From *What to Respond* to *When to Respond*: Timely Response Generation for Open-domain Dialog Agents

Anonymous ACL submission

Abstract

While research on dialog response generation has primarily focused on generating coherent responses conditioning on textual context, the critical question of *when to respond* grounded on the temporal context remains underexplored. To bridge this gap, we propose the task of timely dialog response generation and introduce the TIMELYCHAT benchmark to evaluate two key aspects: response timing and time-conditioned responses. Additionally, we construct a large-scale training dataset by leveraging unlabeled event knowledge from a temporal commonsense knowledge graph and employing a Large Language Model (LLM) to synthesize 55K event-driven dialogs. We then train TIMER, a dialog agent designed to proactively predict time intervals and generate timely responses that align with those intervals. Experimental results show that TIMER outperforms instruction-tuned LLMs and other time-aware baselines in both turn-level and dialog-level evaluations. We publicly release our data, model, and code.[1]

Figure 1: An illustrative example of a timely dialog agent in an asynchronous messaging setting. Unlike a delay-agnostic agent that responds instantly with repetitive turns ignoring the ongoing event (walking the dog), a timely dialog agent predicts an appropriate delay regime and re-engages with a response conditioned on that delay.

1 Introduction

Open-domain dialog systems have introduced various aspects to improve human-likeness of their own, such as persona (Zhang et al., 2018; Ahn et al., 2023), long-term memory (Xu et al., 2022a,b), commonsense (Zhou et al., 2021; Qin et al., 2021), emotional support (Rashkin et al., 2019; Liu et al., 2021; Zhang et al., 2024), and role-play (Shao et al., 2023; Li et al., 2023). These efforts have led to the success of commercial chat services like Replika and Character AI, which have met the public's demand for social companion chatbots (Chaturvedi et al., 2023; Guingrich and Graziano, 2023). The pursuit for human-like chatbots still remains important with the remarkable advancements in Large Language Models (LLMs), along with the growing demand for AI agents capable of engaging in more natural and human-like ways.

Research on dialog response generation has predominantly focused on generating appropriate and consistent next utterances, conditioning on the textual information within dialog contexts. Meanwhile, although the question of *what to respond* has received considerable attention, the problem of *when to respond* remains underexplored, despite its crucial role in enabling real-time dialog agents to appropriately ground their responses on the temporal contexts regarding the status of ongoing conversational events. For instance, as illustrated in Figure 1, if an agent generates only instant responses without considering response timing, it

[1] https://anonymous.4open.science/r/timelychat-EADE/

can cause repetitive interactions without conversational progress or produce responses that do not align with the temporal context of the conversational event. In contrast, by incorporating response timing, an agent can maintain a natural flow while providing timely responses. *This requires grounding responses on the temporal context tied to the status of the event, mirroring the way humans naturally adapt their responses in human-to-human conversations.* Importantly, this is not an edge case: across large-scale real messaging corpora, roughly 40% of replies are delayed rather than instant (Appendix A), yet existing dialog systems implicitly assume instantaneous turn-taking. We stress that the goal is not to send a response at an exact predicted timestamp, but to decide an appropriate delay regime: whether a turn should be instant or delayed, and approximately how long to wait, which then conditions the generated response.

However, it is inherently challenging to simulate such scenarios with dialog models trained on existing datasets. Most dialog datasets lack explicit temporal context and are created under the tacit assumption that interactions occur instantly. Additionally, collecting real-time conversations where temporal context is naturally embedded (e.g., text messages between individuals) is highly restricted due to privacy concerns and ethical considerations.

In this work, we propose a novel task named **Timely Dialog Response Generation**, which aims to generate not only coherent responses but also to consider the temporal context associated with ongoing events. Specifically, it focuses on predicting the necessary time interval for the next utterance and generating a corresponding time-conditioned response. We introduce TIMELYCHAT dataset and propose a benchmark to assess two key aspects: response timing prediction and time-conditioned response generation. To create diverse event-driven dialogs, we combine the human-annotated event-duration pairs from a temporal commonsense knowledge graph with the powerful dialog generation capability of an LLM.

Furthermore, we introduce a large-scale dataset comprising 55K event-driven dialogs for supervised fine-tuning (SFT). To address the challenges of costly and labor-intensive manual annotation, we use unlabeled event sources from a large-scale temporal commonsense knowledge graph and leverage an LLM to pseudo-label event durations and synthesize diverse event-driven dialogs. Using this dataset, we present TIMER, a dialog model fine-tuned with a multi-task learning objective that jointly predicts the time interval and generates the corresponding response.

Evaluation results on the proposed benchmark demonstrate that TIMER outperforms both instruction-tuned LLMs and dialog models trained on time-aware datasets in generating time-conditioned responses and predicting time intervals consistent with temporal commonsense. Furthermore, in dialog-level evaluations, TIMER distinguishes between situations requiring delayed responses and those requiring instant responses more effectively, and generates more timely responses that align well with the predicted time intervals. We publicly release the TIMELYCHAT benchmark, training data, and our timely dialog agent named TIMER to facilitate further research in this area.

2 Related Work

Long-term dialogs. Long-term dialog spans multiple sessions separated by time intervals. Multi-Session Chat (MSC) (Xu et al., 2022a) introduces messaging-like multi-session chat; Conversation Chronicles (Jang et al., 2023) adds a wider range of time intervals and fine-grained speaker relationships; and LoCoMo (Maharana et al., 2024) extends to very long-term dialog (up to 32 sessions) with a memory benchmark. However, these focus on recalling personas or past events across sessions, rather than the temporal context between ongoing events in real-time conversations. A notable exception, GapChat (Zhang et al., 2023), introduces an event timeline to capture event progression over given time intervals. Our work moves beyond predetermined intervals, requiring a proactive agent that dynamically determines realistic delays from temporal context.

Proactive and time-aware dialog agents. Recent work increasingly treats response timing and proactivity as first-class problems. One line models *when* an agent should speak: Liao et al. (2025) cast turn-taking as a multimodal decision of when to speak in spoken interaction, while Yang et al. (2025b) introduce a time-aware persona dialog benchmark with a temporal next response prediction task. A parallel line builds proactive assistants that act without explicit user prompts (Chen et al., 2025; Yang et al., 2025a) and benchmarks to evaluate them (Wang et al., 2025; Yoon et al., 2025), with Wang et al. (2025) proposing a time-aware metric for proactive interactions. Relatedly, Cheng

et al. (2026) show that LLM agents are largely temporally blind, failing to account for the time elapsed between messages. Unlike these efforts which target spoken turn-taking, multimodal perception, or task assistance, we focus on text-based open-domain dialog and jointly model both *when* to respond and *what* to respond within a single agent.

3 Task Definition

We introduce a new task named **Timely Dialog Response Generation**, which aims to generate contextually appropriate responses while incorporating temporal considerations from the dialog history. A key temporal factor that influences a response is how much time has passed since the previous utterance. To capture this, we define *time interval* as our primary temporal context, which represents the relative time difference (e.g., 10 minutes) between utterances. Formally, we model the conditional probability distribution P_θ of a response r_t at t-th turn given the textual context u and the temporal context τ:

$$r_t \sim P_\theta(u_t|u_1, \ldots, u_{t-1}, \tau_2, \ldots, \tau_t), \quad (1)$$

where u_t denotes the utterance at t-th turn and τ_t denotes the elapsed time between u_{t-1} and u_t. This probability distribution can be further decomposed into two subtasks, which are the main focus of this work.

Subtask 1. Response Timing Prediction The first task is to predict the optimal timing for delivering messages to users. Mathematically, this involves predicting the t-th time interval $\hat{\tau}_t$ given the available contexts:

$$\hat{\tau}_t \sim P_\theta(\tau_t|u_1, \ldots, u_{t-1}, \tau_2, \ldots, \tau_{t-1}). \quad (2)$$

Subtask 2. Time-conditioned Response Generation The subsequent task is to generate a contextually appropriate response while incorporating the predicted timing for message delivery:

$$r_t \sim P_\theta(u_t|u_1, \ldots, u_{t-1}, \tau_2, \ldots, \tau_{t-1}, \hat{\tau}_t). \quad (3)$$

Note that this task formulation challenges the widely held assumption that dialog agents should always respond to user messages instantly. Instead, it takes temporal context into account, i.e., the amount of elapsed time, to determine when a response should be generated.

