



ARTICLE

TATA: A Trust-Aware Task-Oriented Agent Framework for Industrial Intelligence Scenarios

Pan Li^{1,2}, Zhi Li³ and Yingyou Wen^{2,*}

¹School of Computer Science and Engineering, Northeastern University, Shenyang, China

²Neusoft Research Institute, Neusoft Group, Shenyang, China

³School of Information Science and Engineering, Shenyang Ligong University, Shenyang, China

*Corresponding Author: Yingyou Wen. Email: wenyinyou@mail.neu.edu.cn

Received: 29 March 2026; Accepted: 14 May 2026; Published: 15 June 2026

ABSTRACT: The rapid advancement of edge intelligence in Industrial Internet of Things (IIoT) is transforming human-computer interaction from conventional “command execution” to complex “human-AI deep collaboration”. Within such safety-critical industrial environments, establishing robust mutual understanding and trust mechanisms becomes a significant prerequisite for decision reliability and efficiency. However, existing industrial interaction systems predominantly focus on task progression and explicit command responses, lacking fine-grained, dynamic tracking of operators’ trust states, cognitive evolution, and behavioral dynamics. Moreover, current LLM-based user simulation in evaluation often exhibit an “over-cooperation” bias, failing to capture the cognitive conflicts and trust crises characteristic of high-pressure, high-risk industrial conditions. To address these challenges, we first propose a trust-aware user behavior model, which utilizes an LLM-parameterized Hidden Markov Model (HMM) to formalize collaborative trust as a dynamic latent variable, thereby structurally characterizing the psychological and behavioral dynamics of operators across multi-turn interactions. Building on this, we introduce TATA, a task-oriented agent framework integrating trust-awareness and cognitive alignment. Through a dual-track state monitoring mechanism and adaptive interaction policy coordination, TATA effectively advances collaborative tasks and fosters relationship maintenance in realistic collaborative environments. Comprehensive evaluations on six industrial task scenarios demonstrates that TATA achieves an optimal balance between collaboration depth and task efficiency, outperforming the strongest baseline by achieving 1.6 to 2.6 times higher collaboration efficiency and an absolute increase of over 15 percentage points in task completion rate. These findings provide valuable insights for developing resilient and adaptive deep human-AI collaboration tailored to IIoT scenarios.

KEYWORDS: Industrial intelligence; agent; human-AI collaboration (HAC)

1 Introduction

The rapid convergence of AI and the Industrial Internet of Things (IIoT) is fundamentally transforming modern industrial systems, enabling intelligent monitoring, predictive maintenance, and data-driven decision-making across smart factories, power grids, and logistics networks [1]. At the center of this transformation, AI agents powered by Large Language Models (LLMs) [2,3] have emerged as a critical operational engine, supporting complex industrial decision-making and task execution through natural language interaction [4]. This shift is driving a transition in human-machine interaction (HMI) paradigms, moving from traditional rule-based command matching toward conversational human-AI collaboration (HAC) systems. Under this emerging interaction mode, operators can engage in multi-turn dialogue with AI

agents to collaboratively address highly specialized industrial tasks through iterative communication rather than one-shot instruction execution [2,3,5].

Nevertheless, the effective operation of current AI-driven industrial collaboration systems fundamentally relies on the assumption that *operators are able to provide clear and sufficiently detailed task specifications*. This prerequisite often fails in dynamic industrial scenarios that require exploratory problem-solving and domain expertise [6]. In practice, when confronted with unexpected situations, industrial operators can typically observe only preliminary symptoms of a problem. Despite possessing extensive tacit experience, they often struggle to fully articulate precise system requirements, potential constraints, and underlying logic at the early stages of a task [7].

Consider a real-world industrial scenario: an operator is troubleshooting an unexpected temperature anomaly in a chemical reactor. Under the traditional command-execution interaction mode, the operator must diagnose the exact cause by themselves and issue explicit commands (e.g., *Check the status of cooling valve V-102.*). This places the entire cognitive burden on the human operator, who may not have access to complete real-time monitoring data. In contrast, in a collaborative dialogue scenario, the AI agent proactively engages in joint problem-solving. It can ask exploratory questions (e.g., *Has the coolant flow rate changed in the past hour?*) and adjust its subsequent dialogue policy based on the operator's responses (e.g., *I'm not sure if that's the issue, but I recall encountering a similar situation last month.*). This collaborative mode reduces the operator's cognitive load and builds problem localization and shared understanding through multi-turn interactions, which are the capabilities that are critically important in industrial edge computing environments with extremely high requirements for safety and real-time performance. If an AI agent engages in "premature execution" or makes misaligned decisions based on incomplete information, it will not only waste the already limited computational and communication resources but also trigger cascading operational risks. Ultimately, this form of ineffective (or even harmful) interaction can rapidly erode operators' trust in AI systems [8,9], thereby directly undermining the adoption and effectiveness of HAC decision-making.

To address these challenges, recent research has begun to explore deep human-AI collaboration mechanisms from a cognitive science perspective [6,7]. Based on the Common Ground theory [10,11], Li et al. proposed the Cognitive Collaborative Dialogue (CCD), a task that emphasizes that *the user's understanding of the problem and the cognition of the solution co-evolve through the interaction process between both parties*. Through collaborative dialogue, AI agents can gradually guide users to jointly construct a shared understanding of the problem space and ultimately generate customized solutions. This insight aligns closely with the requirements of IIoT scenarios. In complex industrial decision-making, it is often necessary to deeply integrate AI-analyzed data insights and collected industrial environment information with the tacit expertise of human experts. This "cognitive alignment" is a gradual process that requires the continuous establishment of mutual trust and consensus, rather than a one-time execution of instructions.

Although the CCD framework has demonstrated effectiveness in collaborative interaction and task completion at a macro level, existing approaches lack comprehensive modeling and fine-grained tracking the dynamic trust and behavioral evolution of both human and agents during dialogue. Specifically, current practices in evaluating multi-turn interaction systems and constructing user simulators predominantly rely on mechanical behavior switching [12], predefined information disclosure strategies [13], or static persona templates [14], making it difficult to capture the psychological dynamics of real human users as they evolve with content and collaboration depth throughout the conversation. More critically, influenced by RLHF in model pretraining, LLM-driven user simulators often exhibit an excessively cooperative attitude [15], tending to accommodate any guidance from the agent, and rarely displaying genuine resistance, skepticism, or emotional fluctuations even when cognitive conflicts or trust decline occur. This idealized simulating behavior obscures the trust crises and collaboration barriers that agents may encounter in real-world

applications. Simultaneously, AI agents themselves lack explicit mechanisms to perceive and respond to cognitive shifts (e.g., *cognitive evolution, adjustment, or conflict*) during each interaction turn. Under such settings, it becomes difficult to simulate authentic HAC processes, ultimately leading to overestimated evaluation results and creating “fake success” in collaborative effectiveness.

To bridge this critical gap, we tackle the aforementioned challenges by focusing on both user modeling and agent design, with a focus on achieving robust HAC in complex industrial decision-making environments. As shown in Fig. 1, we first propose a **trust-aware user behavior model**. This method formalizes users’ “collaborative trust level” during dialogue as dynamic latent variables within an LLM-parameterized Hidden Markov Model (HMM), thereby structurally characterizing the psychological and behavioral dynamics encountered in real-world interactions. Based on these insights, we present **TATA**, a task-oriented agent framework integrating trust-awareness and cognitive alignment. TATA comprises three core modules: a dialogue trajectory planner, a dual-track state monitor, and an adaptive interaction coordinator. The planner forecasts future dialogue trajectories; the monitor jointly assesses task-level cognitive development and relationship-level affective vigilance signals. Guided by a partially observable Markov decision process (POMDP), the coordinator leverages these dual-track observations to dynamically adjust interaction policies, balancing task advancement and trust maintenance throughout the collaboration process. Additionally, we further introduce a novel set of targeted trust dynamics evaluation metrics for more fine-grained assessment of HAC quality throughout the interaction.

In summary, our main contributions are as follows:

- We propose a trust-aware user behavior modeling approach that represents collaborative trust as a dynamic latent variable, establishing a rigorous and practically grounded methodological basis for evaluating HAC in industrial decision-making scenarios.
- We develop the TATA agent framework, which integrates dual-track information fusion and adaptive policy regulation mechanisms to enhance collaboration capabilities in dynamic industrial environments.
- We introduce a process-oriented trust dynamics evaluation system, offering new insights into the deployment of trustworthy AI collaboration systems in industrial intelligence environments.

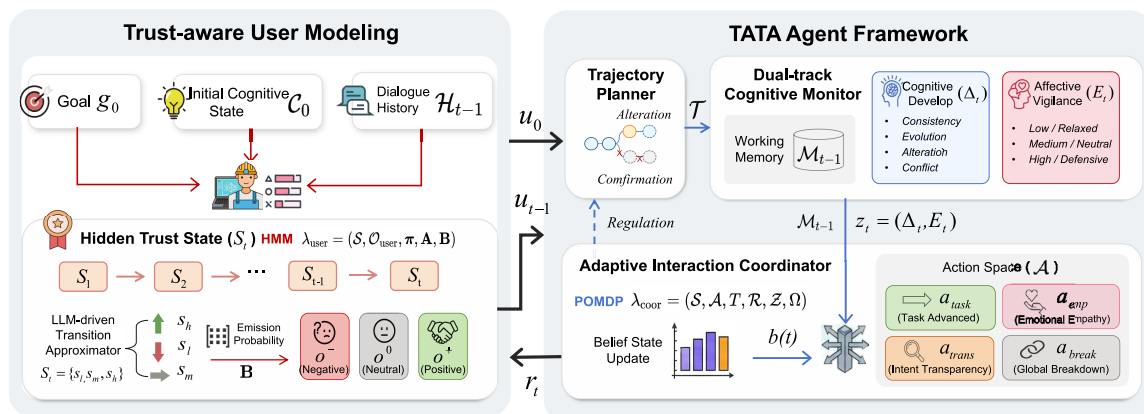


Figure 1: Overview of the whole interaction framework. (1) Trust-aware user modeling (left): using an LLM-driven HMM to model users’ trust dynamics and behavior based on internal state and dialogue history; (2) TATA framework (right): featuring a dual-track state monitor for tracking task and emotional states, and an POMDP-based coordinator to update belief states and select optimal policies for response generation and trajectory regulation.

