WHEN SHOULD AI CHIME IN? BENCHMARKING TIMELY CONTRIBUTION OF PROACTIVE AGENTS IN MULTI-PARTY CHAT

Anonymous authors
Paper under double-blind review

ABSTRACT

Proactive agents can support group collaboration by surfacing useful information without waiting to be asked, but even relevant contributions can be disruptive if poorly timed. In this work, we isolate the timing question: given a fixed candidate contribution such as a clarification or relevant fact, when should an AI agent chime in? We introduce PITCH (**P**roactive **I**ntervention **T**iming in **CH**at), to our knowledge the first benchmark with large-scale human judgments of proactive agent timing collected from participants in the same unscripted multi-party conversations they took part in. To collect high-quality judgments, we design a novel interface in which 2–4 people chat synchronously, author candidate AI contribution messages, place their own and others' contributions into the chat history, and then deliberate as a group about the best placement for each message. PITCH contains 384 candidate contributions and 1,091 judgments, each marking an ideal turn, an acceptable window, or that a contribution should not appear, from 89 participants engaged in 35 conversations. Evaluating 15 models across 28 configurations, we find that while models can place contributions well when given the entire conversation history, they **systematically chime in too early when asked to decide live as the conversation unfolds**: best-turn accuracy falls from 57.8% to 19.0%, with nearly half of placements occurring before every participant's acceptable window. In comparison, lay humans deciding live are more accurate (23.7% vs 19.0%), and fewer of their placements are premature (39.7% vs 46.4%). Model reasoning traces show that 93% of intervention justifications cite relevance, while almost none consider waiting. By providing a human-grounded benchmark, our work reveals gaps in agents' ability to integrate into group settings and points towards proactive agents that can reason explicitly about appropriate timing.

1 INTRODUCTION

Achieving proactivity in AI agents requires both having useful information to contribute and knowing when to chime in with it. Agents are increasingly capable of providing contextualized factual information such as clarifications, examples, or relevant facts that can benefit users Wang et al. (2025), but less attention has focused on timing. Getting the timing right has particularly high stakes in group settings where an agent interjection can negatively disrupt the flow of conversation (Yeo et al., 2026). For instance, in a group chat about travel planning, an agent might have useful information about hotel pricing—but bringing it up after the group has settled on a hotel would be too late, while introducing it before they begin discussing accommodations might distract from the current topic. Appropriate timing therefore involves weighing the benefit of contributing now against the possibility that waiting would produce a better moment (Horvitz, 1999).

Evaluating how well agents perform on timing requires knowing when the participants in a conversation would have actually welcomed a particular contribution (e.g., a message containing useful clarifications). However, most prior efforts to develop proactive agents do not evaluate against this information; instead, they resort to methods that lack ecological or construct validity, including testing on synthetic conversations (Chen et al., 2025), using LLM-as-a-judge to infer timing preference (Liu et al., 2025b), and asking non-participants to validate by reviewing transcripts (Liu et al., 2025a). Similarly, datasets developed explicitly to benchmark agent timing performance also rely

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

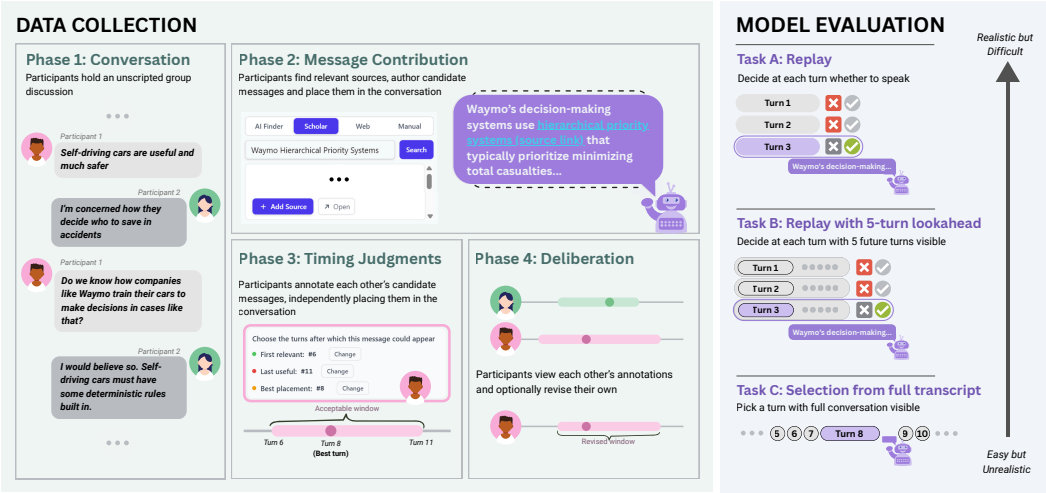


Figure 1: PITCH (**P**roactive **I**ntervention **T**iming in **CH**at) is the first benchmark for proactive intervention timing built from unscripted human conversations and participants' own judgments of when proactive interjections should occur. We collected the benchmark data via a four-phrase protocol: Conversation, Message Contribution, Timing Judgments, and Deliberation. We then evaluated 15 models across 28 configurations against the benchmark in three task setups that differ in their task difficulty and realism for real-world deployment of agents in group conversations.

on synthetic or semi-synthetic conversations with designated intervention points (Patel et al., 2025; Nama et al., 2026). Yet work has shown that synthetic group conversations do not resemble ones between humans (Kamoi et al., 2026). These evaluations also assume that there is only one correct placement for agent interjection and that all participants would agree on the timing. In practice, acceptable timing often spans a window and varies across participants, as our study shows. Finally, participant preferences cannot be inferred from conversation traces alone: a transcript records what participants said, but not when they would have welcomed information that was never introduced.

To address this gap, we introduce PITCH (**P**roactive **I**ntervention **T**iming in **CH**at), which, to our knowledge, is the first benchmark for proactive intervention timing built from **unscripted human conversations and participants' own judgments of when proactive interjections should occur**. PITCH is enabled by a study protocol and interactive interface designed to jointly capture a group's chat conversation, agent contribution messages, and participants' timing preferences. We recruit researchers and students and group each according to shared interests; they then have an unscripted synchronous chat regarding a research topic of their choosing or make shared plans for a hypothetical trip or event, respectively. Next, with the support of search and drafting tools, participants author contributions containing information that could have benefited their discussion. For each contribution, the interface presents the message alongside the transcript and asks each group member to identify an ideal turn and an acceptable window, or indicate that the contribution should not appear at all. Participants make preferences independently before reviewing one another's placements and optionally revising their own, allowing the protocol to capture both shared preferences and disagreements that persist after deliberation.

The resulting benchmark contains 384 human-authored candidate messages from 35 conversations comprising 2,194 turns, with 1,091 timing preferences from 89 participants. For 54% of messages, all ideal turns from participants in the group lie within two turns of one another, while 77% have an acceptable window shared by all participants. Of the preferences revisited during peer deliberation, 85% retain their original ideal turn. We therefore evaluate placements against both individual timing preferences and shared acceptable regions. These preferences provide a participant-grounded reference for the timing an online agent should anticipate, capturing both variability in individual preferences and regions of agreement across participants.

We evaluate 15 proprietary and open-weight models in 28 standard and reasoning configurations under three task setups: *(A) turn-by-turn replay*, *(B) replay with five-turn lookahead*, and *(C) selec-*

tion from the full transcript. When models can inspect the completed conversation before selecting a placement (the *(C) selection from the full transcript* setup), the selected placements almost always come within striking distance of participants' preferences, with 92.7% falling within at least one participant's acceptable window, and 57.8% directly matching at least one participant's ideal turn. However, when models must commit without knowing what follows (the *(A) turn-by-turn replay* setup), the prediction accuracy for ideal turns falls from 57.8% to 19.0%, and 46.4% of predictions occur before any annotator's acceptable window opens, compared to 1.0% after every window closes. We find that while non-participant lay annotators also find the online task challenging, they still perform better than models (23.7% vs 19.0%) and are less likely to make premature placements (39.7% vs 46.4%). When we allows models to see into the near future (the *(B) replay with five-turn lookahead*), accuracy only rises to 22.4%. Inspecting firing justifications made in model reasoning traces, we find that 92.8% argue solely that the contribution is relevant to the current turn, while only 0.14% explicitly weigh waiting for a later moment. Collectively, these results show that recognizing relevance is insufficient for well-timed proactive assistance, and that evaluating models with full conversational hindsight can substantially overstate their ability to intervene appropriately in live collaboration.

We contribute (i) **PITCH, a participant-grounded benchmark of proactive intervention timings** containing human-authored contributions and large-scale timing preferences from participants in unscripted multi-party conversations (Section 2.4); (ii) **an interactive collection protocol** that elicits ideal placements, acceptable alternatives, and persistent differences in participant preferences for fixed contributions (Section 2); (iii) **a formulation and set of evaluation measures for participant-aligned timing** that preserve both individual preferences and shared acceptable regions (Sections 2.2 and 3); and (iv) **empirical findings on model timing alignment**, showing a substantial gap between post hoc and online intervention decisions and models' tendency to treat relevance as sufficient reason to speak even when the timing is not appropriate (Sections 4 and 5).

2 EVALUATING PROACTIVE AGENT TIMING WITH HUMAN CONVERSATIONS AND PARTICIPANT-GROUNDED TIMING JUDGMENTS

PITCH evaluates whether models contribute messages containing helpful factual knowledge at times preferred by the **actual conversation participants**. Because these preferences are participant-specific and cannot be reliably inferred from transcripts or simulations, we designed an interactive protocol that jointly collects real conversations and participants' retrospective timing preferences.

2.1 FOUR PHASES OF THE INTERACTIVE DATA COLLECTION PROTOCOL

Our data collection protocol captures two key pieces of information unavailable from conversation transcripts alone: *what information an AI could have proactively contributed* and *when participants would have wanted to receive it*. Participants first take part in a real group discussion, then author candidate AI contribution messages containing factual information that could have proactively supported the group, such as clarifying a concept, providing evidence, or offering relevant logistical details. Using real discussions also preserves conversational dynamics that synthetic interactions may smooth away (Zhou et al., 2024; Bian et al., 2025; Kamoi et al., 2026).

