

On Generalized Planning for Controlling Opinion Networks: Interpreting Human-AI Dialog States and Beliefs

Bharath Muppasani, Protik Nag, Biplav Srivastava, Vignesh Narayanan

University of South Carolina, Columbia, SC, USA

Abstract

Influencing user opinions through dialog is crucial for conversational agents in applications such as education, healthcare, and customer support. Traditional dialog management systems, while effective for task-oriented interactions, often lack mechanisms to model and control the evolution of user beliefs during conversations. Recent advancements in opinion dynamics—especially those integrating planning techniques for strategic influence within networks—provide new methods to address this challenge. In this paper, we present a formal framework that integrates opinion dynamics into dialog management systems, introducing control mechanisms to model and influence user opinions toward desired targets. We define mechanisms for targeted message propagation, incorporating strategies for belief manipulation within the dialog context. We demonstrate how our framework can effectively guide user opinions during interactions through illustrative examples. This work bridges the gap between opinion dynamics and dialog systems, using AI planning techniques, offering a novel approach to designing conversational agents capable of engaging in belief-driven interactions and strategically influencing user beliefs.

1 Introduction

Developing effective dialog management systems is a central challenge in artificial intelligence, especially for applications that require complex human-computer interactions. Virtual assistants like Siri and Google Assistant need to interpret and respond to a wide range of user queries, ensuring natural and effective communication (Algherairy and Ahmed 2024). Traditional approaches often rely on frameworks like Partially Observable Markov Decision Processes (POMDPs) to handle uncertainties in dialog states by maintaining probabilistic beliefs over possible user intentions and contextual factors (Young et al. 2013). Reinforcement Learning (RL) techniques are used to optimize dialog policies by maximizing expected cumulative rewards, helping to select appropriate system responses based on these probabilistic beliefs (Thomson and Young 2010).

However, current methods struggle with modeling how user beliefs evolve during interactions. This limitation is evident in applications that involve persuasion, negotiation, or

information dissemination, where it’s important to influence and guide user opinions toward desired outcomes. Traditional methods often lack mechanisms to manage the user’s belief state in a way that allows for such influence, limiting their effectiveness in complex conversations.

Integrating principles from opinion dynamics—a field that examines how individual opinions evolve through social interactions—offers a promising avenue to address these challenges. Opinion dynamics models provide mathematical frameworks to simulate the processes by which opinions are formed and altered within networks, capturing the influence of communication and social factors (Sîrbu et al. 2017). Recent research has investigated leveraging planning methodologies to influence user opinions toward targeted outcomes. In (Muppasani et al. 2024c), the authors propose a framework that integrates opinion dynamics with automated planning to simulate strategies for information spread in social networks. Similarly, (Muppasani et al. 2024b) explores intervention planning using learning-based strategies to disseminate accurate information within dynamic opinion networks. By incorporating such models into dialog management systems, it becomes possible to represent and track the temporal evolution of user beliefs more effectively. This enhances the system’s ability to anticipate user responses and adapt its strategies accordingly, leading to more coherent and persuasive dialogs. Additionally, leveraging opinion dynamics can inform the design of reward structures in RL-based dialog systems, aligning them more closely with the goal of achieving desired shifts in user opinions.

In this paper, we propose integrating models from opinion dynamics into dialog management systems. By simulating how opinions evolve during conversations, we aim to improve dialog policy generation and enable systems to maintain credible information narratives during deployment. That is, when the dialog agent is given *official* information for dissemination, it can better anticipate user responses and adapt their strategies. To facilitate understanding, a glossary of key terms used in this paper is provided in Table 1. The remainder of this paper is organized as follows: Section 2 motivates the problem and discusses the foundational concepts of dialog management systems and how planning techniques have been applied to exert control in opinion networks. Section 3 formalizes the problem, introduces mechanisms for opinion updates and topic interrelations, and presents scenar-

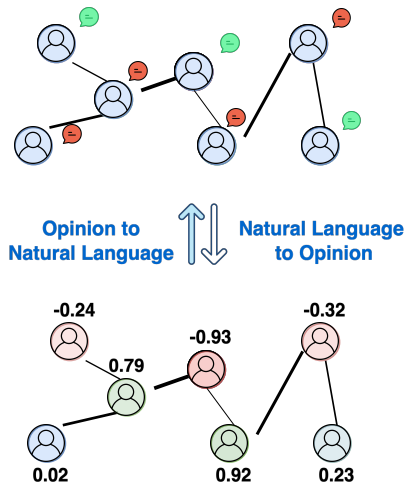


Figure 1: Interpreting dialog states as dynamic opinion networks

ios demonstrating the proposed approach. Finally, Section 4 concludes with a discussion on the broader implications of this work and potential directions for future research.