2 Related Work

2.1 Task-Oriented Dialogue Agents and Industrial Applications

Task-oriented dialogue systems are transitioning from traditional pipeline architectures to LLM-driven agents [16,17]. Current studies primarily leverage LLMs to enhance task planning [18,19], dialogue state tracking [20], and tool invocation [21,22] to improve execution efficiency. However, most such agents typically employ self-iterative unidirectional mechanisms that lack user engagement, rendering them susceptible to failure when user information proves insufficient or ambiguous [23]. Consequently, proactive agents [24] have emerged as a prominent research focus. Moving beyond passive response patterns, these agents optimize interactions by actively clarifying ambiguities, guiding the conversational flow, and preemptively proposing solutions. Representative approaches include specialized model training [25], Chain-of-Thought (CoT) reasoning [26], and knowledge graph inference [27].

Regarding practical applications, the capabilities of LLMs and agents have catalyzed an expansion of task-oriented dialogue agents from everyday domains to complex industrial production [28]. Under the “Human-Centric” paradigm of Industry 5.0, human-machine interaction within IIoT environments is evolving from conventional unidirectional static command parsing toward dynamic, mixed-initiative decision-making and deep collaboration [29,30]. To address the challenges posed by high-noise, highly dynamic environments with stringent fault-tolerance costs, recent research prioritizes intent recognition and mixed-initiative dialogue mechanisms in complex scenarios to enable more flexible task allocation and execution [31–33]. For instance, latest studies have proposed a mixed-initiative dialogue system that dynamically assesses human collaborative willingness based on historical interactions to optimize human-machine task allocation [33]. Additionally, for high-risk industrial collaborative scenarios, researchers have introduced dialogue-based explainable safety mechanisms designed to communicate underlying safety constraints and operational logic transparently to operators, thereby maintaining shared situational awareness during task interruption and recovery [34].

Nevertheless, existing proactive interaction strategies exhibit significant limitations. Most studies still conceptualize humans as static instruction sources or assume that users exhibit perfect cooperation during evaluation [33]. Consequently, the proactivity of current systems remains largely confined to few-turn information acquisition and slot filling, lacking mechanisms for dynamic trust tracking, goal negotiation, and cognitive alignment over extended task durations [35]. These constraints substantially limit their practical applicability in authentic industrial intelligence scenarios.

2.2 User Simulator

User simulation serves as a critical technique for generating interactive dialogue data, widely employed for evaluating and training conversational agents [36]. Compared with conventional user simulators, LLM-based approaches exhibit superior cross-domain adaptability and behavioral diversity [37,38]. LLMs emulate users through persona simulation [14], role-playing [39], and individual behavior modeling [40]. Typical research approaches primarily encompass prompt engineering [41], retrieval-augmented generation [42], efficient fine-tuning [43], and reinforcement learning with preference optimization [44], among others. However, most of these methods rely on static user profiles or are confined to the task objectives themselves, overlooking the dynamic cognitive processes and collaborative participation mechanisms inherent in human-AI interaction. Furthermore, existing LLM-based user simulators may inherently exhibit an “over-cooperation” or “sycophancy” bias [15,45], tending to overly accommodate the other party and avoid cognitive conflicts. In user simulation scenarios for evaluating human-AI collaboration, an overly compliant simulator fails to reproduce the cognitive conflicts, skepticism, or trust crises that may arise during

real collaborative processes, leading to inflated evaluation results, which is the problem that our proposed approach explicitly addresses.

3 Methodology

3.1 Task Formulation

As a significant extension of traditional Task-Oriented Dialogue (TOD) in complex scenarios, **Cognitive Collaborative Dialogue (CCD)** is proposed to address tasks characterized by high exploration, professional expertise, and personalization [6]. Unlike conventional TOD systems that provide predetermined answers or mechanically execute “slot-filling” queries, CCD necessitates multi-turn progressive interactions between agents and users to co-construct shared understanding of underspecified problems [13].

Formally, a CCD task with N turns is defined as: $\mathcal{D}_N = \{g_0, \mathcal{C}_0, \mathcal{T}, \mathcal{H}_N, \mathcal{C}_{\text{final}}\}$, where g_0 is the user’s preliminary goal, \mathcal{C}_0 is the initial cognitive state set representing the user’s initial understanding (e.g., *background context, constraints, preferences*). \mathcal{T} is a serialized collaborative blueprint, defined in [6] as a “cognitive trajectory”, which outlines potential directions and constructs pathways for deeper exploration in subsequent dialogues. $\mathcal{H}_N = \{u_0, r_0, u_1, r_1, \dots, u_{N-1}, r_{\text{final}}\}$ represents the dialogue history, and $\mathcal{C}_{\text{final}}$ is the final evolved state of \mathcal{C}_0 .

However, a notable challenge exists in basic CCD task settings: **users’ internal trust states and cognitive evolution are implicit and unobservable**. As a result, existing frameworks tend to advance tasks in a mechanical manner, lacking mechanisms to perceive and resolve implicit cognitive conflicts. More importantly, most existing evaluation environments rely on static user simulators that may exhibit a tendency toward “over-cooperation” [15], further obscuring the trust crises that could arise in real HAC.

Therefore, to authentically assess and address these issues, it is imperative to first model real user behaviors dynamically within the CCD context (see Section 3.2), and subsequently design agent frameworks capable of operating effectively in such complex and dynamic environments (see Section 3.3). The Overview of the trust-aware human-AI collaborative dialogue framework is shown in Fig. 1.

3.2 Trust-Aware User Behavior Modeling

Effective collaboration relies on the co-evolution of cognitive states between dialogue participants and the establishment of trust in this process. [10,46]. However, conventional user modeling approaches often oversimplify user behavior by employing static persona templates or random information disclosure strategies, neglecting the reality that user actions in real interactions are dynamic manifestations of evolving internal trust and cognitive states [7,13,14].

In highly exploratory collaboration scenarios, “trust” is not a static value but a complex cognitive process that dynamically evolves through multidimensional interactions [47]. Based on Trajectory Epistemic Network Analysis (T-ENA) conducted on real human subjects interacting with conversational agents¹ [48], Li et al. demonstrated that human trust manifests interdependently through analytic processes (e.g., *systematic evaluation and scrutiny*) and affective processes (e.g., *emotional experiences*).

Drawing on these theoretical and empirical foundations, we propose a trust-aware user behavior modeling approach (Fig. 1-Left) that formalizes the user’s “cognition–trust–behavior” process during multi-turn collaboration as an LLM-parameterized Hidden Markov Model (HMM), denoted by $\lambda_{\text{user}} = (\mathcal{S}, \mathcal{O}_{\text{user}}, \boldsymbol{\pi}, \mathbf{A}, \mathbf{B})$. This framework provides a rigorous operational definition for *collaborative trust*.

¹The study recruited 24 participants and collected 1981 lines of dialogue text during the completion of 12 decision-making tasks with a conversational agent.

Notably, distinct from trustworthiness as an intrinsic property of AI models [49], *collaborative trust* in this study is defined as the operator's dynamic, latent cognitive-affective state. This state reflects their willingness to rely on the AI agent, disclose task-relevant information to it, and follow its guidance during multi-turn collaborative problem-solving in industrial scenarios. Collaborative trust is a relational construct that co-evolves through the interaction process. Formally, it is defined as a discrete hidden state variable $S_t \in \mathcal{S} = \{s_l, s_m, s_h\}$, corresponding to Low, Medium, and High trust levels, respectively. This trust state remains unobservable to the agent during each dialogue turn and can only be indirectly inferred through the user's explicit behavior. The observable variable $O_t \in \mathcal{O}_{\text{user}} = \{o^-, o^0, o^+\}$ denotes the user's explicit interaction behavior at t -th turn, categorized into three behavioral modes [48]:

- **Negative Collaboration (o^-):** characterized by high cognitive vigilance and rigorous system scrutiny, with linguistic features such as defensive questioning, topic shifting, and minimal information disclosure, accompanied by negative valence and high arousal emotional expressions;
- **Neutral Collaboration (o^0):** routine interaction with moderate information sharing and emotionally neutral stance;
- **Positive Collaboration (o^+):** exhibiting low cognitive vigilance and high willingness to collaborate, with proactive information sharing, constructive feedback, and focused engagement with the agent.

While the classic HMM assumes a time-invariant state transition matrix (\mathbf{A}), the evolution of user trust is in fact highly dependent upon the agent's response quality, dialogue context, and the user's initial goals. A static matrix cannot capture the rich contextual dependencies inherent in natural dialogue. Therefore, inspired by modeling techniques in [50], we model the state transition as an LLM-driven dynamic transition function. The trust state at turn t (S_t) depends on the previous trust state S_{t-1} , the user's goal g_0 , initial private cognitive status set \mathcal{C}_0 , and the current dialogue history \mathcal{H}_{t-1} , i.e.,

$$S_t \sim P_{\text{LLM}}(S_t | S_{t-1}, g_0, \mathcal{C}_0, \mathcal{H}_{t-1}) \quad (1)$$

This conditional transition function is approximated deterministically by the LLM according to the following criteria:

- **Trust Increase:** The agent accurately uncovers implicit elements within \mathcal{C}_0 or explores new task factors closely related to g_0 , or provides logically sound clarifications to user doubts, thereby facilitating cognitive development and alignment.
- **Trust Decrease:** The agent exhibits forgetfulness (e.g., *repeatedly asking for information already provided*), logical inconsistencies, or makes suggestions that conflict with previously established user preferences.
- **Trust Maintenance:** Routine transitional exchanges or stable information sharing.

