence intervals for the difference.

Q	Difference (Men)	Difference (Women)
Q1	-0.77	-0.51
Q2	-0.62	-0.21
Q3	0.14	0.34
Q4	0.44	0.18

Table 3: Difference in average scores given by the model and humans for both genders. A negative value means that the model gives lower scores than humans.

Q1 and Q2 and slightly higher scores on Questions 3 and 4. Addressing this bias – whether through prompt adjustments, fine-tuning, or mean-correcting the model’s scores – is an important direction for future research. In this study, we use the raw scores. We suggest a mean-correction approach in the full version of the paper, which shows that the model’s performance can be enhanced via this simple correction.

The bias in the model’s scores is more or less consistent across the genders of participants, as in Table 3. However, since our data is over only 15 statements for men and women each, further investigation is required to ensure that the model is not biased.

A common approach to improve noisy ratings is to take the average of multiple individual ratings. We will compare the model’s rating against ratings of individual annotators, as well as pair and triplet ratings (defined by the average of its members’ ratings). The average rating of the remaining human annotators serves as the “golden” rating. Both the model and the human groups are evaluated on how closely their ratings match this golden rating¹.

We determine for each possible individual, pair and triplet whether the model or humans are closer to the golden score for each of the 30 contributions (there are 8-choose-g groups, where g is the group size). We report in Fig. 1 the fraction of data points where the model outperforms the groups of human annotators. Groups of three humans always outperform the model for each question, while single humans never do. In groups of two, humans perform better on Q1, while the model outperforms humans on all other questions. We also aggregate these scores over the 30 questions, and report in Fig. 2 for how many groups the model rates more statements closer to the golden rating than the humans. The same observations apply as in Fig. 1, but the margins

¹We do not consider the golden rating to be the average of *all* human annotators since that will give an unfair advantage to the group of humans since their ratings will be included in the average.

of victory are larger for both humans and the model. The takeaway from this experiment is that the model is a reasonable alternative to human annotators because three annotators would be too expensive for most deliberation events.




Figure 1: Fraction of data points where groups of humans are outperformed by the model. The model ‘wins’ if it rates a statement closer to the golden rating than a group of humans.




Figure 2: Fraction of groups of humans outperformed by the model. The model ‘wins’ if it rates more statements closer to the golden rating than a group of humans.

A second approach to evaluate the performance of the model against human annotators, is to ask other annotators to evaluate the scores and their justifications (See § 4.2). We report the scores of both human and model ratings and justifications in Fig. 3 and the averages in Fig. 4. The evaluators are not informed about the source of the ratings and justification. The model outperforms the humans by a significant margin, as is evident from the 95% confidence intervals (obtained via bootstrapping by drawing 10,000 samples).




Figure 3: Average scores received in the evaluation of rating-justification pairs by individual annotators on the 1-5 Likert scale. S1 and S2 are the contribution sets. The model’s results are marked by stars. Each annotator has a unique color used across all criteria.




Figure 4: Average scores received in evaluations of rating-justification pairs on the 1-5 Likert scale.

5.3 Efficiency

The cost of annotating a 45-minute deliberation event with GPT-4 is under \$4. As of writing, OpenAI charges \$30 per million tokens input and \$60 per million tokens output. A typical 45-minute deliberation event contains 30-50 contributions, and we make four queries for each contribution, where a typical query has between 200-1000 tokens of input and 20 tokens of output. An additional advantage of automated quality estimation is that it can be implemented in real-time (under 2s delays for a query), allowing for platform interventions to be informed by it.

On the other hand, even if we assume a human annotator can consistently estimate quality in 20 seconds per criterion per contribution, at the California minimum wage of \$16 per hour, human annotations would be four times more expensive than this automated approach. Using three human annotators would triple the cost further. This doesn’t even take into account the cost of training the human annotators.