2 Background

In this section, we outline the motivation behind using conversational agents for belief shaping and review relevant literature on dialog management systems, opinion dynamics, and planning techniques, highlighting their integration for enabling strategic influence in conversational systems.

2.1 Motivation

Conversational agents play a growing role not only as information providers but also as systems that guide and influence user beliefs. To illustrate, consider a simple conversation between a user and a chatbot in Table 2. Here, the chatbot provides factual information to correct the user’s misunderstanding about the day of the week. This interaction highlights the basic mechanics of opinion shaping through conversational exchanges, where the chatbot’s factual utterances guide the user toward aligning their beliefs with reality.

Building upon such basic mechanisms, more complex scenarios arise in corporate settings, where businesses use chatbots to promote products or services and engage with customers. A corporate chatbot, typically deployed by organizations to represent their brand or offerings, may aim to persuade users of a product’s benefits by strategically introducing interrelated topics. For example, Table 3 presents a multi-turn interaction where the chatbot guides the user’s opinion about a product’s cost-effectiveness. The chatbot initially recognizes the user’s low confidence in the product’s 5-star energy rating (s_1) and introduces government energy subsidies (s_2) as supporting evidence. Over subsequent turns, it strategically revisits both topics to strengthen the user’s belief in the product’s overall value.

These examples motivate the need for frameworks that can model, understand, and strategically influence user opinions. Unlike traditional dialog management systems that prioritize task completion and intent recognition, scenarios like these introduce the challenge of steering user beliefs dynamically over time. Addressing this challenge requires a principled approach that integrates opinion dynamics, planning, and inter-topic dependencies into dialog systems. Figure 1 illustrates dialog states represented as dynamic opinion networks, where natural language interactions are mapped to quantitative representations through techniques such as sentiment analysis or large language models (LLMs). These numerical representations enable a structured analysis of opinion dynamics within conversational frameworks.

2.2 Dialog Preliminaries

Dialog management systems typically comprise Dialog State Tracking (DST) and Dialog Policy (DP) modules. DST maintains a probabilistic belief state over user intentions, contextual information, and dialog history, while DP maps this state to an appropriate system action. POMDP-based frameworks are commonly used to handle uncertainties, with Reinforcement Learning (RL) employed to optimize dialog policies by maximizing long-term rewards (Young et al. 2013). However, these approaches can face complexity and scalability issues when adapting to evolving user beliefs and more intricate objectives (Brabra et al. 2021).

An alternative line of work employs plan-based dialog management, where dialog strategies are generated by planning algorithms designed to achieve specified goals. While plan-based methods can be more interpretable and flexible, they also require extensive domain knowledge and must manage the complexities of plan execution in dynamic environments (Santos Teixeira and Dragoni 2022).

2.3 Planning for Control in Opinion Networks

Simulation of Opinion Dynamics and Modeling with PDDL Opinion dynamics research provides mathematical and computational frameworks to understand how individual and collective beliefs evolve within social networks. Early models, such as the DeGroot model (DeGroot 1974), examined belief updates as weighted averages of neighbors’ opinions. Over time, more sophisticated approaches incorporated various network factors and behavioral influences (Anderson and Ye 2019), enabling richer simulations of how information spreads and how consensus or polarization emerges.

Building upon these foundations, recent work has explored integrating opinion dynamics with AI planning methodologies. Muppasani et al. (2024c) introduced a numeric PDDL model specifically tailored to opinion dynamics. This model identifies and executes sequences of actions that guide the network’s opinions from a given initial state toward a desired goal state, offering a transparent means of understanding and directing opinion evolution. One of the primary actions introduced in this framework is `spread-info-from` (shown in Listing 1), which models the information propagation process within a network.