The LLM infers the value S_t according to the above criteria. Within this method, the LLM effectively serves as a parameterized, context-dependent transition function approximator. It leverages pretrained semantic knowledge to empirically approximate the state transition process and the conditional distribution $P(S_t | S_{t-1}, g_0, \mathcal{C}_0, \mathcal{H}_{t-1})$. This mechanism overcomes the limitations of fixed transition matrices in capturing the semantic dynamics of conversations.

Furthermore, existing research [48,51,52] indicates a systematic decoupling between latent trust states and explicit behavioral manifestations due to the influence of cognitive load, emotion regulation, or habitual communication patterns. Therefore, we introduce a parametric noise model for the emission process linking the user's intrinsic trust states to their overt behaviors. With probability $1 - 2\epsilon$, the user's explicit behavior aligns with their latent trust state, while each of the two mismatched behaviors occurs with probability ϵ . We establish a natural ordinal correspondence between trust states and behavioral modes by indexing:

$s_1 = s_l \leftrightarrow o_1 = o^-$, $s_2 = s_m \leftrightarrow o_2 = o^0$, $s_3 = s_h \leftrightarrow o_3 = o^+$. Formally, the emission probability matrix \mathbf{B} is defined as:

$$\mathbf{B}_{ij} = P(O_t = o_j | S_t = s_i) = \begin{cases} 1 - 2\epsilon, & \text{if } j = i \\ \epsilon, & \text{otherwise} \end{cases} \quad (2)$$

where $\epsilon \in (0, 0.5)$ is a hyperparameter controlling the degree of behavioral noise; its specific value is determined empirically (see [Section 4.1](#)).

In summary, the complete dialogue process with N turns for user modeling can be formalized as follows:

$$P(S_{1:N-1}, O_{1:N-1}) = P(S_1) \cdot P(O_1 | S_1) \cdot \prod_{t=2}^{N-1} P_{\text{LLM}}(S_t | S_{t-1}, g_0, \mathcal{C}_0, \mathcal{H}_{t-1}) \cdot P(O_t | S_t) \quad (3)$$

This generative process yields a complete user utterance sequence $U_N = \{u_0, (S_1, O_1, u_1), \dots, (S_{N-1}, O_{N-1}, u_{N-1})\}$, where S_t denotes the latent trust state, O_t denotes the sampled behavior category, and u_t is the user utterance generated by the LLM conditioned on the corresponding prompt instruction. The specific algorithm can be represented as Algorithm 1, where $\text{LLM}^{(O_t)}$ in line 10 refers to instruction templates corresponding to different behavioral mode. $\text{isTerminal}(r_t) = \text{True}$ indicates that the agent outputs the final solution and terminates the dialogue.

Algorithm 1: Trust-aware user utterance generation

Require: Initial goal g_0 , cognitive states \mathcal{C}_0 , prior π , noise ϵ , max turns N_{\max}

Ensure: User Utterance Collection U

```

1: Generate  $u_0 \leftarrow \text{LLM}(g_0, \mathcal{C}_0)$ 
2: Obtain agent response  $r_0 = \text{Agent}(u_0)$ 
3: Dialogue history  $\mathcal{H}_0 \leftarrow \{u_0, r_0\}; U \leftarrow \{u_0\}$ 
4: Sample  $S_1 \sim \pi$ 
5: for  $t = 1$  to  $N_{\max} - 1$  do
6:   if  $t \geq 2$  then ▷ (Eq. (1))
7:      $S_t \leftarrow P_{\text{LLM}}(S_t | S_{t-1}, g_0, \mathcal{C}_0, \mathcal{H}_{t-1})$ 
8:   end if
9:   Sample  $O_t \sim \text{Categorical}(\mathbf{B}_{S_t})$  ▷ (Eq. (2))
10:   $u_t \leftarrow \text{LLM}^{(O_t)}(g_0, \mathcal{C}_0, \mathcal{H}_{t-1})$ 
11:   $U \leftarrow U \cup \{(S_t, O_t, u_t)\}$ 
12:  Obtain agent response  $r_t = \text{Agent}(\mathcal{H}_{t-1}, u_t)$ 
13:   $\mathcal{H}_t \leftarrow \mathcal{H}_{t-1} \cup \{u_t, r_t\}$ 
14:  if  $\text{IsTerminal}(r_t)$  then
15:    break
16:  end if
17: end for
18: return  $U$ 

```

3.3 TATA Agent Framework

In the aforementioned complex, dynamic real-world user environments, the asymmetry between “task progression” and “user state” poses significant challenges for conventional agents to perceive and handle cognitive conflicts during interactions, which can easily lead to a breakdown in trust and failure of collaboration. To address this issue, we propose the **TATA agent framework**. This framework is grounded

in the dynamic evolution of user trust observed in authentic human-AI interactions and aims to enable deep collaborative dialogues through mechanisms for bidirectional cognitive alignment and trust maintenance.

TATA comprises three core modules (Fig. 1-Right): **the Trajectory Planner, the Dual-track State Monitor, and the Adaptive Interaction Coordinator**. Among these, the Trajectory Planner is responsible for establishing the high-level task orientation. In this work, we adopt the construction methodology proposed in [6]. Specifically, given the user's initial goal (g_0) and initial cognitive state (\mathcal{C}_0), the planner generates a serialized cognitive trajectory $\mathcal{T} = \{o_i \mid o_i = (i, \text{topic}, \text{desc}, \text{action}, \text{detail}, \text{tip})\}_{i=1}^K$. This trajectory encapsulates potential future discussion points and the available interaction policy space, thereby providing a foundational blueprint for subsequent dialogue exploration. Among this, each element o_i represents an interaction task associated with a specific dialogue stage, including fields such as turn order, topic, description, actions, key details, and tips. These structured elements are stored in JSON format and used to guide the direction of subsequent collaboration. This process can be formally represented as follows:

$$\mathcal{T}_0 = \mathcal{P}^{\text{Init}}(u_0) \quad (4)$$

$$\mathcal{T}_{t+1} = \mathcal{P}^{\text{Regulation}}(u_t, \mathcal{M}_t) \quad (5)$$

Here, \mathcal{T}_t represents the planned trajectory at t -th turn, and $\mathcal{P}(\cdot)$ denotes the prompt template. \mathcal{M}_t is the working memory.

On the other hand, the Cognitive Monitor and Interaction Coordinator operate in tandem, utilizing a real-time dual-track “cognitive & affective” monitoring mechanism and adaptive policy regulation. This enables the agent to dynamically adjust the dialogue trajectory and select appropriate interaction policies, thereby facilitating effective task completion and fostering deep collaboration.

3.3.1 Dual-Track State Monitor

In exploratory cognitive collaboration, a user's internal trust state and cognitive evolution are inherently unobservable. As discussed in Section 3.2, the user's true trust level is modeled as a latent variable, while observable behaviors constitute only noisy and indirect manifestations of this state. If the agent generates responses solely based on the surface semantics of user utterances, it becomes difficult to detect cognitive shifts and trust fluctuations that have already occurred but have not yet been explicitly expressed.

To address this limitation, we design a “dual-track state monitor” that simultaneously extracts signals from real-time working memory and the user utterance at each turn along two dimensions: **cognitive development** (Δ_t) and **affective vigilance** (E_t). These signals provide structured observational inputs from both task-oriented and relational perspectives, enabling the downstream adaptive coordinator to make more informed adjustments. Specifically, Δ_t represents the cognitive alignment with respect to the user's goal, evaluating only the objective informational relationship between the user's current utterance and the historical facts and logic. We define the cognitive development of the current dialogue as: $\Delta_t \in \{\delta_{\text{Cons/Evo}}, \delta_{\text{Alter}}, \delta_{\text{Conf}}\}$, where:

- Consistency or Evolution ($\delta_{\text{Cons/Evo}}$): The user fully endorses the agent's understanding or provides more specific details on existing discussion elements, demonstrating coherent logic and smooth information progression.
- Alteration (δ_{Alter}): The user's preferences or ideas undergo substantive changes or overturn previous assumptions.
- Conflict (δ_{Conf}): The user expresses confusion, skepticism, or resistance toward the agent's recommendations, facts, or reasoning.

Concurrently, the emotional vigilance state E_t evaluates the user's latent emotional experience and cognitive alertness, represented as $E_t \in \{e_l, e_m, e_h\}$, corresponding respectively to a relaxed interaction with low vigilance and positive affect, a task-oriented state with moderate vigilance and neutral emotion, and a highly vigilant state characterized by negative affect, such as irritation, defensiveness, or disengaged responses.

These two states jointly provide a real-time assessment of the current task status and the interpersonal relationship within the dialogue. The above process can be formally expressed as:

$$(\Delta_t, E_t) = \mathcal{P}^{\text{Monitor}}(g_0, u_t, \mathcal{M}_{t-1}) \quad (6)$$

where $\mathcal{P}^{\text{Monitor}}$ denotes the prompt template, g_0 is the user's initial goal, u_t is the current user utterance, and \mathcal{M}_{t-1} is the working memory, which contains a summary of the current history and cognitive elements related to the task details, specifically defined in [6].

3.3.2 Adaptive Interaction Coordinator

Based on the external observations and perceptions of the user's true trust state and dialogue state provided by the Monitor module, the interaction coordinator's policy decision can be essentially formulated as a sequential decision-making problem under uncertainty. To achieve optimal interaction control, we adopt the theoretical perspective of the Partially Observable Markov Decision Process (POMDP) to structure this process. Specifically, we employ the POMDP tuple $\lambda_{\text{coor}} = (\mathcal{S}, \mathcal{A}, T, \mathcal{R}, \mathcal{Z}, \Omega)$ as a conceptual modeling framework to structure the coordinator's decision-making. This formulation provides a principled structural basis that specifies what the agent must track, which actions are available, and why belief-based reasoning is necessitated.