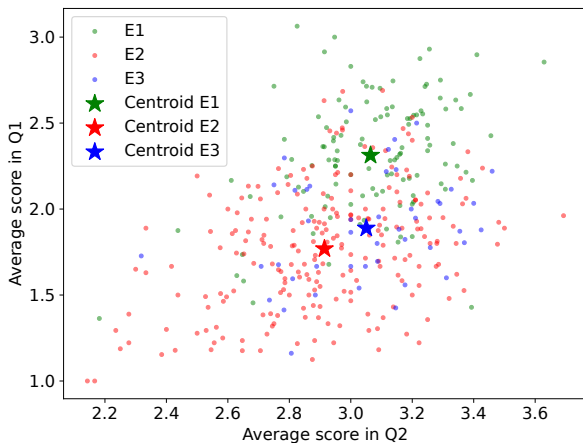


Figure 5: Average room quality on Q1 (justification) and Q2 (novelty) for all rooms in events E1, E2, and E3. The centroid corresponds to the average of the room quality for each room in the event.

5.4 Analysis of Deliberation Events

Since we have established that the model is a reliable annotator, we can use it at scale to inform platform design choices and estimate the average quality of all rooms in deliberation events. For example, in Fig. 5, we plot the average quality of statements in rooms (average over all filtered contributions) for Q1 (justification) and Q2 (novelty). E1 has a higher overall quality on both questions than both E2 and E3 (both using the same agenda design). The figure also shows a large variation between rooms in the same event. This analysis can prove useful when evaluating the event to identify rooms that had a particularly meaningful conversation (for qualitative analysis) or when evaluating the event as a whole.

6 Evaluating Nudges Via Quality Estimates

We can use the quality estimates to evaluate the effect of interventions on the quality of conversation. We will apply this to our speak nudges, which are emitted to a participant after three minutes of not requesting to speak (§4.3). We will first establish the effectiveness of the nudges towards increasing participation. We use dataset E4 here (recall §3). We have data from 27,899 nudges *sent* and 5,222 nudges *skipped*.

6.1 Effectiveness at Prompting to Speak

With the skipped nudges, we have a randomized controlled trial. We can calculate the probability that the participant requests to speak in the 30-second interval following the emission of a nudge (or a skipped nudge) and establish the effect of a nudge by taking the difference in means. Out of the 27,899 nudges, 2,929 (10.4%) resulted in a positive participation outcome, (a request to speak). Out of these, 2625 resulted in an actual contribution. The 5,222 instances where nudges were skipped were followed by 332 requests to speak (6.4%) in 30 seconds, which then led to 301 statements.

This indicates that speak nudges result in a 65.1% increase in the probability of a request to speak and a 63.2%

increase in the probability of a statement. In Fig. 6 we report these probabilities broken down in 5-second intervals, confirming the intuition that the effect of a nudge is most prominent a little after emission.

We report the effectiveness of repeated nudges in the full version of the paper. The first nudge is the most effective (around 15%), with effectiveness dropping quickly. By the fourth nudge, the effectiveness is comparable to a skipped nudge. This shows that sending many nudges is unproductive and that fewer, strategically timed nudges is better.

6.2 Effect of Nudges on Contribution Quality

Now that we have established that nudges increase participation, we are interested in whether those statements are higher or lower quality. We study the effect of nudging on contribution quality in three ways:

Analysis 1. Compare the average quality of contributions following emitted and skipped nudges in dataset E4.

Analysis 2. Compare the average quality of contributions after nudges to all other contributions in E1, E2, and E3.

Analysis 3. Compare the average quality of contributions after nudges to all other contributions for *individual* participants in datasets E1, E2, and E3.

We say a contribution follows a nudge (or a skipped nudge) if the speaking request comes in the following 30 seconds.

Analysis 1. The average quality score for contributions following a nudge is not less than that for contributions following skipped nudges. These are higher by 0.05, 0.07, 0.03, and 0.03 on Q1, Q2, Q3, and Q4, respectively. See Fig. 7 for 95% confidence intervals of the scores. Recall that nudges are sent (or randomly skipped) to participants with lower activity. This result is an important point in favour of targeted nudges. While nudges increase engagement, they do so without decreasing the quality of contributions, thereby generating more ideas in the conversation.

Analysis 2. When compared with all other contributions, we observe that the average quality of nudged contributions is lower on justification (Q1) and novelty (Q2) and higher on constructiveness (Q3), and allowing expansion (Q4) (Fig. 8). This difference in average quality is -0.09 , -0.07 , 0.03 , and 0.02 for Q1, Q2, Q3, and Q4, respectively. This is small compared to the standard deviation of the scores, which are 1.36, 0.96, 0.79, and 0.70 in the same order. This result does not imply causation since nudges are more likely to be sent to less active participants. We give further breakdowns of these results across events and genders in the full version.