Keyword	Sub-discipline	Definition
Belief/Opinion	Opinion Dynamics	A numeric value (e.g., $v_i^T(t) \in [-1, 1]$) representing an agent’s stance on a given topic.
Utterance	Dialog Systems	A spoken or written message by an actor (human or automated agent) exchanged in a dialog conversation.
Intent	Dialog Systems	The underlying purpose or goal conveyed by an utterance. An actor’s intent may span one or more topics (see Topic and Topic Dependency).
Dialog State	Dialog Systems	A representation of the current conversation status, encompassing user intent, and context of previous user-agent interaction .
Topic	Natural Language Processing (NLP)	The theme of a piece of language (text). An intent can be seen as a topic.
Topic Dependency	Ontology/Knowledge Graphs	A structured representation of how different topics are interrelated, often modeled as edges and nodes within an ontology or knowledge graph.
(Dialog) Policy	Dialog Systems, Planning	A set of coordinated actions (plan/ policy) describing how an actor should respond.

Table 1: Glossary of key terms used in the paper and the technical (AI) subdiscipline they are prominently used in.

While this modeling approach is highly expressive and useful for small to medium-sized networks, generating plans for larger networks can become computationally demanding, limiting scalability in certain applications.¹

Listing 1: Numeric PDDL action for opinion propagation (Muppasani et al. 2024c)

```

1 (:action spread-info-from
2 :parameters (?a - agentsource ?t - topic)
3 :precondition (and
4 :effect (and
5   (forall (?a2 - agent)
6     (when (connected-agent ?a ?a2)
7       (increase (have-stance ?a2 ?t)
8         (/ (* (have-trust ?a2 ?a)
9           (- (have-stance ?a ?t) (have-
10             stance ?a2 ?t))))
11   (forall (?a2 - agent)
12     (when (and (connected-agent ?a ?a2)
13       (< (+ (have-stance ?a2 ?t)
14         (/ (* (have-trust ?a2
15           ?a)
16         (- (have-stance ?a ?t)
17           (have-stance ?a2

```

¹Demonstration video at: <https://tinyurl.com/3k7bp99h>

User (Turn 1): “What day is it today?”

Chatbot (Turn 1): “Today is Monday.”

User (Turn 2): “Really? I thought it was Sunday.”

Chatbot (Turn 2): “Actually, yesterday was Sunday.”

User (Turn 3): “Oh, that makes sense now. Thank you!”

Table 2: A simple interaction where the chatbot influences the user’s beliefs using factual information.

User (Turn 1): “I’m not convinced this appliance’s 5-star energy rating (s_1) really makes it cost-effective.”

Chatbot (Turn 1, Topic s_2): “This appliance qualifies for government energy subsidies (s_2), which can reduce your initial purchase costs.”

User (Turn 2): “I wasn’t aware of subsidies (s_2). That does help with the initial cost, but I’m still not sure how the 5-star rating (s_1) affects my bills long-term.”

Chatbot (Turn 2, Topic s_1): “A 5-star energy rating (s_1) indicates that the appliance operates more efficiently, lowering your monthly electricity expenses over time.”

User (Turn 3): “So with subsidies (s_2) reducing the upfront expense and the high energy rating (s_1) cutting ongoing costs, it might actually be cost-effective.”

Chatbot (Turn 3, Topic s_2): “Precisely. By utilizing the available subsidies (s_2), you minimize initial costs and, combined with the appliance’s energy efficiency, achieve a faster return on your investment.”

Table 3: A multi-turn interaction demonstrating how a chatbot uses one topic per utterance to influence the user’s belief across interconnected topics.

```

16   ?t))
17   1)) -1))
18   (assign (have-stance ?a2 ?t) -1))
19 (forall (?a2 - agent)
20   (when (and (connected-agent ?a ?a2)
21     (> (+ (have-stance ?a2 ?t)
22       (/ (* (have-trust ?a2
23         ?a)
24       (- (have-stance ?a ?t)
25         (have-stance ?a2
26           ?t))))
27     1)) 1))
28 (assign (have-stance ?a2 ?t) 1))

```

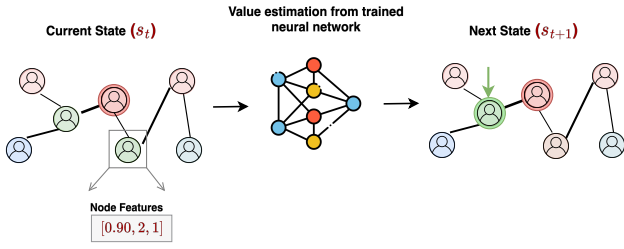


Figure 2: Learning-based framework for intervention planning, utilizing state representations to transition from current state s_t to next state s_{t+1} , enabling strategic interventions (Muppasani et al. 2024b).