Here, the hidden state space $\mathcal{S} = \{s_l, s_m, s_h\}$ represents the set of latent user trust states (Section 3.2). The observation space (\mathcal{Z}) corresponds to the cognitive and affective features perceived by the Monitor, denoted as $z_t = (\Delta_t, E_t)$ (6). The action space (\mathcal{A}) consists of the interaction policies available to the agent. T and Ω denote the state transition function and observation function, respectively. The reward function (\mathcal{R}) is implicitly designed to balance task progression efficiency with the maintenance of collaborative trust.

Since the agent cannot directly access the true hidden state, the Coordinator maintains a belief state distribution b_t at each dialogue turn t . $b_t(s_t)$ represents the posterior probability that the user is in a particular trust state $s_t \in \mathcal{S}$, conditioned on the historical interaction. After executing the previous action a_{t-1} and receiving a new observation z_t , the belief state is, in principle, theoretically updated through a recursive application of Bayes' rule [53]:

$$b_t(s') \propto \Omega(z_t | s', a_{t-1}) \sum_{s \in \mathcal{S}} T(s' | s, a_{t-1}) b_{t-1}(s) \quad (7)$$

Given that TATA operates within an open-ended natural language space where computing high-dimensional integrals is computationally intractable, we employ the powerful contextual reasoning capabilities of LLMs to perform a parametric approximation of this process. Specifically, the LLM takes as input the dialogue history \mathcal{H}_{t-1} , the current dual-track observation z_t , and the previous belief distribution b_{t-1} , and outputs an updated belief distribution b_t through structured prompting. Next, the coordinator selects an action $a_t \in \mathcal{A}$ from the policy space to maximize the expected long-term collaborative reward $\mathbb{E}[\mathcal{R}]$.

In the TATA framework, the reward \mathcal{R} should comprehensively consider the dynamic utility balance between "task completion" and "maintenance of collaborative trust", therefore, we have predefined the following four interaction policies ($a_t \in \mathcal{A} = \{a_{\text{task}}, a_{\text{emp}}, a_{\text{trans}}, a_{\text{break}}\}$):

- Task Advancement (a_{task}): Advances the collaborative task Smoothly along the cognitive trajectory \mathcal{T} . This policy is suitable for routine situations characterized by stable trust and cognitive alignment.
- Emotional Empathy (a_{emp}): Prioritizes emotional responses, soothing and stabilizing the user's emotional state through empathetic expressions and acknowledgments of concern.
- Intent Transparency (a_{trans}): When explicit cognitive conflicts or disagreements arise, the agent proactively introduces external evidence or explains the underlying reasoning logic to rebuild a shared understanding.
- Global Breakdown (a_{break}): When the collaborative state approaches a potential trust breakdown, the agent temporarily suspends the current task process, acknowledges the issue, and renegotiates the direction of collaboration.

Furthermore, the coordinator dynamically selects the most appropriate interaction policy a_t according to the current collaborative context via a policy selection instruction $\mathcal{P}_{\text{Select}}$. The selected policy then guides the generation of the system's response, which can be formulated as:

$$a_t = \mathcal{P}^{\text{Select}}(b_t, z_t, \mathcal{T}, \mathcal{H}_{t-1}) \quad (8)$$

Notably, the proposed framework does not aim to compute an exact solution to the POMDP. Rather, it serves as a structural constraint that guides the coordinator's adaptive decision-making process. It clearly defines the state space the agent needs to track, the boundaries of available policies, and the theoretical foundation for decision optimization. Under these guidance, the coordinator leverages the reasoning capabilities of the LLM as an approximate solver to perform trust maintenance and policy selection.

Through this mechanism, the dialogue agent treats the user as a collaborative partner with complex internal states, enabling dynamic calibration of interaction policies during the conversation and thereby supporting real bidirectional cognitive collaboration.

4 Experiments

4.1 Scenario & User Simulator

Building upon prior research on complex task-oriented dialogue [6,54], this study adopts a dialogue generation method for comparative evaluation and incorporates several competitive baseline methods to address such scenarios. We adapt the scenario selection and user profiling methodologies from CCD to industrial intelligence applications. Drawing upon established researches, we select six representative experimental environments, including: (1) *intelligent logistics route planning*, (2) *equipment fault diagnosis*, (3) *production resource allocation*, (4) *energy management and scheduling*, (5) *quality control and defect detection*, and (6) *supply chain risk assessment* [28,55,56]. Each scenario exhibits high exploratory complexity, domain-specific, and dynamic constraint satisfaction, which require operators and agents to construct mutual understanding through multi-turn dialogue rather than single-turn instruction execution. For each domain, we develop 15 industrial operator profiles representing varying experience levels and cognitive states, yielding 90 instances in total. These profiles are designed to reflect authentic role characteristics and cognitive diversity observed in real-world industrial settings (see [Appendix A](#) for the example profile).

To realistically simulate the psychological dynamics of human operators engaged in HAC at the network edge, and to mitigate the widely observed "over-cooperation" bias associated with using standard LLMs as user simulators, we implement the proposed trust-aware user (operator) behavior model, as detailed in [Section 3.2](#). The specific parameter settings and execution logic are summarized as follows:

Initial Distribution and Emission Probability: During the dialogue initialization phase, we establish the user's prior trust state distribution as $\pi = \{P(s_l) = 0.2, P(s_m) = 0.6, P(s_h) = 0.2\}$, indicating that the

majority of users initially exhibit a neutral-to-cautious psychological disposition. The behavioral noise parameter in the emission matrix \mathbf{B} is set to $\epsilon = 0.1$, implying that the simulated user's explicit behavior perfectly reflects their underlying trust state with probability 0.8, thereby accommodating authentic communication noise.

Behavior Sampling and Dialogue Generation: In each dialogue turn t , the user simulator operates through the following procedure: (1) The LLM evaluates the current dialogue context according to Eq. (1) to determine the updated latent trust state S_t ; (2) The observable behavior O_t is sampled stochastically from the emission distribution \mathbf{B}_{S_t} ; (3) O_t is mapped to the corresponding behavior instruction template, guiding the LLM to generate a user utterance u_t with the appropriate response characteristics.

To further enhance the authenticity and exploratory capacity of the interaction, we introduce additional “alteration noise” during dialogues. Specifically, when the user's trust state $S_t \in \{S_m, S_h\}$, the simulator triggers an exploratory shift event with a probability 0.2. In such events, the user proactively overturns previous task details or introduces new preference shifts (e.g., “*Actually, the priority of the logistics fleet has just changed; we need to reroute through Zone B*”), thereby assessing the dialogue agent's adaptive capacity and strategic modulation within dynamic cognitive environments. The maximum number of dialogue turns is set to 30 to reflect the typical strict latency and interaction constraints found in industrial settings.

4.2 Baselines

We compare our method against five representative baselines: (1) ReAct [57], an agent design paradigm enabling autonomous “think-act-observe” cycles that integrate logical reasoning with tool utilization to accomplish complex tasks; (2) Proactive [26], an agent framework employing proactive methodologies that affords two alternative actions—seeking clarification or maintaining inaction; (3) ProCoT [26], a strategy enhancing LLM proactivity by incorporating reasoning and planning steps into prompts, thereby generating descriptive deliberations requisite for decision-making and response generation; (4) Direct-Prompting [6], a robust baseline wherein the LLM is directly prompted to engage in collaborative dialogue; and (5) CoCo-Agent [6], an advanced cognitive collaboration framework featuring trajectory planning, interaction coordination, and monitoring regulation capabilities.

In the experiments, all aforementioned methods are evaluated through interactions with the proposed trust-aware user simulator. Furthermore, to assess the performance of TATA under conventional user modeling settings, we additionally employed the “hybrid user simulator” introduced in [6]. In each dialogue turn, this simulator randomly selects either an proactive or passive behavior mode with equal probability to generate responses, thereby serving as a robustness stress test to emulate highly dynamic and unpredictable conversational environments. Consequently, this simulator is also employed as a comparative baseline.

4.3 Implementation Details

In this study, we strictly follow the setup of [6] for both the user simulator and the evaluation judge, employing GPT-4.1 as the backbone model in each case (temperature = 0). This choice intentionally aligns with the original CCD experimental protocol to control confounding factors, ensuring that any performance gains observed are attributable to the proposed framework rather than differences in the underlying LLMs.

Although using the same backbone model for simulation and evaluation may theoretically introduce self-preference bias, we mitigate this potential dependency and ensure generalization through three mechanisms: (1) **Multi-variant evaluation:** We evaluate all methods using Qwen-2.5-72B [58] and GPT-4.1 [59] as backbone models with temperature set to 0.6; (2) **Human verification:** We manually audit a randomly 5% subset of GPT-4.1 evaluation outputs that pass validity checks to ensure scoring reliability; (3) **Statistical**

robustness: Dialogues are generated following the procedure in Algorithm 1 and §4.1. Specifically, a user profile is randomly sampled from the 90 available user profiles and input to the user simulator, which then interacts with the agent until the agent produces a final solution or the dialogue reaches 30 turns. Using this procedure, we generate 200 dialogues, from which 100 are randomly sampled, and all metrics are averaged over three iterations. This generation process is consistent with the settings in [6].

Notably, the trust state (S_t) inside the user simulator is strictly invisible to all methods. Each method receives only the natural language utterance u_t generated by the simulator. This design ensures that all compared methods operate under identical observation conditions, thereby providing a controlled and consistent basis for comparative evaluation.