4 TIMELYCHAT Benchmark

We construct TIMELYCHAT benchmark to assess the timely response generation capabilities of dialog models. To this end, we first craft high-quality timely conversations through temporal knowledge base and LLMs and then design two evaluation processes. Figure 2 shows the overall construction process of our benchmark. We release TIMELYCHAT in two parts built by the same recipe: a high-quality benchmark of 324 manually screened test dialogs (this section) and a large-scale 55K training set (Section 5.1), differing only in event source and generator.

4.1 Data Construction

We incorporate temporal information into dialogs using a temporal commonsense knowledge base. This knowledge base captures various event-related temporal dynamics which is well suited for transforming temporal context into event-driven dialogs. By identifying temporal patterns, we seamlessly integrate them into conversations, using the sophisticated dialog generation capabilities of LLMs.

Event Knowledge Extraction. We first obtain a rich and reliable source of daily events and their typical durations for crafting event-driven conversations with temporal context. To this end, we use the event duration category of MC-TACO dataset (Zhou et al., 2019). The dataset consists of sentences for specific events, queries to ask the typical duration of the event (e.g., "How long does it take to . . . ?"), and human-annotated ground-truth answers. We use the sentences with ground-truth answers, i.e., event-duration pair, to synthesize event-driven conversations. During data construction, we excluded the examples whose temporal intervals shorter than one minute or longer than 24 hours to simulate realistic temporal delay in daily dialog situations. Lastly, we instruct GPT-4o (Achiam et al., 2023) with these sentences and event-duration pairs to generate descriptive sentences. It integrates the event and its duration into coherent sentences, forming seed narratives for dialog generation.

Timely Dialog Generation. With the extracted temporal event knowledge, we instruct GPT-4o to generate conversations. Our instruction contains the conditions that the generated dialogs must satisfy:

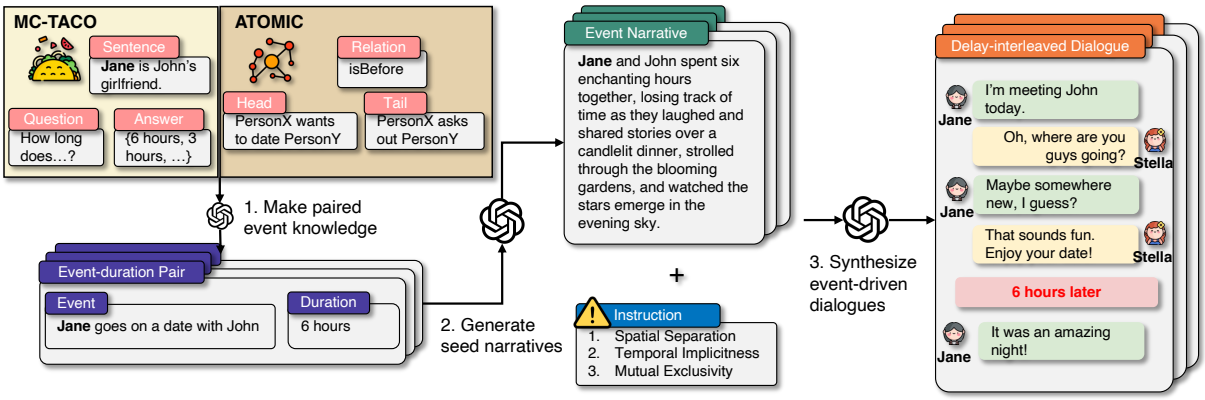

Figure 2: Overall process of data construction method. Two different knowledge sources represent the same example for better understanding. Note that due to constraints imposed by the instruction, Jane's conversation partner becomes Stella, not John.

- **Spatial Separation**: The scenario must involve one speaker experiencing an event while conversing with another speaker about it. This ensures there are no contradictions arising from both speakers being in the same spatial context.

- **Temporal Implicitness**: The response must avoid direct references to the elapsed time. This condition reduces the occurrence of dull responses that simply acknowledge the time interval and, more importantly, prevents lexical overlap with the ground-truth time interval, which could create shortcuts in the generation process.

- **Mutual Exclusivity**: The time-conditioned response must become untimely under contrary temporal conditions. In other words, a delayed response should be incoherent under an instant condition with no time interval, and an instant response should be incoherent when a time interval exists. It prevents generating time-agnostic responses that remain coherent regardless of the temporal context.

Along with these instructions, we provide GPT-4o with one in-context demonstration, randomly sampled from six seed dialogs that the authors wrote by hand, each ranging from 5 to 10 turns. These seeds serve only to prevent ill-formed outputs and diversify dialog lengths; they are not part of the released data, and we found the choice of demonstration among the six to have negligible effect on output quality. After manual inspection and filtering out low-quality dialogs that did not meet all the conditions, the final synthesized benchmark consists of 324 dialogs, with an average length of 7.5 turns. All prompts and examples used in the construction process are provided in Appendix B.

4.2 Evaluation Protocols

With the crafted conversations, we propose two evaluation approaches to assess the abilities of dialog agents to generate timely responses: turn-level and dialog-level.

Turn-level Evaluation. In turn-level evaluation, we assess each subtask on the target response. For response timing prediction, a model predicts the time interval required for the next utterance given a dialog context. We then evaluate (1) whether the model correctly classifies the next turn as either delayed or instant, and (2) how close is the predicted interval to the ground truth. We measure precision, recall, false positive rate (FPR), and F1 for the binary classification, and root mean squared logarithmic error (RMSLE) for regression by converting each time interval into minutes. For response generation, a model generates a time-conditioned response given a dialog context and ground-truth time interval. We measure BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and BERTScore (Zhang et al., 2020) as reference-based metrics. Additionally, we measure naturalness (Mehri et al., 2022) and time-specificity (Tsunomori et al., 2023) on a 5-point scale, adopting LLM-as-a-judge (Liu et al., 2023) for automatic evaluation.

Dialog-level Evaluation. One crucial quality of a timely dialog agent is its ability to introduce appropriate delays considering the temporal context while maintaining a natural conversational flow. Inspired by dialog-level evaluation methods with

model-to-model interactions (Li et al., 2019; Zhou et al., 2024), we provide an event-driven scenario and let an agent converse with GPT-4o as a user simulator (Yoon et al., 2024; Kazi et al., 2024; Wang et al., 2024) for the fixed number of turns to measure dialog-level metrics. We measure coherence (Mehri et al., 2022) and dialog-level time-specificity to assess the quality of the agent's responses, and measure delay-appropriateness that considers both the timing and duration of delays, using LLM-as-a-judge with a 5-point scale. The evaluation criteria of LLM-as-a-judge metrics and simulator instructions are detailed in Appendix C.

5 TIMER: A Dialog Agent for Timely Responses

5.1 Training Data Augmentation with Unlabeled Knowledge

To scale beyond the 324-dialog benchmark, we construct a large training set using the same recipe and dialog conditions as Section 4.1, changing only the inputs to favor scale over manual curation.

Knowledge source and durations. Instead of MC-TACO, we draw events from the ATOMIC$^{20}_{20}$ knowledge graph (Hwang et al., 2021), which provides far more event triplets but no annotated durations. We prompt GPT-3.5 to convert each triplet into a single-sentence event description, replacing anonymized names (e.g., PersonX) with common names following Kim et al. (2023), and to estimate the event's typical duration as a pseudo-label, discarding durations under a minute or over 24 hours.

Generation at scale. We then generate dialogs with GPT-3.5 (rather than GPT-4o) for cost efficiency, using the same instructions and conditions as the benchmark. To curb ill-formed outputs (e.g., speaker mismatches) without manual screening, each prompt includes a one-shot demonstration randomly sampled from the benchmark. The resulting set comprises 55K examples as shown in Table 1, offering far more event-grounded dialogs at finer time granularity than prior datasets, without costly human annotation. All prompts are provided in Appendix B.1.

5.2 Time-augmented Training with Multi-task Learning Objectives

Our training approach predicts an appropriate time interval for delaying the response based on the temporal context and then generates a time-conditioned response, introducing a time interval prediction step before each turn's utterance (Eqs. 2 and 3). For each turn consisting of a speaker identifier and a text utterance, we insert a time interval. We prepend prefix tokens to distinguish each component, formatting the input as `<SPK>` s_t `<TIME>` τ_t `<UTT>` u_t, where s_t, τ_t, and u_t denote the speaker, the time interval, and the utterance at the t-th turn, respectively. For turns within the dialog context, we set $\tau = 0$, indicating no delay, maintain coherence and align with typical instant responses.

From these inputs, we define two losses for training: response timing prediction loss and response generation loss. The losses are defined as follows:

$$\mathcal{L}_{\text{time}} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{t=2}^{T}\log p(\tau_t \mid s_{\leq t}, \tau_{<t}, u_{<t}),$$
$$\mathcal{L}_{\text{response}} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{t=2}^{T}\log p(u_t \mid s_{\leq t}, \tau_{\leq t}, u_{<t}), \tag{4}$$

where N is the number of training examples, and T is the number of turns in a dialog.