25)
26)

Control and Intervention in Opinion Networks Beyond simulation, recent research has focused on developing strategies to actively control and influence dynamic opinion networks. Muppasani et al. (2024b) explored learning-based techniques to plan interventions aimed at disseminating accurate information. They introduced a ranking algorithm to generate labeled data for training supervised neural network classifiers that identify key nodes for disseminating accurate information. Additionally, they developed a reinforcement learning-based centralized dynamic planning framework, utilizing graph convolutional networks (GCNs) to generate scalable plans. Their findings demonstrated that learning-based planners can effectively control misinformation spread by strategically selecting nodes for intervention. These methods determine where and when to apply interventions, ultimately guiding opinions toward more accurate or desired beliefs. Such control-oriented strategies represent a significant step forward in leveraging opinion dynamics not merely for analysis, but for purposeful, goal-driven influence within complex social environments.

These advances suggest that dialog systems can be extended to proactively shape user beliefs. By modeling user opinions as dynamic states influenced by dialog choices, it becomes possible for a dialog manager to not only respond to user queries but also steer user beliefs strategically. In this manner, frameworks that combine opinion dynamics modeling with planning and learning can support the design of conversational agents that guide users toward specific opinion states, going beyond traditional task-oriented dialogs.

3 Problem Formulation

This section formalizes the integration of opinion dynamics concepts into dialog management systems. The objective is to model how user opinions on specific topics evolve over the course of interactions. By incorporating formal opinion update mechanisms and inter-topic dependencies, this framework facilitates the strategic selection of utterances for guiding user beliefs toward specified target states.

3.1 Notation and Basic Setup

Consider a network $G = (V, E)$, where V is a set of agents and $E \subseteq V \times V$ is a set of directed edges. Each edge $(v_i, v_j) \in E$ is associated with a weight $w_{ij} \in [0, 1]$ reflecting the extent to which agent v_j influences agent v_i . Let \mathcal{T} denote a set of topics, each represented by a lowercase element $s \in \mathcal{T}$.

The set V includes human agents v_i^h and chatbot agents v_j^c . The opinion of agent i on topic s at time-step t is defined as $v_i^s(t) \in [-1, 1]$. Interpretations are as follows:

- If $v_i^s(t) > 0$, agent i is inclined to regard topic s as true.
- If $v_i^s(t) < 0$, agent i is inclined to regard topic s as false.
- The magnitude $|v_i^s(t)|$ represents the strength of the belief.

Time is discretized into steps $t = 0, 1, \dots, t_f$, where t_f denotes a final time-step. Figure 2 provides a visual depiction of such an opinion network state, where state representations are used to model transitions from the current state s_t to the next state s_{t+1} . This framework supports the design of strategic interventions to influence opinion dynamics within the network.

3.2 Mechanisms of Opinion Updates

The evolution of opinions is modeled using two primary mechanisms. These mechanisms characterize how a chatbot’s utterances influence a human user’s beliefs.

Mechanism 1: Targeted Message Propagation This mechanism employs a linear adjustment model commonly adopted in opinion dynamics. Suppose the chatbot v_j^c produces an utterance $u_j^s(t)$ at time-step t that conveys its fixed target opinion $v_j^{c,s}(t)$ on topic s . The influence on a human agent v_i^h is captured by:

$$v_i^{h,s}(t+1) = v_i^{h,s}(t) + w_{ij}(v_j^{c,s}(t) - v_i^{h,s}(t)), \quad (1)$$

where:

- $v_i^{h,s}(t)$ is the current opinion of the human agent i on topic s .
- $v_j^{c,s}(t)$ is the chatbot’s target opinion on topic s .
- w_{ij} is the weight indicating how susceptible agent i is to agent j .

