



A Review of Plan-Based Approaches for Dialogue Management

Milene Santos Teixeira¹ · Mauro Dragoni²

Received: 12 May 2021 / Accepted: 13 January 2022
© The Author(s) 2022

Abstract

Dialogue systems deliver a more natural mean of communication between humans and machines when compared to traditional systems. Beyond input/output components that understand and generate natural language utterances, the core of a dialogue system is the *dialogue manager*. The aim of the dialogue manager is to mimic all cognitive aspects related to a natural conversation and it is responsible for identifying the current state of the dialogue and for deciding the next action to be taken by a dialogue system. Artificial intelligence (AI) planning is one of the techniques available in the literature for dialogue management. In a dialogue system, AI planning deals with the action selection problem by treating each utterance as an action and by choosing the actions that get closer to the dialogue goal. This work aims to provide a systematic literature review (SLR) that investigates recent contributions to plan-based dialogue management. This SLR aims at answering research questions concerning: (i) the types of AI planning exploited for dialogue management; (ii) the planning characteristics that justify its adoption in dialogue system; (iii) and, the challenges posed on the development of plan-based dialogue managers. The present SLR was performed by querying four scientific repositories, followed by a manual search on works from the most eminent authors in the field. Further works that were cited by the retrieved papers were also considered for inclusion. Our final corpus is composed of forty works, including only works published since 2014. The results indicate that AI planning is still an emerging strategy for dialogue management. Although AI planning can offer a strong contribution to dialogue systems, especially to those that require predictability, some relevant challenges might still limit its adoption. Our results contributed to discussions in the field and they highlight some research gaps to be addressed in future studies.

Introduction

Conversational agents or dialogue systems have been built with several purposes; from informal chit-chat conversations [1] to more sophisticated dialogue systems that support patients on health treatments [2, 3]. These systems have in common the fact that they try to reproduce the way humans communicate by mimicking cognitive aspects that are part of human' conversations. Aspects such as empathy [4], intention awareness [5], emotional management [6], and trust-building [7] are frequently exploited to deliver interactions that are more pleasant to the end-user if compared to traditional systems.

Dialogue systems that conduct interactions aiming at accomplishing a task or a goal in a given domain are known as *task-oriented* dialogue systems. One of the main components of this kind of systems is the dialogue manager [8], which is responsible for the flow of the conversation. The dialogue manager first identifies the current state of the conversation and, according to the recognized state, it selects the next action to be executed by the system.

Current commercial dialogue systems usually implement a *finite-state* or a *frame-based* dialogue manager [9–11]. These types of dialogue managers can be very efficient to simple dialogues whose goal is to fill a few slots. However, they are highly domain-specific and heavily based on hand-crafted rules and features, requiring high human effort to update or to design a new dialogue system. In addition, such dialogue managers are not capable of processing complex dialogues that require actions beyond slot-filling (e.g., *acknowledge*, *propose*, *assert*). More recent alternatives to building complex dialogues rely on data-driven techniques [12–14]. Data-driven approaches can be very efficient at learning policies from training data and they are

✉ Milene Santos Teixeira
msantosteixeira@fbk.eu

Mauro Dragoni
dragoni@fbk.eu

¹ University of Trento and Fondazione Bruno Kessler, Trento, Italy

² Fondazione Bruno Kessler, Trento, Italy

capable of outperforming hand-crafted dialogue managers [15]. However, they also can be limited to the quality of the training data and to the capabilities of the simulators that are commonly used during their training. Another important limitation of these approaches is that they remove some of the control on the policy, not being fully capable of providing explanations on their decisions, a behavior that may not be desirable in some conversational domains such as healthcare [16].

The drawbacks imposed by these approaches bring to light the need for alternatives for dialogue management that are reliable and that can be significantly easy to generate. A technique that automates action selection (avoiding the handcrafting of the dialogue policy) provides verifiable agents and offers policy control without limiting the flexibility of the dialogue is *automated planning*.

Artificial Intelligence (AI) planning [17] consists of a process that selects and structures domain actions in some rational order aiming at achieving a final pre-stated goal. Considering that a conversation also aims to achieve a goal and that dialogue managers deal with the problem of action selection, AI planning is a promising alternative for dialogue management [18–20]. Planning for human–machine dialogue offers the possibility of reducing the cost of the dialogue by anticipating the outcome of each dialogue action to identify the states that could be achieved. This way, it is possible to choose the best path that leads to the goal of the dialogue.

However, the development of methods that implement AI planning to structure and manipulate effective human–machine dialogue is still in its early stages. Indeed, when compared to other approaches for dialogue management, planning has been used on a quite smaller scale. Several may be the reasons for such low adoption; yet, no detailed analysis of the contributions and challenges of AI planning for dialogue management is available in recent literature.