4.4 Evaluation Metrics

To comprehensively evaluate the performance of agents in dynamic cognitive collaboration, we adopt the integrated evaluation framework of CCD tasks, which assesses both the quality of cognitive collaboration and task completion capability. Additionally, to capture the dynamic evolution of user trust and cognitive states in real-world interactions, we introduce a new set of joint trust and efficiency metrics. Specifically, our evaluation encompasses the following aspects:

- **Cognitive Collaboration Quality:**
 - **Cognition Coverage Index (CCI):** Proportion of the user’s initial cognitive states accurately identified by the agent.
 - **Cognition Gain Index (CGI):** Degree of new cognitive elements discovered or existing elements enhanced during interaction.
 - **CoCo-FI:** Harmonic mean of *CCI* and *CGI*, reflecting overall collaborative effectiveness.
See [Appendix B.1](#) for detailed formulation.
- **Task Completion Capability:**
 - **Inform*:** Similarity between the agent’s final solution and the user’s initial goal (normalized to [0, 1]).
 - **Success*:** Proportion of cognitive elements successfully addressed in the final solution.
See [Appendix B.2](#) for detailed formulation.
- **Trust Maintenance Capability:** To quantitatively evaluate the agent’s ability to maintain and repair trust, we map the latent trust state S_t to a numerical value via a mapping function: $v(s_l) = -1$, $v(s_m) = 0$, and $v(s_h) = 1$. Furthermore, to capture the cumulative nature of trust dynamics, where successive negative (or positive) experiences generate a “snowball effect [60]”, we define an inertia-based *trust momentum* (V_t):

$$V_1 = v(S_1), \quad V_t = \gamma \cdot V_{t-1} + v(S_t), \quad t \geq 2 \quad (9)$$

where $\gamma \in [0,1)$ denotes the *trust inertia coefficient*. Based on this formulation, we introduce three evaluation metrics.

- **Average Trust Momentum (ATM):** Measure the mean accumulated trust momentum over the entire dialogue:

$$\text{ATM} = \frac{1}{N-1} \sum_{t=1}^{N-1} V_t \quad (10)$$

A higher ATM indicates more persistent and stable collaborative trust. The exponential accumulation mechanism ensures that sustained deterioration in trust is penalized far more heavily than occasional fluctuations.

– **Trust Recovery Rate (TRR)**: The proportion of low-trust events after which the agent achieves a significant recovery in *trust momentum* within the subsequent K turns. Let $\mathcal{E} = \{t : S_t = s_l\}$ denote the set of low-trust events, we define:

$$\text{TRR}^{(w=K)} = \frac{1}{|\mathcal{E}|} \sum_{t \in \mathcal{E}} \mathbb{I} \left[\max_{1 \leq w \leq K} (V_{t+w} - V_t) > \eta \right] \quad (11)$$

where w is the recovery observation window and $\eta > 0$ is the minimum recovery magnitude threshold

– **Trust Collapse Rate (TCR)**: The probability of low-trust states during the dialogue:

$$\text{TCR} = \frac{1}{N-1} \sum_{t=1}^{N-1} \mathbb{I}[S_t = s_l] \quad (12)$$

where $\mathbb{I}[\cdot]$ is the indicator function that equals 1 when the condition holds and 0 otherwise.

4.5 Results and Analysis

For the aforementioned metrics of collaborative quality and task completion, we adopt a hybrid evaluation framework that integrates LLM-assisted processing with direct numerical computation [61–63].

Specifically, for the cognitive collaboration metrics, following the approach of Li et al. [6], we treat the user’s initial utterance u_0 as the initial goal g_0 and employ GPT-4.1 to extract task-relevant details C_t from the dialogue history \mathcal{H}_t to compute the metrics regarding *CCI*, *CGI*, *CoCo-F1* and *Success**. For the task completion metric *Inform**, we utilize the GPT-4.1 model to evaluate the relevance of the final response on a scale from 1 to 10, and then normalize the resulting score.

In addition, for the three trust-related collaboration metrics proposed in this work, we compute them by directly tracking the state of S_t within the user simulator and applying Eqs. (9)–(12). In our experiments, the parameters are set to $\gamma = 0.5$ and $\eta = 1$.

4.5.1 Analysis of Collaboration Quality and Task Completion

Table 1 compares the performance of different methods under both *hybrid user simulators* and *trust-aware user simulators*, covering multiple evaluation metrics, including average dialogue turns, collaboration quality (*CCI*, *CGI*, *CoCo-F1*), and task completion rate (*Inform**, *Success**).

As shown in Table 1, for both user types, traditional baseline models (e.g., ReAct, Proactive, ProCoT and Direct) exhibit minimal capacity for sustaining multi-turn interactions, with an average turn count (Avg. Turn) ranging merely from 1.98 to 2.98. Although these methods demonstrate acceptable performance in fulfilling the user’s initial task goals, their collaboration quality remains limited, as reflected by consistently low cognitive gain and overall collaboration scores (typically *CoCo-F1* < 0.40). The tendency to prematurely conclude dialogue without sufficient interactions may result in solutions that address only superficial user requirements, failing to facilitate deep, bidirectional collaboration.

Table 1: Evaluation results of collaborative quality and task completion (Sample = 100, iteration = 3). The standard deviations for collaborative quality and task completion metrics are consistently small (<0.025), indicating the high stability of evaluation results.

Methods	Avg. Turn	Collaborative Quality			Task Completion	
		CCI	CGI	CoCo-FI	Inform*	Success*
Hybrid User Simulator [6]						
ReAct (GPT-4.1) [57]	2.72	0.847	0.255	0.392	0.962	0.783
Proactive (GPT-4.1) [26]	1.93	0.601	0.112	0.189	0.915	0.650
ProCoT (GPT-4.1) [26]	2.20	0.778	0.098	0.174	0.975	0.787
Direct (GPT-4.1) [6]	2.97	0.860	0.260	0.399	0.967	0.761
CoCoAgent						
(Qwen2.5-72B-Instruct) [6]	10.28	0.881	0.372	0.523	0.984	0.832
CoCoAgent (GPT-4.1) [6]	18.05	0.890	0.481	0.624	0.981	0.712
TATA (Qwen2.5-72B-Instruct)						
TATA (GPT-4.1)	5.75	0.901	0.326	0.478	0.979	0.787
TATA (GPT-4.1)	5.82	0.888	0.372	0.525	0.988	0.902
Trust-Aware User Simulator (Ours)						
ReAct (GPT-4.1) [57]	2.79	0.856	0.203	0.328	0.899	0.592
Proactive (GPT-4.1) [26]	1.98	0.728	0.135	0.228	0.884	0.585
ProCoT (GPT-4.1) [26]	2.14	0.817	0.114	0.200	0.912	0.522
Direct (GPT-4.1) [6]	2.98	0.884	0.206	0.334	0.957	0.711
CoCoAgent						
(Qwen2.5-72B-Instruct) [6]	23.78	0.950	0.496	0.651	0.949	0.657
CoCoAgent (GPT-4.1) [6]	29.66	0.948	0.478	0.635	0.971	0.683
TATA (Qwen2.5-72B-Instruct)						
TATA (GPT-4.1)	9.39	0.954	0.411	0.574	0.977	0.805
TATA (GPT-4.1)	10.66	0.940	0.419	0.580	0.985	0.838

Note: The * indicates that these metrics differ from the traditional Inform and Success metrics as defined in [64].

On the other hand, CoCoAgent, which is specifically designed for collaborative dialogue, achieves the highest cognitive collaboration scores under both user settings (with $CoCo-FI = 0.651$ on the 72B variant under the trust-aware user setting), indicating exceptional exploratory capabilities. However, these impressive collaboration metrics come at the expense of conversation efficiency and task focus. Under the hybrid user setting, the average dialogue length increases substantially, while under the trust-aware user setting, the number of turns approaches the predefined upper limit (23.78 in 72B variant and 29.66 in GPT-4.1 variant). Meanwhile, the core task completion metric ($Success^*$) shows a noticeable decline (only 0.657 and 0.683, respectively).

Considering the “trust-aware user modeling” and “alteration noise” mechanisms introduced in this work, CoCo-Agent lacks effective task state awareness. When confronted with potential user ideas shifts or cognitive conflicts, it exhibits an overly compliant, appeasing behavior by continuously adjusting the dialogue trajectory to accommodate the user’s divergent thinking and expand the scope of discussion, ultimately resulting in topic drift and solution non-convergence. In industrial environments where safety and latency constraints are stringent, such uncontrolled exploration without effective convergence control may introduce substantial operational risks. Furthermore, the performance discrepancy observed between the two user settings provides additional evidence that conventional static simulators, lacking psychological state modeling, can mask trust-related breakdowns and collaboration barriers that may arise in real HAC, thereby artificially inflating evaluation outcomes.

By comparison, the proposed TATA demonstrates consistently balanced and robust performance across all experimental settings. Under the trust-aware user scenario, The TATA (GPT-4.1) variant maintains high-quality cognitive collaboration ($CoCo-F1 = 0.580$) while attaining excellent initial goal alignment and task completion ($Inform^* = 0.985$, $Success^* = 0.838$). Under the hybrid user simulator, TATA (GPT-4.1) achieves the highest $Inform^*$ (0.988) and $Success^*$ (0.902) among all methods with an average dialogue length of 5.82 turns, while still preserving solid collaboration quality ($CoCo-F1 = 0.525$).

This demonstrates that TATA’s adaptive interaction coordinator effectively determines when to engage in divergent cognitive exploration (using a_{trans} or a_{break} policies) and when to converge topics for task advancement (a_{task}), thereby effectively solving users’ practical problems while ensuring the depth of collaboration. This mechanism further explains TATA’s superior dialogue efficiency, with a collaboration efficiency ratio ($CoCo-F1/turns$) that is 1.6–2.6 times higher than that of CoCo-Agent. Such efficiency advantages are particularly critical in industrial edge computing environments, where computational and communication resources are severely constrained.