This mechanism shifts the human agent’s opinion linearly toward the chatbot’s target opinion.

Mechanism 2: Utterance-Based Opinion Updates In addition to direct linear adjustments, we also model opinion updates as a function of the utterances exchanges during the conversation. Let $U_{ij}^s(t)$ be the collection of utterances pertaining to topic s exchanged between v_i^h and v_j^c at time step t .

Define:

$$f^s(U_{ij}^s(t))$$

as a function that determines the opinion shift induced by these utterances. This function may incorporate factors such as argument quality or sentiment. Assuming a conversation

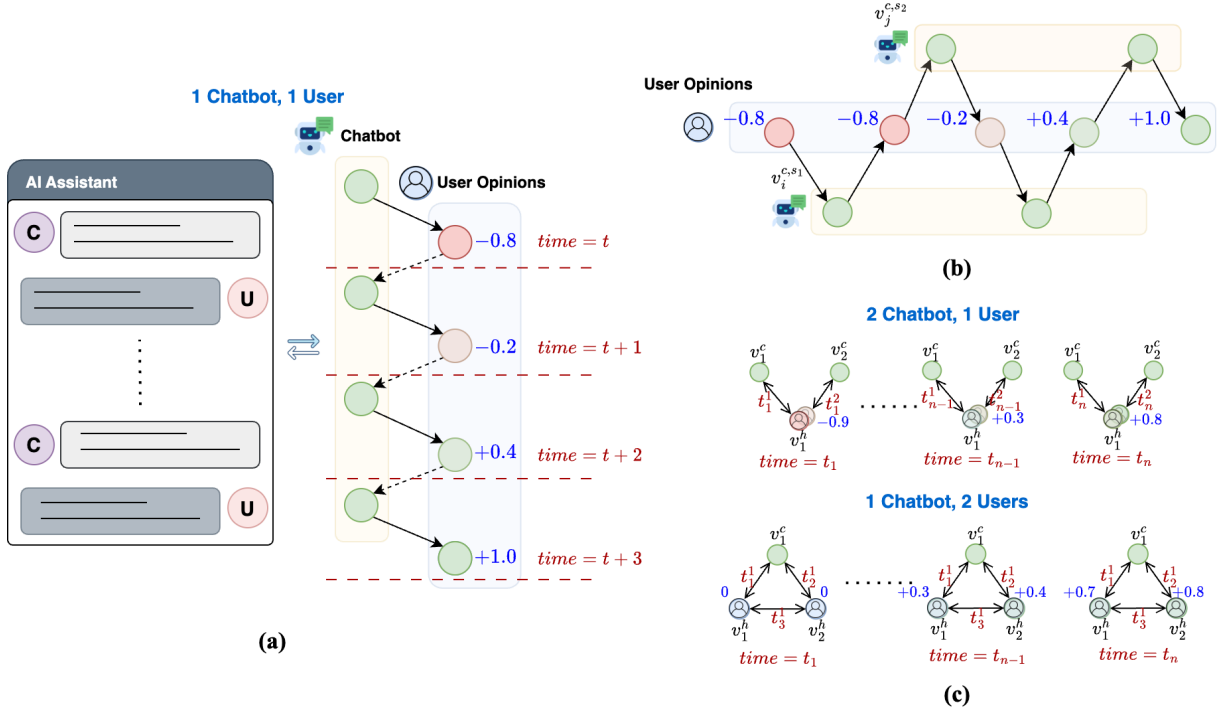


Figure 3: Illustration of opinion dynamics in dialog systems. (a) Induced opinion network observed in the interaction between a chatbot and a user over time. (b) The evolution of user opinions is modeled as a dynamic process influenced by inter-topic dependencies and chatbot interventions. (c) Scenarios involving multiple agents, including one chatbot interacting with multiple users and multiple chatbots influencing a single user.

impact factor $\gamma \geq 0$ to control the overall effect, the update rule becomes:

$$v_i^{h,s}(t+1) = v_i^{h,s}(t) + \gamma f^s(U_{ij}^s(t)). \quad (2)$$

Here:

- γ scales the influence of the utterances.
- $f^s(U_{ij}^s(t))$ determines how the exchanged utterances shift the human agent's opinion.