Concretely, the data collection protocol consists of four phases (shown in Figure 1):

- **Phase 1: Conversation.** Groups hold real-time, unscripted typed conversations in two settings: *open research discussion*, where research literature and scientific evidence support *expert* information needs, and *collaborative planning* of trips or events, where prices, logistics, and options support *everyday* information needs. See example topics in Appendix A.
- **Phase 2: Message contribution.** Participants review the conversation and author factual messages that could have helped the group, with optional search and AI support.
- **Phase 3: Timing judgments.** Group members independently judge others' contributions, selecting an ideal placement, an acceptable timing window, or that the message should not appear.
- **Phase 4: Deliberation.** Participants view peers' judgments and may revise their own, letting us measure whether timing preferences persist after reconsideration without requiring consensus.

This design offers two benefits. First, participant-authored contributions reflect information needs that arose in the actual discussion, making them more likely to be genuinely useful. Second, it

avoids restricting the benchmark to information that current models can already identify, retrieve, synthesize, and contextualize from an ongoing group conversation.

2.2 Task Definition: Proactive Agent Contribution Timing

We define a **contribution** as a message that introduces factual information intended to help the conversation, such as by clarifying a concept, providing evidence or an example, or offering relevant logistical information. In a deployed proactive-agent setting, this message would be generated by the AI agent. Given a fixed contribution, PITCH evaluates *when, if ever, it should be introduced*. Fixing the message isolates timing from upstream challenges such as identifying information needs, retrieving relevant information, and generating the contribution itself.

Formally, a **conversation** $C = (u_1, \ldots, u_T)$ is a sequence of turns from an unscripted multi-party human discussion, and m is a participant-authored candidate contribution. The model chooses when, if ever, to insert m into this fixed trace.

For each contribution, participant i provides a **timing judgment** consisting of an *ideal turn* t_i, the best turn after which to insert the message, and an *acceptable window* $[s_i, e_i]$, spanning the earliest point at which the message becomes relevant to the latest point at which it remains useful. Participants may instead select *do-not-surface* if they believe the message should not appear anywhere in the conversation. Candidate contributions were reviewed for coherence and factual plausibility.

2.3 Participant Recruitment and Study Procedures

Participant recruitment. We recruited 89 participants through social media, personal networks, and online and physical advertisements at several large research universities. Participants received a $35 gift card for a 90-minute session. They reported topic interests in *open research discussion* and *collaborative planning* in advance and were grouped by shared preferences.

Study procedure. Participants first held a 15–25 minute typed conversation in groups of 2–4 using our study platform (Appendix C). Afterward, they reviewed the transcript and authored factual messages they believed could have helped the group, with optional search and AI drafting support. Participants were instructed to write messages that an ideal participant or agent could plausibly have contributed during the conversation. Each contribution was then evaluated by all group members, who independently selected an ideal turn and acceptable timing window or marked it *do-not-surface*. This yielded 1,091 timing judgments: 384 from message authors and 707 from their 1–3 peers. Finally, participants could revisit their judgments after seeing peers' placements and optionally revise them. Of the preferences revisited during peer deliberation, 85% retain their original ideal turn.

2.4 Pitch Benchmark Composition

The resulting dataset comprises 35 conversations and 2,194 turns from 89 participants (96 group memberships), yielding 384 candidate messages and 1,091 timing preferences. (Table 1). Each candidate message received a median of 3 preferences and at least 2 preferences (Table 2).

	All	**Collective Planning**	**Open Research Discussion**
Number of conversations	35	14	21
Conversation length in turns (median, range)	52 (16-227)	79.5 (30-227)	42 (16-92)
Participant group size (median, range)	3 (2-4)	2 (2-4)	3 (2-4)
Turn length in words (median)	10	7	15
Number of candidate messages (count)	384	158	226
Candidate message length in words (median)	67	61	71

Table 1: Statistics of conversation and candidate messages of PITCH.

3 Experiments

To measure how well model preferences align with those of human participants from the conversation group, we define the following metrics and task setups.

Human timing labels	All	Collective Planning	Open Research Discussion
Number of participants (unique participants)	96 (83)	35 (30)	61 (59)
Timing judgment	1,091	415	676
Distinct participants per message (median, mean)	3, 2.8	3, 2.6	3, 3.0
Placements / do-not-surface	1,047 / 44	404 / 11	643 / 33
Gap between earliest and latest ideal turn (median / p90)	2 / 19	1.5 / 14.6	3 / 20
Messages with a unanimous ideal turn	114 (30%)	58 (37%)	56 (25%)
Individual window size (median)	7	6	7
Anyone's best window size	11	9	12
unanimous window size	3	3.5	3

Table 2: Statistics of human judgments collected in PITCH.

Ground truth of human preferences. For a candidate message judged by N participants, each annotator either names a placement $(t_i, [s_i, e_i])$ or marks do-not-surface (DNS). Let $\mathcal{P}$ be the participants who place the message. The ground truth is the pair $\mathcal{B} = \{ t_i : i \in \mathcal{P} \}$ and $\mathcal{W} = \bigcup_{i \in \mathcal{P}} [s_i, e_i]$. where $\mathcal{B}$ is the set of ideal turns they named, and $\mathcal{W}$ is the set of turns inside any of their acceptable windows. We define the ground truth as sets of turns instead of a single point in the conversation because of the inherent ambiguity of timing judgment. For a given message, appropriate turn-taking timing could vary with participants' knowledge, goals, and interpretation of the discussion Cox et al. (2025); Templeton et al. (2023), so reducing these preferences to one turn would discard valid timing preferences.

Task setups. We evaluated 15 models from five providers, using standard and, where supported, maximum-reasoning configurations, for 28 configurations in total. Each configuration was evaluated across three task setups that vary in how much of the conversation is revealed before the model must commit, and consequently in how closely they reflect realistic deployment settings.

(A) Turn-by-turn replay (*replay*): This is the most realistic deployment setting. The conversation is revealed one turn at a time, and the model decides at each turn whether to intervene. Once it answers "yes," that turn becomes its final placement and cannot be revised.
(B) Replay with five-turn lookahead (*delayed replay*): This is a less constrained replay setting in which the model can see the next five turns at each yes-or-no decision.
(C) Selection from the full transcript (*pick*): This serves as an upper bound rather than a realistic deployment scenario. The agent sees the entire conversation and selects one turn.

Metrics. We quantify performance across task setups using metrics defined against human timing preferences.

Metric	Definition
Anyone's Ideal Turn	The predicted turn is some annotator's ideal turn.
Anyone's Window	The predicted turn falls inside some annotator's window.
Unanimous Window	The predicted turn falls inside every annotator's window. Defined on 296 of 384 messages; NA elsewhere.
Timing Drift	Signed turns from the nearest ideal turn ($+$ late, $-$ early).
Endorsement Share	Fraction of participants who endorse the model's output as a ideal turn, an acceptable turn, or do-not-surface. Endorsement share allows us to compare model performance against a human baseline: model share is $\text{endorsers}(x)/N$ while human share is $(\text{endorsers}(t_i) - 1)/(N - 1)$, the annotator's own vote removed from both sides.

4 RQ1: HOW DO PARTICIPANTS JUDGE GOOD TIMING?

Participants agree on an acceptable window but disagree on ideal turn. An individual annotator's acceptable window has a median width of 6 turns, covering 10.6% of the corresponding conversation after normalizing each window by conversation length. When every annotator shares at least one acceptable turn, the unanimous window has median width 3 turns. However, participants routinely disagree on the exact ideal turn placement (Figure 2). The best-turn spread has median 2 turns and p90 19 turns, and only 30% of messages have a unanimous ideal turn. To further quantify the disagreement among annotations, we measure each annotator against the rest of their own group

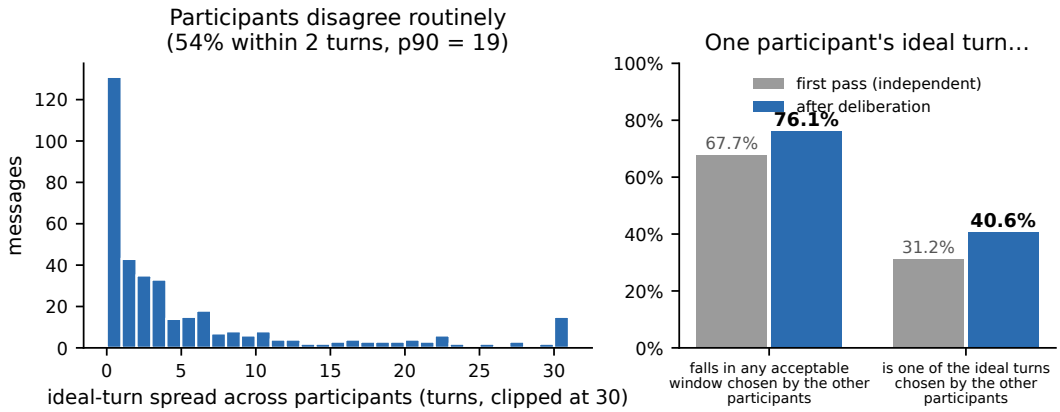


Figure 2: Left: spread between participants' ideal turns. Right: participants scored against the rest of their own group.

Task Setup/Baseline	Anyone's ideal turn	Anyone's Window	Unanimous Window†	Median arm	Best arm
(A) turn-by-turn replay (*replay*)	19.0% [16.5, 21.7]	50.5% [46.8, 54.3]	24.2%	19.0%	32.8%
(B) replay with five-turn lookahead *(delayed replay)*	22.4% [19.3, 25.5]	52.2% [48.2, 56.1]	27.8%	22.1%	33.9%
(C) selection from the full transcript (*pick*)	57.8% [54.3, 61.2]	92.7% [90.9, 94.3]	73.3%	58.7%	64.8%
Human *replay* baseline	23.7%	53.7%	33.2%	—	—
Random turn policy	3.7%	25.5%	7.6%	—	—

Table 3: Model rows are pooled over 28 configurations, with 95% bootstrap intervals resampling items. Best arm is `gemini-3.6-flash` in both *replay* task setups and `gpt-5.6-sol` in *pick*. †Unanimous Window is defined on the 296 of 384 messages with a non-empty intersection across participants and is not comparable to the other columns.

(i.e., *endorsement share*). An annotator's chosen turn falls within others' acceptable windows 76.1% of the time, but exactly matches their ideal turn only 40.6%.