4.5.2 Analysis of Dynamic Trust Maintenance

For the evaluation of trust maintenance capability, we computed the relevant metrics for 200 dialogue samples under each method according to Eqs. (9)–(12). Table 2 and Fig. 2-left show the average trust momentum (ATM), trust collapse rate (TCR), and trust recovery rate (TRR, window size = [1, 5] turns) recorded by the user simulator for all comparison methods.

Table 2: Trust metric comparison across methods.

Methods	ATM (↑)	TCR (↓)	TRR ^(w=2) (↑)	TRR ^(w=3) (↑)	TRR ^(w=4) (↑)	TRR ^(w=5) (↑)
ReAct (GPT-4.1)	0.398	0.096	0.267	0.333	0.333	0.333
Proactive (GPT-4.1)	0.350	0.081	0.180	0.222	0.222	0.222
ProCoT (GPT-4.1)	0.192	0.167	0.000	0.000	0.000	0.000
Direct (GPT-4.1)	0.249	0.156	0.425	0.512	0.512	0.512
CoCoAgent (Qwen2.5-72B)	1.171	0.048	0.472	0.655	0.775	0.864
CoCoAgent (GPT-4.1)	1.627	0.035	0.465	0.689	0.851	0.900
TATA (Qwen2.5-72B)	1.335	0.051	0.521	0.631	0.751	0.772
TATA (GPT-4.1)	1.015	0.040	0.614	0.767	0.800	0.800

As shown in the table, traditional baseline models exhibit extremely low average trust scores ($ATM < 0.4$), with trust collapse rates (TCR) are mostly around 10%. ProCoT (GPT-4.1) even reaches as high as 16.7%. The absence of trust-aware mechanisms renders these approaches incapable of addressing cognitive conflicts and emotional fluctuations during task execution, resulting in rapid trust deterioration and limited recovery. Conversely, collaborative-capable methods, benefiting from extended dialogue turns, afford greater opportunities for cognitive collaboration and build trust. In terms of ATM and TCR metrics, CoCo-Agent (GPT-4.1) achieved optimal performance ($ATM = 1.627$, $TCR = 3.5\%$), while our proposed TATA also demonstrates strong competitiveness, maintaining high trust momentum and a low collapse rate ($ATM > 1.0$, TCR is around 5%), far outperforming standard baselines.

Furthermore, as shown in Fig. 2-left, when trust drops to a low level, methods with collaborative capabilities demonstrate a significant improvement in recovery performance over the subsequent 1–5 turns

as the window size increases, whereas the recovery rates of other methods are mostly below 0.6 or even 0. Notably, at the minimal window size ($w = 2$), TATA (GPT-4.1) achieved the peak TRR of 0.614, surpassing CoCo-Agent (GPT-4.1) ($\text{TRR}^{(w=2)} = 0.465$) by 32.0%. This outcome directly validates the efficacy of TATA’s dual-track monitoring and adaptive policy coordination. By simultaneously detecting cognitive conflict and emotional vigilance, the agent can rapidly adjust targeted recovery policies to contain the situation before trust degradation escalates. Although CoCo-Agent eventually narrows the performance gap over extended windows via trajectory adjustment, it does so at the expense of conversational efficiency. The introduction of significant informational redundancy ultimately compromises overall task completion (as shown in Table 1).

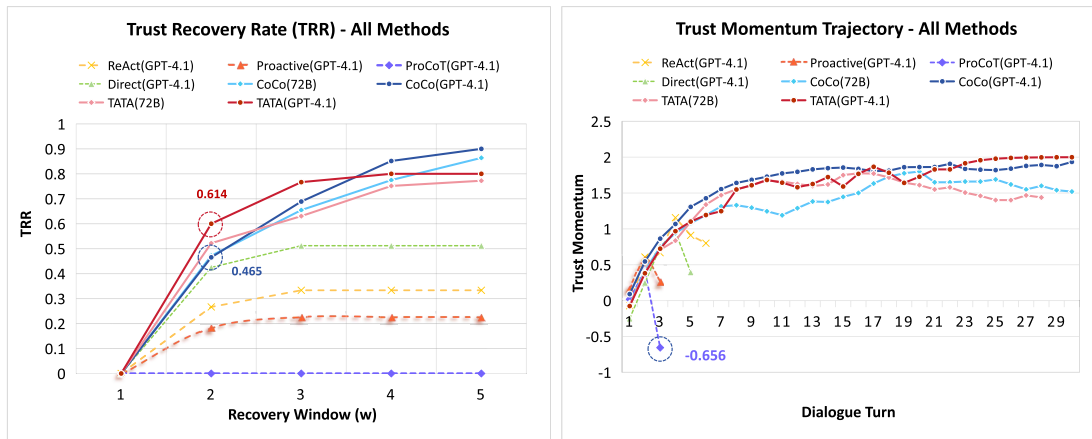


Figure 2: The trust recovery rate (TRR) across window size (left) and the trust momentum trajectory across the dialogue (right) for all methods.

To visualize trust dynamics across multi-turn interactions, we tracked the average trust momentum trajectories of each method throughout the entire dialogue (Fig. 2-right). The analysis further confirms TATA’s mechanistic advantages in dynamic interactions:

(1) Traditional baseline methods typically end the dialogue prematurely within 3–5 turns, accompanied by a sharp decline in trust momentum (e.g., ProCoT drops to -0.656 when $turn = 3$). This indicates that in challenging dynamic trust simulation environments, simple, unperceptive interactions and task progression easily trigger users’ defense mechanisms, leading to collaboration failure. (2) The trust momentum of the two CoCo-Agent variants steadily rises to around 1.9 during the middle to late stages of the dialogue and remains stable. However, considering the extremely long dialogue turns and low task success rates, it is evident that CoCo-Agent tends to adopt a “compliant” and conservative interaction policy to avoid conflict. While this maintains a high level of superficial trust with the user, it fails to achieve effective cognitive fusion and collaboration. (3) In contrast, the trust trajectory of TATA (GPT-4.1) exhibits an evolution pattern consistent with human cognitive processes. It rapidly establishes initial trust, maintains dynamic stability with reasonable fluctuations during mid-dialogue substantive exploration, and efficiently converges upon task completion. This pattern exemplifies robust, authentic scenarios rather than mere conflict avoidance, showcasing more resilient collaborative capabilities.

4.6 Ablation Study

We employ Qwen2.5-72B as the backbone model and conduct a series of ablation studies to validate the efficacy of the proposed method. Given the tight input-output dependencies among the Planner, Monitor, and Coordinator, direct module removal would disrupt system operation. Therefore, we construct degraded

variants by preserving the overall interaction skeleton while disabling only the core functionality of the target module, and compare these against the full model.

Specifically, we evaluate: (1) **Full TATA**, retaining all modules; (2) **TATA w/o Monitor**, preserving the working memory and task execution framework but disabling dual-track state monitor with fix *cognitive development* Δ_t and *emotional vigilance state* E_t to cognitively consistent ($\delta_{Cons/Evol}$) and emotionally neutral states e_m ; and (3) **TATA w/o Coordinator**, retaining trajectory planning and monitoring but removing the POMDP-based belief update and policy selection mechanisms, instead generating responses based on the current subtask and proceeding sequentially. All models utilize the same trust-aware user simulators, test domains, and evaluation protocols to ensure comparability.

Results in [Tables 3](#) and [4](#) demonstrate that removing any core module precipitates significant performance degradation. When dual-track state monitoring is ablated (**w/o Monitor**), the agent loses real-time perception of implicit cognitive shifts and emotional fluctuations. This not only diminishes cognitive collaboration quality (CoCo-F1) and task success rate (Success*), but also causes average trust momentum (ATM) to plummet from 1.335 to 0.796 due to the inability to perceive user states, representing a 20.3% decline in short-term trust repair capacity ($TRR^{(w=2)}$).

Similarly, removing the adaptive Coordinator (**w/o Coordinator**) reduces the agent to a sequential trajectory executor incapable of flexibly invoking strategies to advance collaboration depth or maintain relationships (CoCo-F1 and Success* dropping to minimum values of 0.432 and 0.712, respectively). This results in an almost doubled trust collapse rate (TCR: 0.094) and a 26.8% reduction in short-term trust repair capacity ($TRR^{(w=2)}$) compared to the full version. These findings substantiate that the dynamic state observation provided by the Monitor and the interaction policy flexibility afforded by the Coordinator are highly complementary, which together constituting the core mechanisms for sustaining dynamic collaborative trust and enabling robust industrial HAC.

Table 3: Ablation experimental results of collaborative quality and task completion (Sample = 100, iteration = 3). The standard deviations for collaborative quality and task completion metrics are <0.027.

Methods	Avg. Turn	Collaborative Quality			Task Completion	
		CCI	CGI	CoCo-F1	Inform*	Success*
TATA (Qwen2.5-72B)	9.39	0.954	0.411	0.574	0.977	0.805
TATA w/o Monitor	10.78	0.917	0.355	0.512	0.985	0.767
TATA w/o Coordinator	5.00	0.933	0.281	0.432	0.989	0.712

Note: The * indicates that these metrics differ from the traditional Inform and Success metrics as defined in [64].

Table 4: Ablation experimental results of trust metric comparison across methods.

Methods	ATM (↑)	TCR (↓)	$TRR^{(w=2)}$ (↑)	$TRR^{(w=3)}$ (↑)	$TRR^{(w=4)}$ (↑)	$TRR^{(w=5)}$ (↑)
TATA (Qwen2.5-72B)	1.335	0.051	0.521	0.631	0.751	0.772
TATA w/o Monitor	0.796	0.084	0.415	0.518	0.600	0.633
TATA w/o Coordinator	0.595	0.094	0.381	0.629	0.677	0.677

5 Conclusions

This study investigates the critical challenges of task-level cognitive alignment and relationship-level trust maintenance within HAC systems operating in highly dynamic and complex industrial intelligence scenarios.