This mechanism enables modeling of more context-dependent and content-sensitive opinion changes.

Topic Interrelations In practice, agents often hold opinions on multiple, potentially interrelated topics. Such interdependencies can arise from logical, causal, or semantic relationships among concepts. Ontologies and knowledge graphs, such as ConceptNet (Liu and Singh 2004) and WordNet (Fellbaum 2000), provide structured ways to represent and reason about these semantic relations. By leveraging these knowledge-rich resources, one could capture complex, context-dependent influences between topics in a principled manner. However, integrating such ontological structures directly into the opinion update process can be non-trivial, requiring appropriate abstraction and scaling methods.

In this work, we use a numerical representation to incorporate inter-topic dependencies presented by Anderson and Ye (2019). We define a topic interrelation matrix S , where

each element $S_{ss'} \in [-1, 1]$ quantifies how belief in topic s' influences belief in topic s :

$$S_{ss'} = \begin{cases} > 0 & \text{if belief in } s' \text{ positively influences belief in } s, \\ < 0 & \text{if belief in } s' \text{ negatively influences belief in } s, \\ 0 & \text{if there is no direct influence.} \end{cases}$$

Each agent v_i maintains a belief vector $\mathbf{v}_i(t) = [v_i^{s_1}(t), v_i^{s_2}(t), \dots, v_i^{s_n}(t)]$, where $\mathcal{T} = \{s_1, \dots, s_n\}$.

Considering these interrelations, the opinion update rule for agent i on topic s becomes:

$$v_i^s(t+1) = v_i^s(t) + w_{ij}(v_j^s(t) - v_i^s(t)) + \sum_{s' \in \mathcal{T}} S_{ss'}(v_i^{s'}(t) - v_i^s(t)). \quad (3)$$

In this formulation, the final term models how other topics s' within the same agent's belief vector shape its opinion on topic s .

3.3 Scenarios

The above mechanisms and interrelations can be applied to various scenarios. We outline two illustrative cases.

Scenario 1: Single Chatbot Influencing a Single Human Agent Over Multiple Topics Consider a scenario involving a single chatbot v_j^c interacting with a single human agent

v_i^h across multiple topics $\epsilon \mathcal{T}$. At each time-step t , the chatbot selects utterances $\{u_j^s(t)\}_{s \in \mathcal{T}}$ intended to influence the human agent’s opinions.

The opinion update for each topic s considering direct influence and inter topic dependencies is:

$$v_i^{h,s}(t+1) = v_i^{h,s}(t) + w_{ij}(v_j^{c,s}(t) - v_i^{h,s}(t)) + \sum_{s' \in \mathcal{T}} S_{ss'}(v_i^{s'}(t) - v_i^s(t)). \quad (4)$$

The chatbot’s objective is to select utterances $\{u_j^s(t)\}$ to minimize the deviation of the human’s final opinions from its own target opinions:

$$\min_{\{u_j^s(t)\}} \sum_{s \in \mathcal{T}} |v_i^{h,s}(t_f) - v_j^{c,s}(t_f)|, \quad (5)$$

where t_f is the final time-step. This objective formalizes a planning or optimization problem in which the chatbot attempts to steer the user’s opinions toward a desired configuration.

Scenario 2: Multiple Chatbots Competing to Influence a Network of Human Agents Consider multiple chatbots $\{v_k^c\}$ and human agents $\{v_i^h\}$. Each chatbot has its own target opinions and aims to influence multiple human agents. The opinion update for a human agent v_i^h on topic s is then:

$$v_i^{h,s}(t+1) = v_i^{h,s}(t) + \sum_k w_{ik}(v_k^{c,s}(t) - v_i^{h,s}(t)) + \sum_{s' \in \mathcal{T}} S_{ss'}(v_i^{h,s'}(t) - v_i^{h,s}(t)). \quad (6)$$

Here, each chatbot v_k^c selects utterances $\{u_k^s(t)\}$ to influence a network of human agents. An analogous objective for each chatbot is:

$$\min_{\{u_k^s(t)\}} \sum_{v_i^h \in V_h} \sum_{s \in \mathcal{T}} |v_i^{h,s}(t_f) - v_k^{c,s}(t_f)|, \quad (7)$$

where $V_h \subseteq V$ denotes the set of human agents in the network. In this setting, multiple chatbots attempt to concurrently shape the network’s opinion distribution, potentially leading to strategic interactions.