Placements are spread through the conversation. Participants' placements of candidate messages were evenly spanned across the conversation. Across 1,047 placements, the mean relative position is 0.493, and the median is 0.486, with the middle half between 0.27 and 0.73. This broad distribution indicates that there is no universally good time to place a candidate message. This distribution also helps explain some annotator disagreement. A message can both reasonably reopen a thread the group dropped or support a claim after the group returns to the topic. We therefore treat disagreement in annotation data as a signal about the range of recipient preferences and define the model's target as a set of acceptable placements, rather than forcing the model to match a single averaged placement that may reflect no annotator's judgment.

5 RQ2: Can Prompted Models Match Human Timing preferences?

Models can often identify plausible placements when the full conversation is visible, but the same models struggle when they must decide online (Table 3).

Models remain far below the participant-grounded target online. In *replay*, model placements match some annotator's ideal turn on 19.0% of messages and fall inside some annotator's window on 50.5%, leaving nearly half of model placements outside every participant-endorsed window. Five-turn lookahead in the *delayed replay* task setup only marginally improves performance, reaching 22.4% on Anyone's ideal turn and 52.2% on Anyone's Window. There is also considerable variance in accuracy across configurations, from 6.2% for `gpt-oss-120b` with reasoning to 32.8% for `gemini-3.6-flash` with reasoning, measured by ideal turn accuracy. Full-conversation context substantially improves models' performance. With full hindsight in the *pick* task setup, models'

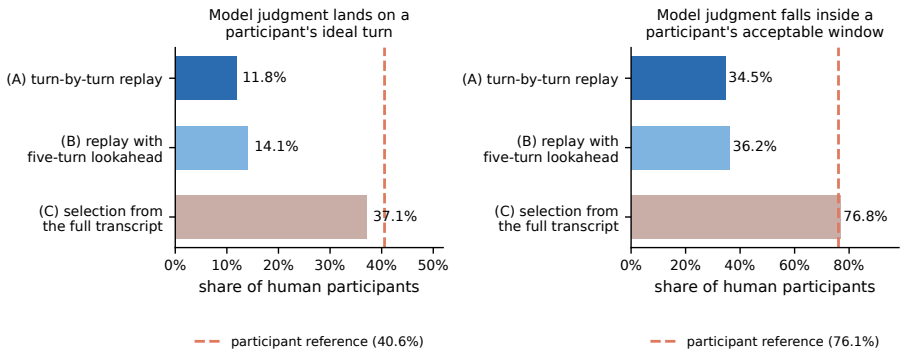


Figure 3: Share of the annotator group a single placement satisfies, against the human reference (dashed). *Pick* reaches the human line; both *replay* task setups fall far short of it.

average accuracy rises from 50.5% to 92.7% on Anyone's Window, but still misses nearly half of the placements (57.8%), judging by Anyone's ideal turn.

Online models satisfy fewer in-group preferences than participants' own placements. Endorsement share lets us compare humans and models using the same metric: the share of the original conversation group a single placement satisfies (Figure 3). After deliberation, a participant's placement on average satisfies 40.6% of the rest of the group on the best-turn criterion and 76.1% on the window criterion; it satisfies the entire group 28.6% and 63.2% of the time. Models reach participant-level endorsement only in the *pick* task setup with hindsight of full conversation context (Appendix Table 6). In *replay*, models are 28.1 points below the deliberated participant reference on best-turn share ([-32.4, -24.0]) and 40.6 points below it on window share ([-44.4, -36.7]) (Appendix Figure 10). Under online evaluation, models achieve less than one-third of the human best-turn share and less than half of the window share.

Models also underperform independent lay annotators in *replay* mode. The original participants provide the benchmark reference but their retrospective placements do not measure how people perform under the online constraint. We therefore ran a separate study with independent Prolific annotators to set a layperson baseline for the *replay* task (study design is detailed in Appendix D). Of 81 recruited annotators, 56 annotators contributed 192 placements. We find lay annotators matched a participant's ideal turn on 23.7% of placements and fell within a participant's union window on 53.7% of placements; an improvement compared to 19.0% and 50.5% for models, respectively. Laypeople still fall short, suggesting the importance of both the additional context that participants bring and the inherent difficulty of anticipating where a discussion will head.

Errors in the online task setups are systematically premature. The dominant error in the online task setup is firing before any annotator's window starts (Figure 4). In *replay*, 46.4% of predictions are issued before any annotator's window has opened, compared with 1.0% issued after every window has closed. This pattern indicates that models tend to place a given message before any of the acceptable windows have begun. In *replay*, the average timing shift is -11.1 (-15.2 conditional on being early) (Figure 9 in Appendix). Participants place messages on average 49.3% of the way into a full conversation, while models chime in on average 26.9% of the way (Figure 5). Considered against a median conversation length of 52 turns, this distribution shows that the average online prediction arrives approximately one fifth of a conversation before the point at which the group judged it appropriate. Compared to models, we also find that lay annotators are less likely to chime in prematurely (39.7% vs 46.4%).

Justifications at the point of commitment are most often about relevance, not timing. To examine what models attend to when they commit, we used a `deepseek-v4-pro` judge to classify 4,349 *replay* trajectories from the 13 models with inspectable reasoning traces, withholding model identity and prediction correctness (Appendix G.4). In *replay*, 92.8% of firing justifications argue that the content fits the current point in the conversation, 7.0% argue that the present moment is appropriate, and 0.14% explicitly weigh waiting. These labels suggest that when asked whether a message should be delivered at the current turn, models often answer whether the message is relevant now, rather than explicitly reasoning about whether the timing could be better in a later moment.

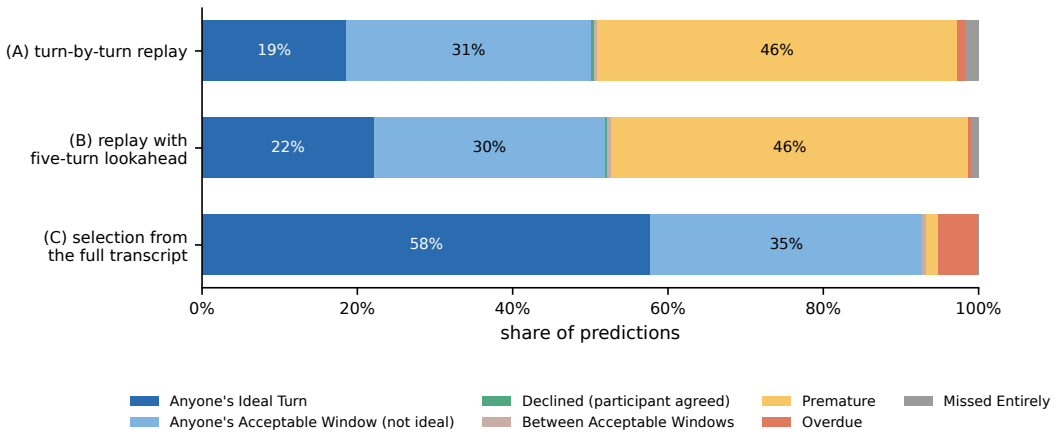


Figure 4: Where predictions land. Online, the single dominant outcome is firing before any annotator's window opened; with hindsight it nearly vanishes and the residual error becomes lateness.

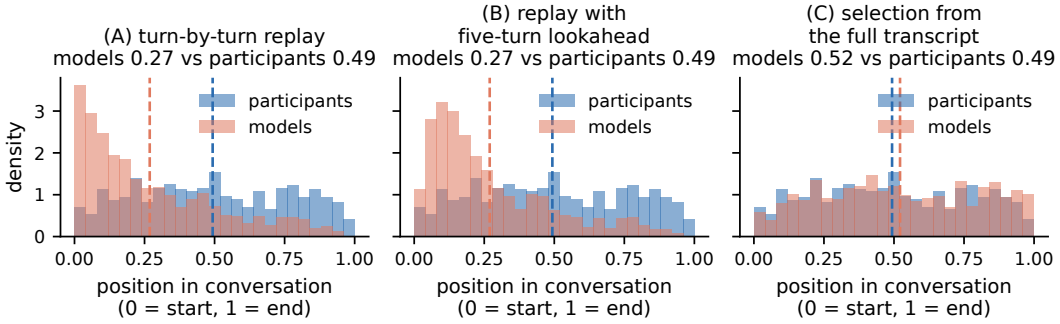


Figure 5: Where models place messages against where people do. The distributions barely overlap online but then nearly coincide with hindsight.

6 Discussion, Limitations, and Future Work

The same proactive message can be appropriate at multiple distinct points in a conversation. Half the messages in Pitch have ideal placements more than two turns apart even after deliberation, and some cluster around multiple distinct points in a conversation. For example, a snack suggestion was placed when food was first raised and again when the topic returned much later. This motivates metrics and ground truth that preserve multiple valid timings rather than force a single gold turn. We found that participants tend to agree when a message has a clear conversational anchor, such as when a participant directly asks about the cost of a bus fare. They are more likely to disagree when the same message could be useful at multiple points in the conversation. For example, a comment about medical evaluation could either revive an earlier thread or support a related claim made twenty turns later. These cases suggest that some disagreements that remain after deliberation reflect different valid uses of the same message, rather than disagreement about the facts.

What does the retrospective reference measure? The reference in Pitch is retrospective: participants see the full conversation and may revise placements after peers' choices, whereas replay models observe only prior turns. Alignment therefore measures how well models anticipate participants' timing preferences, rather than how they compare with humans making the same online decision. These retrospective preferences still provide a meaningful target for when recipients consider an intervention appropriate, but they cannot establish whether a live intervention would improve the discussion, since it may alter what follows.

Our data also captures preferences for only a single proactive interjection, and thus cannot model how prior interventions or intervention frequency shape later timing preferences. Future work could

study these sequential dynamics through richer human data collection or simulation-based benchmarks as models become better at reproducing realistic group conversations.

Under *replay*, model placements receive 34.5% endorsement versus 76.1% for participants' retrospective placements. Together with frequent premature interventions, this gap motivates policies that explicitly weigh acting now against waiting. Prompt ablations and an online human baseline could further separate model-specific limitations from decision framing and task difficulty. Similar challenges arise in proactive coaching, where whether, when, and how to intervene are distinct decisions Srinivas et al. (2025).