To address the over-idealization of existing evaluation environments, we first develop a trust-aware user behavior model grounded in theoretical and empirical research from cognitive science and trust dynamics. By formalizing collaborative trust as a dynamic latent variable within a HMM, the proposed model characterizes operators' psychological fluctuations and cognitive evolution across multi-turn interactions, thereby providing a realistically grounded and reproducible testbed for evaluating the reliability of industrial collaborative intelligence systems. Building upon this, we propose the TATA agent framework. Beyond macro-level task planning, TATA introduces a dual-track cognitive/emotional state monitor and dynamically calibrates the policy space under the guidance of an adaptive interaction coordinator. This enables AI agents to conduct targeted interactions that enhance collaborative efficiency while reducing redundant interaction turns. Furthermore, we introduce a novel set of process-oriented trust evaluation metrics and conduct comparative experiments across six representative complex industrial decision-making scenarios. Extensive experimental results demonstrate that TATA achieves an effective balance among task completion capability, collaborative decision efficiency, and relationship stability in HAC scenarios.

We anticipate that this work will provide new insights into human-in-the-loop (HITL) intelligent systems for industrial applications. Future research may further explore extending the TATA framework to multi-agent or multi-user collaborative settings within cloud edge device coordinated architectures.

6 Limitations

Although the trust-aware user modeling approach and the proposed TATA agent framework demonstrate significant improvements in collaborative efficiency across industrial task domains, several limitations remain for further investigation: (1) Regarding user behavior simulation, although a structured HMM framework and noise mechanisms were introduced based on empirical observations, the LLM-based simulator still cannot fully replicate the authentic behavioral patterns of real industrial operators under complex conditions, such as varying levels of professional expertise and psychological stress tolerance. (2) In terms of mechanism design, while employing an LLM as an approximator for Bayesian belief updating and POMDP policy selection circumvents the challenge of high-dimensional computation, it introduces inherent uncontrollability of LLMs, such as sensitivity to prompt design and potential hallucination risks. (3) The current framework lacks deep integration with structured industrial knowledge bases, which limits the precision of agent recommendations in highly specialized scenarios. Furthermore, due to constraints imposed by the current simulation-based validation environment, this study has not yet conducted benchmarking of inference latency or computational overhead when deploying the system on real industrial edge devices. Future work will focus on human-in-the-loop (HITL) dataset collection, the integration of domain knowledge and sensor data, and the exploration of lightweight neural trust tracking methods combined with model quantization techniques, aiming to achieve a balanced trade-off between computational cost and collaboration depth.

Acknowledgement: Not applicable.

Funding Statement: The authors received no specific funding for this study.

Author Contributions: Pan Li: Conceptualization, methodology and writing—original draft preparation; Zhi Li: Investigation, conceptualization and validation; Yingyou Wen: Supervision and project administration. All authors reviewed and approved the final version of the manuscript.

Availability of Data and Materials: The raw data supporting the conclusions of this article will be made available by the authors on reasonable request.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A Examples of Industrial Scenarios and User Profiles

The following is an example of a specific operator profile for the “Equipment Failure Diagnosis” domain:

- “user_id”: “equipment_23”
- “occupation”: “Equipment maintenance engineer”
- “background”: “Mechanical engineering major”
- “experience”: “17 years of industrial field experience”
- “initial goal (g_0)”: “Abnormal temperature fluctuations detected in the reactor, need to locate the problem quickly to avoid production shutdown.”
- “initial cognitive state (C_0)”:
 - (1) “known symptoms”: “Temperature occasionally exceeds setpoint”
 - (2) “constraints”: “Cannot immediately shut down for comprehensive inspection”
 - (3) “preferences”: “Prefer to check common fault points first”
 - (4) “knowledge”: “Encountered a similar issue last month caused by a stuck cooling valve”

Appendix B Evaluation Metric Details

Appendix B.1 Cognitive Collaboration Quality Metric Details

- **Initial Cognition Coverage Index (CCI):** Evaluate the ability to elicit and identify the user’s cognitive starting points during multi-turn interactions.

$$CCI = \frac{|\mathcal{K}_0 \cap \mathcal{K}|}{|\mathcal{K}_0|} \quad (A1)$$

Here, \mathcal{K}_0 is the predefined initial set of cognitive element keys and \mathcal{K} denotes the set of cognitive element keys extracted after the dialogue. $|\mathcal{K}_0 \cap \mathcal{K}|$ represents the number of *common* elements keys that the agent’s success in eliciting from the existing ones. A higher CCI indicates more comprehensive coverage of the user’s preliminary cognitive state.

- **Cognition Gain Index (CGI):** Evaluates the incremental cognitive gains achieved during the interaction, including the discovery of new elements and the semantic evolution of existing ones.

$$CGI = \frac{|\mathcal{K} \setminus \mathcal{K}_0| + \sum_{k \in \mathcal{K}_0 \cap \mathcal{K}} (1 - \text{BScore}_{\text{F1}}(C_0[k], C[k]))}{|\mathcal{K}|} \quad (A2)$$

where $k \in \mathcal{K}_0 \cap \mathcal{K}$ and $|\mathcal{K} \setminus \mathcal{K}_0|$ denotes the *newly* identified cognitive element set. $\text{BScore}_{\text{F1}}(C_0[k], C[k])$ is the BERTScore F1 value measuring the semantic change between the initial state ($C_0[k]$) and final state ($C[k]$) of element k . A higher CGI indicates greater effectiveness in guiding users to discover new

cognitive elements or deepen existing ones (For example, users may adjust their preliminary preferences in case of agent recommendations and persuasion.).

- **Cognitive Collaboration F1 (CoCo-F1):** The harmonic mean (F1-Score) of CCI and CGI as the core composite metric:

$$\text{CoCo-F1} = 2 \cdot \frac{CCI \cdot CGI}{CCI + CGI} \quad (\text{A3})$$

Higher scores indicate a more in-depth and comprehensive collaboration, which quantitatively reflects the process of understanding co-construction and evolution.

Appendix B.2 Task Completion Capability Metric Details

- **Success***: Calculates the proportion of collected cognitive elements that are successfully incorporated into the agent's final response.

$$\text{Success}^* = \frac{|\mathcal{K}_{\text{final}} \cap \mathcal{K}|}{|\mathcal{K}|} \quad (\text{A4})$$

where $\mathcal{K}_{\text{final}}$ is the key set of cognitive state extracted in r_{final} .

References

1. Joshi B, Singh A, Kumar N, Rautela S. Fuzzy-deep learning-based artificial intelligence for edge computing and real-time decision-making in uncertain IoT environments. In: Proceedings of the 2025 First International Conference on Advances in Computer Science, Electrical, Electronics, and Communication Technologies (CE2CT); 2025 Feb 21–22; Bhimtal, Nainital, India. p. 1301–6.
2. ManusAI. Leave it to manus. 2025 [cited 2025 Oct 9]. Available from: <https://manus.im/>.
3. Yang H, Yue S, He Y. Auto-GPT for online decision making: benchmarks and additional opinions. arXiv:2306.02224. 2023.
4. Cheng J, Kang H, Shao Y, Li N, Chen P, Wang R, et al. Survey on efficient large language models: principles, algorithms, applications, and open issues. IEEE Trans Neural Netw Learn Syst. 2026;37(5):2025–45. doi:10.1109/TNNLS.2025.3628671.
5. Yuan S, Song K, Chen J, Tan X, Shen Y, Ren K, et al. Easytool: enhancing LLM-based agents with concise tool instruction. In: Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers); 2025 Apr 29–May 4; Albuquerque, New Mexico. p. 951–72.
6. Li P, Wang J, Yang Q, Guo T, Liu Y, Wen Y. From answering to discussing: advancing human-AI cognitive collaboration in dialogue agents. Inf Process Manag. 2026;63(5):104711.
7. Kim TS, Lee Y, Yu J, Chung JY, Kim J. DiscoverLLM: from executing intents to discovering them. arXiv:2602.03429. 2026.
8. Zhang Z, Zhao H. Advances in multi-turn dialogue comprehension: a survey. arXiv:2103.03125. 2021.
9. Becker J. Multi-agent large language models for conversational task-solving. arXiv:2410.22932. 2024.
10. Clark HH. Using language. Cambridge, UK: Cambridge University Press; 1996.
11. Tolzin A, Janson A. Uncovering the mechanisms of common ground in human-agent interaction: review and future directions for conversational agent research. Internet Res. 2026;36(1):292–315. doi:10.1108/intr-06-2023-0514.
12. Dong W, Chen S, Yang Y. Protod: proactive task-oriented dialogue system based on large language model. In: Proceedings of the 31st International Conference on Computational Linguistics; 2025 Jan 19–24; Abu Dhabi, United Arab Emirates. p. 9147–64.
13. Laban P, Hayashi H, Zhou Y, Neville J. LLMs get lost in multi-turn conversation. arXiv:2505.06120. 2025.

