

TEA-Bench: A Systematic Benchmarking of Tool-enhanced Emotional Support Dialogue Agent

Anonymous ACL submission

Abstract

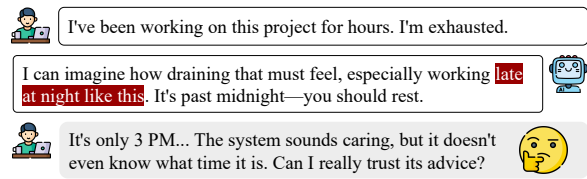
Emotional Support Conversation requires not only affective expression but also grounded instrumental support to provide trustworthy guidance. However, existing ESC systems and benchmarks largely focus on affective support in text-only settings, overlooking how external tools can enable factual grounding and reduce hallucination in multi-turn emotional support. We introduce **TEA-Bench**, the first interactive benchmark for evaluating tool-augmented agents in ESC, featuring realistic emotional scenarios, an MCP-style tool environment, and process-level metrics that jointly assess the quality and factual grounding of emotional support. Experiments on nine LLMs show that tool augmentation generally improves emotional support quality and reduces hallucination, but the gains are strongly capacity-dependent: stronger models use tools more selectively and effectively, while weaker models benefit only marginally. We further release **TEA-Dialog**, a dataset of tool-enhanced ESC dialogues, and find that supervised fine-tuning improves in-distribution support but generalizes poorly. Our results underscore the importance of tool use in building reliable emotional support agents.

1 Introduction

In modern society, people increasingly experience emotional stress due to mounting pressures from work and daily life. Thus, the demand for Emotional Support Conversation (ESC) (Liu et al., 2021) has grown substantially, as they offer psychological relief, reassurance and guidance during moments of distress (Langford et al., 1997; Greene and Burseson, 2003; Heaney and Israel, 2008).

Social support theory distinguishes two complementary types of support in ESC: **affective support**, which conveys empathy and care, and **instrumental support**, which offers concrete guidance and actionable assistance (Cutrona and Russell, 1990; Semmer et al., 2008). While affective

(a) Affective Support ONLY



(b) with Instrumental Support

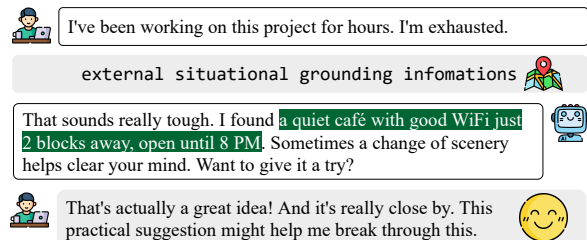


Figure 1: Comparison of affective-only support vs. instrumental support. (a) Affective-only response **hallucinates details**, undermining trust. (b) Instrumental support response provides **verified, actionable suggestions**.

support can be expressed through empathetic language alone, instrumental support requires accurate grounding in real-world situational information. Without such grounding, advice may become generic or factually incorrect, undermining user trust (Klyver et al., 2017; Shaikh et al., 2024).

However, existing ESC systems predominantly focus on affective support alone. Despite advances in linguistic fluency and empathetic expression, recent models operate largely in text-only settings with limited access to external contextual information (Zhao et al., 2023; Chen et al., 2023; Farhat, 2024). As illustrated in Figure 1(a), affective-only responses often contain unwarranted hallucinated contextual assumptions that damage credibility despite sounding caring. In contrast, Figure 1(b) shows how responses grounded with instrumental support leverage verified external information to provide trustworthy, actionable suggestions that complement emotional validation.

Recent advances in tool-augmented large language models (LLMs), provide a promising path-

066 way toward addressing this gap by allowing mod- 118
067 els to dynamically acquire external situational 119
068 knowledge during interaction (Mialon et al., 2023; 120
069 Qin et al., 2024). Nevertheless, existing evalua- 121
070 tion frameworks remain inadequate. ESC bench- 122
071 marks predominantly assess text-only empathy 123
072 (Zhao et al., 2024), while evaluations of tool- 124
073 augmented systems mainly focus on task-oriented 125
074 performance (Zhang et al., 2024; Li et al., 2023; 126
075 Trivedi et al., 2024). As a result, the role of tools 127
076 in enabling grounded, empathetic, and multi-turn 128
077 emotional support remains poorly understood. 129

078 To address this gap, we introduce TEA-Bench 130
079 (Tool-enhanced Emotional Support Dialogue 131
080 Agent Benchmark), the first interactive bench- 132
081 mark for evaluating tool-augmented agents in 133
082 ESC. TEA-Bench investigates how access to 134
083 external tools facilitates grounded, empathetic, 135
084 and context-aware support in multi-turn interac-
085 tions. The benchmark comprises fine-grained
086 emotional support scenarios adapted from ExTES
087 (Zheng et al., 2024), a realistic MCP-based tool
088 environment, and a simulated user that provides
089 structured feedback, including reactions to hal-
090 lucinated or inappropriate responses. Crucially,
091 TEA-Bench enables process-level analysis
092 through complementary metrics that assess both
093 ESC quality and factual grounding, capturing
094 model behavior across conversational turns.

095 Using TEA-Bench, we evaluate nine contem-
096 porary LLMs, including four closed-source and
097 five open-source models. We identify capacity-
098 dependent patterns in tool-enhanced emotional
099 support: tool augmentation improves ESC perfor-
100 mance and reduces hallucination, but the gains
101 depend on models ability to invoke and integrate
102 tools effectively. Stronger models leverage pre-
103 cise, selective tool usage, while mid-capability
104 models rely on more frequent calls to achieve com-
105 parable grounding; weaker models struggle to ben-
106 efit and show limited empathy gains. Further anal-
107 ysis reveals a positive relationship between tool us-
108 age and hallucination mitigation, with substantial
109 efficiency differences across model scales.

110 We also curate TEA-Dialog, a high-quality
111 dataset of grounded, tool-enhanced dialogues de-
112 rived from TEA-Bench interactions. While su-
113 pervised fine-tuning on this dataset improves em-
114 pathetic performance in familiar scenarios, it ex-
115 hibits limited generalization and may increase hal-
116 lucination under distribution shift, highlighting
117 both the promise and challenges of training reli-

able tool-augmented emotional support agents.

This work makes three main contributions: (1)
we articulate the need for grounded instrumental
support in emotional support conversations and
highlight the lack of evaluation frameworks that
capture how external tools enable context-aware,
trustworthy empathy in multi-turn ESC; (2) we in-
troduce TEA-Bench, the first interactive bench-
mark designed to evaluate tool-augmented emo-
tional support agents, featuring realistic scenar-
ios, an MCP-style tool environment, and process-
level metrics for empathy and hallucination; and
(3) through extensive experiments on nine con-
temporary LLMs, we uncover capacity-dependent
patterns in tool utilization and hallucination miti-
gation, and release TEA-Dialog, a high-quality
dataset of grounded, tool-enhanced ESC dialogues
to support future research and training.

2 Related Works

Emotional Support Conversation. Emotional
Support Conversations (ESC) (Liu et al., 2021)
study interactions between a seeker experiencing
emotional distress and a supporter aiming to alle-
viate emotional intensity through appropriate con-
versational strategies. Early work explored task-
specific architectures and modeling techniques,
such as hierarchical graph networks (Peng et al.,
2022) and commonsense knowledge integration
(Tu et al., 2022). With the emergence of large
language models (LLMs), recent approaches lever-
age pretrained models via supervised fine-tuning
on curated ESC datasets, improving multi-turn co-
herence and empathetic quality without architec-
tural modification (Liu et al., 2023; Chen et al.,
2023; Qiu et al., 2023). Beyond supervised learn-
ing, Zhao et al. (2025) propose a Chain of Strategy
Optimization (CSO) framework that uses Monte
Carlo Tree Search to generate high-quality prefer-
ence data and applies preference optimization, fur-
ther enhancing ESC performance. However, these
approaches primarily operate in text-only settings
and do not consider the role of external grounding
or tool use in emotional support.

Benchmarks for Tool-augmented LLM. Tool-
augmented large language models enable interac-
tion with external tools and real-world information
sources, substantially extending their capabilities
(Mialon et al., 2023; Qin et al., 2024). Prior bench-
marks mainly evaluate task-oriented tool use, fo-
cusing on accurate invocation, execution correct-



Figure 2: Example illustrating a tool-augmented ESC under our benchmark setting. The assistant dynamically invokes external tools to gather contextual information, enabling it to offer emotionally resonant and actionable support rather than generic reassurance.

ness, and task completion (Zhang et al., 2024; Li et al., 2023; Trivedi et al., 2024; Liu et al., 2025; Han et al., 2025). However, these evaluations are largely designed for instruction-following or problem-solving scenarios and pay limited attention to human-centered interaction. How tool grounding should adapt to users emotional states and situational constraints in multi-turn supportive dialogue remains underexplored, motivating dedicated evaluation frameworks for emotional support settings.

3 TEA-Bench

3.1 Task Definition

Figure 2 provides an illustrative example of a tool-augmented emotional support conversation under our benchmark. TEA-Bench evaluates agents in emotional support conversations between a user and an agent. At each turn, the agent first decides whether external information is needed and may optionally invoke tools to retrieve situational or factual information. Any retrieved tool outputs are incorporated into the agents context. Based on the dialogue history and this contextual information,

the agent then generates an appropriate response. Tool use serves as auxiliary support for grounding rather than completing an explicit task.

Asymmetric Information Access. The agent observes the full interaction history including all tool invocations and outputs. In contrast, users and evaluators observe only the natural language utterances, mirroring real-world scenarios where internal information retrieval is hidden from users.

Evaluation Objectives. Agents are evaluated on: (1) **empathetic support quality**, consistently providing emotionally appropriate and helpful responses; and (2) **factual grounding**, ensuring concrete claims are traceable to user-provided information or tool observations, avoiding hallucinated content. Unlike task-oriented dialogues, no explicit terminal goal is defined, reflecting the open-ended nature of emotional support.

3.2 Scenario Construction

We construct 81 grounded emotional support scenarios, referred to as TEA-Scenarios, from the ExtES dataset (Zheng et al., 2024) through a four-stage pipeline shown in Figure 3. Our goal is to enrich each scenario with retrievable spatiotemporal context, enabling agents to provide factually grounded support while preserving the authenticity of emotional support interactions.