Current recruitment limitations and opportunities for broader participation. Our current participant pool reflects the practical demands of synchronous group sessions, with many participants recruited from university communities and through word of mouth. Future work could broaden the demographic range and examine how timing preferences relate to participants' roles and perceptions within a conversation, such as topic expertise, relative influence, or personality, as well as whether models align differently with different participants. The research discussions also center largely on AI, CS, HCI, and related computational fields; extending the protocol to natural sciences, humanities, and other domains could reveal additional conversational dynamics and timing patterns.

Human baselines and the potential for "superhuman" proactive agents. The considerable gap between retrospective participant judgments and the layperson online baseline suggests opportunities to go beyond matching online human decisions. With richer participant context and better models of how the conversation may unfold, an agent could potentially anticipate preferred intervention points more accurately than a layperson deciding in real time. Another direction is to explicitly trade off placement quality against latency by allowing the agent to wait briefly before intervening. In threaded text chat, interface design allows for this tradeoff by enabling an agent to insert a contribution into the recent conversational history, making slightly delayed but better-placed interventions feasible. Indeed, our *delayed replay* task setup probes this possibility, showing some limited improvements in placement accuracy with up to a 5-turn delay in timing.

7 Related Work

Proactive AI agents. The question of when an agent should act dates back to mixed-initiative systems, where assistance must balance the value of acting against the cost of interruption (Horvitz, 1999). Recent proactive agents increasingly model timing explicitly, often with a gate that decides whether the agent should speak (Mao et al., 2024; Eckhaus et al., 2025; Chen et al., 2025; Yang et al., 2026; Huang et al., 2026). Related benchmarks evaluate this through per-turn SPEAK/SILENT prediction (Patel et al., 2025; Chen et al., 2025; Nama et al., 2026), but typically designate a single correct intervention point, often in synthetic or semi-synthetic conversations. In contrast, PITCH uses unscripted human conversations and timing judgments from the actual participants, allowing multiple acceptable placements and treating persistent disagreement as meaningful variation (Aroyo & Welty, 2015; Plank, 2022; Abdulhai et al., 2026).

Human–AI collaborative agents. A parallel body of work studies AI agents in human conversations and collaborative group settings (Do et al., 2022; Houde et al., 2025; Yeo et al., 2026). In groups, proactive timing is especially consequential because an intervention enters a shared space and draws multiple participants' attention (Yeo et al., 2026). Prior work often evaluates downstream outcomes or predefined intervention policies rather than whether timing aligns with participants' preferences. PITCH complements this literature by isolating the timing of a fixed contribution and evaluating it against participant-grounded judgments.

8 Conclusion

We introduce PITCH, a benchmark for evaluating when proactive agents should intervene in multi-party conversations. By jointly collecting unscripted human conversations, participant-authored messages, and participants' timing judgments, PITCH provides a human-grounded testbed for studying proactive timing. Our results show that current models often intervene too early and that appropriate timing can vary across participants and conversational contexts. We hope PITCH enables future work on agents that better understand not only what to contribute, but when to contribute it.

ETHICS STATEMENT

The study was reviewed and approved by our institution's IRB. Participants gave informed consent to the public release of the conversation they held and the annotations they produced, were compensated $35 for the session, and could withdraw at any point, including after the fact. Participants' names are removed from the released dataset. Due to constraints in the data collection procedures, the corpus has limitations: it is English-language text among small, broadly egalitarian groups drawn largely from research and social networks in the United States, so a system calibrated to it encodes one set of conversational norms.

REPRODUCIBILITY STATEMENT

Section 2 gives the collection protocol. Section 3 defines every metric. Section 3 specifies the three task setups and the decoding settings; Appendix G gives the prompts. Appendix E lists the 15 models and 28 configurations together with coverage and cost. Appendix E gives the construction of the no-model reference random policy. The dataset and the analysis notebooks, and the figure generation script are included in the supplementary material. We will release the code of the annotation systems upon acceptance of the paper.

AI USE STATEMENT

We used generative AI as part of the research methodology to classify model reasoning traces, as described in Section 5. Second, we used AI tools to assist with paper revision and language editing. Third, we used coding agents to help develop the data collection interface and supporting system, and to assist in writing data analysis code. The authors reviewed and verified all AI-assisted writing and code, and the reported analyses were produced by running the verified analysis code. Generative AI was not used to make final research conclusions or replace author review and interpretation. Other uses of generative AI were not applicable.

REFERENCES

Marwa Abdulhai, Ryan Cheng, Donovan Clay, Tim Althoff, Sergey Levine, and Natasha Jaques. Consistently simulating human personas with multi-turn reinforcement learning. *Advances in Neural Information Processing Systems*, 38:52920–52957, 2026.

Lora Aroyo and Chris Welty. Truth is a lie: Crowd truth and the seven myths of human annotation. *AI Magazine*, 36(1):15–24, 2015.

Ning Bian, Xianpei Han, Hongyu Lin, Baolei Wu, and Jun Wang. Social simulations with large language model risk utopian illusion. *arXiv preprint arXiv:2510.21180*, 2025.

Tuochao Chen, Nicholas Scott Batchelder, Alisa Liu, Noah A Smith, and Shyamnath Gollakota. Llamapie: Proactive in-ear conversation assistants. In *Findings of the Association for Computational Linguistics: ACL 2025*, pp. 13801–13824, 2025.

Christopher Cox, Riccardo Fusaroli, Yngwie A Nielsen, Sunghye Cho, Roberta Rocca, Arndis Simonsen, Azia Knox, Meg Lyons, Mark Liberman, Christopher Cieri, et al. Social context matters for turn-taking dynamics: A comparative study of autistic and typically developing children. *Cognitive Science*, 49(10):e70124, 2025.

Hyo Jin Do, Ha-Kyung Kong, Jaewook Lee, and Brian P Bailey. How should the agent communicate to the group? communication strategies of a conversational agent in group chat discussions. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–23, 2022.

Niv Eckhaus, Uri Berger, and Gabriel Stanovsky. Time to talk: Llm agents for asynchronous group communication in mafia games. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing (Findings)*, pp. 11356–11368, 2025.

Eric Horvitz. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 159–166, 1999.

Stephanie Houde, Kristina Brimijoin, Michael Muller, Steven I Ross, Dario Andres Silva Moran, Gabriel Enrique Gonzalez, Siya Kunde, Morgan A Foreman, and Justin D Weisz. Controlling ai agent participation in group conversations: A human-centered approach. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*, pp. 390–408, 2025.

Tao Huang, Jun-Wei Zhang, Pengju Yan, Rui Yang, and Xiao-Lin Li. Chatagents: Enhancing autonomous decision-making in multi-agent conversations via intrinsic motivation. *2026 International Conference on Artificial Intelligence and Materials (ICAIM)*, pp. 1–6, 2026. URL https://api.semanticscholar.org/CorpusID:290254458.

Ryo Kamoi, Ameya Godbole, Longqi Yang, Rui Zhang, Mengting Wan, and Pei Zhou. Evaluating llm-simulated conversations in modeling inconsistent and uncollaborative behaviors in human social interaction. *arXiv preprint arXiv:2603.17094*, 2026.

Xingyu Bruce Liu, Shitao Fang, Weiyan Shi, Chien-Sheng Wu, Takeo Igarashi, and Xiang 'Anthony' Chen. Proactive conversational agents with inner thoughts. In Naomi Yamashita, Vanessa Evers, Koji Yatani, Sharon Xianghua Ding, Bongshin Lee, Marshini Chetty, and Phoebe O. Toups Dugas (eds.), *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems, CHI 2025, YokohamaJapan, 26 April 2025- 1 May 2025*, pp. 184:1–184:19. ACM, 2025a. doi: 10.1145/3706598.3713760. URL https://doi.org/10.1145/3706598.3713760.

Ziyi Liu, Bahareh Sarrafzadeh, Pei Zhou, Longqi Yang, Jieyu Zhao, and Ashish Sharma. Promediate: A socio-cognitive framework for evaluating proactive agents in multi-party negotiation. *CoRR*, abs/2510.25224, 2025b. doi: 10.48550/ARXIV.2510.25224. URL https://doi.org/10.48550/arXiv.2510.25224.

Manqing Mao, Paishun Ting, Yijian Xiang, Mingyang Xu, Julia Chen, and Jianzhe Lin. Multi-user chat assistant (MUCA): a framework using llms to facilitate group conversations. *CoRR*, abs/2401.04883, 2024. doi: 10.48550/ARXIV.2401.04883. URL https://doi.org/10.48550/arXiv.2401.04883.

Vihaan Nama, Shreya Mendi, Zian Ye, and Brinnae Bent. When2speak: A dataset for temporal participation and turn-taking in multi-party conversations for large language models. *CoRR*, abs/2605.05626, 2026. doi: 10.48550/ARXIV.2605.05626. URL https://doi.org/10.48550/arXiv.2605.05626.

Deep Anil Patel, Iain Melvin, Christopher Malon, and Martin Renqiang Min. Discussllm: Teaching large language models when to speak. *CoRR*, abs/2508.18167, 2025. doi: 10.48550/ARXIV.2508.18167. URL https://doi.org/10.48550/arXiv.2508.18167.

Barbara Plank. The "problem" of human label variation: On ground truth in data, modeling and evaluation. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pp. 10671–10682. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022.EMNLP-MAIN.731. URL https://doi.org/10.18653/v1/2022.emnlp-main.731.

Vidya Srinivas, Xuhai Xu, Xin Liu, Ayush Kumar, Isaac Galatzer-Levy, Shwetak Patel, Daniel McDuff, and Tim Althoff. Substance over style: Evaluating proactive conversational coaching agents. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 20848–20880, 2025.

Emma M Templeton, Luke J Chang, Elizabeth A Reynolds, Marie D Cone LeBeaumont, and Thalia Wheatley. Long gaps between turns are awkward for strangers but not for friends. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1875):20210471, 2023.

Ruotong Wang, Xinyi Zhou, Lin Qiu, Joseph Chee Chang, Jonathan Bragg, and Amy X Zhang. Social-rag: Retrieving from group interactions to socially ground ai generation. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pp. 1–25, 2025.

Shu Yang, Difei Xu, Jiaxin Pei, and Di Wang. Proact: Towards breakdown-aware proactive agent in multi-user collaboration. *arXiv preprint arXiv:2607.03730*, 2026.