14. Zhao Z, Vania C, Kayal S, Khan N, Cohen SB, Yilmaz E. PersonaLens: a benchmark for personalization evaluation in conversational AI assistants. *arXiv:2506.09902*. 2025.
15. Sharma M, Tong M, Korbak T, Duvenaud D, Askill A, Bowman SR, et al. Towards understanding sycophancy in language models. *arXiv:2310.13548*. 2023.
16. Yi Z, Ouyang J, Xu Z, Liu Y, Liao T, Luo H, et al. A survey on recent advances in LLM-based multi-turn dialogue systems. *ACM Comput Surv*. 2025;58(6):1–38. doi:10.1145/3771090.
17. Luo J, Zhang W, Yuan Y, Zhao Y, Yang J, Gu Y, et al. Large language model agent: a survey on methodology, applications and challenges. *arXiv:2503.21460*. 2025.
18. Wang L, Xu W, Lan Y, Hu Z, Lan Y, Lee RKW, et al. Plan-and-solve prompting: improving zero-shot chain-of-thought reasoning by large language models. *arXiv:2305.04091*. 2023.
19. Zhang D, Zhoubian S, Hu Z, Yue Y, Dong Y, Tang J. ReST-MCTS*: LLM self-training via process reward guided tree search. *Adv Neural Inf Process Syst*. 2024;37:64735–72. doi:10.52202/079017-2066.
20. Feng Y, Rahmani HA, Lipani A, Yilmaz E. Towards asking clarification questions for information seeking on task-oriented dialogues. *arXiv:2305.13690*. 2023.
21. Qin Y, Liang S, Ye Y, Zhu K, Yan L, Lu Y, et al. ToolLLM: facilitating large language models to master 16,000+ real-world APIS. *arXiv:2307.16789*. 2023.
22. Hudecek V, Dušek O. Are LLMs all you need for task-oriented dialogue. *arXiv:2304.06556*. 2023.
23. Qian C, He B, Zhuang Z, Deng J, Qin Y, Cong X, et al. Tell me more! Towards implicit user intention understanding of language model driven agents. *arXiv:2402.09205*. 2024.
24. Deng Y, Lei W, Lam W, Chua TS. A survey on proactive dialogue systems: problems, methods, and prospects. *arXiv:2305.02750*. 2023.
25. Zhang X, Deng Y, Ren Z, Ng SK, Chua TS. Ask-before-plan: proactive language agents for real-world planning. In: *Proceedings of the Findings of the Association for Computational Linguistics: EMNLP 2024*; 2024 Nov 12–16; Miami, FL, USA. p. 10836–63.
26. Deng Y, Liao L, Chen L, Wang H, Lei W, Chua TS. Prompting and evaluating large language models for proactive dialogues: clarification, target-guided, and non-collaboration. *arXiv:2305.13626*. 2023.
27. Besta M, Blach N, Kubicek A, Gerstenberger R, Podstawski M, Gianinazzi L, et al. Graph of thoughts: solving elaborate problems with large language models. *Proc AAAI Conf Artif Intell*. 2024;38:17682–90.
28. Fernandez C, Fernández I, Aceta C. LAMIA: an LLM approach for task-oriented dialogue systems in industry 5.0. In: *Proceedings of the 15th International Workshop on Spoken Dialogue Systems Technology*; 2025 May 27–30; Bilbao, Spain. p. 27–30.
29. Adel A. Future of industry 5.0 in society: human-centric solutions, challenges and prospective research areas. *J Cloud Comput*. 2022;11(1):40. doi:10.1186/s13677-022-00314-5.
30. Tóth A, Nagy L, Kennedy R, Bohuš B, Abonyi J, Ruppert T. The human-centric industry 5.0 collaboration architecture. *MethodsX*. 2023;11(16):102260. doi:10.1016/j.mex.2023.102260.
31. Kraus M, Wagner N, Minker W. Modelling and predicting trust for developing proactive dialogue strategies in mixed-initiative interaction. In: *Proceedings of the 2021 International Conference on Multimodal Interaction*; 2021 Oct 18–22; Montréal, QC, Canada. p. 131–40.
32. Peng J, Kimmig A, Wang D, Niu Z, Tao X, Ovtcharova J. Intention recognition-based human-machine interaction for mixed flow assembly. *J Manuf Syst*. 2024;72(4):229–44. doi:10.1016/j.jmsy.2023.11.021.
33. Yu A, Li C, Macesanu L, Balaji A, Ray R, Mooney R, et al. Mixed-initiative dialog for human-robot collaborative manipulation. *arXiv:2508.05535*. 2025.
34. Xu Y, Zhan X, Kaltungo AY, Ng MS, Ishizawa T, Fujimoto K, et al. Dialogue based interactive explanations for safety decisions in human robot collaboration. *arXiv:2604.05896*. 2026.
35. Xi Z, Chen W, Guo X, He W, Ding Y, Hong B, et al. The rise and potential of large language model based agents: a survey. *Sci China Inf Sci*. 2025;68(2):121101. doi:10.1007/s11432-024-4222-0.
36. Kong C, Fan Y, Wan X, Jiang F, Wang B. PlatoLM: teaching LLMs in multi-round dialogue via a user simulator. In: *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*; 2024 Aug 11–16; Bangkok, Thailand. p. 7841–63.

37. Ni B, Wang Y, Wang L, Kveton B, Derroncourt F, Xia Y, et al. A survey on llm-based conversational user simulation. In: Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers); 2026 Mar 24–29; Rabat, Morocco. p. 4266–301.
38. Ahmad A, Hillmann S, Möller S. Simulating user diversity in task-oriented dialogue systems using large language models. arXiv:2502.12813. 2025.
39. Wang N, Peng Z, Que H, Liu J, Zhou W, Wu Y, et al. RoleLLM: benchmarking, eliciting, and enhancing role-playing abilities of large language models. In: Findings of the Association for Computational Linguistics: ACL 2024; 2024 Aug 11–16; Bangkok, Thailand. p. 14743–77.
40. Lin J, Tomlin N, Andreas J, Eisner J. Decision-oriented dialogue for human-AI collaboration. *Trans Assoc Comput Linguist.* 2024;12(1):892–911. doi:10.1162/tacl_a_00679.
41. Luo X, Tang Z, Wang J, Zhang X. DuetSim: building user simulator with dual large language models for task-oriented dialogues. arXiv:2405.13028. 2024.
42. Wang X, Sen P, Li R, Yilmaz E. Adaptive retrieval-augmented generation for conversational systems. In: Findings of the Association for Computational Linguistics: NAACL 2025; 2025 Apr 29–May 4; Albuquerque, New Mexico. p. 491–503.
43. Naous T, Laban P, Xu W, Neville J. Flipping the dialogue: training and evaluating user language models. arXiv:2510.06552. 2025.
44. Rafailov R, Sharma A, Mitchell E, Manning CD, Ermon S, Finn C. Direct preference optimization: your language model is secretly a reward model. *Adv Neural Inf Process Syst.* 2023;36:53728–41.
45. Cheng M, Yu S, Lee C, Khadpe P, Ibrahim L, Jurafsky D. ELEPHANT: measuring and understanding social sycophancy in LLMs. arXiv:2505.13995. 2025.
46. Lee JD, See KA. Trust in automation: designing for appropriate reliance. *Hum Factors.* 2004;46(1):50–80.
47. Li M, Erickson IM, Cross EV, Lee JD. Estimating trust in conversational agent with lexical and acoustic features. *Proc Hum Factors Ergon Soc Annu Meet.* 2022;66(1):544–8. doi:10.1177/1071181322661147.
48. Li M, Kamaraj AV, Lee JD. Modeling trust dimensions and dynamics in human-agent conversation: a trajectory epistemic network analysis approach. *Int J Hum Comput Interact.* 2024;40(14):3571–82.
49. Huang Y, Sun L, Wang H, Wu S, Zhang Q, Li Y, et al. TrustLLM: trustworthiness in large language models. arXiv:2401.05561. 2024.
50. Liu J, Tan YK, Fu B, Lim KH. From intents to conversations: generating intent-driven dialogues with contrastive learning for multi-turn classification. arXiv:2411.14252. 2025.
51. Gross JJ. Emotion regulation: affective, cognitive, and social consequences. *Psychophysiology.* 2002;39(3):281–91.
52. Hoff KA, Bashir M. Trust in automation: integrating empirical evidence on factors that influence trust. *Hum Factors.* 2015;57(3):407–34. doi:10.1177/0018720814547570.
53. Spaan MT. Partially observable Markov decision processes. In: Reinforcement learning: state-of-the-art. Berlin/Heidelberg, Germany: Springer; 2012. p. 387–414.
54. Li P, Yang Q, Xu S, Li X, Li Z, Wang C, et al. Adaptive-TOD: an LLM-driven and adaptive agent for diverse interaction modes. *Neurocomputing.* 2025;652:130991.
55. Kumar N, Kumar RR. Human-AI collaboration in operations and supply chain management: a systematic literature review. *Manag Rev Q.* 2025;24(3):691. doi:10.1007/s11301-025-00575-9.
56. Wang K, Du N. Real-time monitoring and energy consumption management strategy of cold chain logistics based on the internet of things. *Energy Inform.* 2025;8(1):34. doi:10.1186/s42162-025-00493-w.
57. Yao S, Zhao J, Yu D, Du N, Shafran I, Narasimhan K, et al. React: synergizing reasoning and acting in language models. In: Proceedings of the 11th International Conference on Learning Representations (ICLR); 2023 May 1–5; Kigali, Rwanda.
58. Yang A, Yang B, Hui B, Zheng B, Yu B, Zhou C, et al. Qwen2 technical report. arXiv:2407.10671. 2024.
59. OpenAI, Achiam J, Adler S, Agarwal S, Ahmad L, Akkaya I, et al. GPT-4 technical report. arXiv:2303.08774. 2024.
60. Zhang M, Press O, Merrill W, Liu A, Smith NA. How language model hallucinations can snowball. arXiv:2305.13534. 2023.

61. Qian K, Beirami A, Kottur S, Shayandeh S, Crook P, Geramifard A, et al. Database search results disambiguation for task-oriented dialog systems. arXiv:2112.08351. 2021.
62. Liu Y, Iyer D, Xu Y, Wang S, Xu R, Zhu C. G-eval: NLG evaluation using GPT-4 with better human alignment. arXiv:2303.16634. 2023.
63. Zheng L, Chiang WL, Sheng Y, Zhuang S, Wu Z, Zhuang Y, et al. Judging LLM-as-a-judge with MT-bench and chatbot arena. *Adv Neural Inf Process Syst*. 2023;36:46595–623. doi:10.52202/075280-2020.
64. Budzianowski P, Wen TH, Tseng BH, Casanueva I, Ultes S, Ramadan O, et al. Multiwoz—a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*; 2018 Oct 31–Nov 4; Brussels, Belgium.