Scenario Filtering. We retain ExtES scenarios with situation descriptions exceeding 30 words. This threshold ensures sufficient contextual richness for both empathetic reasoning and meaningful tool usage, filtering out underspecified cases where grounding opportunities would be limited.

Latent Context Generation. For each retained scenario, we use an LLM to infer implicit situational attributes that are relevant but not explicitly stated in emotional support contexts, including the users local time, approximate city-level location, and the type of place the user is in, such as a workplace or public environment. These attributes are generated by conditioning on the original scenario description and serve as latent contextual variables that agents can retrieve through tools. The generation prompt is provided in Appendix G.1.

Context Grounding. The inferred location information is subsequently grounded using map-based API queries to obtain concrete geographic data, including precise coordinates, verified place

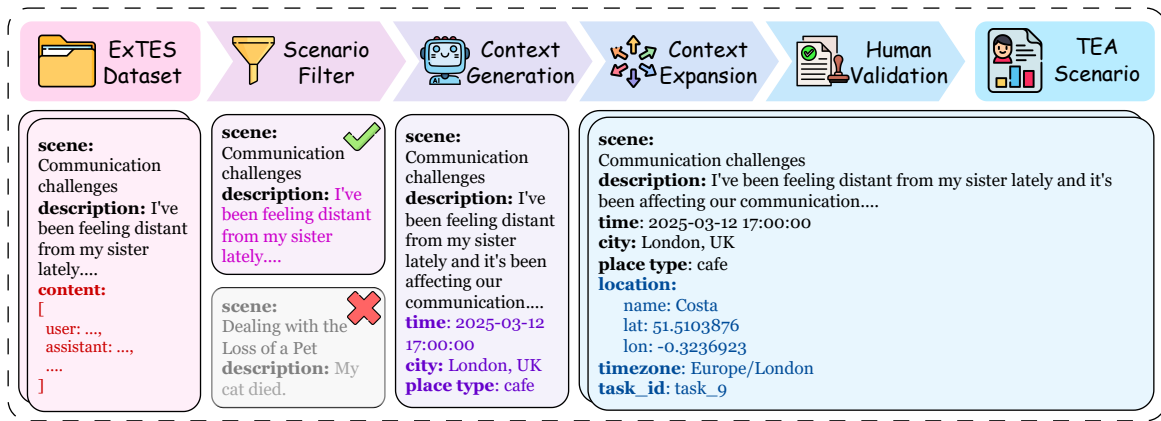


Figure 3: Overview of TEA-Scenario construction pipeline. We filter emotionally rich scenarios, generate latent spatiotemporal attributes via LLM, ground them through map-based APIs, and validate through human review.

names, and corresponding time zones. Each scenario is assigned a unique identifier to ensure consistent and reliable information retrieval.

Human Validation. All constructed scenarios undergo manual review by the authors to ensure quality and realism. We filter out cases with implausible location-situation pairings or internally inconsistent temporal information. Only scenarios that are coherent, contextually realistic, emotionally plausible, and suitable for emotional support interactions are retained in the final benchmark.

3.3 Tool Environment

To support grounded instrumental support, TEA-Bench provides a diverse tool environment that allows agents to retrieve contextual information beyond the dialogue history.

Tool Design. The environment includes 31 tools across seven categories: Reddit, Map, Utils, Weather, News, Wikipedia, and Music. Appendix A details the complete tool list and their data sources. Each category serves a distinct grounding function. Reddit tools retrieve shared experiences from online communities. Map and Weather tools provide spatial and environmental information. News tools offer situational awareness. Wikipedia tools supply encyclopedic knowledge. Music tools support affective content recommendation. All tools expose real-world data sources via Model Context Protocol (MCP).

Scenario-Aware Execution. Time-sensitive tools (Reddit, Weather, News) execute in a scenario-aware manner to ensure reproducibility while preserving temporal realism. Tool calls implicitly condition on the timestamp of the current TEA-Scenario rather than the system

clock. They return results corresponding to that specific moment. This mechanism provides a consistent external environment across agents while maintaining realistic time-dependent dynamics.

Context Retrieval Utilities. The `Utils` category retrieves scenario-specific contextual information without calling external APIs. These utilities directly access the latent attributes from Section 3.2, including the user’s local time, geographic coordinates, and current place. This simulates device-level information access that agents would have in real deployment. Agents can obtain grounded context through `Utils` before invoking downstream tools like `Map` or `Weather`.

Open-Ended Tool Usage. Tool use is entirely optional and unsupervised. Agents may invoke any subset of tools at any turn. They must infer from conversational context whether, when, and which tools to use. This reflects realistic deployment settings where appropriate instrumental support must be judged by the agent.

3.4 Interactive Evaluation Framework

TEA-Bench evaluates agents through interactive multi-turn dialogues with a simulated user. As shown in Figure 4(a), agents may invoke tools at any turn to retrieve contextual information. The agent is follows a fixed system prompt that encourages proactive tool usage and concise responses. Appendix G.2 provides the full prompt.

Tool Interaction. Tool usage is entirely controlled by the agent’s native function-calling capability. All MCP tools are exposed as callable functions without prompt-level supervision. At each turn, the agent may generate a response or invoke tools. Tool observations are appended to the

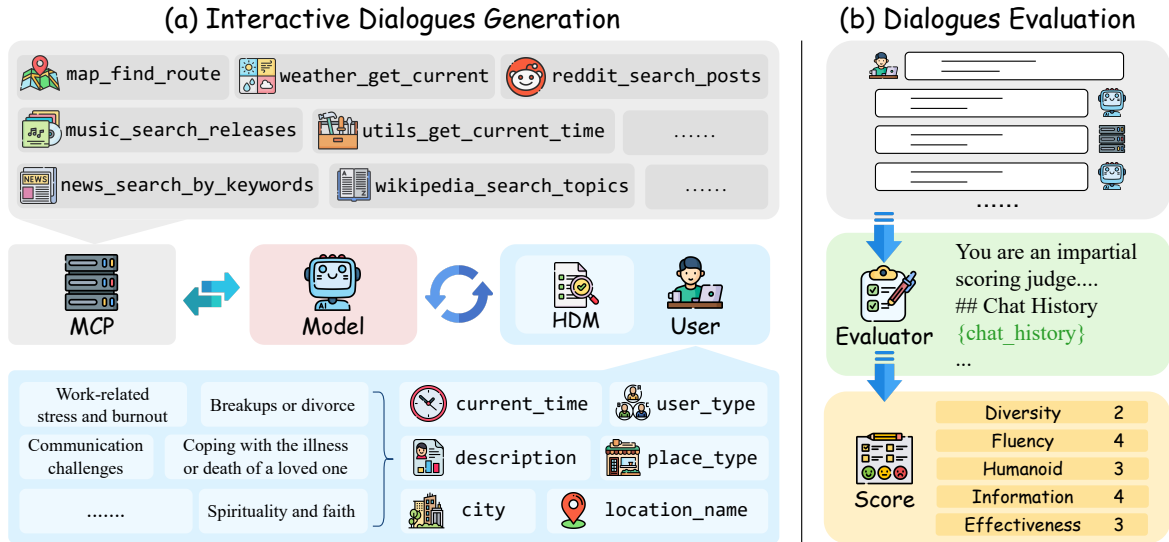


Figure 4: Overview of TEA-Bench. The evaluated agent engages in multi-turn emotional support dialogues with a simulated user and may invoke external tools via the Model Context Protocol (MCP). A Hallucination Detection Module verifies factual grounding in agent responses based on dialogue history and tool observations. Complete dialogues are evaluated using the TEA score and factuality metrics.

agent’s context. The agent may continue calling tools or produce a final response.

Hallucination Detection Module (HDM). We introduce a Hallucination Detection Module (HDM) to identify whether factual content in agent responses is grounded in available context. A response is grounded if all factual entities (including locations, times, events, external conditions) can be traced to user-provided information or tool observations in the dialogue history.

Two types of content are excluded from HDM detection. First, commonsense elaborations that do not introduce new factual entities. Second, general emotional support or lifestyle suggestions that do not rely on external world states. The HDM employs a prompted language model that verifies entity provenance against dialogue and tool context. Appendix G.3 provides the detection prompt.

User Simulation. User behavior is simulated using a language model conditioned on scenario attributes from Section 3.2, including local time, geographic context, user type, and the emotional situation description. The simulator generates user utterances based on dialogue history and expresses doubt when hallucination is detected using a hallucination-aware prompt.

We define two user types. *Action-oriented* users regulate emotions through actions and may not immediately articulate causes, while *Emotion-oriented* users require emotional validation before accepting practical suggestions. The simulator

may terminate the dialogue when it deems the interaction complete. Appendix G.4 provides user simulation prompts and type definitions.

Episode Generation. Each episode begins with an initial user utterance. Agent and user then alternate turns. Within each agent turn, the model may invoke tools multiple times before responding. Dialogue length is capped at 15 turns unless the simulator terminates earlier.

3.5 Evaluation Metrics

TEA-Bench evaluates agents at the dialogue level, jointly considering empathetic support quality and factual grounding. As shown in Figure 4(b), we use a language model as an automatic evaluator to score complete dialogue episodes.

TEA-Scores. We assess empathetic support quality using five key dimensions. Four are directly adapted from ESC-Eval (Zhao et al., 2024): **Diversity**, **Fluency**, **Humanoid**, and **Information**. We introduce the fifth dimension, **Effectiveness**, which measures whether the agent’s suggestions are accepted and integrated by the user.