The goal of this work is to provide a systematic literature review (SLR) in recent contributions of plan-based dialogue systems, with a focus on dialogue management. We intend to provide the reader with a broad picture of the current state of the art in this field. To achieve our goal, we identify the most recent works on the topic and extract data that help us to identify the following aspects: the types of AI planning exploited for dialogue management; the planning characteristics that justify its adoption; and the challenges posed on the development of plan-based dialogue managers. Our research reveals that planning can contribute to richer dialogues on very different application domains, but it still lacks research to overcome some relevant challenges (e.g., evaluation strategies).

The remainder of this paper is organized as follows: first, we discuss the background of the main elements of

this SLR, i.e., automated planning and dialogue systems (Background). Next, we present the method used to select the papers and list our research questions (Method). Results discusses the process conducted to extract the works to compose our final corpus. In Answering the Research Questions, we discuss the answers to the research questions. Then, on Discussion, we provide some further discussion on the topic and, finally, Conclusions concludes this SLR.

Background

This section introduces some fundamental notions from the theory of automated planning and dialogue systems, which constitute the background of this survey.

Automated Planning

Automated planning [17], also known as AI planning, can be viewed as a search problem that describes the sequence of actions that must be applied in the world in order to achieve a goal. Planning considers that the world has different *states*. A state can be changed to a new one by the so-called *actions*, which have *preconditions* that must be met to be applicable to that state; and present *effects* that describe what changes in the world after its application, i.e., add or delete values.

Besides the description of the available actions, a planning problem requires the specification of the *initial* and the *goal* state.

This description must be provided to an *automated planner* in a common representation language like the Planning Domain Definition Language (PDDL). Some examples of automated planners are the Planner for Relevant Policies (PRP) [21], Model Based Planner (MBP) [22], and the Planning with Knowledge and Sensing (PKS) [23]. Starting from the initial state, the automated planner investigates all possible states that can result from the application of the available actions until reaching the goal. This way, a *plan* is a path (sequence of actions) that takes the initial state to the goal state.

Dialogue Systems

Conversational agents have been built with different purposes and they fall into two categories according to the purpose of the dialogue [1]. *Chatbots* are dialogue systems whose only goal is to informally chit-chat with a human on any topic for which training data is available. Chatbots usually cover multi-turn interactions that imitate a human–human conversation rather than participate and contributing to some purpose. *Task-oriented dialogue systems*, on the other hand, conduct single or multi-turn dialogue interactions with the user in order to accomplish a task in

the given domain (e.g., booking a hotel). To retrieve the necessary information to accomplish the given task, these systems usually conduct a *slot-filling* dialogue.

Most task-oriented dialogue systems follow a standard “NL pipeline” design, which is composed of three main components [24]: the (i) *Natural Language Understanding* (NLU) component receives the user input and recognizes an *intent*; the (ii) *Dialogue Manager* receives this intent and tracks the current state of the dialogue to decide the next dialogue action; the (iii) *Natural Language Generation* (NLG) component converts this action to a human-understandable response that is given by the system to the end-user.

In this study, we keep our focus on the dialogue manager component, as the main contribution of AI planning is on action selection. Both the NLU and NLG components are out of the scope of this work and we refer the reader to [25] and [26] for surveys in those topics, respectively.

The dialogue manager [8] is responsible for the flow of the conversation in a dialogue system. Computationally, a dialogue can be seen as a series of communicative acts [27] (dialogue actions) that are used to accomplish particular purposes (goals). Therefore, in a dialogue manager, each utterance can be treated as an action and the dialogue manager deals with the problem of selecting an action to execute in a given *state*. States are used to understand and track the current status of the conversation, i.e., to identify what has already been discussed and what should come next. AI planning is capable of identifying the possible states that can be achieved within a dialogue and, for each state, deciding which action should be taken aiming to reach the goal of the dialogue. In this review, we keep our focus on works that have employed AI planning in the dialogue manager component.

Method

This work follows recognized guidelines for conducting a SLR. We refer mainly to the ones discussed by [28] and by [29]. In this section, we detail the steps that were conducted to accomplish the SLR.

Motivation and Study Goals

As discussed in the introduction, AI planning has been reported as an effective and efficient technique for dialogue management [18–20]. Nevertheless, it still has not gained due attention, presenting low adoption among current dialogue systems. A broad view of how planning strategies can be integrated within dialogue systems is missing in recent literature and the lack of such study led to an underestimation of the potentiality of such techniques. First, the lack of a comprehensive overview of what has been produced so far

might limit the evolution of already defined approaches. This might result in the design of redundant strategies to handle issues that have been successfully covered by previous literature. Another consequence is the low adoption of this technique since a shared knowledge of its characteristics and advantages for dialogue management is not available. Such comprehension would make it possible for dialogue engineers to consider the adoption of AI planning on scenarios that require features such as predictability, which are more effective on knowledge-based approaches when compared to data-driven ones. Finally, not knowing which are the challenges and open problems on plan-based dialogue management might lead to unexpected barriers to the development of plan-based dialogue systems.