Analysis 3. For each participant, we calculated the difference in the average quality of their nudged and other contributions. We then averaged this difference across all participants with at least one nudged and one non-nudged contribution. This is 607 participants over E1, E2, and E3. The results are shown in (Table 4). The small negative effect seen in Analysis 2 is now reduced to near zero!

The intuition that the rating difference between nudged and other contributions is driven by the fact that nudges are sent to less-active individuals is further corroborated by the




Figure 6: Probability of a nudge leading to a speaking request within 30 seconds in dataset E4. Red bars show the control with skipped nudges, and blue bars show the case with nudges. Error bars represent 95% confidence intervals for each 5-second interval.

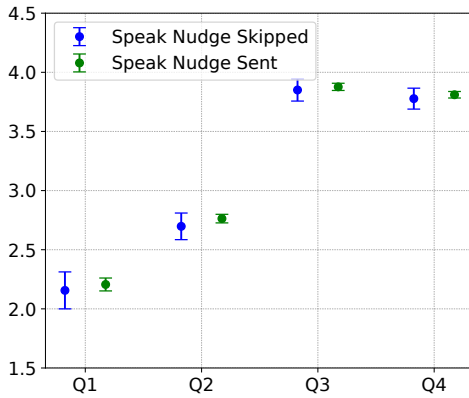


Figure 7: Average quality of statements for which the associated speaking request came after a skipped nudge or a sent nudge with 95% confidence intervals in dataset E4.

following observation. The Pearson Correlation Coefficient between the number of contributions made by a participant and their average quality scores are: Q1: 0.21, Q2: 0.24, Q3: 0.13, and Q4: 0.17, all nontrivial positive numbers, indicating a weak correlation between the average activity level and the quality of contributions. Analysis 3 further reinforces the takeaway of Analysis 1 that nudged statements have the same average quality as non-nudged statements.

7 Conclusions and Future Work

In this work, we demonstrate that LLMs are valuable tools for rating deliberative contributions, providing consistent and scalable quality assessments that significantly reduce the resource demands associated with human annotations. This approach aids in evaluating platform interventions, such as targeted nudges, and overall deliberation quality.

One significant limitation is the unpredictability of LLMs (Perković, Drobnjak, and Botički 2024). In our tests on 30




Figure 8: Average quality of nudged and other statements in events E1, E2, and E3, with 95% confidence intervals.

Q	Nudge Effect	Nudge Effect 95% CI
Q1	-0.012	[-0.112, 0.089]
Q2	0.022	[-0.046, 0.090]
Q3	-0.008	[-0.062, 0.043]
Q4	0.021	[-0.024, 0.065]

Table 4: Nudge effect on individual-level quality, averaged over all participants in events E1, E2, and E3. Negative numbers indicate that the quality is lower for nudged statements.

randomly selected statements across 5 trials, we observed low average variances of 0.13, 0.08, 0.02, and 0.01 in the model’s ratings on Q1-Q4, respectively, indicating good consistency in the model’s ratings. Ensuring this consistency across different contexts remains crucial. Additionally, using AI in this context raises issues of transparency and accountability. LLM outputs are highly sensitive to prompt design (Sclar et al. 2023; Errica et al. 2024), and the rapid development of LLMs presents the risk that current queries may not work as expected with future models. Analysts must remain vigilant about potential biases and may need to adjust queries over time. Evaluating AI-generated justifications can aid explainability, but ongoing human quality control may still be necessary. While developing a self-hosted LLM could mitigate some risks, it would also introduce additional costs and complexity.

Further research includes improving the model through fine-tuning and mean correction to enhance the accuracy of quality assessments. Another key direction is studying the relationship between discussion quality and opinion change, measured through pre- and post-event surveys, to understand how discourse quality influences participant perspectives. Designing platform interventions based on real-time quality assessments is an exciting avenue for future work, potentially allowing dynamic adjustments to enhance the effectiveness of deliberative processes. Integrating automated assessments with other evaluation metrics could provide a more comprehensive understanding of deliberative quality.