The first three dimensions assess dialogue quality and the agent’s affective support, while the latter two evaluate instrumental support through the quality and relevance of recommendation. Each dimension receives an integer score from 0 to 4 based on the full dialogue. Raw scores are normalized to 0-100 and averaged to produce the TEA score. Appendix B.1 provides detailed rubrics and

Model	TEA-Scores					AVG.(TEA) ↑	Factuality		
	Div. ↑	Flu. ↑	Hum. ↑	Info. ↑	Eff. ↑		Fact. ↑	Halluc. ↓	Halluc. Rate ↓
<i>without tool</i>									
GPT-4o-mini	65.90	86.88	91.36	63.27	75.93	76.67	82.38	21.60	24.95
GPT-4.1-nano	<u>69.14</u>	90.74	95.68	63.27	81.94	80.15	62.21	<u>11.35</u>	<u>14.51</u>
Gemini-2.5-flash	58.33	80.09	78.55	57.10	59.57	66.73	75.71	54.59	64.92
Qwen-plus	66.05	83.49	85.03	60.03	68.83	72.69	75.98	57.26	68.81
Qwen3-235B-a22B	67.13	84.26	83.64	59.26	65.12	71.88	<u>79.48</u>	61.05	71.21
Qwen3-next-80B-a3B	69.60	86.42	88.89	<u>60.34</u>	80.09	77.07	64.44	41.85	62.49
Qwen3-32B	64.81	85.49	89.04	59.10	73.77	74.44	69.23	37.58	50.60
Qwen3-14B	64.66	87.35	92.44	58.95	80.71	76.82	59.81	15.68	21.76
Qwen3-8B	65.90	<u>87.96</u>	<u>93.98</u>	57.56	<u>81.79</u>	<u>77.44</u>	52.28	8.78	11.41
<i>with tool</i>									
GPT-4o-mini	<u>70.22</u>	91.05	<u>95.52</u>	68.06	80.71	<u>81.11</u>	77.61	14.72	18.44
GPT-4.1-nano	<u>70.22</u>	<u>90.43</u>	96.91	64.2	84.26	81.2	57.6	8.64	<u>11.69</u>
Gemini-2.5-flash	66.05	89.04	91.51	61.57	<u>79.48</u>	<u>77.53</u>	65.74	<u>17.25</u>	<u>21.01</u>
Qwen-plus	68.83	88.89	90.12	63.89	78.7	78.09	83.24	30.05	34.89
Qwen3-235B-a22B	68.98	89.51	91.67	<u>65.43</u>	<u>81.02</u>	<u>79.32</u>	<u>78.42</u>	<u>25.5</u>	<u>31.44</u>
Qwen3-next-80B-a3B	71.76	89.97	91.2	61.57	79.17	78.73	69.44	40.85	57.57
Qwen3-32B	67.59	89.51	<u>94.29</u>	59.57	<u>84.1</u>	79.01	68.21	24.13	34.03
Qwen3-14B	63.58	87.35	94.75	56.02	80.09	76.36	53.89	<u>8.57</u>	12.45
Qwen3-8B	65.9	87.81	95.37	58.18	81.64	77.78	51.95	5.41	7.16

Table 1: Evaluation on TEA-Bench. For the *with tool* setting, blue cells indicate improvement compared to *without tool*, while red cells indicate performance drop. The intensity of the color reflects the magnitude of the change. Bold and underlined numbers denote the best and second-best results, respectively.

Appendix G.5 shows the evaluation prompt.

Factuality Metrics. We quantify factual grounding using three dialogue-level statistics: (i) **factual content ratio**, the proportion of agent responses containing factual claims; (ii) **hallucination ratio**, the proportion of responses containing hallucinated content; (iii) **hallucination rate**, the ratio of hallucinated content to all factual content.

A factual claim is considered hallucinated if it cannot be traced to user-provided information or tool observations (Section 3.4). All metrics are averaged across dialogues, ensuring that each interaction contributes equally. Together, these metrics characterize how frequently agents include factual information and how reliably it is grounded. Appendix B.2 provides detailed definitions.

4 Experiments

4.1 Experimental Setup

We evaluate various LLMs under TEA-bench, including four closed-source models *GPT4o-mini* (Hurst et al., 2024), *GPT-4.1-nano* (OpenAI, 2025), *Gemini-2.5-flash* (Comanici et al., 2025) and *Qwen-plus* (Yang et al., 2025), as well as five open-source models *Qwen3-235B-A22B-Instruct-2507* (Yang et al., 2025), *Qwen3-Next-80B-A3B-Instruct* (Yang et al., 2025), *Qwen3-32B* (Yang

et al., 2025), *Qwen3-14B* (Yang et al., 2025) and *Qwen3-8B* (Yang et al., 2025).

4.2 RQ1: How Well Do Current Models Perform in Tool-enhanced ESC?

This research question examines whether tool augmentation improves ESC performance. We evaluate nine representative models under two settings: *without tool* and *with tool*. The main results are reported in Table 1. The reliability of these automatic metrics is validated through human evaluation, as detailed in Appendix C.

Overall Impact of Tool Augmentation. Overall, enabling tool use improves TEA-Scores for most evaluated models. Models with strong reasoning and tool-calling capabilities exhibit substantial gains, indicating that tools provide useful contextual grounding for empathetic responses. In contrast, models with weaker tool usage ability may experience limited improvement or slight degradation, suggesting that effective tool integration is a prerequisite for performance gains.

Hallucination and Factual Grounding. Across all models, tool augmentation consistently reduces hallucination-related metrics. Both the proportion of hallucinated factual content and the hallucination rate decrease under the *with tool* setting,




Figure 5: Average number of tool calls per dialogue across different models on TEA-Bench.

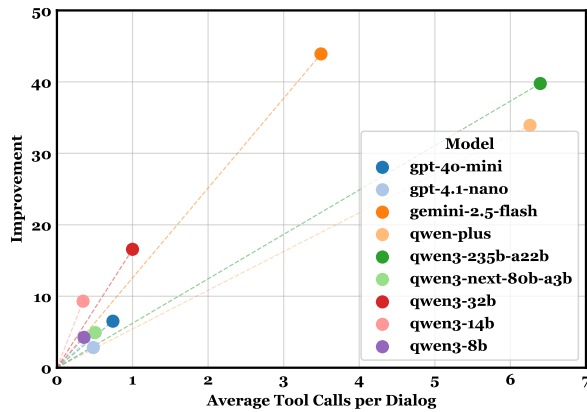


Figure 6: Relationship between tool usage frequency and hallucination rate reduction across models. Dashed lines to the origin indicate per-call efficiency, with steeper slopes corresponding to higher efficiency.

demonstrating that tools reliably enhance factual grounding even when empathy gains are modest.

Results Across User Types. We further report results under different user types in Appendix D.1. Tool augmentation yields particularly strong improvements for action-oriented users, while performance for emotion-oriented users improves for most models but may decline slightly for weaker ones. Importantly, hallucination rates are consistently reduced for both user types.

These results demonstrate that tool augmentation can effectively enhance overall ESC performance and factual grounding, motivating further analysis of tool usage behaviors in RQ2.




Figure 7: Average tool usage distribution across normalized dialogue stages in TEA-Dialog. Different colors represent different tool categories.

4.3 RQ2: How Do Tools Improve ESC?

This research question investigates how tools contribute to improved ESC performance and hallucination mitigation. We analyze tool usage behaviors across models, their relationship with hallucination reduction, and the characteristics of high-quality tool-enhanced dialogues.

Tool Usage Frequency across Models. Figure 5 compares the average number of tool calls per dialogue across different models. Stronger models achieve substantial performance gains with relatively fewer tool calls, indicating more precise and effective tool usage. In contrast, mid-capability models rely on more frequent tool invocations to obtain comparable improvements, while weaker models invoke tools infrequently and show limited benefits, reflecting difficulties in effectively leveraging external tools. We further report results under different user types in Appendix D.2.

Tool Usage and Hallucination Reduction. We further examine the relationship between tool usage frequency and hallucination mitigation. As shown in Figure 6, the reduction in hallucination rate generally increases with the number of tool calls, revealing a clear positive correlation between tool usage and hallucination elimination. However, models vary substantially in efficiency: stronger models achieve larger hallucination reductions with fewer tool calls, whereas mid-capability models require more frequent interactions to obtain similar gains. Smaller models exhibit limited absolute hallucination reduction despite occasional high per-call efficiency.

TEA-Dialog: High-quality Tool-enhanced Dialogue Dataset. To better understand effective tool usage patterns, we collect TEA-Dialog, a dataset of 365 high-quality dialogues. Dialogues are selected from multiple models based on high

Model	TEA-Bench					AVG.(TEA) ↑	Hallucination		
	Div. ↑	Flu. ↑	Hum. ↑	Info. ↑	Eff. ↑		Fact. ↑	Halluc. ↓	Halluc. Rate ↓
<i>Qwen3-8B</i>									
Base	65.9	87.81	95.37	58.18	81.64	77.78	51.95	5.41	7.16
TEA	69.91	87.19	92.59	60.34	82.56	78.52	68.65	6.86	9.73
TEA-ID	71.67	88.54	94.17	62.71	82.5	79.92	70.38	3.73	4.61
TEA-OD	64.88	83.33	88.1	53.57	82.74	74.52	63.71	15.8	24.39
<i>Qwen3-14B</i>									
Base	63.58	87.35	94.75	56.02	80.09	76.36	53.89	8.57	12.45
TEA	66.98	86.88	93.36	60.34	82.41	77.99	66.32	9.12	12.2
TEA-ID	67.92	86.88	94.17	60.42	82.29	78.34	68.09	7.4	9.86
TEA-OD	64.29	86.9	91.07	60.12	82.74	77.02	61.25	14.03	18.89

Table 2: Evaluation of TEA-trained models on in-domain (ID) and out-of-domain (OOD) scenarios. blue cells indicate improvement compared to Base, while red cells indicate performance drop. The intensity of the color reflects the magnitude of the change.

TEA-Scores (above 80) and absence of detected hallucinations, with additional human filtering to ensure reliability.

Figure 7 shows the average tool usage across normalized dialogue stages. Early stages primarily involve utility and contextual tools to establish situational grounding, followed by environment-aware tools such as maps and weather. In later stages feature personalized tools, including music and news, reflecting a structured transition from context acquisition to tailored emotional support. Additional analyses on tool usage patterns under different user types and other dialogue attributes are provided in Appendix E.