ShunYi Yeo, Tianyi Zhang, Scott Bateman, Gary Hsieh, Young-Ho Kim, Simon Tangi Perrault, Jiannan Li, and Anthony Tang. Group conversational agents: A review of designs that support and shape group interaction. In *Proceedings of the 2026 Designing Interactive Systems Conference*, pp. 2765–2779, 2026.

Xuhui Zhou, Zhe Su, Tiwalayo Eisape, Hyunwoo Kim, and Maarten Sap. Is this the real life? is this just fantasy? the misleading success of simulating social interactions with llms. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 21692–21714, 2024.

A CONVERSATION TOPICS

The benchmark evaluates knowledge placement within ongoing conversations spanning a diverse set of topics and conversational goals. Each conversation begins with a prompt developed from topic interests the users provided in a pre-study survey. These are designed to cover a range of domains and information needs.

Below, we provide two representative prompts in full, followed by a list of keywords describing the broader prompt set to illustrated the breadth of topics and activities represented in the benchmark. Candidate knowledge messages associated with these conversations are described separately in Appendix B.

A.1 REPRESENTATIVE TOPICS

Collaborative Planning. *"You're planning a 3-day trip together to a US city for Labor Day weekend. Plan their itinerary together, a mix of must-see spots and a couple of local, less touristy picks, plus where you should eat and how to get around, given the time you have. Consider the hours, current prices, transportation, etc. By the end, land on: a city, a day-by-day itinerary, and a rough per-person budget."*

This prompt represents the collaborative planning subset of the benchmark, in which participants may develop travel itineraries or event plans together and encounter information needs over the course of the conversation.

Scholarly Discussion. *"What Does HCI Owe AI Development Right Now? The dominant framing in AI research is capabilities: can the model do the task? HCI has decades of knowledge about interaction patterns, user agency, and sociotechnical systems, but is that knowledge actually reaching the people building agents? What concrete contributions should HCI researchers be pushing into AI development today, and what's stopping them? Is formalization the bottleneck, or is it access and incentives?"*

This prompt represents the scholarly discussion subset, where the conversation develops around an open-ended research question rather than a concrete planning task.

A.2 TOPIC BREADTH

Across the benchmark, 25 prompts span a range of travel, event organization, learning, and budget-constrained planning as well as AI in education and research, evaluation, design, and accessibility.

Collaborative planning

- Plan a two- to three-day trip
- Plan a reunion weekend
- Organize a small creative hackathon
- Organize an esports competition
- Plan a lab group retreat
- Learn a skill together
- Run a budget-constrained event

Research discussions

- Explore AI tools for student learning
- Consider AI for research assistance
- Discuss existing AI benchmarks and progress measurement
- Examine generative AI for user simulation
- Consider AI beyond the productivity paradigm
- Discuss who should determine AI agent behavior control

- Examine conversational AI voice and tone
- Consider LLM agents for accessibility
- Discuss LLM assumptions across diverse users
- Explore unresolved issues around autonomous driving technology
- Examine how HCI knowledge can inform AI development

B CANDIDATE KNOWLEDGE MESSAGES

For each conversation, participants produce candidate knowledge messages that they believe appropriate to introduce at a particular point in their dialogue. We provide five representative candidate messages below, illustrating the range of knowledge being evaluated.

Example 1:

"If you're looking for more budget-friendly seafood options beyond the upscale spots, there are some great affordable alternatives. [Yankee Lobster Co.](https://www.busybeingshasha.com/blog/yankeelobster) gets consistently recommended for reasonably priced lobster rolls, and [Pauli's](https://www.tripadvisor.com/Restaurants-g60745-c33-zfp16-Boston_Massachusetts.html) shows up on multiple "best cheap seafood" lists. For a true local experience, check out Burke's - it's described as the definition of "hole in the wall" seafood on [r/boston](https://www.reddit.com/r/boston/comments/14kyu1q/best_hole_in_the_wall_cheap_seafood), praised for being professional and clean but not fancy. You could definitely do a good seafood meal for much less than the "splurge" places while still getting that authentic Boston experience."

This example corresponds to the trip-planning prompt in Appendix A.

Example 2:

AI researchers may be undervaluing interface design, a concern that is receiving attention from design leaders. Jakob Nielsen recently argued we might be heading toward ["No More User Interface"](https://jakobnielsenphd.substack.com/p/no-more-ui) as AI becomes the primary interaction layer. However, there's a counter-argument that [AI is exposing what design really is](https://uxdesign.cc/what-ai-exposes-about-design-319029d48441) - the strategic thinking behind interfaces, not just the visual output. Some companies are replacing UX research with [AI-generated user testing](https://smart-interface-design-patterns.com/articles/ai-user-research), which could further marginalize HCI contributions if the field doesn't articulate its unique value in understanding human-AI interaction patterns.

This example corresponds to the scholarly-discussion prompt in Appendix A.

Example 3:

That's a really important point about dialects. Research shows that reward models used in LLM training exhibit systematic biases against non-standard dialects like African American Language [**Rejected Dialects: Biases Against African American Language in Reward Models**](https://www.semanticscholar.org/paper/f400b0d5724203633f56299cb19f34bf18fd31b9). So longer context windows might not help if the underlying training biases cause models to misinterpret dialectal variations.

Example 4:

Boston's on the Northeast Corridor, which is actually Amtrak's strongest route! The Northeast Regional runs frequently between Boston and NYC(https://www.timeout.com/usa/news/a-defense-of-the-northeast-regional-amtraks-best-east-coast-train-111424) (about 3.5 hours) and continues down to DC. It's much more reliable than most west coast Amtrak routes. Plus you could potentially take the Acela for the premium experience - that gets you to NYC in under 3 hours.

C STUDY PLATFORM FOR THE INTERACTIVE DATA COLLECTION PROTOCOL

We developed a web-based data collection platform to collect conversations while supporting both candidate knowledge message creation and cross-annotation. It hosts the full study workflow.

Conversation and contribution. Participants first use the platform to engage in a synchronous group conversation. After the conversation, they review the transcript and create candidate knowledge items they believe could have been useful during the discussion. The message contribution interface (Figure 6) requires participants to identify where in the conversation the item would be relevant, including ideal turn and a range of acceptable turns. Integrated tools help with resource search and retrieval as well as message authoring.

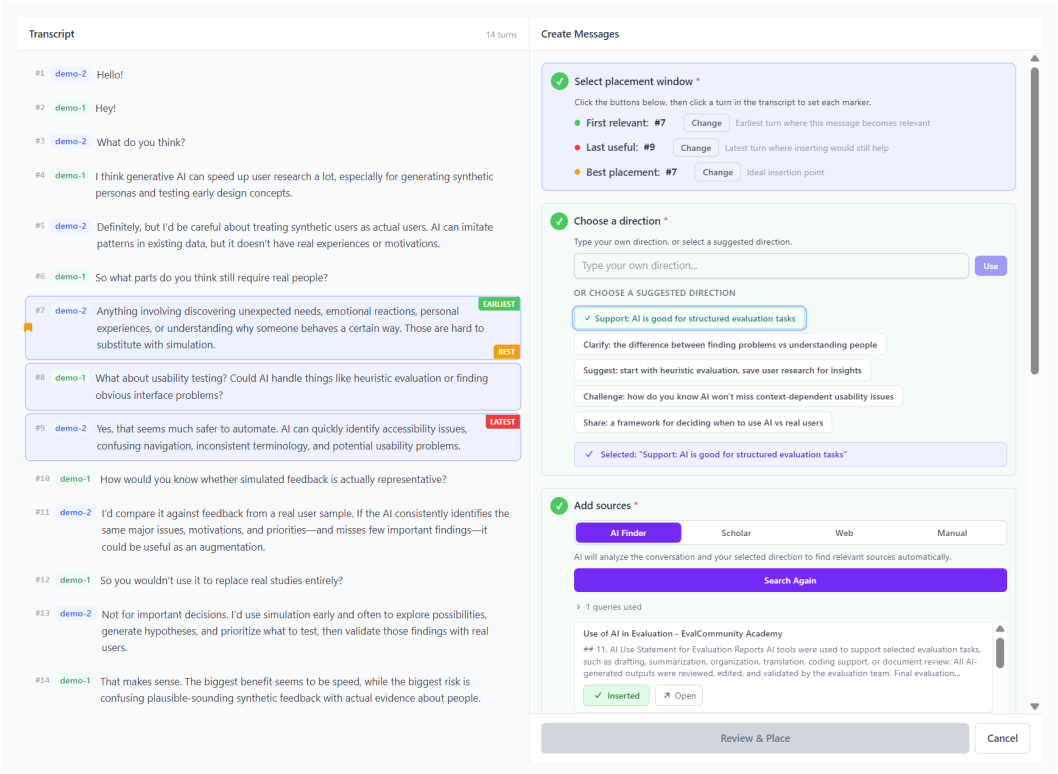


Figure 6: Message contribution interface. The full transcript sits in a scrollable pane on the left and the candidate message and placement controls on the right.

Timing judgments. Post-construction, participants play the role of participants, reviewing the neutralized (Appendix G.3) contributed messages created by each other. The interface allows them to similarly indicate where these should appear within the conversation, if at all, and optionally providing brief rationale for their decisions (Figure 7).

Deliberation. A subsequent deliberation stage presents each item from the set of created knowledge messages with every participant's placement, allowing participants to review and revise their own annotations (Figure 8).

Researchers manage study group formation, review knowledge items, configure annotation assignments, and export the resulting data through an admin interface. The platform also includes an interactive tutorial to familiarize users with each stage of the study before they begin the actual task.

The platform was implemented as a full-stack web application using React and TypeScript with a FastAPI PostgreSQL backend.