Acknowledgments

We thank James Fishkin, Lovish Chopra, and Shoaib Mohammed for their valuable discussions. We are grateful to Tenzin Kartsang for coordinating the human annotators, as well as to the annotators themselves. We also extend our thanks to our deliberation partners, particularly those at Meta, for their assistance in data collection. Additionally, we acknowledge the use of ChatGPT from OpenAI for enhancing the clarity of parts of this paper.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bakker, M.; Chadwick, M.; Sheahan, H.; Tessler, M.; Campbell-Gillingham, L.; Balaguer, J.; McAleese, N.; Glaese, A.; Aslanides, J.; Botvinick, M.; et al. 2022. Fine-tuning language models to find agreement among humans with diverse preferences. *Advances in Neural Information Processing Systems*, 35: 38176–38189.
- Beyene, A.; Domingo, N.; Duncan, I.; Graciano, T.; Hein, T.; Hwang, E.; Ladabaum, I.; Martinez, S.; Muir, L.; Ramaswami, A.; Soh, S.; Xiong, J.; Wohlford Kidd, K.; Wright Garnett, L.; Siu, A.; and Terra, L. 2021. Shaping Our Future - Final Report.
- Black, L. W.; Burkhalter, S.; Gastil, J.; and Stromer-Galley, J. 2014. Methods for analyzing and measuring group deliberation. In *Sourcebook for political communication research*, 323–345. Routledge.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Ding, S.; and Ito, T. 2023. Self-agreement: a framework for fine-tuning language models to find agreement among diverse opinions. In *Pacific Rim International Conference on Artificial Intelligence*, 298–309. Springer.
- Dutwin, D. 2003. The Character of Deliberation: Equality, Argument, and the Formation of Public Opinion. *International Journal of Public Opinion Research*, 15(3): 239–264.
- Errica, F.; Siracusano, G.; Sanvito, D.; and Bifulco, R. 2024. What Did I Do Wrong? Quantifying LLMs’ Sensitivity and Consistency to Prompt Engineering. *arXiv preprint arXiv:2406.12334*.
- Esau, K.; Friess, D.; and Eilders, C. 2017. Design Matters! An Empirical Analysis of Online Deliberation on Different News Platforms. *Policy & Internet*, 9(3): 321–342.
- Fish, S.; Gözl, P.; Parkes, D. C.; Procaccia, A. D.; Rusak, G.; Shapira, I.; and Wüthrich, M. 2023. Generative social choice. *arXiv preprint arXiv:2309.01291*.
- Fishkin, J. 2009. *When the people speak: Deliberative democracy and public consultation*. Oup Oxford.
- Fishkin, J.; Siu, A.; Diamond, L.; and Bradburn, N. 2021. Is Deliberation an Antidote to Extreme Partisan Polarization? Reflections on “America in One Room”. *American Political Science Review*, 115(4): 1464–1481.
- Fishkin, J. S. 1997. *The voice of the people: Public opinion and democracy*. Yale university press.
- Fishkin, J. S.; Siu, A.; Chang, S.; Ciesla, E.; and Kartsang, T. 2023. public_meta_community_forum_final_report._stanford_ddl.1(2).pdf. Result Analysis, Stanford Deliberative Democracy Lab, Stanford CA.
- Friess, D.; and Eilders, C. 2015. A systematic review of online deliberation research. *Policy & Internet*, 7(3): 319–339.
- Gelauff, L.; Nikolenko, L.; Sakshuwong, S.; Fishkin, J.; Goel, A.; Munagala, K.; and Siu, A. 2023. Achieving parity with human moderators: A self-moderating platform for online deliberation 1. In *The Routledge Handbook of Collective Intelligence for Democracy and Governance*, 202–221. Routledge.
- Gerhards, J. 1997. Diskursive versus liberale Öffentlichkeit: Eine empirische Auseinandersetzung mit Jürgen Habermas. *Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 49(1): 1–34.
- Goel, A.; Gueta, A.; Gilon, O.; Liu, C.; Erell, S.; Nguyen, L. H.; Hao, X.; Jaber, B.; Reddy, S.; Kartha, R.; et al. 2023. Llm accelerate annotation for medical information extraction. In *Machine Learning for Health (ML4H)*, 82–100. PMLR.
- Gold, V.; El-Assady, M.; Bögel, T.; Rohrdantz, C.; Butt, M.; Holzinger, K.; and Keim, D. 2015. Visual Linguistic Analysis of Political Discussions: Measuring Deliberative Quality. *Digital Scholarship in the Humanities*, fqv033.
- Goodin, R. E.; and Dryzek, J. S. 2006. Deliberative Impacts: The Macro-Political Uptake of Mini-Publics. *Politics & Society*, 34(2): 219–244.
- Henri, F. 1992. Computer Conferencing and Content Analysis. In *Collaborative learning through computer conferencing: The Najaden papers*, 117–136. ISBN 978-3-642-77684-7.
- James, L. R.; Demaree, R. G.; and Wolf, G. 1993. rwg: An assessment of within-group interrater agreement. *Journal of applied psychology*, 78(2): 306.
- LeBreton, J. M.; and Senter, J. L. 2008. Answers to 20 questions about interrater reliability and interrater agreement. *Organizational research methods*, 11(4): 815–852.
- Li, R.; Patel, T.; and Du, X. 2023. Prd: Peer rank and discussion improve large language model based evaluations. *arXiv preprint arXiv:2307.02762*.
- Luskin, R. C.; Iyengar, S.; and Fishkin, J. S. 2014. Online versus Face-to-Face Deliberation: Evidence from Parallel Deliberative Polls.
- Mason, R. 1992. Evaluation Methodologies for Computer Conferencing Applications. In Kaye, A. R., ed., *Collaborative Learning Through Computer Conferencing*, 105–116.

Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-77686-1 978-3-642-77684-7.

Newman, D. R. 1995. A content analysis method to measure critical thinking in face-to-face and computer supported group learning. *Interpersonal Computing and Technology Journal*, 3(2): 56–77.

O’Neill, T. A. 2017. An overview of interrater agreement on Likert scales for researchers and practitioners. *Frontiers in psychology*, 8: 264983.

Park, J.; Wisniewski, P.; and Singh, V. 2024. Leveraging Large Language Models (LLMs) to Support Collaborative Human-AI Online Risk Data Annotation. *arXiv preprint arXiv:2404.07926*.

Perković, G.; Drobnjak, A.; and Botički, I. 2024. Hallucinations in LLMs: Understanding and Addressing Challenges. In *2024 47th MIPRO ICT and Electronics Convention (MIPRO)*, 2084–2088. IEEE.

Refuel AI. 2023. LLM Labeling Technical Report. <https://www.refuel.ai/blog-posts/llm-labeling-technical-report>. Accessed: May 9, 2024.

Sclar, M.; Choi, Y.; Tsvetkov, Y.; and Suhr, A. 2023. Quantifying Language Models’ Sensitivity to Spurious Features in Prompt Design or: How I learned to start worrying about prompt formatting. *arXiv preprint arXiv:2310.11324*.

Small, C. T.; Vendrov, I.; Durmus, E.; Homaei, H.; Barry, E.; Cornebise, J.; Suzman, T.; Ganguli, D.; and Megill, C. 2023. Opportunities and risks of LLMs for scalable deliberation with Polis. *arXiv preprint arXiv:2306.11932*.

Smith, G.; and Wales, C. 2000. Citizens’ Juries and Deliberative Democracy. *Political Studies*, 48(1): 51–65.

Steenbergen, M. R.; Bächtiger, A.; Spörndli, M.; and Steiner, J. 2003. Measuring Political Deliberation: A Discourse Quality Index. *Comparative European Politics*, 1(1): 21–48.

Van Der Does, R.; and Jacquet, V. 2023. Small-Scale Deliberation and Mass Democracy: A Systematic Review of the Spillover Effects of Deliberative Minipublics. *Political Studies*, 71(1): 218–237.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wu, N.; Gong, M.; Shou, L.; Liang, S.; and Jiang, D. 2023. Large language models are diverse role-players for summarization evaluation. In *CCF International Conference on Natural Language Processing and Chinese Computing*, 695–707. Springer.