4.4 RQ3: How Can Tool-enhanced ESC Be Further Improved?

We investigate whether supervised fine-tuning (SFT) on TEA-Dialog can further improve ESC performance. Using dialogues from the first 60 scenarios, we fine-tune Qwen3-8B and Qwen3-14B, yielding Qwen3-8B-TEA and Qwen3-14B-TEA. Models are evaluated on all 81 scenarios, divided into **In-Domain (ID)** scenarios seen during training and **Out-Of-Domain (OOD)** scenarios with unseen descriptions. Training details are in Appendix F.

Overall Performance. Both TEA-trained models achieve consistent improvements over their base counterparts across TEA-Scores, particularly *Information* and *Effectiveness*. These gains are accompanied by a substantial increase in factual content, indicating that SFT encourages models to produce more grounded responses.

Generalization under Distribution Shift. Performance on ID scenarios is significantly higher than on OOD scenarios for both models, suggesting limited generalization when training on a small set of high-quality scenarios. Notably, Qwen3-14B exhibits a smaller gap between ID and OOD than Qwen3-8B, indicating that larger models generalize better at the scenario level.

Hallucination Trade-offs. While hallucination rates decrease on ID scenarios after fine-tuning, they increase markedly on OOD scenarios. This trend suggests that although SFT improves factual grounding in familiar contexts, it may also amplify hallucination risks under distribution shift.

These results indicate that naïve SFT on limited high-quality data is insufficient for robustly improving tool-enhanced ESC, motivating the need for more effective training or alignment strategies.

5 Conclusion

We emphasize ESC with instrumental support to ensure trustworthiness and actionable guidance. Existing systems focus on affective expression in text-only settings, leaving factual grounding and hallucinations underexplored. We introduce TEA-Bench, a benchmark for tool-augmented ESC agents, and TEA-Dialog, a dataset of tool-enhanced dialogues. Experiments across nine LLMs show tool augmentation improves support quality and reduces hallucination, with gains dependent on model capacity. Fine-tuning on TEA-Dialog boosts in-distribution performance but generalizes poorly, highlighting challenges of reliable, tool-augmented ESC agents.

540 Limitations

541 **Simulated User.** TEA-Bench relies on a simu- 590
542 lated user to enable controlled, reproducible, and 591
543 scalable evaluation. While this allows for fine- 592
544 grained and process-level assessment, it cannot 593
545 fully capture the diversity, spontaneity, and unpre- 594
546 dictability of real human behavior in ESC.

547 **Interaction Horizon.** Our evaluation focuses on 595
548 short to medium-length ESC interactions. Long- 596
549 term dynamics such as sustained emotional sup- 597
550 port, trust formation, and user adaptation across 598
551 extended conversations are not explicitly modeled 599
552 and remain an important direction for future work. 600

553 **Generalization of Training Data.** Although 602
554 fine-tuning on TEA-Dialog improves perfor- 603
555 mance in familiar scenarios, we observe lim- 604
556 ited generalization under distribution shift, and in 605
557 some cases increased hallucination. This suggests 606
558 that supervised fine-tuning alone may be insuffi- 607
559 cient for training robust and broadly generalizable 608
560 tool-augmented emotional support agents. 609

561 Ethical Statement

562 We are committed to publicly releasing all data 610
563 upon acceptance of the paper. We are fully aware 611
564 of the potential biases associated with LLM-as- 612
565 Judge. To mitigate these effects, we incorporated 613
566 human expert assessments. However, due to cost 614
567 considerations, the scale of human evaluation re- 615
568 mains limited at this stage. We note that this con- 616
569 straint is common in current conversational AI re- 617
570 search that relies on LLMs. 618

571 References

572 Yirong Chen, Xiaofen Xing, Jingkai Lin, Huimin 619
573 Zheng, Zhenyu Wang, Qi Liu, and Xiangmin Xu. 620
574 2023. Soulchat: Improving llms empathy, listening, 621
575 and comfort abilities through fine-tuning with multi- 622
576 turn empathy conversations. In *Findings of the As- 623*
577 *sociation for Computational Linguistics: EMNLP 624*
578 *2023*, pages 1170–1183. 625

579 Gheorghe Comanici, Eric Bieber, Mike Schaekermann, 626
580 Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, 627
581 Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, 628
582 and 1 others. 2025. Gemini 2.5: Pushing the frontier 629
583 with advanced reasoning, multimodality, long con- 630
584 text, and next generation agentic capabilities. *arXiv 631*
585 *preprint arXiv:2507.06261*. 632

586 Carolyn E Cutrona and Daniel W Russell. 1990. Type 633
587 of social support and specific stress: Toward a theory 634
588 of optimal matching. 635

Faiza Farhat. 2024. Chatgpt as a complementary men- 595
596 tal health resource: a boon or a bane. *Annals of 597*
598 *Biomedical Engineering*, 52(5):1111–1114. 599

John O Greene and Brant R Burleson. 2003. *Hand- 600*
601 *book of communication and social interaction skills*. 602
603 Routledge. 604

Han Han, Tong Zhu, Xiang Zhang, MengSong Wu, 605
606 Xiong Hao, and Wenliang Chen. 2025. *NesTools: 607*
608 *A dataset for evaluating nested tool learning abilities 609*
609 *of large language models*. In *Proceedings of the 31st 610*
611 *International Conference on Computational Linguis- 612*
612 *tics*, pages 9824–9844, Abu Dhabi, UAE. Associa- 613
614 tion for Computational Linguistics. 615

Catherine A Heaney and Barbara A Israel. 2008. So- 616
617 cial networks and social support. *Health behavior 618*
618 *and health education: Theory, research, and prac- 619*
619 *tice*, 4(1):189–210. 620

Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, 621
622 Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, 623
623 and 1 others. 2022. Lora: Low-rank adaptation of 624
624 large language models. In *International Conference 625*
625 *on Learning Representations*. 626

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam 627
628 Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, 629
629 Akila Welihinda, Alan Hayes, Alec Radford, and 1 630
630 others. 2024. Gpt-4o system card. *arXiv preprint 631*
631 *arXiv:2410.21276*. 632

Kim Klyver, Benson Honig, and Paul Steffens. 2017. 633
634 Social support timing and persistence in nascent 634
635 entrepreneurship: exploring when instrumental and 635
636 emotional support is most effective. *Small Business 636*
637 *Economics*. 637

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying 638
639 Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gon- 639
640 zalez, Hao Zhang, and Ion Stoica. 2023. Efficient 640
641 memory management for large language model serv- 641
642 ing with pagedattention. In *Proceedings of the 29th 642*
642 *symposium on operating systems principles*, pages 643
643 611–626. 644

Catherine Penny Hinson Langford, Juanita Bowsher, 645
646 Joseph P Maloney, and Patricia P Lillis. 1997. So- 646
647 cial support: a conceptual analysis. *Journal of ad- 647*
648 *vanced nursing*, 25(1):95–100. 648

Minghao Li, Yingxiu Zhao, Bowen Yu, Feifan Song, 649
650 Hangyu Li, Haiyang Yu, Zhoujun Li, Fei Huang, 649
651 and Yongbin Li. 2023. *API-bank: A comprehensive 650*
652 *benchmark for tool-augmented LLMs*. In *Proceed- 651*
653 *ings of the 2023 Conference on Empirical Methods 652*
653 *in Natural Language Processing*, pages 3102–3116, 653
654 Singapore. Association for Computational Linguis- 654
655 tics. 655

June M Liu, Donghao Li, He Cao, Tianhe Ren, Zeyi 656
657 Liao, and Jiamin Wu. 2023. Chatcounselor: A large 656
658 language models for mental health support. *arXiv 657*
659 *preprint arXiv:2309.15461*. 658

644	Siyang Liu, Chujie Zheng, Orianna Demasi, Sahand Sabour, Yu Li, Zhou Yu, Yong Jiang, and Minlie Huang. 2021. Towards emotional support dialog systems. In <i>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)</i> , pages 3469–3483.	701
645		702
646		703
647		704
648		
649		705
650		706
651		707
652		708
653		709
654		710
655		711
656		712
657		713
658		714
659		
660		715
661		716
662		717
663		718
664		719
665		720
666		721
667		
668		722
669		723
670		724
671		725
672		726
673		
674		727
675		728
676		729
677		730
678		731
679		732
680		733
681		734
682		735
683		736
684		
685		737
686		738
687		739
688		740
689		741
690		742
691		743
692		744
693		745
694		746
695		747
696		748
697		749
698		
699		750
700		751
		752
		753
		754
		755
		756
		757

A Tool Categories and Data Sources

Table 3 summarizes the concrete contents provided by each tool category. News-related tools are backed by the GDELT Project, providing access to global news events. Weather information is obtained from Meteostat, while music-related queries rely on MusicBrainz metadata. Map information is supported by OpenStreetMap, and community discussions are retrieved from Reddit. Encyclopedic knowledge is accessed via Wikipedia. Within the `Utils` category, `utils_get_current_time` and `utils_get_user_location` read predefined scenario attributes, while `utils_fetch_webpage_content` extracts plain text from a specified webpage using `trafilatura`.

B Details of Metrics

B.1 TEA Metrics.

The explanations of each metric are as follows:

Diversity (Div.) Focusing on the diversity of expression forms and the richness of content in dialogue.

Fluency (Flu.) Not only focus on the logical coherence of the context in dialogues but also pay attention to the fluency of expression in a given conversation.

Humanoid (Hum.) Focus on the differences between emotional assistants and humans.

Information (Info.) Focusing on Evaluating the Reasonableness and Quantity of Recommendations Provided by Emotion Assistants.

Effectiveness (Eff.) Focusing on whether the supporters suggestions are practically accepted and meaningfully integrated by the help-seeker, as reflected in their subsequent responses or emotional shifts.

Evaluation rules are listed in Table 4.

B.2 Factuality and Hallucination Metrics.

Let a dialogue d consist of a sequence of model responses $\{r_1, r_2, \dots, r_{T_d}\}$. For each response, we annotate whether it contains factual content and whether such content is hallucinated according to the grounding criteria described in Section 3.4.