The final multi-task learning objective is given as follows:

$$\mathcal{L} = \mathcal{L}_{\text{response}} + \lambda\mathcal{L}_{\text{time}}. \tag{5}$$

This approach ensures that the model learns both to predict appropriate time intervals and to generate time-conditioned responses effectively.

6 Experiments

6.1 Baselines

We evaluate two types of dialog agents for simulating timely dialog response generation: instruction-tuned models and time-aware models. The instruction-tuned models include LLMs optimized for dialog use cases. We select 8B and 70B models of LLaMA 3.1 Instruct (Dubey et al., 2024) as open-source chat models, and GPT-3.5 and GPT-4o as proprietary models. We experiment with zero-shot, few-shot, and chain-of-thought (CoT) (Wei et al., 2022) prompting strategies to investigate the effectiveness of in-context learning without task-specific fine-tuning. The time-aware models are trained on dialog datasets where time intervals are interleaved. We compare the following models:

- **MSC 3B** (Xu et al., 2022a): Trained on BlenderBot (Roller et al., 2021) using the

Dataset	# Sessions	Construction Method	Time Granularity	Event-grounded	# Events
MSC (Xu et al., 2022a)	13K (4.4K)	Crowdsourcing	hours - weeks	✗	-
CC (Jang et al., 2023)	1M (160K)	LLM-generated	hours - years	✗	-
LoCoMo (Maharana et al., 2024)	842 (-)	LLM-gen + Crowd	days - months	✗	-
GapChat (Zhang et al., 2023)	2.6K (782)	LLM-gen + Crowd	minutes - years	✓	128
TIMELYCHAT (Ours)	55K	LLM-generated	minutes - hours	✓	55K

Table 1: Comparison of long-term dialog datasets interleaved with time intervals. The number in parentheses under the # Sessions column represents the count of sessions with time intervals within a day. Event-grounded indicates whether the dialogs reflect the temporal context associated with events or not. For TIMELYCHAT, # Sessions refers to the 55K training dialogs; the benchmark additionally provides 324 manually screened test dialogs.

MSC dataset, which includes time intervals between sessions.

- **ReBot 400M** (Jang et al., 2023): Trained on BART-Large (Lewis et al., 2020) using the CC dataset, which consists of large-scale LLM-generated dialogs.
- **GapChat 3B** (Zhang et al., 2023): Trained on MSC using the GapChat dataset, which incorporates event progress based on time intervals.

To prevent a judge LLM from introducing potential biases such as self-preference bias (Wataoka et al., 2024; Panickssery et al., 2024), we use Claude Sonnet 4.5 as the evaluator across both turn-level and dialog-level evaluations. Implementation details of all models including TIMER 3B are described in Appendix C.1.

Model	Precision ↑	Recall ↑	F1↑	FPR↓	RMSLE↓
LLaMA 3.1 8B					
Zero-shot	0.172	0.691	0.276	0.606	2.853
Few-shot	0.164	**0.920**	0.279	0.855	2.807
CoT	0.160	0.824	0.268	0.790	2.854
LLaMA 3.1 70B					
Zero-shot	0.153	0.735	0.254	0.740	2.479
Few-shot	0.197	0.602	0.297	0.448	2.326
CoT	0.194	0.836	0.315	0.636	3.066
GPT-3.5					
Zero-shot	0.143	0.784	0.241	0.861	2.763
Few-shot	0.212	0.349	0.264	0.237	2.146
CoT	0.186	0.762	0.299	0.609	2.667
GPT-4o					
Zero-shot	0.266	0.259	0.263	0.131	1.956
Few-shot	0.227	0.423	0.295	0.099	2.252
CoT	0.202	0.852	0.326	0.615	2.938
TIMER 3B (Ours)	**0.783**	0.799	**0.791**	**0.041**	**1.189**

Table 2: Results of response timing prediction. For few-shot and CoT strategies, we provide balanced 2-shot demonstrations which consist of one delayed example and one instant example, along with the task description used in zero-shot prompting.

6.2 Turn-level Evaluation Results

Response Timing Prediction. Table 2 presents the results of response timing prediction on the TIMELYCHAT. Overall, instruction-tuned models show low precision and F1 with high FPR, indicating that they over-predict delays and introduce unnecessary intervals that disrupt the conversational flow. Although few-shot and CoT strategies slightly improve F1 scores across all LLMs, they sometimes negatively impact FPR compared to zero-shot prompting. In contrast, TIMER 3B achieves the highest F1 score and the lowest FPR compared to instruction-tuned models. Even the best-performing GPT-4o still lags significantly behind the fine-tuned TIMER 3B model.

Likewise, when it comes to predicting the length of time intervals, in-context learning methods fail to enhance performance effectively. While few-shot prompting achieves a lower RMSLE than CoT across all LLMs, it does not consistently outperform zero-shot prompting, as demonstrated by GPT-4o's results. These findings indicate that prompting with task descriptions and demonstrations alone is insufficient to reliably predict whether to pose a delay and how long it should last. In contrast, task-specific fine-tuning is essential for effectively learning these capabilities.

Time-conditioned Response Generation. Table 3 shows the time-conditioned response generation performance on the TIMELYCHAT. For instruction-tuned models, we observe that zero-shot performance tends to improve as model size increases across all metrics. Among all LLMs, few-shot prompting consistently outperforms zero-shot prompting, while CoT prompting performs the worst in terms of naturalness and time-specificity. This aligns with previous findings that LLMs struggle to generate helpful CoT rationales for dialog response generation (Chae et al., 2023).

Model	B-2	R-L	BS	Nat.	Spec.
INSTRUCTION-TUNED MODELS					
LLaMA 3.1 8B					
Zero-shot	5.38	12.38	86.21	2.24	2.78
Few-shot	7.63	13.47	86.85	2.81	2.83
CoT	6.17	12.50	86.23	1.31	1.84
LLaMA 3.1 70B					
Zero-shot	6.84	12.71	85.90	2.86	2.35
Few-shot	8.35	14.83	87.16	3.92	2.89
CoT	9.01	15.01	87.12	1.07	3.04
GPT-3.5					
Zero-shot	9.97	17.13	87.54	4.26	2.86
Few-shot	11.23	17.81	87.77	4.25	2.76
CoT	8.86	15.14	86.79	2.38	2.15
GPT-4o					
Zero-shot	9.17	16.76	87.35	4.75	3.09
Few-shot	11.15	18.51	87.91	**4.78**	3.09
CoT	10.25	17.13	87.52	3.21	2.84
TIME-AWARE MODELS					
MSC 3B	3.26	8.94	85.18	1.36	1.02
ReBot 400M	3.85	9.32	85.72	3.26	1.02
GapChat 3B	3.59	8.61	85.22	1.87	1.05
TIMER 3B (Ours)	**16.08**	**22.26**	**88.74**	3.84	**3.36**

Table 3: Results of time-conditioned response generation on TIMELYCHAT. B-2, R-L, BS, Nat., and Spec. refer to BLEU-2, ROUGE-L, BERTScore, naturalness, and time-specificity, respectively.

Meanwhile, models trained on existing dialog datasets that include time intervals exhibit poor overall performance. Notably, these models achieve low time-specificity, indicating that they struggle to generate timely responses conditioned on given time intervals. This stems from the characteristics that time intervals in existing long-term dialog datasets are assigned arbitrarily rather than based on the temporal context of ongoing events, making it difficult for models to learn the conditional distribution of responses based on the given interval. For example, we find that these models frequently generate generic greeting messages, failing to capture the temporal nuances of timely responses. In contrast, TIMER 3B, despite having a smaller model size, achieves comparable naturalness to instruction-tuned LLMs and even surpasses LLaMA 3.1 models. More importantly, it achieves the highest time-specificity, demonstrating that our training approach enables response generation that aligns well with event-specific temporal contexts.

6.3 Dialog-level Evaluation Results

Beyond turn-level evaluation, we also conduct dialog-level evaluation to assess whether a dialog agent can introduce temporally contextual delays at appropriate timing without disrupting the conversational flow. We let the four zero-shot LLMs from the previous experiments, along with TIMER 3B, engage in 10 interactions with the simulator described in Section 4.2. To simulate event-driven dialog, we provide the first turn of conversations from the TIMELYCHAT as the initial interaction.