We consider two topics T_1 and T_2 , with the user’s opinions on these topics at time t denoted by $v_i^{h,T_1}(t)$ and $v_i^{h,T_2}(t)$, respectively. At $t = 0$, the user’s initial opinions are $v_i^{h,T_1}(0) = -0.5$ (user believes it is not Monday) and $v_i^{h,T_2}(0) = 0.0$ (neutral about T_2). The conversation impact factor is set as $\gamma = 0.8$, and $f^T(u_j^c(t), u_i^h(t))$ represents the impact of utterances on topic T .

3.4 Illustrative Example: Opinion Update

Consider the conversation presented in Table 2. The topics involved are: *Topic* s_1 - “Today is Monday.” and *Topic* s_2 - “Yesterday was Sunday.”. Assume a strong positive interrelation $S_{s_1 s_2} = 1.0$, indicating that belief in s_2 strongly supports belief in T_1 . A human agent v_i^h interacts with a chatbot v_j^c over time. The utterances at each time step t are represented as $u_i^h(t)$ for the human and $u_j^c(t)$ for the chatbot.

Opinion Updates We consider two topics s_1 and s_2 , with the user’s opinions on these topics at time t denoted by $v_i^{h,s_1}(t)$ and $v_i^{h,s_2}(t)$, respectively. At $t = 0$, the user’s initial opinions are $v_i^{h,s_1}(0) = -0.5$ (user believes it is not Monday) and $v_i^{h,s_2}(0) = 0.0$ (neutral about s_2). The conversation impact factor is set as $\gamma = 0.8$, and $f^s(u_j^c(t), u_i^h(t))$ represents the impact of utterances on topic s .

For this example:

- At $t = 0$: $f^{s_1} = +1$, $f^{s_2} = 0$ (chatbot provides factual information about s_1).
- At $t = 1$: $f^{s_1} = 0$, $f^{s_2} = +1$ (chatbot reinforces s_2 by clarifying it was Sunday).

The opinion updates follow equation 4.

At $t = 0$:

- $v_i^{h,s_2}(1) = 0.0$ (no change in s_2).
- $v_i^{h,s_1}(1) = -0.5 + 0.8 = 0.3$ (chatbot shifts belief on s_1).

At $t = 1$:

- $v_i^{h,s_2}(2) = 0.8$ (positive reinforcement of s_2).
- $v_i^{h,s_1}(2) = 0.3 + 0.8 = 1.0$ (bounded at 1.0).

At $t = 2$: No new information, so opinions remain:

- $v_i^{h,s_2}(3) = 0.8$.
- $v_i^{h,s_1}(3) = 1.0$.

By providing fact-based information and leveraging strong inter-topic relationships, the chatbot can influence user opinions more effectively, even when multiple topics are intertwined. Simple factual utterances may not only shift beliefs directly but also exert indirect effects through related topics.

This formalism provides a mathematical foundation for integrating opinion dynamics into dialog management. By representing opinions, their evolution, and intertopic dependencies, a dialog manager can systematically guide users toward targeted belief states. Existing planning algorithms can then be employed to find optimal sequences of utterances that drive these shifts. This process is further illustrated in Figure 3, which visualizes the concept of chat unrolling. The figure captures how sequential interactions between a chatbot and a user influence opinion states over time, highlighting incremental updates in user beliefs with each dialog turn. Each step of the unrolling process reflects how the dialog manager’s strategies and the user’s responses interact to progressively adjust opinions, ultimately converging toward the desired belief states. By providing a step-by-step visualization of the temporal evolution of opinions, Figure 3 highlights the underlying interaction mechanisms through two distinct scenarios involving multiple users and chatbots.

4 Discussion and Conclusion

This paper proposes a formal framework that integrates opinion dynamics into dialog management systems. The key idea is to conceptualize user opinions as dynamic states shaped by direct interactions and interrelations among multiple topics. By representing these interdependencies through

a matrix-based structure, the approach enables strategic interventions that guide user beliefs toward targeted goals.