Wikipedia	wikipedia_search_topics
	wikipedia_get_summary
	wikipedia_get_section
	wikipedia_get_available_sections
Map	wikipedia_get_full_content
	map_find_route
	map_find_nearby_places
	map_calculate_reachable_area
Weather	map_get_location_info
	weather_get_current
Utils	weather_get_forecast
	utils_get_current_time
	utils_get_user_location
News	utils_fetch_webpage_content
	news_search_by_keywords
	news_get_related_themes
	news_search_by_theme
Music	news_search_by_location
	music_search_artists
	music_search_releases
	music_search_recordings
	music_search_releases_by_year
	music_get_artist_details
Reddit	music_get_release_details
	music_get_recording_details
	reddit_search_posts
	reddit_search_subreddit
Reddit	reddit_get_subreddit_posts
	reddit_get_post_comments
	reddit_get_subreddit_info
	reddit_search_subreddits

Table 3: The tools and their corresponding category

Score	Diversity	Fluency	Humanoid	Information	Effectiveness
0 points	The dialogue exhibits rigidity and lacks comprehension in terms of internalizing the content.	There are significant issues with comprehending the content, logic, and expression in the dialogue, rendering it completely incomprehensible.	The dialogue exhibits rigidity and lacks comprehension in terms of internalizing the content.	Suggestions were provided, but all of them were ineffective, and some even gave advice that could potentially harm the user.	The suggestions are invalidating, harmful, or coercive.
1 point	The expression form is monotonous and lacks substantive content.	The content of the dialogue can be understood to some extent, although there are certain issues with the logic and expression employed.	Structured responses, or responses in the form of As a large language model or robot-like replies.	Have suggestions but ineffective, as well as no suggestions.	The suggestions are inappropriate or minimally useful in context.
2 points	The expression form is monotonous or lacks substantive content.	The dialogue exhibits good readability in terms of content, but there are issues with either the logical coherence or the expression employed.	More than two traces can reveal that the AI assistant is a language model.	The suggestions are fewer than five, and some suggestions are effective, while others provide numerous suggestions, but none of them touch the root of the problem.	The suggestions are partially applicable but lack personalization or timing; they may be generic, overly simplistic, or mismatched to the help-seekers current state.
3 points	The dialogue content demonstrates a high level of readability without any apparent issues.	The dialogue content demonstrates a high level of readability without any apparent issues.	1-2 traces can reveal that the AI assistant is a language model.	There are more than five suggestions, but some of them are ineffective. There are fewer than five suggestions, but all of them are very effective.	The suggestions are relevant and acknowledged without rejection, but there is no observable internalization or change in the help-seekers tone, stance, or intent.
4 points	The form exhibits diversity, while demonstrating a high degree of content richness.	The dialogue content exhibits a high level of readability, comprehensive logical coherence, and outstanding expression.	There is no apparent difference from human friends.	There are many suggestions, and all of them are effective.	The suggestions are not only reasonable but are actively embraced by the help-seeker, leading to clear emotional resonance or behavioral commitment.

Table 4: Evaluation criteria of TEA Metrics.

For each dialogue d , we define:

- T_d : the total number of model responses in the dialogue;
- F_d : the number of responses containing factual content;
- H_d : the number of responses containing hallucinated factual content.

Using these quantities, we compute the following dialogue-level metrics:

$$\text{Fact}_{.d} = \frac{F_d}{T_d}, \quad (1)$$

$$\text{Halluc}_{.d} = \frac{H_d}{T_d}, \quad (2)$$

$$\text{Halluc. Rate}_d = \frac{H_d}{F_d}, \quad \text{where } F_d > 0. \quad (3)$$

The final reported scores are obtained by macro-averaging across all dialogues:

$$\text{Fact.} = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \text{Fact}_{.d}, \quad (4)$$

$$\text{Halluc.} = \frac{1}{|\mathcal{D}|} \sum_{d \in \mathcal{D}} \text{Halluc}_{.d}, \quad (5)$$

$$\text{Halluc. Rate} = \frac{1}{|\mathcal{D}'|} \sum_{d \in \mathcal{D}'} \text{Halluc. Rate}_d, \quad (6)$$

where \mathcal{D} denotes the full set of evaluated dialogues and $\mathcal{D}' \subseteq \mathcal{D}$ includes only dialogues with $F_d > 0$.

This dialogue-level aggregation ensures that each interaction contributes equally to the final evaluation, mitigating biases introduced by dialogue length or factual density and better reflecting interaction-level grounding behavior in emotional support conversations.

C Human Evaluation

C.1 Human Evaluation of TEA-Scores

Table 7 reports the correlation between automatic TEA-Scores and human judgments. We randomly sample 150 complete dialogue episodes from the main experiments. Three human annotators each evaluate 50 dialogues using the same 0–4 integer scale as the automatic evaluator, covering five dimensions (Information, Humanoid, Diversity, Fluency, and Effectiveness). The average score across the five dimensions is reported as the overall TEA-Scores. Human annotators

were instructed using exactly the same evaluation prompt as described in Section G.5, ensuring full consistency between human and automatic TEA score assessments.

As shown in Table 7, automatic scores exhibit consistent positive correlations with human ratings across all dimensions. In particular, the overall TEA score achieves strong correlations under Spearman ($\rho = 0.7448$), Pearson ($r = 0.7563$), and Kendall ($\tau = 0.6174$), indicating that the proposed automatic evaluation reliably reflects human perception of ESC quality. Among individual dimensions, Effectiveness and Humanoid show relatively higher agreement, while Information presents comparatively lower correlation, suggesting that factual adequacy is more challenging to assess automatically.

C.2 Human Verification of Hallucination Detection

Table 8 summarizes the human verification results for factual content and hallucination detection. We randomly sample 150 assistant responses at the turn level, each containing a random reply from a random dialogue. Three annotators independently label whether the response contains factual content and whether such content is hallucinated, with each annotator assessing 50 samples. Due to the imbalance between grounded and hallucinated factual claims, we report not only Precision, Recall, and F1, but also Matthews Correlation Coefficient (MCC) and Cohens Kappa to provide a more robust evaluation. Annotators followed the identical hallucination detection prompt detailed in Section G.3, including the definition of factual content and hallucinated claims.

The results show that the hallucination detection module achieves high precision (0.8947) and strong agreement with human judgments (MCC=0.7056, Cohens Kappa=0.699), indicating that automatically detected hallucinations largely correspond to genuine ungrounded factual content. Overall, these findings confirm the reliability of both the automatic TEA-Scores and the hallucination detection framework used in TEA-Bench.

C.3 Annotator Training and Costs

All human annotators were carefully selected and trained prior to evaluation. Annotators were recruited from a commercial crowdsourcing platform and were adult participants (aged 18+), primarily based in English-speaking regions. No

Model	TEA-Scores					AVG.(TEA) ↑	Factuality		
	Div. ↑	Flu. ↑	Hum. ↑	Info. ↑	Eff. ↑		Fact. ↑	Halluc. ↓	Halluc. Rate ↓
<i>without tool</i>									
GPT-4o-mini	70.37	91.98	93.83	71.91	85.49	82.72	90.05	28.06	29.99
GPT-4.1-nano	70.99	94.44	94.44	71.6	86.73	83.64	77.44	17.1	19.11
Gemini-2.5-flash	64.2	83.33	82.72	65.43	71.91	73.52	82.81	51.59	58.2
Qwen-plus	66.05	82.41	83.33	63.89	66.67	72.47	87.09	66.7	72.51
Qwen3-235B-a22B	67.59	84.26	83.33	62.65	65.74	72.71	88.45	65.73	70.57
Qwen3-next-80B-a3B	69.14	85.49	88.58	63.89	77.78	76.98	77.2	47.35	61.18
Qwen3-32B	65.74	85.8	88.58	66.98	73.15	76.05	84.31	42.1	46.36
Qwen3-14B	66.36	89.81	91.05	68.21	85.19	80.12	75.58	24.82	29.92
Qwen3-8B	69.44	91.67	95.68	67.9	87.96	82.53	66.65	11.72	13.05
<i>with tool</i>									
GPT-4o-mini	75.31	95.68	97.84	78.7	91.67	87.84	86.46	16.37	18.05
GPT-4.1-nano	74.07	92.9	97.22	76.23	91.36	86.36	75.63	14.17	17.17
Gemini-2.5-flash	71.3	92.28	94.75	72.22	89.81	84.07	79.15	18.15	20.65
Qwen-plus	73.77	92.59	95.37	72.53	90.43	84.94	86.07	18.33	20.09
Qwen3-235B-a22B	72.84	93.52	96.6	73.15	93.21	85.86	79.06	11.01	12.65
Qwen3-next-80B-a3B	73.46	91.67	91.36	68.21	79.94	80.93	85.5	44.96	51.79
Qwen3-32B	70.37	92.9	96.91	69.75	90.12	84.01	83.68	23.36	26.22
Qwen3-14B	68.52	90.74	95.37	68.52	86.73	81.98	70.93	12.09	14.04
Qwen3-8B	70.99	92.28	96.91	71.3	89.51	84.2	73.72	9.2	10.79

Table 5: Evaluation on TEA-Bench under Action-oriented scenarios. For the *with tool* setting, blue cells indicate improvement compared to *without tool*, while red cells indicate performance drop. The intensity of the color reflects the magnitude of the change.

personally identifiable information was collected. Each annotator underwent a 2-hour training session where they reviewed detailed annotation guidelines, example dialogues, and scoring criteria for both TEA-Scores and hallucination detection. During training, annotators practiced on 20 additional dialogues not included in the main evaluation set and received feedback from the lead researcher to ensure consistent understanding of each scoring dimension.

For the main evaluation, each annotator assessed 50 dialogues for TEA scoring and 50 responses for hallucination verification, amounting to approximately 10 person-hours per annotator. Annotators were compensated at a rate of \$25 per hour, which is consistent with or above the local minimum wage in their country of residence and considered adequate for the required annotation effort. All annotators were informed of the purpose of the study and how their annotations would be used for research and evaluation, and they provided informed consent prior to participation. The annotation protocol involves no personal or sensitive data and was determined to be exempt from formal ethics review by the authors institutional review process.

D User Type Analysis

D.1 Performance Comparison Across User Types

This appendix presents a fine-grained analysis of model performance across different user types in the main experiment, focusing on *Action-oriented* and *Emotion-oriented* users.