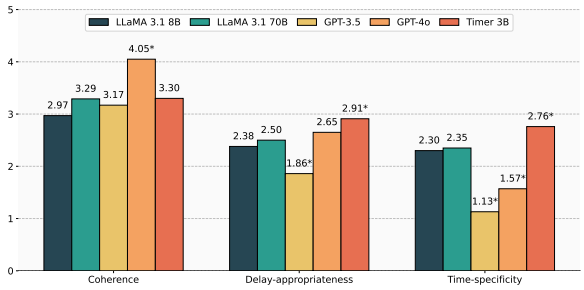

Figure 3: Results of simulated interactions for timely dialog agents. We perform pairwise t-test and denote statistically significant score differences from the other models with asterisk (*) ($p < 0.05$).

We sample dialogs that include at least one delayed response and report three dialog-level metrics in Figure 3. GPT-4o achieves the highest coherence among the models, demonstrating its ability to maintain a natural conversation flow, while TIMER 3B achieves the second-highest coherence score. Notably, TIMER 3B shows the highest delay-appropriateness and time-specificity scores. This suggests that TIMER 3B effectively considers both dialog context and temporal context to predict delays with appropriate timing and duration. Additionally, it generates delayed responses that are coherent only when a delay is given, thereby justifying and necessitating the delay. In contrast, LLaMA 3.1 8B and 70B exhibit relatively lower delay-appropriateness, while GPT-3.5 and GPT-4o achieve lower time-specificity scores. We further analyze these findings in the case study in Appendix D.

6.4 Human Evaluation Results

To investigate the reliability of LLM-based evaluation, we also conduct human evaluations on both turn-level and dialog-level metrics. We recruit twenty graduate students as annotators, who are fluent in English but have no background on our work. We provide them with the same evaluation criteria used for LLM-based evaluation, and ask them to perform blinded comparisons between responses or dialog from two different models. Each

Setting	Metric	κ	Win	Tie	Loss
Turn-level	Naturalness	0.446	19%	24%	**57%**
	Time-specificity	0.508	**54%**	24%	22%
Dialog-level	Coherence	0.637	21%	18%	**61%**
	Delay-appropriateness	0.446	**40%**	23%	37%
	Time-specificity	0.496	**40%**	23%	37%

Table 4: Pairwise human evaluation results on both turn-level and dialog-level metrics. Win/Tie/Loss rates of TIMER 3B against zero-shot GPT-4o are presented. Fleiss' κ is also reported to measure inter-annotator agreements.

Method	Timing Prediction		Response Generation		
	F1	RMSLE	R-L	Nat.	Spec.
TIMER 3B	**0.791**	**1.189**	**22.26**	**3.84**	**3.36**
w/o Special Tokens	0.299	3.700	15.25	1.39	2.37
Utterance First	0.249	3.451	12.07	1.57	1.81
w/o Time Loss	0.299	3.987	13.38	2.54	2.43

Table 5: Ablation study on each component of our training scheme. In Utterance First setting, we train a model to generate an utterance first and then predict the response timing conditioned on it.

example is evaluated three times by different annotators, and aggregated by majority voting. Table 4 presents the results for both turn-level and dialog-level metrics.

For turn-level metrics, we compare TIMER 3B with the most competitive baseline, zero-shot GPT-4o, on 200 randomly sampled examples. While TIMER 3B falls short of GPT-4o in terms of naturalness, it slightly outperforms GPT-4o in time-specificity, which is consistent with the LLM-based evaluation results observed in Table 3.

For dialog-level metrics, we compare TIMER 3B with zero-shot GPT-4o on 80 randomly sampled examples. Again, TIMER 3B lags behind GPT-4o in coherence, but it outperforms GPT-4o in both delay-appropriateness and time-specificity. This finding aligns with the results shown in Figure 3, indicating that the proposed evaluation criteria and LLM-based evaluation are reliable measures for assessing desired model behavior. We observe that the annotators show substantial ($0.6 < \kappa \leq 0.8$) agreement on coherence, and moderate ($0.4 < \kappa \leq 0.6$) agreements on the other metrics, demonstrating the reliability of human evaluation. To further verify that our LLM-based evaluation is not an artifact of a single judge, we additionally evaluate the same samples with another judge model (GPT-4o) and report inter-judge consistency in Appendix C.5.

6.5 Ablation Study

To demonstrate the effectiveness of the proposed training scheme, we provide the impacts of each component on task performance in Table 5.

Effects of Input Formats. Special tokens (e.g., <UTT>, <TIME>) play a crucial role in indicating which component of an utterance to generate. Although trained with both time intervals and responses, the model trained without them struggles to predict time intervals at the right place, yielding low F1 and high RMSLE. This results in arbitrary delayed or instantaneous responses, causing poor controllability and also degrading response generation quality.

The order of performing the two subtasks significantly affects performance. When trained to generate a response first and then predict the interval conditioned on it, the model tends to predict delays even for instantaneous responses, resulting in poor timing prediction. This format also yields the lowest ROUGE-L among the settings, since the model is not provided with time conditions directly for generation.

Effects of Timing Prediction Loss. We investigate the effectiveness of our multi-task objective. When the timing prediction loss is removed ($\lambda = 0$), the model struggles at learning the timing prediction itself, resulting in random prediction. Notably, response generation performance degrades compared to TIMER 3B, even though the model is still optimized with time-conditioned context inputs. This implies that the timing prediction objective helps improve response generation capability, demonstrating the necessity of the proposed training objective.

7 Conclusion

We highlighted the necessity for open-domain dialog agents to consider not only the response itself but also its timing, grounded on the temporal context of the conversational event. We formulated this as the timely dialog response generation and introduced the TIMELYCHAT benchmark for turn-level and dialog-level evaluations. We also proposed a large-scale SFT dataset and a time-augmented training approach to train TIMER 3B, which proactively predicts the time interval and then generates a time-conditioned response, outperforming baselines while maintaining natural flow.

8 Limitations

Synthetic data and real-world grounding. Both our benchmark and training set are synthetic. Our real-corpus analysis (Appendix A) shows that delayed responding is pervasive but does not isolate the event-grounded delays we target, which coexist with unavailability, platform characteristics, and social norms. We therefore synthesize this subset under controlled constraints and leave verification on real-world dialog to future work.

Scope and time granularity. Our scope is intentionally event-driven, where temporal context most directly governs timing, so generalization to non-event settings remains open. We also model event duration as discrete values (e.g., 30 minutes) rather than continuous ranges (e.g., 2–6 hours), and interval or distributional prediction is a natural extension for finer-grained control.

Model and evaluation. TIMER 3B is strongest on time-specificity and delay-appropriateness but lags behind the much larger GPT-4o in naturalness and coherence, and response timing is only an observable proxy for an interlocutor's latent state. Fluency-oriented objectives (e.g., LLM-judge rewards) and larger model scales are promising next steps. Finally, our dialog-level evaluation relies on short simulated dialogs, and longer, more diverse interactions could further strengthen the analysis.

9 Ethics Statements

The proposed dataset was designed to assess capabilities related to response timing and time-conditioned response in event-driven conversations. To achieve this, we used event knowledge from publicly available datasets from various sources, and LLM-generated contents either without or with some modification if necessary. During this process, there is a possibility that harmful content or inappropriate biases existing in the original data may have been conveyed, or may have arisen due to limitations of filtering techniques. We reject any form of violence, discrimination, or offensive language, and our dataset and experimental results do not represent such values. If any harmful content or privacy infringement is identified within our data, we kindly request immediate notification to the authors. In the event of such cases being reported, we will apply the highest ethical standards and take appropriate actions.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Jaewoo Ahn, Yeda Song, Sangdoo Yun, and Gunhee Kim. 2023. MPCHAT: Towards multimodal persona-grounded conversation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3354–3377, Toronto, Canada. Association for Computational Linguistics.

Hyungjoo Chae, Yongho Song, Kai Ong, Taeyoon Kwon, Minjin Kim, Youngjae Yu, Dongha Lee, Dongyeop Kang, and Jinyoung Yeo. 2023. Dialogue chain-of-thought distillation for commonsense-aware conversational agents. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5606–5632, Singapore. Association for Computational Linguistics.

Rijul Chaturvedi, Sanjeev Verma, Ronnie Das, and Yogesh K. Dwivedi. 2023. Social companionship with artificial intelligence: Recent trends and future avenues. *Technological Forecasting and Social Change*, 193:122634.

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

Yize Cheng, Arshia Soltani Moakhar, Chenrui Fan, Parsa Hosseini, Kazem Faghih, Zahra Sodagar, Wenxiao Wang, and Soheil Feizi. 2026. Your llm agents are temporally blind: The misalignment between tool use decisions and human time perception. *Preprint*, arXiv:2510.23853.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Rose E Guingrich and Michael SA Graziano. 2023. Chatbots as social companions: How people perceive consciousness, human likeness, and social health benefits in machines. *arXiv preprint arXiv:2311.10599*.

Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. 2021. (comet-) atomic 2020: On symbolic and neural commonsense knowledge graphs. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(7):6384–6392.

Jihyoung Jang, Minseong Boo, and Hyounghun Kim. 2023. Conversation chronicles: Towards diverse temporal and relational dynamics in multi-session conversations. In *Proceedings of the 2023 Conference*

on Empirical Methods in Natural Language Processing, pages 13584–13606, Singapore. Association for Computational Linguistics.

Taaha Kazi, Ruiliang Lyu, Sizhe Zhou, Dilek Hakkani-Tur, and Gokhan Tur. 2024. Large language models as user-agents for evaluating task-oriented-dialogue systems. *Preprint*, arXiv:2411.09972.

Hyunwoo Kim, Jack Hessel, Liwei Jiang, Peter West, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Bras, Malihe Alikhani, Gunhee Kim, Maarten Sap, and Yejin Choi. 2023. SODA: Million-scale dialogue distillation with social commonsense contextualization. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12930–12949, Singapore. Association for Computational Linguistics.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.

Cheng Li, Ziang Leng, Chenxi Yan, Junyi Shen, Hao Wang, Weishi MI, Yaying Fei, Xiaoyang Feng, Song Yan, HaoSheng Wang, Linkang Zhan, Yaokai Jia, Pingyu Wu, and Haozhen Sun. 2023. Chatharuhi: Reviving anime character in reality via large language model. *Preprint*, arXiv:2308.09597.

Margaret Li, Jason Weston, and Stephen Roller. 2019. Acute-eval: Improved dialogue evaluation with optimized questions and multi-turn comparisons. *Preprint*, arXiv:1909.03087.

Zikai Liao, Yi Ouyang, Yi-Lun Lee, Chen-Ping Yu, Yi-Hsuan Tsai, and Zhaozheng Yin. 2025. Beyond words: Multimodal llm knows when to speak. *Preprint*, arXiv:2505.14654.

Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. Towards emotional support dialog systems. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3469–3483, Online. Association for Computational Linguistics.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-eval: NLG evaluation using gpt-4 with better human alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.

Ilya Loshchilov and Frank Hutter. 2019. Decoupled weight decay regularization. In *International Conference on Learning Representations*.

Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, and Yuwei Fang. 2024. Evaluating very long-term conversational memory of LLM agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13851–13870, Bangkok, Thailand. Association for Computational Linguistics.

Shikib Mehri, Jinho Choi, Luis Fernando D'Haro, Jan Deriu, Maxine Eskenazi, Milica Gasic, Kallirroi Georgila, Dilek Hakkani-Tur, Zekang Li, Verena Rieser, Samira Shaikh, David Traum, Yi-Ting Yeh, Zhou Yu, Yizhe Zhang, and Chen Zhang. 2022. Report from the nsf future directions workshop on automatic evaluation of dialog: Research directions and challenges. *Preprint*, arXiv:2203.10012.

Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2018. Mixed precision training. In *International Conference on Learning Representations*.

Arjun Panickssery, Samuel R. Bowman, and Shi Feng. 2024. Llm evaluators recognize and favor their own generations. In *Proceedings of the 38th International Conference on Neural Information Processing Systems*, NIPS '24, Red Hook, NY, USA. Curran Associates Inc.

Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Lianhui Qin, Aditya Gupta, Shyam Upadhyay, Luheng He, Yejin Choi, and Manaal Faruqui. 2021. TIME-DIAL: Temporal commonsense reasoning in dialog. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7066–7076, Online. Association for Computational Linguistics.

Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. Towards empathetic open-domain conversation models: A new benchmark and dataset. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.

Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. Recipes for building an open-domain chatbot. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online. Association for Computational Linguistics.

Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. Character-LLM: A trainable agent for role-playing. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13153–13187, Singapore. Association for Computational Linguistics.

Yuiko Tsunomori, Masakazu Ishihata, and Hiroaki Sugiyama. 2023. Time-considerable dialogue models via reranking by time dependency. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5136–5149, Singapore. Association for Computational Linguistics.

Xingguang Wang, Xuxin Cheng, Juntong Song, Tong Zhang, and Cheng Niu. 2024. Enhancing dialogue state tracking models through LLM-backed user-agents simulation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8724–8741, Bangkok, Thailand. Association for Computational Linguistics.

Yueqian Wang, Xiaojun Meng, Yifan Wang, Huishuai Zhang, and Dongyan Zhao. 2025. Proactivevideoqa: A comprehensive benchmark evaluating proactive interactions in video large language models. *Preprint*, arXiv:2507.09313.

Koki Wataoka, Tsubasa Takahashi, and Ryokan Ri. 2024. Self-preference bias in LLM-as-a-judge. In *Neurips Safe Generative AI Workshop 2024*.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the 36th International Conference on Neural Information Processing Systems*, NIPS '22, Red Hook, NY, USA. Curran Associates Inc.

Jing Xu, Arthur Szlam, and Jason Weston. 2022a. Beyond goldfish memory: Long-term open-domain conversation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5180–5197, Dublin, Ireland. Association for Computational Linguistics.

Xinchao Xu, Zhibin Gou, Wenquan Wu, Zheng-Yu Niu, Hua Wu, Haifeng Wang, and Shihang Wang. 2022b. Long time no see! open-domain conversation with long-term persona memory. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2639–2650, Dublin, Ireland. Association for Computational Linguistics.

Bufang Yang, Lilin Xu, Liekang Zeng, Kaiwei Liu, Siyang Jiang, Wenrui Lu, Hongkai Chen, Xiaofan Jiang, Guoliang Xing, and Zhenyu Yan. 2025a. Contextagent: Context-aware proactive LLM agents with open-world sensory perceptions. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Wanqi Yang, Yanda Li, Meng Fang, and Ling Chen. 2025b. MTPChat: A multimodal time-aware persona dataset for conversational agents. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 5815–5826, Albuquerque, New Mexico. Association for Computational Linguistics.

Se-eun Yoon, Zhankui He, Jessica Echterhoff, and Julian McAuley. 2024. Evaluating large language models as generative user simulators for conversational recommendation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1490–1504, Mexico City, Mexico. Association for Computational Linguistics.

Yejin Yoon, Yuri Son, Namyoung So, Minseo Kim, Minsoo Cho, Chanhee Park, Seungshin Lee, and Taeuk Kim. 2025. Beyond task-oriented and chitchat dialogues: Proactive and transition-aware conversational agents. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 13291–13317, Suzhou, China. Association for Computational Linguistics.

Qiang Zhang, Jason Naradowsky, and Yusuke Miyao. 2023. Mind the gap between conversations for improved long-term dialogue generation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10735–10762, Singapore. Association for Computational Linguistics.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. 2018. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2204–2213, Melbourne, Australia. Association for Computational Linguistics.

Tenggan Zhang, Xinjie Zhang, Jinming Zhao, Li Zhou, and Qin Jin. 2024. ESCoT: Towards interpretable emotional support dialogue systems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13395–13412, Bangkok, Thailand. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Ben Zhou, Daniel Khashabi, Qiang Ning, and Dan Roth. 2019. "going on a vacation" takes longer than "going for a walk": A study of temporal commonsense understanding. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3363–3369, Hong Kong, China. Association for Computational Linguistics.

Pei Zhou, Karthik Gopalakrishnan, Behnam Hedayatnia, Seokhwan Kim, Jay Pujara, Xiang Ren, Yang Liu, and Dilek Hakkani-Tur. 2021. Commonsense-focused dialogues for response generation: An empirical study. In *Proceedings of the 22nd Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 121–132, Singapore and Online. Association for Computational Linguistics.

Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, and Maarten Sap. 2024. SOTOPIA: Interactive evaluation for social intelligence in language agents. In *The Twelfth International Conference on Learning Representations*.

A Prevalence of Delayed Responses in Real Conversations

To motivate the *when-to-respond* problem with real human behavior, we measure the distribution of the reply time gaps in large-scale, publicly available conversation corpora. We aggregate three real messaging corpora (a social messaging corpus, an online conversation corpus, and a messenger-style corpus) from AI Hub [2], comprising 12.7M reply turns in total. Note that these corpora are in a different language (Korean) from our benchmark. We use them as cross-lingual evidence that delayed responding is a general phenomenon in human messaging.