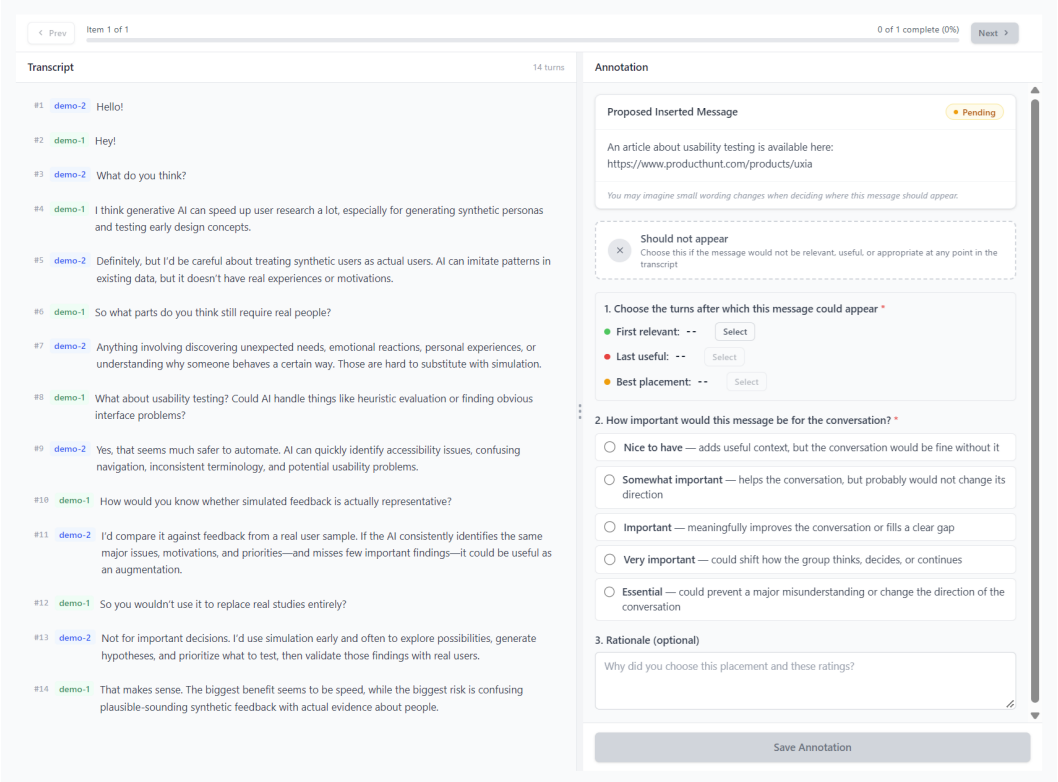


Figure 7: Timing judgment interface. Users navigate the transcript in the left pane while making annotation decisions in the right pane for each message contributed by other participants in the group.

D Human Baseline for Intervention Timing

To contextualize model performance, we conducted a human baseline study measuring how lay users perform the same replay intervention-timing task. The goal was to determine whether models' difficulty in selecting an appropriate intervention point and tendency to speak early reflects the difficulty of the task or a gap between model behavior and human preferences.

D.1 Study Design

Participants were recruited through Prolific and completed the replay task (Section 3) on a web-based annotation platform. We collected data across both scholar and trip conversation scenarios which allows us to measure human baseline performance across the full range of contexts used in model evaluation.

Our ground truth is derived from the data collected for PITCH. For each candidate contribution, we considered two measures: ideal turn (the set of turns selected as ideal intervention points in our ground truth) and union windows (span of earliest to latest acceptable intervention points across participants). These measures allow direct comparison with model performance.

D.2 Annotation Platform and Protocol

The annotation interface presented participants with a conversation transcript alongside three corresponding candidate contributions, de-anonymized and pulled from the PITCH dataset. The interface simulates a natural flow of conversation by allowing participants to progress turn by turn, deciding at each point whether any of the three candidate contributions should be sent. This reflects the replay task we used for model evaluation.

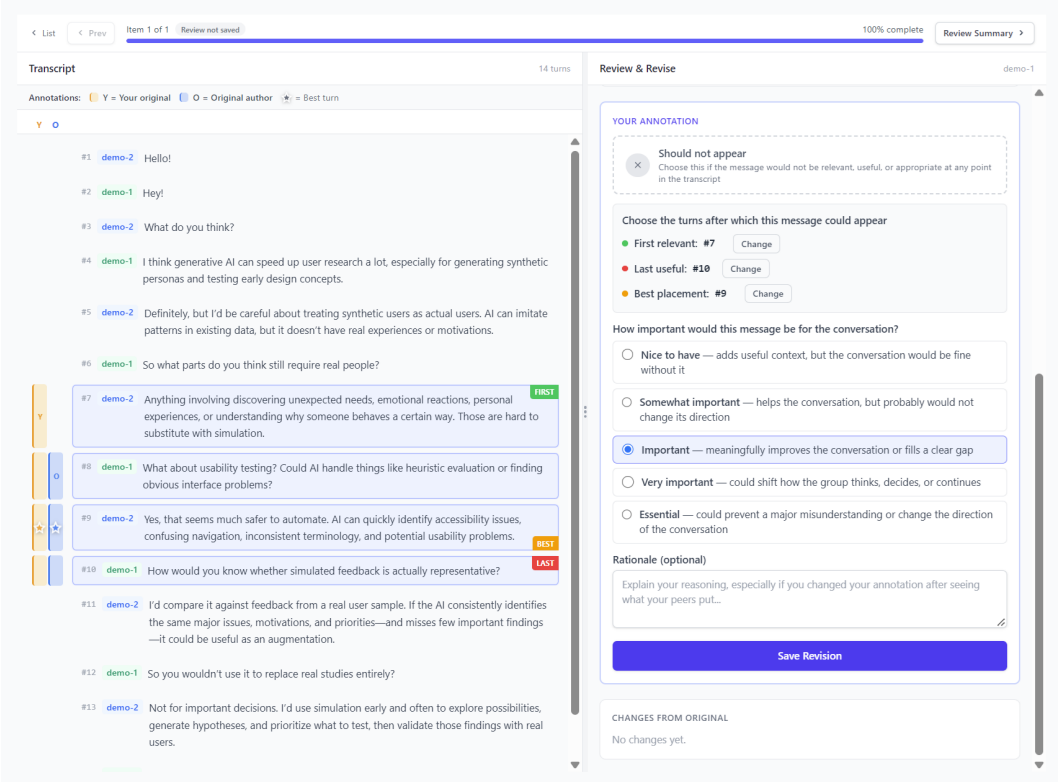


Figure 8: Deliberation interface. Users are presented with the conversation transcript once again, alongside timing judgments made on a contribution message by all other participants. The right pane allows users to optionally revise and update their judgments.

At each turn in the conversation, participants could see:

- The conversation history up to that point
- The three candidate contributions
- A "Send" button for each contribution that had not yet been placed
- A "Continue" button to advance to the next turn without placing any contribution

If participants reached the end of the conversation with contributions still unplaced, they were given the option to mark those remaining contributions as "not relevant"–indicating the message did not fit anywhere in the conversation.

Each session consisted of placing three candidate contributions within a single conversation and participants could complete multiple sessions across different conversations.

Participants received a base compensation of $1.00 USD, with up to an additional $2.00 bonus for placing contributions within the ground truth union windows. This incentive structure encouraged accuracy while acknowledging that reasonable people may disagree on exact timing.

D.3 Participants and Quality Control

We recruited 109 participants, of which 59 completed selections for collaborative-planning scenarios and 50 for scholarly-discussion scenarios. After manual quality control, 331 data points were retained for analysis.

Quality control consisted of filtering for suspicious completion patterns: users who completed items in implausibly short times or who selected the first available turns regardless of conversation content suggested that they did not read the instructions or conversations carefully.

Table 4: Lay-user and model performance on intervention timing. Human performance is reported separately for scholarly and collaborative scenarios, as well as in aggregate (N=331 placements).

	Human (Lay Users)			Models
Metric	**Scholarly**	**Collaborative**	**Aggregate**	
Anyone's ideal turn	22.4%	24.6%	23.7%	19.0%
Within anyone's window	55.2%	52.6%	53.7%	50.5%
Within unanimous window	34.3%	32.5%	33.2%	24.2%
Earlier than anyone's window	37.6%	41.1%	39.7%	46.4%
Mean timing shift (turns)	−4.3	−11.8	−8.7	−11.1

Provider	**Models**
Google	`gemini-3.1-pro`, `gemini-3.6-flash`
OpenAI	`gpt-5.6-sol`, `gpt-5.6-luna`
Anthropic	`opus-4-6`, `sonnet-4-6`, `haiku-4-5`
Open-weight (Together AI)	`deepseek-v4-pro`, `glm-5.2`, `kimi-k3`, `minimax-m3` `qwen-3.7-max`, `gpt-oss-120b`, `inkling` `nemotron-3-ultra-550b-a55b`

Table 5: The 15 models evaluated, each in standard and (where supported) reasoning mode.

D.4 Results

The resulting lay-user baseline substantially exceeds the performance of models evaluated on the same task (Table 4).

Both lay users and models exhibit a consistent early placement bias. Participants placed contributions earlier than anyone's window in 39.7% of cases, with a mean timing shift of −8.69 turns from the nearest acceptable intervention point. However, this bias is notably less severe than models' tendency to intervene early.

E Models, Coverage, and Cost

Coverage and input. Models receive the neutralized candidate message and a system prompt describing the proactive-agent task, with conversation visibility determined by the framing (Section 3). The prompt is included in Appendix G. Standard mode uses temperature 1.

Scale and cost. The run comprises 77,755 API calls at a total cost of $507.83: $262.89 for replay against $18.05 for pick.

Random Baseline A random turn drawn uniformly over the conversation.

F Additional Results

F.1 Model Performance by Conversation Scenario

Table 7 separates the full-corpus model results in Table 3 into scholarly discussions and trip-planning conversations. The scenario split is smaller than the difference among task framings, and its direction changes with the framing. In replay, trip-planning accuracy is 2.0 points higher on Anyone's ideal turn, while Anyone's Window is nearly unchanged (50.8% against 50.3%). With five-turn lookahead, scholarly discussions are 2.5 points higher on Anyone's ideal turn. The largest scenario gap appears in pick: scholarly discussions are 4.9 points higher on Anyone's ideal turn, 3.4 points higher on Anyone's Window, and 5.8 points higher on Unanimous Window.