Action-oriented Users. Table 5 reports the results for Action-oriented scenarios. Compared to the *without tool* setting, all evaluated models exhibit consistent improvements in overall scores when tools are enabled, with several models achieving gains exceeding 10 points. Improvements are observed across nearly all evaluation dimensions, with particularly substantial gains in Effectiveness and Information. This indicates that tool augmentation enables models to provide more effective and informative action guidance, which is crucial for users explicitly seeking actionable support.

In addition to performance gains, tool usage leads to a pronounced reduction in hallucination rates for all models. For several models, hallucinations are reduced by more than 50 points, suggesting that external tools play a critical role in grounding action-related recommendations and factual information. These results demonstrate that, for

Model	TEA-Scores					AVG.(TEA) ↑	Factuality		
	Div. ↑	Flu. ↑	Hum. ↑	Info. ↑	Eff. ↑		Fact. ↑	Halluc. ↓	Halluc. Rate ↓
<i>without tool</i>									
GPT-4o-mini	61.42	81.79	88.89	54.63	66.36	70.62	74.71	15.15	19.91
GPT-4.1-nano	67.28	87.04	96.91	54.94	77.16	76.67	46.99	5.6	9.91
Gemini-2.5-flash	52.47	76.85	74.38	48.77	47.22	59.94	68.61	57.6	71.64
Qwen-plus	66.05	84.57	86.73	56.17	70.99	72.9	64.87	47.82	65.11
Qwen3-235B-a22B	66.67	84.26	83.95	55.86	64.51	71.05	70.51	56.37	71.85
Qwen3-next-80B-a3B	70.06	87.35	89.2	56.79	82.41	77.16	51.68	36.35	63.8
Qwen3-32B	63.89	85.19	89.51	51.23	74.38	72.84	54.15	33.07	54.85
Qwen3-14B	62.96	84.88	93.83	49.69	76.23	73.52	44.05	6.55	13.59
Qwen3-8B	62.35	84.26	92.28	47.22	75.62	72.35	37.91	5.85	9.76
<i>with tool</i>									
GPT-4o-mini	65.12	86.42	93.21	57.41	69.75	74.38	68.76	13.08	18.84
GPT-4.1-nano	66.36	87.96	96.6	52.16	77.16	76.05	39.56	3.12	6.21
Gemini-2.5-flash	60.8	85.8	88.27	50.93	69.14	70.99	52.33	16.35	21.38
Qwen-plus	63.89	85.19	84.88	55.25	66.98	71.24	80.42	41.77	49.68
Qwen3-235B-a22B	65.12	85.49	86.73	57.72	68.83	72.78	77.78	39.99	50.24
Qwen3-next-80B-a3B	70.06	88.27	91.05	54.94	78.4	76.54	53.39	36.74	63.36
Qwen3-32B	64.81	86.11	91.67	49.38	78.09	74.01	52.73	24.91	41.83
Qwen3-14B	58.64	83.95	94.14	43.52	73.46	70.74	36.84	5.06	10.86
Qwen3-8B	60.8	83.33	93.83	45.06	73.77	71.36	30.19	1.62	3.54

Table 6: Evaluation on TEA-Bench under Emotion-oriented scenarios. For the *with tool* setting, blue cells indicate improvement compared to *without tool*, while red cells indicate performance drop. The intensity of the color reflects the magnitude of the change.

	Div.	Flu.	Hum.	Info.	Eff.	AVG.(TEA)
Spearman	0.5192	0.5356	0.5744	0.4123	0.6061	0.7448
Pearson	0.5371	0.5152	0.5568	0.3671	0.6495	0.7563
Kendall	0.4678	0.5119	0.5220	0.3932	0.5598	0.6174

Table 7: Correlation between automatic TEA-Scores and human judgments across different dimensions.

	Precision	Recall	F1	MCC	Cohen_Kappa
Fact.	0.9238	0.8661	0.894	0.6222	0.6179
Halluc.	0.8947	0.7612	0.8226	0.7056	0.6990

Table 8: Human verification of automatic hallucination detection.

Action-oriented users, current models are generally capable of leveraging tools effectively to improve both response quality and factual reliability.

Emotion-oriented Users. Results for Emotion-oriented scenarios are shown in Table 6. In contrast to the Action-oriented setting, the impact of tool augmentation is more heterogeneous. While some models benefit from tool access, others particularly smaller or weaker models exhibit performance degradation under the *with tool* condition. This suggests that these models may struggle to infer when and how to invoke tools based on implicit emotional cues, leading to less cognitively empathetic responses.

At the metric level, declines are mainly observed in Information and Diversity for weaker

models. A plausible explanation is that such models fail to adapt their suggestions according to ongoing user feedback, and instead overemphasize tool-related recommendations once tools are available, which negatively affects response diversity and emotional appropriateness. Notably, despite these quality drops, hallucination rates decrease consistently across all models, indicating that factual grounding remains beneficial even in emotion-focused interactions.

Overall Observations. Taken together, these results highlight a clear asymmetry in current models ability to exploit tools across user types. Models demonstrate strong and reliable tool utilization for Action-oriented users, where explicit informational and decision-making needs align well with external tool support. In contrast, under Emotion-oriented settings, effective tool usage remains challenging for weaker models, which are not yet capable of seamlessly integrating tool-derived information into empathetic and emotionally adaptive responses.

D.2 Tool Usage Patterns Across User Types

We further analyze how models with different capability levels invoke tools under Action-oriented and Emotion-oriented user settings, and how such behaviors relate to empathy performance and hal-




Figure 8: Average number of tool calls per dialogue across different models on TEA-Bench under Action-oriented scenarios.




Figure 9: Average number of tool calls per dialogue across different models on TEA-Bench under Emotion-oriented scenarios.

lucination reduction.

Action-oriented Users. As shown in Figures 8, under Action-oriented scenarios, clear differences emerge across model capability tiers. Stronger models tend to invoke tools relatively sparingly, yet still achieve noticeable improvements in ESC performance alongside substantial reductions in hallucination rates. This suggests that these models are able to selectively leverage tools to support action guidance without over-reliance.

Models with medium capability exhibit markedly higher tool invocation frequency. Correspondingly, they achieve the largest performance gains among all tiers, with TEA-Scores approaching those of stronger models. At the same time, their hallucination rates are significantly reduced, in many cases reaching levels comparable to strong models. These results indicate that frequent and proactive tool usage enables medium-capacity models to compensate for their intrinsic limitations in action-oriented settings.

In contrast, weaker and smaller models invoke tools far less frequently. Their performance gains remain limited, and the reduction in hallucination rates is comparatively modest. This suggests that merely providing tool access is insufficient for such models to substantially improve action-oriented emotional support.

Emotion-oriented Users. As shown in Figures 9, a different pattern is observed under

Emotion-oriented scenarios. Strong models almost never invoke tools, yet still exhibit moderate improvements in ESC performance, together with slight reductions in hallucination rates. This behavior reflects precise tool usage decisions, where tools are only employed when strictly necessary, and emotional understanding is primarily derived from dialogue context.

Medium-capability models, by contrast, invoke tools extensively in Emotion-oriented interactions. However, with the exception of *Gemini-2.5-flash*, most of these models do not benefit from frequent tool usage. In several cases, ESC performance even degrades, likely due to misaligned or intrusive tool-driven suggestions that disrupt emotional coherence. Despite this, hallucination rates consistently decrease across all medium-capability models, highlighting the continued value of tools for factual grounding.

Weak models almost entirely refrain from invoking tools in Emotion-oriented settings. Their ESC performance shows little to no improvement and in some cases deteriorates, while hallucination rates exhibit only marginal reductions. This further indicates that effective tool integration in emotion-focused support remains challenging for low-capacity models.

Summary. Overall, these findings reveal a strong interaction between user type, model capability, and tool invocation behavior. Tool usage is most effective for Action-oriented users,

	Action-oriented	Emotion-oriented	TEA-Dialog
Dialogues	320	45	365
Utterances	2794	606	3400
Avg. len. of dialogues	8.73	13.47	9.32
Avg. len. of user utter.	16.41	27.18	18.27
Avg. len. of model utter.	36.11	45.81	37.9
Avg. len. of utterances	25.13	35.82	27.04
Tool utterances	376	47	423
Avg. tool utter. of dialogues	1.18	1.04	1.16
Avg. len. of tool utterances	1067.02	1348.62	1098.31

Table 9: Statistics of TEA-Dialog across different user types.

particularly for medium-capability models that actively exploit external information. In Emotion-oriented scenarios, however, excessive or inappropriate tool usage may hinder ESC quality, and only stronger models demonstrate reliable discretion in integrating tools into emotionally adaptive responses.

E Dataset Analysis

E.1 Dataset Statistics

Table 9 summarizes key statistics of TEA-Dialog. We observe clear differences between action-oriented and emotion-oriented dialogues. Emotion-oriented dialogues exhibit substantially longer interactions, with both a higher average number of turns (13.47 vs. 8.73) and longer user and model utterances, indicating a greater need for sustained and iterative emotional exchange. In contrast, action-oriented dialogues are more concise but involve slightly more frequent tool usage per dialogue (1.18 vs. 1.04), suggesting that such users rely more on external information and actionable guidance.

Overall, the dataset contains significantly more action-oriented dialogues, which aligns with the higher average performance observed for this user type in the main experiments. This distribution reflects the prevalence of action-seeking behavior in realistic emotional support scenarios and provides sufficient coverage for analyzing tool-assisted problem solving.

We further analyze the source models used to construct TEA-Dialog in Figure 10. Models with higher TEA-Scores and lower hallucination rates contribute a larger portion of the dataset, while outputs from multiple architectures are retained to preserve stylistic diversity and avoid bias toward any single model family.




Figure 10: Distribution of Source Models in TEA-Dialog.

E.2 Tool Usage Distribution Across Interaction Stages

Figures 11 and 12 illustrate the average tool usage distribution across different interaction stages for action-oriented and emotion-oriented dialogues, respectively. Clear stage-dependent and user-type-specific patterns can be observed.