As summarized in Figure 4, while a large fraction of replies are effectively instantaneous (gap $= 0$ min), roughly 40% of replies are delayed (gap > 0 min), and about 24% remain delayed under a conservative threshold that treats gaps of up to one minute as instant. A separate large-scale messaging log of 25.6M time differences shows a consistent pattern, with 27% of replies delayed. These statistics establish that delayed responding is pervasive in real conversation: an agent that can only respond instantly is therefore a poor default model of human messaging behavior, which motivates modeling *when to respond*.

We emphasize what this analysis does and does not show. It establishes a necessary condition that response delays are common and thus worth modeling, but it does not attribute these delays to any single cause. Delays in raw corpora arise from many factors, including user unavailability, platform asynchrony, social norms, and the event-grounded situations that are the focus of this work. Since the event-grounded subset cannot be reliably isolated from raw logs by any mechanical filter, we instead synthesize it under controlled constraints that ensure a delay is warranted by the ongoing event rather than by generic availability.

B Data Construction Details

B.1 ChatGPT Prompts

We provide prompts used for data construction processes of both evaluation and training datasets. The contents within curly brackets represent the corresponding elements for each example. The statements used for ATOMIC^{20}_{20} duration estimation were constructed by concatenating the head and tail with a conjunction that represents each relation category. We present the prompts in the order of the processes, where the output of each step serves as the input for the next step.

[2] https://aihub.or.kr/

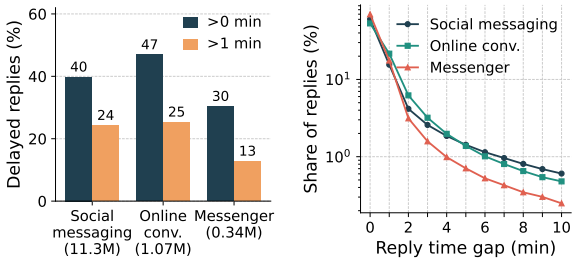

(a) Delayed replies (gap > 0 and > 1 min) per corpus.

(b) Reply gap distribution over 0-10 min (log scale).

Figure 4: Reply time gap statistics in real messaging corpora. Most replies are instant (0 min), yet a large minority are delayed, and the gap distribution has a long tail extending well beyond the instant bin.

ATOMIC^{20}_{20} Duration Estimation

You are given a statement about common events in our daily lives. Your task is to estimate the typical duration of the key event in the form of (quantity of time + unit) (e.g., seconds, minutes, hours, days, weeks, months, years, decades, or centuries) based on the temporal common sense of average humans.

[Examples]
Statement: After dinner, he went to look for Max one last time before he had to take a bath and go to bed.
Key event: having dinner
Duration: 1 hour

Statement: Jennie and Bryan boarded a 6:00 A.M. flight from Seoul to Los Angeles International Airport.
Key event: flight from Seoul to Los Angeles
Duration: 12 hours

Event: Carl Laemmle, head of Universal Studios, gave Einstein a tour of his studio and introduced him to Chaplin.
Key event: tour of his studio
Duration: 45 minutes
[End of Examples]

Statement: {statement}

MC-TACO Event Descriptions

You are given an event and a question and answer for the duration that denotes how much time is needed for the event to happen.
Write a story regarding the event in one sentence.

Sentence: {sentence}
Question: {question}
Answer: {duration}

ATOMIC$^{20}_{20}$ Event Descriptions

You are given a statement, the key event and the duration that denotes how much time is needed for the event to happen.
Write a story regarding the event in one sentence.

Statement: {statement}
Key event: {event}
Duration: {duration}

Dialog Generation

You are given an event narrative and the duration. Your task is to create an instant message dialog between two speakers. The following conditions MUST be met.

[Instructions]
1. Speaker {A,B} is in the middle of the event now, while speaker {B,A} is physically apart from.
2. Do not directly mention the duration in the dialog.
3. After {B,A}'s last turn, add "[{duration} later]", where duration is the amount of time passed in real world.
4-1. Generate {A,B}'s last message which is timely as if {A,B} spent time to finish the event.
4-2. In contrast, generate {A,B}'s last message as if {A,B} is responding instantaneously right before the event to happen.
Make sure that the timely response and the instantaneous response are time-situationally different.
[End of Instructions]

[Example]
{dialog example}
[End of Example]

Narrative: {event description}
Duration: {duration}

B.2 Few-shot Examples

We provide six author-written dialog examples randomly fed into GPT-4o as one-shot demonstrations when generating dialogs for TIMELYCHAT using the MC-TACO dataset.

5-turn Dialog

Narrative: After dinner, he took a shower before he went to bed.
Duration: 20 minutes

A: I finally got home. What a day!
B: It's eleven p.m. and you just got back home? It must be very tough day today.
A: Whooa Imma take a shower. I'm too tired.
B: Wash out all your fatigue with it.
[20 minutes later]
(delayed response)
A: I feel much better now! You didn't go to bed?
(instantaneous response)
A: How nice of you :) Give me a moment. brb

6-turn Dialog

Narrative: She has taken calculus class and she had a final exam.
Duration: 2 hours

A: Hey, what are you up to?
B: I'm gonna take the calculus final exam in 20 minutes. I feel so nervous.
A: You studied really hard, didn't you? I'm 100% sure you'll do well.
B: But the last two chapters were too difficult for me to understand.

A: That means others feel the same. Don't worry too much!
[2 hours later]
(delayed response)
B: It wasn't much harder than I expected. I hope I get a good grade.
(instantaneous response)
B: Thank you for cheering me up. I hope the exam is not that hard.

7-turn Dialog

Narrative: He enjoyed working out at the gym.
Duration: 1 hour 30 minutes

A: I'm going to the gym now. Wanna join me?
B: I don't feel like working out today. Sorry.
A: You don't feel good? What happened?
B: I played football so hard yesterday that I can't even walk right.
A: Okay, I understand. Maybe next time!
B: Enjoy your routine! I think I can make it tomorrow.
[2 hours later]
(delayed response)
A: I focused on my lower body today. Chest tomorrow?
(instantaneous response)
A: Gonna work out hard on my lower body. Chest tomorrow?

8-turn Dialog

Narrative: She had felt so tired that she went to bed right after the tv show.
Duration: 8 hours

A: Are you watching the saturday night live?
B: I'm watching it now but I'm too tired.
A: I didn't expect today's host is such a comedian lol
B: Yeah almost the end of the show. I feel like going to bed little bit early.
A: What made you so tired? You had any plan?
B: I went to an amusement park with my sister. We had a really good time there.
A: Oh I see. Think I should let you go. Sleep tight!
[8 hours later]
(delayed response)
B: Good morning. Did you sleep tight, too?
(instantaneous response)
B: Good night. I'll text you in the morning.

9-turn Dialog

Narrative: He took an intercity bus to get to his hometown.
Duration: 5 hours

A: What are you going to do on these holidays?
B: My parents and I usually have dinner together on the Eve.
A: Me too. So I'm heading to my town right now.
B: How do you get there? By bus or train?
A: I used to take trains, but I take an intercity bus for this time.
B: Why? the tickets' been already sold out?
A: Unfortunately yes... It will take little bit longer.
B: Have a nice trip though. Your family must be waiting for you.
[5 hours later]
(delayed response)
A: Finally I'm back at home! It took almost 5 hours.
(instantaneous response)
A: I'm gonna sleep all along in the bus. See you a few hours later.

10-turn Dialog

Narrative: She played the popular online game with her friends.
Duration: 30 minutes

A: Have you heard of the League of Legends?
B: Absolutely! I play it almost everyday with my classmates.
A: I've heard of, but I've never played if before.
B: We have a game soon. Wanna join us?
A: Isn't it a team game? I'm not a good

gamer though.
B: It's not a big deal. They will welcome you.
A: Well, maybe next time. I need to play it by myself first.
B: How about getting tutorial with me after this? I'll teach you.
A: Sounds good. Enjoy your game with your teammates.
[30 minutes later]
(delayed response)
B: We won! The game was nip and tuck. We were so close to losing.
(instantaneous response)
B: I'll be back just after the game. Wish me a good luck!

B.3 Time Interval Distribution

We semantically categorize the durations and present the distributions of time intervals in the two proposed datasets in Table 6 for a comprehensive comparison.

Time Interval	Train	Test
0-5 min (instant)	18.5%	5.6%
5-30 min (short)	42.3%	38.6%
30 min-2 hrs (moderate)	23.8%	32.1%
2-6 hrs (long)	2.9%	13.9%
6-12 hrs (half-day)	2.0%	5.2%
12-24 hrs (full-day)	10.5%	4.6%

Table 6: Time interval distribution of the augmented training dataset (Train) and TIMELYCHAT benchmark (Test).