These are descriptive subgroup results rather than a controlled scenario effect. Scenario is omitted from the model prompt, but the two subsets differ in both conversation structure and message

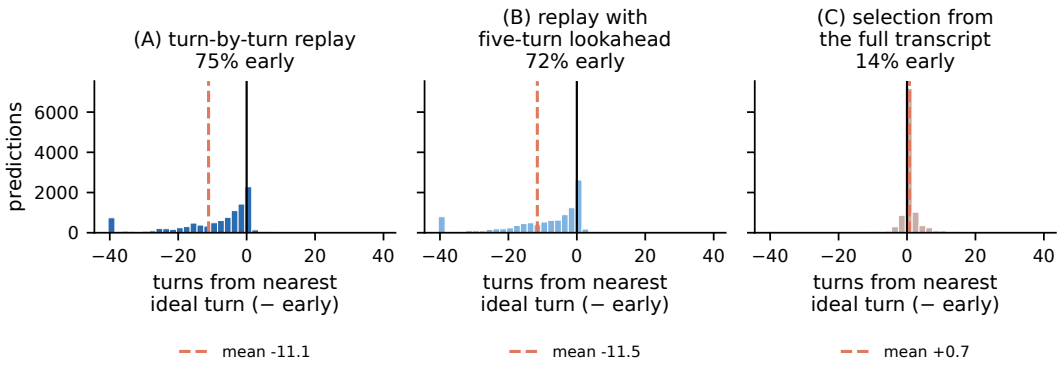


Figure 9: Timing drift by framing. Averaging the online and hindsight distributions would report a mild early bias that describes neither, which is why the framings are never pooled.

	Best-turn criterion		Window criterion	
	mean share	satisfied all	mean share	satisfied all
Human annotator, after deliberation	40.6%	28.6%	76.1%	63.2%
Human annotator, first pass	31.3%	18.0%	67.7%	53.2%
(A) turn-by-turn replay	37.1%	18.1%	76.8%	56.5%
(B) replay with five-turn lookahead	14.1%	6.7%	36.2%	21.4%
(C) selection from the full transcript	11.8%	5.4%	34.5%	18.7%

Table 6: Share of the annotator group a single placement satisfies, the only quantity on which humans and models are scored by an identical definition. Bold marks the two human reference rows. Both human rows exclude the annotator's own vote from numerator and denominator; a model has no vote in the group, which is what makes the rows comparable. The first-pass row is each annotator's independent judgment before seeing anyone else's. Paired differences reported in the text are computed on the *n* items where every arm produced a placement, so they will not equal differences of these marginal columns.

composition. Scholarly discussions have a median of 42 turns and three participants, whereas trip-planning conversations have a median of 79.5 turns and two participants (Table **??**). The bootstrap intervals for Scholar and Trip also overlap in every framing. The split therefore does not support attributing the observed differences to the scenario label itself.

G Prompts

The strings below are the prompts used to draft candidate messages, strip placement cues from them, elicit timing decisions, and label firing justifications. A token in braces is a runtime field. Every other character is fixed.

G.1 Timing

The same system message and user-message scaffold are used for pick, replay, and delayed replay. The reported standard and native-reasoning configurations both instruct the model to leave the reasoning field empty; native reasoning, where a provider supports it, is the configuration marked (R) in Figure 11.

System message.

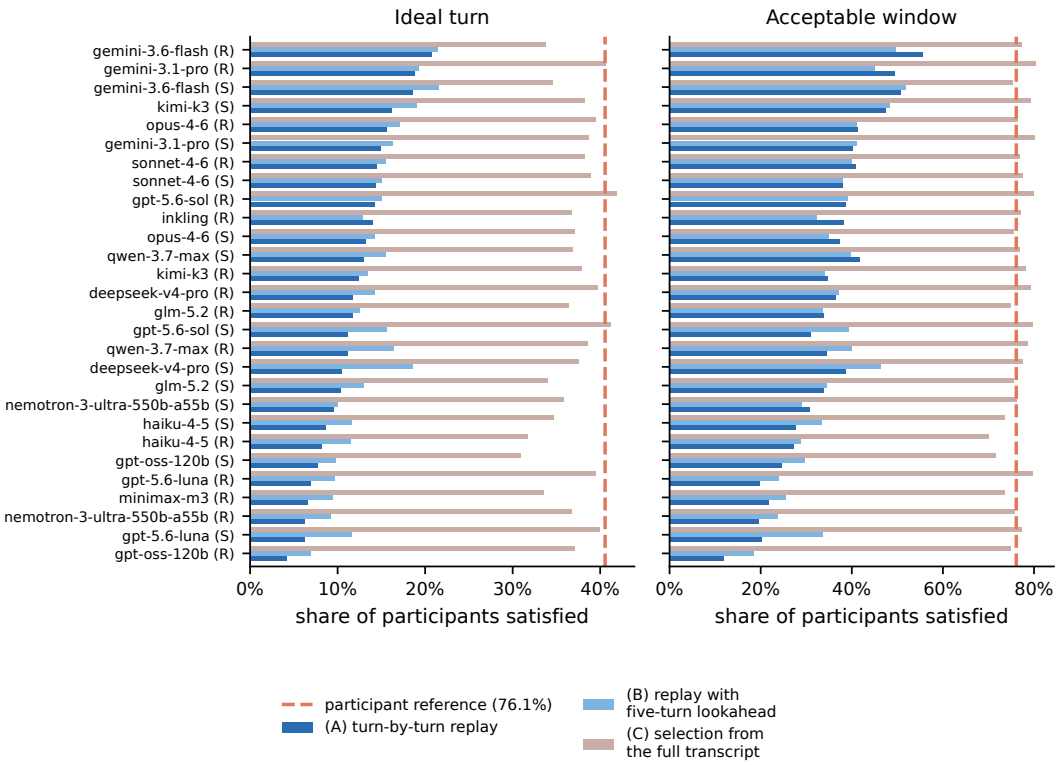


Figure 10: Every configuration against the human reference (dashed). In pick, three arms clear the human best-turn bar and eighteen clear the window bar; in both online framings, none do. (R) means reasoning mode and (S) means standard model configuration

Task Setups	Subset	Anyone's ideal turn	Anyone's Window	Unanimous Window†	Median arm	Best arm
(A) turn-by-turn replay	Full	19.0% [16.5, 21.7]	50.5% [46.8, 54.3]	24.2%	19.0%	32.8%
	Scholar	18.2% [14.8, 21.9]	50.3% [45.6, 55.0]	23.6%	18.1%	28.3%
	Trip	20.2% [16.2, 24.2]	50.8% [45.0, 56.2]	25.0%	20.6%	39.9%
(B) replay with five-turn lookahead (eased)	Full	22.4% [19.3, 25.5]	52.2% [48.2, 56.1]	27.8%	22.1%	33.9%
	Scholar	23.4% [19.3, 27.6]	53.7% [48.3, 58.8]	30.2%	24.6%	32.3%
	Trip	20.9% [16.6, 25.4]	50.1% [44.2, 55.9]	24.9%	21.5%	37.3%
(C) selection from the full transcript (upper bound)	Full	57.8% [54.3, 61.2]	92.7% [90.9, 94.3]	73.3%	58.7%	64.8%
	Scholar	59.8% [54.9, 64.3]	94.1% [92.3, 95.8]	75.9%	60.2%	66.8%
	Trip	54.9% [49.8, 60.0]	90.7% [87.4, 93.6]	70.1%	55.7%	63.3%

Table 7: Model performance on the full corpus and by conversation scenario, pooled over 28 configurations. Full contains 384 messages, Scholar 226, and Trip 158. Brackets are 95% bootstrap intervals resampling items. The best configuration differs across subsets: `gemini-3.6-flash` is best in replay, except that its standard arm leads Scholar and its reasoning arm leads Full and Trip; `kimi-k3` standard leads Scholar in delayed replay, while `gemini-3.6-flash` reasoning leads Full and Trip; and `gpt-5.6-sol` reasoning leads Full and Scholar in pick, while its standard arm leads Trip. †Unanimous Window is defined on 296 of 384 Full, 160 of 226 Scholar, and 136 of 158 Trip messages.

```
You are observing a naturally unfolding small-group conversation.
↪  Conversations typically span roughly 20{200 turns, averaging about
↪  60. You hold information that could help the group's conversation
↪  become more informative, productive, and smooth, which you would
↪  contribute in the form of a message inserted into the
↪  conversation.
```

User-message scaffold.

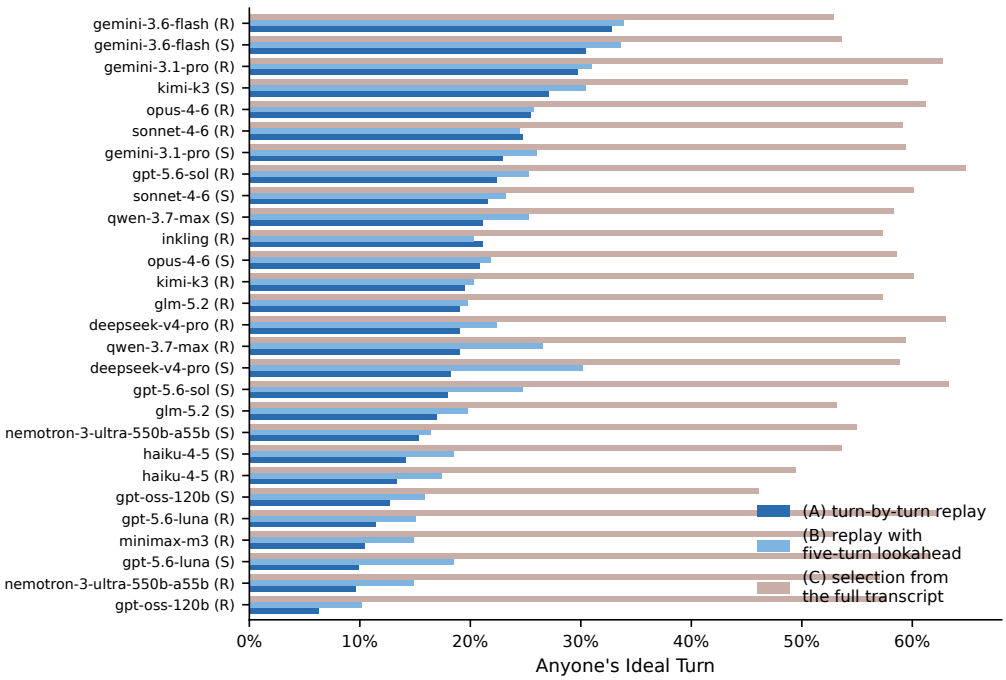


Figure 11: Anyone's ideal turn for all 28 configurations. The gap between the pick bar and the replay bar is wider than the entire spread across models within either framing, and it is positive for every configuration. (R) indicates models in native reasoning mode with max effort; (S) indicates models in lowest reasoning effort.

```
Visible conversation:
{formatted visible transcript}

Candidate messages to insert (each message appears once):
{candidate_messages_json}

{task_instruction}

{batch_instruction_if_any}

{reasoning_instruction}

Return exactly one valid JSON object matching this shape:
{answer_template_json}

{yes_example_if_replay_or_delayed_replay}

Do not use Markdown fences or include any text outside the JSON
↪  object.
```