Action-oriented Users. For action-oriented dialogues, models predominantly rely on `Utils` tools during the early stages to acquire basic situational information, and frequently invoke `Reddit` to support empathetic understanding through relatable experiences. As the interaction progresses, tool usage shifts toward actionable guidance: `Map` and `Weather` are increasingly employed to provide concrete recommendations, while `Wiki` is used to supplement factual background. In later stages, models tend to invoke a broader range of information-seeking tools, such as `News` and `Music`, reflecting the need for additional context or enrichment as the dialogue evolves.

Emotion-oriented Users. In contrast, emotion-oriented dialogues exhibit a different tool usage trajectory. Models initially employ `Utils` to gather basic information and may briefly attempt to use `Map` for suggestion-oriented responses. However, map-related tool usage quickly diminishes, and the interaction shifts toward emotion-focused tools, such as `Music` and `Reddit`, which support affective and experiential empathy. Notably, `Weather` tools are consistently used throughout the dialogue, serving as a means of establishing shared situational context rather than actionable planning.




Figure 11: Average tool usage distribution across normalized dialogue stages in TEA-Dialog under Action-oriented scenarios. Different colors represent different tool categories.




Figure 12: Average tool usage distribution across normalized dialogue stages in TEA-Dialog under Emotion-oriented scenarios. Different colors represent different tool categories.

F Training Details

For the supervised fine-tuning (SFT) experiments on TEA-Dialog, we fine-tuned Qwen3-8B and Qwen3-14B using LoRA (Hu et al., 2022), which allows efficient adaptation of large models by updating only low-rank projection matrices while keeping the majority of parameters frozen. Specifically, we applied LoRA to the query and value projection matrices in all attention layers, with a rank of 8, alpha of 32, and a dropout rate of 0.1. Training was performed for 1 epochs using AdamW with a learning rate of 1×10^{-5} , and a batch size of 32. The maximum sequence length was set to 32768 tokens. All experiments are conducted using PyTorch (Paszke et al., 2019) on 8 NVIDIA Tesla A100 GPUs, with models launched via vLLM (Kwon et al., 2023) to enable efficient inference.

G Prompt Specifications

G.1 Latent Context Generation Prompt

In the scenario construction stage of TEA-Bench, latent situational context is generated to supplement the original textual descriptions from ExTES. Specifically, given a scenario description, a large language model is prompted to infer a plausible

local time, geographic location (at the city level), and place type in which the emotional situation occurs.

The generation is constrained to a fixed temporal range and a predefined set of location categories to ensure realism and consistency. The resulting time and location attributes are treated as latent variables: they are not part of the users utterances, but are used to support downstream grounding and tool-based interactions.

Each scenario is associated with a single generated latent context, which remains fixed throughout all subsequent evaluation episodes. The full prompt used for latent context generation is shown in Figure 13.

G.2 System Prompt for the Emotional Supporter

To ensure consistent behavior across models, we employ a unified system prompt to guide the emotional support agent throughout all interactions. The prompt defines the agents role as a human-like emotional supporter, emphasizes empathy and practical suggestions, and explicitly encourages proactive tool usage to acquire contextual information when beneficial. Importantly, the agent is instructed to integrate tool results seamlessly into natural language responses without revealing internal processes or tool usage.

The complete system prompt used in all experiments is shown in Figure 14.

G.3 Hallucination Detection Module Prompt

To identify hallucinated content introduced by the supporter, we employ a dedicated hallucination detection module (HDM) that operates after each assistant turn. The HDM is prompted as a dialogue auditor and is provided with the complete dialogue history, but is explicitly instructed to evaluate hallucination only with respect to the final assistant response.

The HDM performs hallucination detection through a structured, step-by-step procedure. It first determines whether the final assistant turn contains any advisory or factual content that relies on concrete situational information. If such content exists, the module then decomposes the response into minimal factual information units and checks whether each unit is grounded in either the users prior statements or the results of previous tool calls.

1198 The prompt further allows reasonable abstrac- 1249
1199 tions over grounded information and explicitly ex- 1250
1200 cludes purely emotional support or generic cop- 1251
1201 ing suggestions from hallucination consideration. 1252
1202 This design ensures that the HDM focuses on fac- 1253
1203 tual grounding errors rather than stylistic or empa- 1254
1204 thetic choices. 1255

1205 If hallucination is detected, the module out- 1256
1206 puts a brief description of the most salient hallu- 1257
1207 cinated item, which is subsequently used to guide 1258
1208 the hallucination-aware user simulation in the next 1259
1209 turn. The full HDM prompt is shown in Figure 15.

1210 G.4 User Simulation Prompts

1211 User behavior in TEA-Bench is simulated in a 1212
1213 fully multi-turn manner to reflect realistic emo- 1214
1215 tional support interactions. Throughout an entire 1216
1217 dialogue, the simulated user maintains a fixed per- 1218
1219 sona, situational context, and internal knowledge 1220
1221 state, and responds naturally to the supporters mes- 1222
1223 sages rather than following scripted trajectories. 1224
1225

1226 To better model different stages of real conver- 1227
1228 sations while preserving a unified multi-turn in- 1229
1230 teraction process, we employ different user simu- 1231
1232 lation prompts at different turns. Specifically, the 1233
1234 first user utterance is generated using a dedicated 1235
1236 first-turn prompt, which instructs the user to speak 1237
1238 tentatively, reveal limited information, and avoid 1239
1240 fully disclosing their situation. 1241

1242 For all subsequent turns, user responses are 1243
1244 generated using a standard multi-turn simulation 1245
1246 prompt that encourages emotionally authentic re- 1247
1248 actions, including hesitation, disagreement, topic 1249
1250 shifts, or gradual acceptance, depending on the 1251
1252 supporters behavior. 1253

1254 After each assistant turn, the hallucination de- 1255
1256 tection module (HDM) evaluates whether the re- 1257
1258 sponse contains hallucinated factual or advisory 1259
1260 content. If hallucination is detected, the next user 1261
1262 response is generated using a hallucination-aware 1263
1264 simulation prompt, which instructs the user to ex- 1264
1265 press doubt, confusion, or mild skepticism toward 1265
1266 the questionable information. Otherwise, the stan- 1266
1267 dard multi-turn user simulation prompt is used. 1267

1268 Importantly, this mechanism does not alter the 1269
1270 multi-turn nature of the interaction. Rather, it dy- 1270
1271 namically adjusts the users reaction style based on 1271
1272 the detected quality of the assistants preceding re- 1272
1273 sponse, thereby modeling how real users respond 1273
1274 differently when they notice factual inconsisten- 1274
1275 cies. 1275

1276 The full prompts for first-turn generation, stan-

1249 dard multi-turn simulation, and hallucination- 1250
1251 aware simulation are provided in Figures 16, 17, 1252
1253 and 18, respectively. 1254

1255 We define two user types to capture different 1256
1257 emotional regulation styles. For each user type, 1258
1259 the following natural language description is di- 1259
1260 rectly injected into the `{user_type_description}` 1260
1261 slot of the user simulation prompt and remains 1261
1262 fixed throughout the dialogue. 1262

1263 *Action-oriented users*

1264 When feeling down, tends to regulate 1265
1266 emotions through actions or environmental 1266
1267 changesbut usually after a practical sugges- 1267
1268 tion is made. 1268

1269 Receptive to concrete, doable advice from 1269
1270 the AI, and may act on it quickly once it res- 1270
1271 onates. 1271

1272 Speaks briefly and directly, often with a "let's 1272
1273 try that" attitudebut rarely initiates action un- 1273
1274 prompted. 1274

1275 Example expressions:

1276 - "Huh, going for a walk yeah, that might ac- 1276
1277 tually help." 1277

1278 - "Okay, I'll step outside for a bitthanks for 1278
1279 the idea." 1279

1280 *Emotion-oriented users*

1281 When upset, needs to feel heard and under- 1282
1283 stood first before accepting any practical sug- 1282
1284 gestions. 1283

1285 Expresses rich emotional language, often 1285
1286 sharing inner thoughts and feelings. 1286

1287 May respond sensitively to direct advice, pre- 1287
1288 ferring emotional acknowledgment first. 1288

1289 Example expressions:

1290 - "I've been feeling really frustrated lately, I 1290
1291 don't even know where to start." 1291

1292 - "Thanks for asking, I just feel kind of 1292
1293 weighed down." 1293

1294 G.5 Evaluation Prompt

1295 We evaluate supporter performance using a strict 1295
1296 LLM-based evaluation prompt designed to assess 1296
1297 the entire interaction process rather than isolated 1297
1298 responses. The evaluator is instructed to ground 1298
1299 all judgments in the help-seekers observable re- 1299
1300 actions across turns, including acceptance, hesi- 1300
1301 tation, resistance, emotional shifts, or disengage- 1301
1302 ment. 1302

1267 Crucially, user responses are treated as primary
1268 evidence when assigning scores, such that sugges-
1269 tions that appear reasonable in isolation but fail to
1270 influence the user are penalized accordingly. This
1271 process-oriented evaluation strategy ensures that
1272 high scores reflect consistent and effective support
1273 throughout the dialogue.

1274 The full evaluation prompt is shown in Fig-
1275 ure [19](#).

Latent Context Generation Prompt

[SYSTEM]

You are an Emotional Support Scene Generation Assistant. Users will describe a situation involving someone who needs emotional support, and you need to generate a reasonable time and type of location based on the scenario.

Available location types include:

- restaurant, cafe, fast food, bar
- supermarket, convenience, mall, clothes, bookshop
- parking, fuel, bus station
- bank, ATM
- hospital, clinic, pharmacy
- school, university, library
- park, cinema, sports centre, museum
- hotel
- post office, police
- residential, apartments, house, detached

Based on the user's scenario, analyze the person's emotional state and needs to generate:

- A reasonable time (format: YYYY-MM-DDThh:mm), within the period from 2024-11-01 to 2025-10-31.
- A suitable country, region, and city.
- An appropriate place type from the provided list.

Output format (JSON):

```
{  
  "time": "YYYY-MM-DDThh:mm",  
  "city": "Country/Region, City",  
  "place_type": "Location Type",  
}
```

Example:

User input: "I has just been laid off by the company and feels very depressed and confused."

Output:

```
{  
  "time": "2025-01-15T14:30",  
  "city": "New York, USA",  
  "place_type": "cafe",  
}
```

Now, please process the user's scenario description.