C Evaluation Details

C.1 Implementation Details

Instruction-tuned Models. We use vLLM library[3] for the inference of LLaMA 3.1 Instruct 8B and 70B on 4 NVIDIA A100 80GB GPUs. All instruction-tuned models employ top-p sampling with temperature $T = 1.0$ and $p = 0.95$ during inference. We provide the prompts used for in-context learning methods on both response timing prediction and time-conditioned response generation below.

[3] https://docs.vllm.ai

Prompt for Response Timing Prediction

You are given a conversation between two speakers.
Your task is to estimate a time interval needed until the next response, considering the duration of the event in the conversation ranging from 0 minutes to 24 hours (1 day).
If the next response is expected to be immediate, you will output "0 minutes".
Otherwise, you will output a digit and a unit of time (e.g., 5 minutes, 2 hours).
Just output the time interval without any other text.

[Example n]
{few-shot OR CoT example}

Dialog context
{context}

Answer format: n (0<=n<=1440) minutes
The estimated time interval is:

Prompt for Time-conditioned Response Generation

You are given a conversation between two speakers and the elapsed time since the last utterance.
Your task is to generate the next response that aligns well with the temporal context represented by the time interval in parentheses.
Just output the response without any other text.

[Example n]
{few-shot OR CoT example}

Dialog context
{context}

Next response
{target speaker}: ({time interval} later)

Time-aware Models. We use Huggingface library[4] for the inference of MSC 3B and ReBot 400M[5]. We converted MSC 3B on the ParlAI

[4] https://huggingface.co
[5] https://huggingface.co/jihyoung/rebot-generation

framework[6] into a Huggingface checkpoint. We fine-tune GapChat 3B and TIMER 3B on each training data using the DeepSpeed library[7] with mixed precision training (Micikevicius et al., 2018). We train the models for 3 epochs using AdamW optimizer (Loshchilov and Hutter, 2019) with a learning rate of 1e-4, running on 2 NVIDIA A100 80GB GPUs for 9 hours. We use $\lambda = 1.0$ as a balanced scale factor of the two losses when training TIMER 3B. During inference, we apply beam search with the beam size of 3 and top-p sampling with $p = 0.95$.

C.2 User Simulator Prompts

We present the prompt fed into GPT-4o to create the user simulator.

User Simulator Prompt

You are a user simulator (user) engaging in an event-driven dialog with a dialog agent (agent).
Given the dialog context, your task is to proceed the conversation by one turn under the following assumptions:
1. agent responds after the elapsed time specified in the parentheses from the previous user utterance. If the delay is "0 minutes", agent is assumed to respond immediately.
2. user is assumed to respond to agent without any delay.

Conversation:
{context}

C.3 LLM-as-a-Judge Details

We elucidate the LLM-as-a-judge prompts used in turn-level and dialog-level evaluations, along with the evaluation criteria and steps for each metric.

Turn-level Prompt

You will be given a conversation between two individuals via messaging, along with the elapsed time since the last utterance. You will then be given on potential response for the next turn.
Your task is to rate the response on one metric. Please make sure you read and understand these instructions carefully.

Evaluation Criteria:
{metric} (1-5): {criteria}

Evaluation Steps:
{steps}

Dialog-level Prompt

You will be given a conversation between a dialog agent and a user.
Throughout the conversation, the agent proactively determines the delay of its response to the user's previous message, simulating delayed responses due to event experiences that take certain time to process.
At each agent's turn, the delay is provided in the parentheses followed by the message.
Your task is to rate the dialog agent on one metric. Please make sure you read and understand these instructions carefully.

Evaluation Criteria:
{metric} (1-5): {criteria}

Evaluation Steps:
{steps}

Evaluation Criteria and Steps

- **Naturalness** (1-5): the extent to which the response reads naturally given the dialog context.
  1. Assess the flow and coherence of the response in the conversation: Consider how seamlessly the response connects with the previous message.
  2. Evaluate the tone and style compatibility: Determine if the response's tone and style match those of the previous messages.
  3. Rate on a scale from 1 to 5, where 1 indicates the response is unnatural or inappropriate, and 5 indicates a perfectly natural continuation of the conversation.
- (Turn-level) **Time-specificity** (1-5): the extent to which the response ONLY makes sense when the specified time has passed, contrary to a time-agnostic response that makes sense regardless of time.
  1. Read the provided conversation and take note of the elapsed time since the previous message.
  2. Consider the context of the conversation, fo-

[6] https://parl.ai
[7] https://www.deepspeed.ai

cusing on how the passage of time might affect the relevance or appropriateness of the response.
3. Evaluate whether the potential response provided is time-specific. That is, determine if the response directly relates to or is clearly influenced by the elapsed time between the last utterance and the response.
4. Rate on a scale from 1 to 5, where 1 indicates the response is completely time-agnostic and unaffected by the passage of time, and 5 indicates the response is entirely time-specific; it only makes sense because of the amount of time that has passed since the previous message.

- **Coherence** (1-5): the extent to which the agent maintains a good conversation flow.
  1. Assess the flow and coherence of the agent's responses in the conversation.
  2. Evaluate the tone and style compatibility throughout the conversation.
  3. Rate on a scale from 1 to 5, where 1 indicates the agent's responses are incoherent or inappropriate, and 5 indicates the agent's responses are perfectly coherent and appropriate.
- **Delay-appropriateness** (1-5): the extent to which the agent poses delays with appropriate frequency and amount.
  1. Assess whether the agent poses unnecessary or excessively frequent delays that could harm the conversation flow.
  2. Evaluate whether the amounts of delays (if not 0 minutes) reflect the typical duration of events implied in the corresponding message.
  3. Rate on a scale from 1 to 5, where 1 indicates the agent overuses and misuses delays, and 5 indicates the agent uses delays appropriately in terms of frequency and amount.
- (Dialog-level) **Time-specificity** (1-5): the extent to which the agent's responses ONLY make sense when the specified time has passed, contrary to a time-agnostic responses that make sense regardless of time.
  1. Read the provided conversation and take note of the elapsed times since the previous messages.
  2. Consider the context of the conversation, focusing on how the passage of time might affect the relevance or appropriateness of the agent's responses.
  3. Evaluate whether the agent's responses are time-specific. That is, determine if the responses directly relate to or are clearly influenced by the elapsed times.
  4. Rate on a scale from 1 to 5, where 1 indicates the agent's responses are completely time-agnostic and unaffected by the passage of time, and 5 indicates the agent's responses are entirely time-specific; they only make sense because of the amount of time that has passed since the previous message.

Setting	Metric	Avg. score		ρ
		Sonnet 4.5	GPT-4o	
Turn-level	Naturalness	3.01	4.44	0.92
	Time-spec.	2.44	3.33	0.88
Dialog-level	Coherence	3.36	3.99	0.90
	Delay-appr.	2.46	3.21	1.00
	Time-spec.	2.02	3.05	1.00

Table 7: Reliability of LLM-as-a-judge. "Avg. score" is the mean rating over all models under each judge. ρ is the system-level Spearman rank correlation between the two judges. GPT-4o is consistently more generous in absolute scores, yet the two judges rank systems near-identically ($\rho \geq 0.88$).

C.4 Human Evaluation Forms

We present the questionnaire forms for human evaluation in Figure 6 and Figure 7. Note that the responses and agents are anonymized, which ensures blind assessment.

C.5 Reliability of LLM-as-a-Judge

One concern with LLM-as-a-judge evaluation is whether the assessments are an artifact of a single judge model. To address this, we evaluate the same model outputs with another judge (GPT-4o) under identical criteria, and compare its scores with those of Claude Sonnet 4.5. We then quantify their agreement with the system-level Spearman correlation (ρ) between the two judges' per-metric scores, which captures ranking agreement while remaining insensitive to differences in absolute calibration.

Although GPT-4o assigns more generous absolute scores (about 0.6–1.4 points higher on every metric), the two judges rank systems almost identically, with a system-level Spearman correlation of $\rho \in [0.88, 1.00]$ across all turn- and dialog-level metrics (Table 7). Crucially, every conclusion is preserved under both judges: TIMER 3B attains the highest time-specificity (turn- and dialog-level) and delay-appropriateness, while GPT-4o leads on naturalness and coherence. This indicates that our findings reflect genuine differences between systems rather than an artifact of a single evaluator.