The proposed mechanisms—ranging from targeted messaging to utterance-based opinion adjustments—demonstrate how a dialog system can influence user beliefs in a principled manner. Unlike traditional dialog frameworks that focus on intent recognition and task completion, the model presented here broadens the scope to include actively shaping user beliefs over time. This capability holds particular relevance in contexts where persuasion, negotiation, or the dissemination of accurate information is desired.

To realize the full potential of this approach, future work should explore ways to refine and scale the underlying planning algorithms. Additionally, it will be important to develop guidelines that ensure responsible opinion shaping while preserving user autonomy and trust. Evaluating these methods in real-world settings will be crucial to validate their effectiveness, efficiency, and acceptability. The role of planning knowledge formulated as concepts and relationships about domains, problems, plans and planners, formalized in an ontology, can be crucial. Such an ontology can be used for many applications for better dialogs like planner selection, generating explanations and discovering action hierarchies. An example of this is (Muppasani et al. 2024a).

In summary, this research advances the integration of opinion dynamics and dialog management, offering a foundation for developing more adaptive, context-aware, and influential conversational agents. By accounting for the complexity of interrelated topics and the evolution of user beliefs over time, such systems can transcend mere information exchange and serve as more persuasive, informative, and responsive tools in domains such as education, healthcare, and customer support. With planning and generalization techniques playing a crucial role in policy generation for dialogs, explicitly modeling the induced opinion network and controlling opinions or topics within them can further enhance the trustworthiness and effectiveness of conversational agents. In this manner, the presented work aims to extend the capabilities of current dialog systems beyond traditional intent-response paradigms, enabling a more strategic and controlled influence over user opinions.

References

Algherairy, A.; and Ahmed, M. 2024. A review of dialogue systems: current trends and future directions. *Neural Computing and Applications*, 36(12): 6325–6351.

Anderson, B. D.; and Ye, M. 2019. Recent advances in the modelling and analysis of opinion dynamics on influence networks. *International Journal of Automation and Computing*, 16(2): 129–149.

Brabra, H.; Báez, M.; Benatallah, B.; Gaaloul, W.; Bouguelia, S.; and Zamanirad, S. 2021. Dialogue management in conversational systems: a review of approaches, challenges, and opportunities. *IEEE Transactions on Cognitive and Developmental Systems*, 14(3): 783–798.

DeGroot, M. H. 1974. Reaching a consensus. *Journal of the American Statistical Association*, 69(345): 118–121.

Fellbaum, C. 2000. WordNet : an electronic lexical database. *Language*, 76: 706.

Liu, H.; and Singh, P. 2004. ConceptNet — A Practical Commonsense Reasoning Tool-Kit. *BT Technology Journal*, 22: 211–226.

Muppasani, B.; Gupta, N.; Pallagani, V.; Srivastava, B.; Mutharaju, R.; Huhns, M. N.; and Narayanan, V. 2024a. Building a Plan Ontology to Represent and Exploit Planning Knowledge and Its Applications. In *Eighth International Conference on Data Science and Management of Data (CODS-COMAD '24)*, India.

Muppasani, B.; Nag, P.; Narayanan, V.; Srivastava, B.; and Huhns, M. N. 2024b. Towards Effective Planning Strategies for Dynamic Opinion Networks. *arXiv preprint arXiv:2410.14091*.

Muppasani, B.; Narayanan, V.; Srivastava, B.; and Huhns, M. N. 2024c. Expressive and Flexible Simulation of Information Spread Strategies in Social Networks Using Planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 23820–23822.

Santos Teixeira, M.; and Dragoni, M. 2022. A review of plan-based approaches for dialogue management. *Cognitive Computation*, 14(3): 1019–1038.

Sîrbu, A.; Loreto, V.; Servedio, V. D.; and Tria, F. 2017. Opinion dynamics: models, extensions and external effects. *Participatory sensing, opinions and collective awareness*, 363–401.

Thomson, B.; and Young, S. 2010. Bayesian update of dialogue state: A POMDP framework for spoken dialogue systems. *Computer Speech & Language*, 24(4): 562–588.

Young, S.; Gašić, M.; Thomson, B.; and Williams, J. D. 2013. Pomdp-based statistical spoken dialog systems: A review. *Proceedings of the IEEE*, 101(5): 1160–1179.