Each visible turn is written as below, with a blank line between turns:

```
Turn {turn_number} ({speaker}): {text}
```

Candidate messages use request-local labels, not UUIDs:

```
[
  {
    "message_id": "item_1",
    "message": "{neutralized candidate message text}"
  },
  {
    "message_id": "item_2",
```

```
    "message": "{neutralized candidate message text}"
  }
]
```

Task instruction, (A) turn-by-turn replay.

```
At checkpoint {checkpoint}, decide whether to insert each message now,
↪  immediately after turn {checkpoint}. Use decision "YES" or "NO";
↪  YES must select turn {checkpoint}, and NO must set selected_turn
↪  to null.
```

Task instruction, (B) replay with five-turn lookahead

```
You are now at turn {visible}. Decide whether to insert each message
↪  at one of turns {start} through {end} - the {span} most recent
↪  turns. The message would be inserted immediately after the turn
↪  you select. Use decision "YES" or "NO"; YES must select a turn
↪  between {start} and {end}, and NO must set selected_turn to null.
```

Task instruction, (C) selection from the full transcript .

```
For each message, select the single ideal turn in the full
↪  conversation to insert it. The message would be inserted
↪  immediately after the turn you select. Use decision "PICK" and its
↪  selected_turn.
```

Batched-message instruction. Inserted for {batch_instruction_if_any} when a request contains more than one candidate:

```
Judge each message completely independently, as if it were the only
↪  message you were given. Do not rank or compare the messages
↪  against one another, and do not let your decision about one
↪  message influence your decision about any other. Two messages may
↪  be inserted at the same turn, and any number of them may be
↪  inserted or not inserted.
```

Answer template, (A) turn-by-turn replay or (B) replay with five-turn lookahead

```
{
  "reasoning": null,
  "final_answer": [
    {
      "message_id": "item_1",
      "decision": "NO",
      "selected_turn": null
    }
  ]
}
```

Answer template, (C) selection from the full transcript .

```
{
  "reasoning": null,
  "final_answer": [
    {
      "message_id": "item_1",
      "decision": "PICK",
      "selected_turn": 1
    }
  ]
}
```

G.2 CANDIDATE-MESSAGE DRAFTING

Participants author messages with search and drafting support (Section 2). Optional blocks are included only when the corresponding input is present.

System message.

```
You are an expert AI assistant embedded in a group conversation
↪  annotation tool. Your task is to generate a concrete, helpful
↪  agent utterance that could be inserted into the conversation to
↪  surface relevant external knowledge for the participants. The
↪  utterance should be natural, informative, and directly useful to
↪  the ongoing discussion. It should sound like a knowledgeable
↪  colleague contributing a fact, reference, clarification, or
↪  suggestion | not like a chatbot.
```

User-message template.

```
## Scenario
{scenario}

## Conversation transcript
{transcript_text}

{optional_user_selected_resources_or_search_results}

{optional_response_requirements}

{optional_current_draft_block}

{final_instruction}
```

If participant-selected resources are available, the resources block is:

```
## User-selected resources (use these as primary sources)
1. **{title}**{optional_manual_source_note}
   {snippet up to 500 characters}
   URL: {url}
```

If only search results are available, the resources block is:

```
## Relevant search results
1. **{title}**
   {snippet_or_abstract up to 500 characters}
   URL: {url}
```

Optional contribution-type guidance:

```
## Response requirements
**Contribution type:** {type_guidance}
**Importance level:** {importance_label}
```

Contribution-type guidance values:

```
heads_up: Introduce a useful fact, paper, example, or context about a
↪  topic the conversation has touched on but not yet considered.
pushback: Challenge or question something someone said in the
↪  conversation.
support: Confirm or back up something someone said in the conversation
↪  with evidence.
suggestion: Recommend what the group could do, consider, read,
↪  compare, or decide next.
```

Importance labels:

```
importance >= 0.75: Essential or very important
importance >= 0.5: Important
importance >= 0.25: Somewhat important
otherwise: Nice to have
```

Final instruction without an existing draft:

```
Based on the conversation and any search results above, generate a
↪  single concrete and helpful agent utterance that would benefit the
↪  group discussion. Keep the response under 100 words. Be concise
↪  and direct. Include inline citations for the sources provided
↪  above. For sources with URLs, use markdown hyperlinks: e.g.
↪  'According to [Smith et al. (2023)](https://example.com/paper),
↪  ...'. For manual sources without URLs, cite them by title in bold:
↪  e.g. 'As noted in **Introduction to Algorithms**, ...'. Only cite
↪  sources that were provided above. Respond with only the utterance
↪  text, no extra commentary.
```

Final instruction with an existing draft:

```
The user has already started writing the response shown above. Refine
↪  their draft into a concise agent utterance that preserves the
↪  original intent. Keep the response under 100 words. Be concise and
↪  direct. Include inline citations for the sources provided above.
↪  For sources with URLs, use markdown hyperlinks: e.g. 'According to
↪  [Smith et al. (2023)](https://example.com/paper), ...'. For manual
↪  sources without URLs, cite them by title in bold: e.g. 'As noted
↪  in **Introduction to Algorithms**, ...'. Only cite sources that
↪  were provided above. Respond with only the utterance text, no
↪  extra commentary.
```

G.3 NEUTRALIZATION

Timing runs use the neutralized message (Section E).

System message.

```
You are a text neutralization assistant. Your task is to rewrite an AI
↪  agent utterance so that it preserves all factual and informational
↪  content but removes any conversational cues that tie it to a
↪  specific moment in a conversation.

Specifically, you must remove or rephrase:
- Discourse markers (e.g. 'Actually...', 'By the way...', 'So...')
- Turn-specific references (e.g. 'As you mentioned earlier...', 'Going
↪  back to what Alice said...')
- Response-to-question patterns (e.g. 'To answer your question...',
↪  'That's a great point, but...')
- Conversation participant name references (e.g. 'As Bob noted...',
↪  '@Alice'). Only remove names of conversation participants, NOT
↪  names that appear in citations or references.
- Hedging or politeness tied to the conversation flow
- Turn numbers or placement references (e.g. 'responding to turn 5',
↪  'after turn 3', 'at this point in the conversation'). Remove ALL
↪  references to specific turn numbers as they leak placement
↪  information.

INSTRUCTION-STYLE DETECTION: If the text is written as behavioral
↪  instructions for an AI (e.g. 'The AI should fact-check...', 'When
↪  the user says X, respond with...', 'Fact-check what was said')
↪  rather than a direct message, you MUST rewrite it as the concrete
↪  utterance the AI would actually send. Convert the instruction into
↪  the actual helpful message that would appear in the conversation.
```

```
IMPORTANT: You must preserve all academic citations (e.g. 'Smith et
↪  al., 2020', 'doi:10.1234/...'), bibliographic references,
↪  institution names (e.g. 'MIT', 'World Health Organization'),
↪  author names within citations, markdown hyperlinks (e.g.
↪  '[text](url)'), and any other sourcing information. These are
↪  factual content, not conversational cues.

The output should read as a standalone informative statement that
↪  could appear in any context. Preserve the exact factual claims,
↪  data, references, citations, and recommendations. Do not add or
↪  remove information.

Respond with ONLY the neutralized text, no commentary.
```

User message.

```
Please neutralize the following agent utterance:

---
{raw_text}
---
```

G.4 FIRING-JUSTIFICATION JUDGE

The replay reasoning trace analysis in Section 5 uses the following judge.

System message.

```
You label the reasoning behind a decision to add a message to a group
↪  conversation AT A PARTICULAR MOMENT.

The reasoner was shown a conversation up to one turn and asked a
↪  yes/no question: should this message be inserted here, now? It
↪  could not see what comes next. It said yes. You are given its
↪  justification.

The distinction that matters is between arguing that the CONTENT FITS
↪  and arguing that THIS MOMENT is right. They are different claims,
↪  and a justification can make one without the other.

LABELS:
- timing_argued: Argues that THIS MOMENT is the right one - that
↪  something just happened which makes now, rather than earlier or
↪  later, the time.
- relevance_only: Argues the message is useful, on-topic or a good fit
↪  here, without arguing that this moment is better than a later one.
- waiting_considered: Explicitly weighs holding off - names a later
↪  moment that might be better, or says the topic may develop
↪  further.
- urgency: Argues the window is closing: the conversation is about to
↪  move on, or a decision is about to harden.
- prerequisite_met: Argues a condition has now been satisfied that was
↪  not satisfied earlier - the topic has been raised, the question
↪  asked, the claim made.
- no_reason: States the decision with no justification.

Rules:
- "rationales" must contain ONLY exact label keys from the list above,
↪  verbatim and lowercase, e.g. "relevance_only". Never put a
↪  sentence, an explanation, or a quote in that array. Explanations
↪  belong in "evidence".
- Assign every label that applies. `relevance_only` and
↪  `timing_argued` are the important pair and are mutually exclusive.
```

```
- Assign `relevance_only` when the justification only establishes that
↪  the message is useful, on-topic, or answers something - without
↪  arguing that later would be worse.
- Assign `timing_argued` when it gives a reason THIS turn beats a
↪  later one.
- `no_reason` is exclusive.
- "evidence" is at most 20 words quoted from the justification.

Example of the required shape:
{"rationales": ["relevance_only"], "primary": "relevance_only",
↪  "argues_timing": false, "evidence": "directly answers the question
↪  about price"}

Return ONLY JSON:
{"rationales": ["..."], "primary": "...", "argues_timing": true,
↪  "evidence": "short quote"}
```

User message.

```
CONVERSATION, UP TO AND INCLUDING TURN {checkpoint}:
{conversation_excerpt}

MESSAGE TO ADD:
{message_text}

DECISION: insert it here, after turn {checkpoint}

JUSTIFICATION:
{justification}
```