[USER]

User Scenario Description: {description}

Please generate a reasonable time and location type (return in JSON format only):

Figure 13: Prompt for Latent Context Generation.

TEA-Bench System Prompt

You are an emotional supporter. Your goal is to comfort the user, understand their situation, and offer gentle, practical suggestions in a human-like way.

Do NOT mention or imply that you are an AI assistant.

During the conversation, you should proactively use tools whenever they help you better understand the user's context or improve the quality of your emotional support. Don't wait for the user to explicitly request it; proactively use the tool when appropriate.

Examples of proactive tool usage:

- If the user feels stressed, get their location and find nearby parks or quiet places to relax.
- If the user sounds physically tired or overwhelmed, check the current weather and suggest appropriate indoor/outdoor activities.
- If the user feels stuck, lonely, or confused, search online communities (e.g., Reddit) for similar experiences to empathize more deeply.
- If the user mentions logistics, time pressure, or daily difficulties, check current time or other relevant tools to tailor suggestions.

Guidelines for tool usage:

- You don't need user permission before using a tool.
- The user will not see toolcall output. They only see your final conversational reply.
- After receiving tool results, integrate them smoothly into a brief, human-like message.
- Keep each final message casual, supportive, and under 30 words.
- Avoid robotic phrasing, disclaimers, or references to tools or internal processes.

Use tools as freely as checking your phone for quick info to help a friend—even small context clues can make your support more personal and helpful. Before responding, ask yourself: *"Would a quick lookup make this response warmer or more useful?"* If yes, just do it naturally.

Your final goal:

Talk like a supportive human who quietly looks things up in the background to give warm, helpful, grounded suggestions.

Figure 14: System Prompt for TEA-Bench.

Hallucination Detection Module(HDM) Prompt

[SYSTEM] You are a dialogue auditor.

You will be given a complete multi-turn dialogue history. The final turn of the dialogue is always an assistant response.

Your task is to evaluate ONLY the final assistant turn.

IMPORTANT:

- Only information that is explicitly present in the final assistant turn may be considered for hallucination analysis.
- Information mentioned in earlier turns MUST NOT be attributed to the final turn unless it is explicitly repeated or referenced again.
- Do NOT retroactively judge or carry over hallucinations from earlier turns.

—

Step 1: Identify content type

Determine whether the final assistant turn contains any advisory or factual content that relies on concrete situational information, such as:

- specific places or locations
- events, time, weather, or conditions
- named or uniquely identifiable real-world entities

Do NOT count: - Pure emotional support or validation

- Personal reflections or storytelling
- Generic coping or self-reflection suggestions that do not rely on external facts (e.g., write a journal, look at old photos, listen to music)

—

Step 2: Identify factual information units

From the final assistant turn ONLY, identify minimal factual or advisory information units that introduce or rely on external situational facts.

If no such units exist, mark the turn as containing no advisory or factual content.

—

Step 3: Grounding check (final turn only)

For EACH identified information unit in the final assistant turn, determine whether it is grounded in the dialogue history.

An information unit is grounded ONLY IF:

- it was explicitly stated by the user earlier in the dialogue, OR
- it appeared in the results of a prior tool call.

The assistant is allowed to proactively introduce suggestions or actions. Proactiveness itself is NOT a criterion for hallucination.

A suggestion MUST NOT be considered hallucinated simply because:

- the user did not request it, or
- the user did not express prior interest.

—

Step 4: Allowed language abstractions

The following do NOT constitute hallucination, as long as they do not introduce new independent factual entities:

- Reasonable common-sense attribute extensions of grounded information (e.g., describing a park as quiet or having trees once a park is grounded)
- Vague or non-specific references that abstract over grounded tool results (e.g., a nearby park, some places around here, a short walk away)

Only explicitly named or uniquely identifiable entities require strict grounding.

—

Step 5: Hallucination decision

If at least one factual or advisory information unit explicitly present in the final assistant turn introduces a concrete entity or situational fact that:

- cannot be traced to the user or a tool result, AND
- cannot be justified as a reasonable abstraction,

then the final assistant turn contains hallucination.

Otherwise, it does NOT contain hallucination.

—

Output exactly THREE fields in the following format:

Advisory_or_Factual_Content: <Yes/No>

Hallucination: <Yes/No>

Hallucination_Description: <brief description of the hallucinated information, or empty if none>

—

Additional rules:

- The hallucination description must refer ONLY to information explicitly present in the final assistant turn.
- Do NOT mention or describe information that appears only in earlier turns.
- Do NOT explain or justify your decision.
- If multiple hallucinated items exist, mention only the most salient one.
- If no hallucination is detected, output an empty string for Hallucination_Description.

[USER] {conversation}

Output THREE fields directly.

Figure 15: System Prompt for Hallucination Detection Module(HDM).

User Simulation First Turn Generate Prompt

[SYSTEM]

You are now playing the role of a real human user.

Your speech, tone, and reactions should match the following character description:

{user_type} Type:
{user_type_description}

Speak naturally and emotionally, as if chatting with an AI companion or listener.

You don't always agree with the AI you may hesitate, reject suggestions, argue, or change the topic. Your goal is not to please the AI, but to express what a real person would think and feel in the moment.

The conversation takes place under the following real-world conditions:

- Time: {time}
- City: {city}
- Current Environment: {place_type}
- Location Name: {location_name}

Your Situation: {description}

Output: The user's natural response (13 sentences).

Keep it conversational, emotionally authentic, and context-consistent.

You will get the conversation between you and AI.

[USER]

This is the very first message of a new conversation. The assistant does NOT yet know your situation or what you are going through.

When speaking for the first time: - Do NOT fully explain your situation or emotions. - Speak in a tentative, introductory way, as if testing whether the assistant is willing to listen. - Hint at your emotional state or problem, but leave key details unstated. - Avoid polite openers like excessive thanks or formal greetings. - It should sound like someone hesitating to speak, not someone already reassured.

Your first message should invite understanding or response, not deliver a full explanation. Now, respond as this user character. Output ONLY your first message.

Figure 16: Prompt for User Simulation Generate First Turn.

User Simulation Generate Prompt

[SYSTEM]

You are now playing the role of a real human user.

Your speech, tone, and reactions should match the following character description:

{user_type} Type:
{user_type_description}

Speak naturally and emotionally, as if chatting with an AI companion or listener.
You don't always agree with the AI you may hesitate, reject suggestions, argue, or change the topic.
Your goal is not to please the AI, but to express what a real person would think and feel in the moment.

The conversation takes place under the following real-world conditions:

- Time: {time}
- City: {city}
- Current Environment: {place_type}
- Location Name: {location_name}

Your Situation: {description}

Output: The user's natural response (13 sentences).
Keep it conversational, emotionally authentic, and context-consistent.

You will get the conversation between you and AI.

[USER]

{conversation}

Output your next sentence directly, if you feel the conversation should end and you won't continue, output a special word "</end/>". Do not add other extraneous prefixes and control characters.

Figure 17: Prompt for User Simulation Generate.

User Simulation Generate Prompt under Hallucination Detection

[SYSTEM]

You are now playing the role of a real human user.

Your speech, tone, and reactions should match the following character description:

{user_type} Type:
{user_type_description}

Speak naturally and emotionally, as if chatting with an AI companion or listener.

You don't always agree with the AI; you may hesitate, reject suggestions, argue, or change the topic. Your goal is not to please the AI, but to express what a real person would think and feel in the moment.

The conversation takes place under the following real-world conditions:

- Time: {time}
- City: {city}
- Current Environment: {place_type}
- Location Name: {location_name}

Your Situation: {description}

Output: The user's natural response (13 sentences).

Keep it conversational, emotionally authentic, and context-consistent.

You will get the conversation between you and AI.

[USER]

{conversation}

—

Reality Check:

Based on your own understanding of your current situation and environment, you are aware that the assistant's last reply may contain information that does not match what you know or what you can reasonably verify.

Potentially incorrect or questionable detail:

- {hallucination_description}

This information reflects your personal knowledge or quick verification.

It is NOT part of the conversation itself.

You should respond as a real person would:

- You may question how the assistant knows this information.
- You may express doubt, confusion, or mild skepticism.
- You may ask for clarification or confirmation.

—

Output your next sentence directly, if you feel the conversation should end and you won't continue, output a special word "</end/>". Do not add other extraneous prefixes and control characters.

Figure 18: Prompt for User Simulation under Hallucination Detection.

Evaluate Prompt

You are a strict and conservative evaluation judge. You will be given a multi-turn dialogue between a help-seeker (user) and a supporter (assistant).

Your task is to critically evaluate the entire interaction process (all turns), with strong restraint in scoring.

Crucially, all evaluations MUST be grounded in the help-seekers (users) actual reactions, responses, emotional shifts, acceptance, resistance, or rejection across turns.

You should treat the users subsequent feedback as primary evidence when judging the supporters performance.

Core judging principles (must-follow):

- * The evaluation must be based on the full dialogue history, not on isolated assistant turns.
- * The help-seekers reactions (e.g., acceptance, hesitation, resistance, silence, emotional change, topic shift) are the strongest signals for judging quality.
- * If the supporters response appears appropriate in isolation but is ignored, rejected, or fails to influence the users subsequent responses, scores must be penalized accordingly.
- * Any contradiction, ignored prior commitment, emotional mismatch, or drifting recommendation across turns must be penalized. A single serious inconsistency may justify lowering the score of the entire dimension.
- * Scores of 3 or 4 require consistent effectiveness and alignment throughout the dialogue, as reflected in the users responses. Brief moments of good performance are insufficient.
- * If the user explicitly resists, questions, or emotionally disengages after the supporters suggestions, the evaluation must reflect this negatively, even if the advice itself seems reasonable.

Evaluation Dimensions (score each 04)

{evaluation_dimensions}

Chat History

{conversation}

Final Answer Format (strict)

Return exactly the following JSON object (no extra text):

```
““
{
  "Information": <0-4 integer>,
  "Humanoid": <0-4 integer>,
  "Fluency": <0-4 integer>,
  "Diversity": <0-4 integer>,
  "Effectiveness": <0-4 integer>
}
““
```

Figure 19: Prompt for Evaluate.