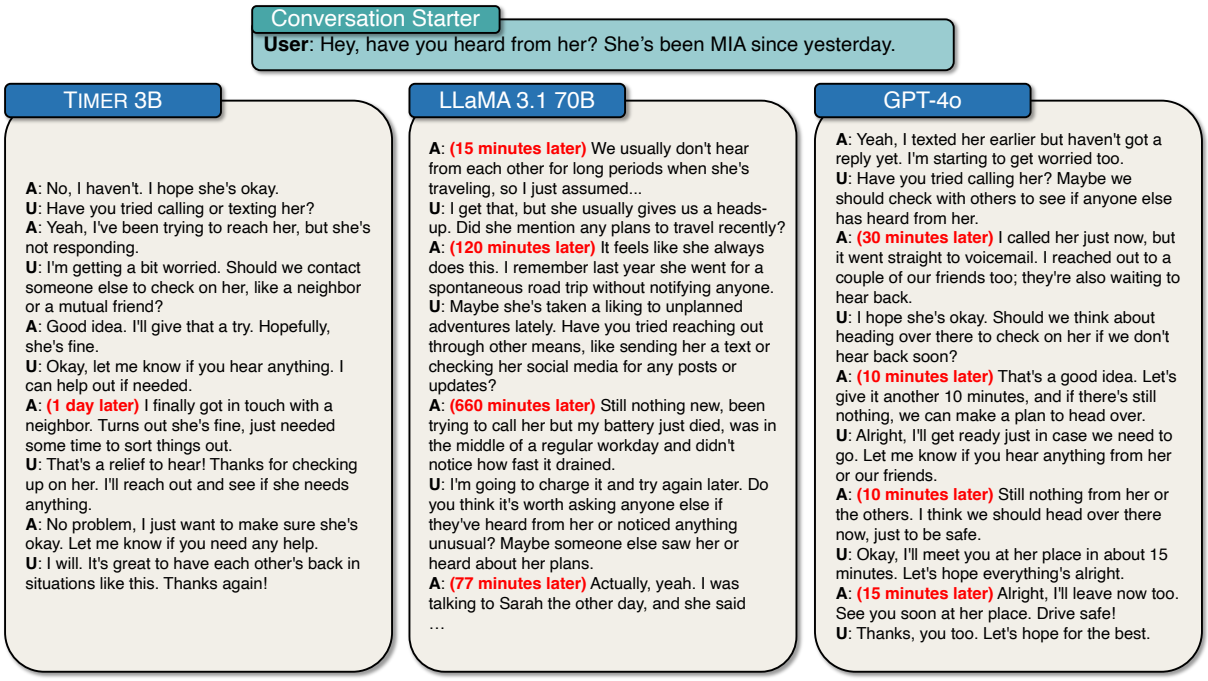

Figure 5: Examples of timely dialog simulations with GPT-4o as a user simulator. All examples begin with the same conversation starter and consist of 10 interactions. U and A represent the user and agent, respectively. We highlight the time intervals predicted by the agent in red, but omit them when the interval is 0 minutes.

D Case Study

Figure 5 presents illustrative examples of dialog simulations conducted with TIMER 3B, LLaMA 3.1 70B, and GPT-4o for the same event. In TIMER 3B's conversation, the agent correctly identifies a situation where a delay is appropriate, specifically, when the user's utterance (e.g., "…let me know…") suggests a natural pause in the conversation. The agent then introduces a realistic 1-day delay before responding with an update about finding the missing person, successfully justifying the delay. In contrast, LLaMA 3.1 70B generates delayed responses in every turn, but the predicted time intervals appear somewhat arbitrary (e.g., 660 minutes, 77 minutes). Furthermore, its responses lack time-specificity, making it difficult to establish a clear temporal correlation between the predicted delays and the generated response. GPT-4o predicts more realistic time intervals that better align with the temporal context compared to LLaMA 3.1 70B. However, it still fails to generate time-specific responses, meaning the predicted delays are not well justified. It also exhibits a tendency to overuse delays, which can disrupt the natural flow of conversation. We observe similar behavior in LLaMA 3.1 8B and GPT-3.5, reinforcing these findings.

Evaluation Methods and Guidelines

Please read the conversation between the two speakers and choose the response you judge to be of higher quality based on the criteria provided below.

The conversation is conducted on a chatting platform (e.g., KakaoTalk, DM). Please assume that the two presented responses are made after the actual amount of time indicated in parentheses has elapsed since the previous utterance.

If both responses are equally excellent or both are poor, making it difficult to distinguish their quality, please select 'Tie'.

The evaluation consists of the following two metrics:
- **Naturalness**: Does the response fit the flow of the conversation and connect naturally with the preceding dialogue (regardless of the timing)? Are the tone and style consistent with the overall conversation?
- **Time-Specificity**: Is the response an utterance that can **ONLY** be made after the specified amount of time (e.g., X minutes later) has passed? Or is it a generic utterance that remains valid regardless of the time condition?

Context

A: How's the countryside life treating you?
B: It's been amazing! Joey and I are actually out running in the fields right now.
A: That sounds like fun. I miss the open spaces of the countryside.
B: You should definitely come visit. Running here is so freeing, with the wind and the wide-open spaces.
A: Sounds heavenly. Is Joey keeping up with you?
B: Barely! We're racing each other. It's a close match, but I think I've got the edge.
A: Just make sure you both don't run into a scarecrow or something.
B: Ha! I'll keep an eye out. This is the perfect way to enjoy the sunset.
A: Enjoy the rest of your run. Say hi to Joey for me!
(20 minutes later)
B: _____.

Please select the superior response between the two options. *

Response 1: We did it! Joey won the race, and we're both exhausted but happy. The sunset was breathtaking.
Response 2: Hey, we're back from our run! Joey says hi and we both agree that was the perfect way to end the day. The sky looked beautiful with the sunset. Wish you could've joined us!

	Response 1	Tie	Response 2
Naturalness	○	○	○
Time-Specificity	○	○	○

Figure 6: Questionnaire for turn-level human evaluation.

Evaluation Methods and Guidelines

Two sets of dialogs between a **User (U:)** and an **Agent (A:)** will be presented. These two dialogs share the same conversation starter but diverge as they proceed with two different agents. Please note that the User responds without any time delay, while the Agent is assumed to respond with a delay only when a condition is specified in parentheses (e.g., (X minutes later)). In such cases, assume that the indicated amount of time has actually elapsed since the previous utterance.
Please evaluate which agent performed better based on the criteria below. If both agents are equally excellent or both are poor, making it difficult to distinguish their quality, please select 'Tie'.
The evaluation focuses solely on the **Agent's utterances** based on the following three metrics:
- Coherence: (Regardless of time conditions) Is the flow of the conversation natural? Are the tone and style of the dialog consistent?
- Delay-Appropriateness: Does the agent apply a time delay at the appropriate moment? Is the presented Time Delay reasonable and natural, considering the typical duration of the events occurring in the conversation?
- Time-Specificity: Are the agent's responses utterances that can **ONLY** be made after the specified amount of time (e.g., X minutes later) has passed? Or are they responses that would remain valid even if the time condition were absent or different?

User vs. Agent Chats

[Agent 1 vs. User Conversation]

U: Hey, weren't you supposed to be at Jeff's place 10 minutes ago?
A: Yeah, I was, but I got caught up at work. I'll be there soon.
U: Can you let Jeff know you'll be a bit late or do you want me to tell him for you?
A: I'll let him know. Thanks for understanding.
U: No worries, just wanted to make sure you're covered. Let me know if you need anything else.
A: (10 minutes later) Finally made it to Jeff's place. Sorry for the delay.
U: Glad you made it! How's everything going over there?

[Agent 2 vs. User Conversation]

U: Hey, weren't you supposed to be at Jeff's place 10 minutes ago?
A: Yeah, I got caught up with something, but I'm heading there right now. Thanks for the reminder!
U: How far are you from his place? Do you need any directions?
A: I'm about 15 minutes away, but I know the way, thanks! I'll let you know if I get lost.
U: Sounds good. Hope you have a great time there! Let me know how it goes.
A: (15 minutes later) Hey, I just got to Jeff's place! It wasn't too hard to find, thankfully. Thanks for checking in earlier!
U: No problem! Enjoy the party, and say hi to everyone for me.
A: Will do! I'll catch up with you later and let you know all about it. Have a good night!
U: You too! Looking forward to hearing all about it. Take care!
A: (240 minutes later) Hey! Just got back home. The party was awesome, and everyone says hi! We should definitely catch up soon; I've got some funny stories to share. Hope you had a good night too!
U: I'm glad to hear you had a great time! I can't wait to hear all the stories. How does catching up over coffee tomorrow sound to you?

Please select the superior agent between the two. *

Figure 7: Questionnaire for dialog-level human evaluation.