The goal of this work is to provide a broad overview of recent contributions to plan-based dialogue management. As part of our goal, we aim at extracting the planning techniques that have offered a great contribution to dialogue systems. We also aim to identify the main characteristics of planning that justify its adoption in dialogue management. We expect that the comprehension of these characteristics will instigate a further investigation of plan-based dialogue systems. Lastly, we also have as our goal the identification of the main challenges on the topic. The finding of such challenges is crucial to decide on the suitability of AI planning for a given dialogue problem, but also to support future work on investigating and overcoming them.

[Research Questions](#) presents the research questions that were elaborated to motivate our investigation.

Research Questions

Starting from the goals described in [Motivation and Study Goals](#), we formulated three research questions to be answered within this review:

- **RQ1.** What types of planning have been used to support dialogue management in recent plan-based dialogue systems?
- **RQ2.** What are the most eminent planning characteristics that motivate the adoption of planning for dialogue management?
- **RQ3.** What challenges can be identified in recent plan-based dialogue systems?

The general questions above motivate discussions on the state of the art of plan-based dialogue management. *RQ1* seeks to identify which are the planning techniques that were exploited in the retrieved works, motivating a discussion on the reasons for their adoption. By describing some planning characteristics that are expected on efficient dialogue systems, *RQ2* helps us to justify the application and benefits of AI planning to support dialogue management. Finally,

RQ3 is devoted to the identification of the main challenges that are still encountered in the development of plan-based dialogue managers.

Search Strategy

To help answer our research questions, we retrieved works by using three strategies: (i) applying a search query to available scientific databases; (ii) conducting a manual search, i.e., by starting from the results of point (i), we extracted the most cited authors and checked if all their relevant works were already included; (iii) including additional works that did not appear in the results of the two previous strategies, but that were mentioned by the works retrieved by them.

To query scientific databases, four sources have been selected: *scopus*¹, *acm*², *dblp*³, and *ieee*⁴. The decision on the selection of these repositories was made given their extensive coverage of well-established scientific literature in computer science. In addition, these repositories include works that have been published in peer-reviewed conference and workshop proceedings. As a search strategy, a query was formulated including keywords related to the purpose of this work and the research questions listed in [Research Questions](#). The query is shown below:

```
("planning" OR "plan-based") AND
("dialogue" OR "dialog" OR "conversation" OR "conversational agents")
```

The resulting query aimed at focusing on a specific portion of works that employ AI planning on dialogue systems. The conjunction “AND” has been employed to avoid the retrieval of works that exploit different dialogue techniques, such as data-driven approaches. The authors have considered expanding the query to include further terms related to dialogue, such as *argumentation* or *interaction*. However, the inclusion of these terms resulted in a huge set of non-related works, revealing a deviation from the focus of our search.

From the authors’ previous knowledge of the works of some researchers in this field, several works that have exploited plan-based dialogue did not include the terms present in our query either in their title or abstract (e.g., [30, 31]). This way, we also opted to conduct a manual search on the bibliography of the most cited authors elicited in the previous step in order to include these and similar works.

Finally, relevant works mentioned within the selected corpus were also considered for inclusion, as long as they

Table 1 Exclusion Criteria

EC 1	The paper is not available.
EC 2	The paper is not in English.
EC 3	The paper is a technical report.
EC 4	The paper is shorter than 4 pages.
EC 5	The paper was published before 2014.
EC 6	The paper does not cover either dialogue or AI planning.
EC 7	The paper does not address dialogue management.

meet our eligibility criteria described in [Article Selection](#) and [Quality Assessment](#).

Article Selection

To filter the retrieved works and to focus only on the ones that meet our main objectives, we have defined some exclusion criteria (EC), which are listed in [Table 1](#).

Papers were excluded if they satisfied at least one of the exclusion criteria between EC1 and EC7. These criteria aim to exclude papers that are not available online, are written in some language rather than English, are not scientific contributions, or that are too short to provide a complete approach. To keep the focus on the most recent contributions, we also excluded works that have been published before 2014. Papers that do not address either dialogue systems or automated planning are also excluded as they cannot contribute to the purpose of this review. To direct this review to dialogue management, we excluded works whose main focus is on natural language processing. Finally, we opted to not exclude works that referred to the same approach as those, eventually, presented slightly different planning techniques.

Quality Assessment

To assess the quality of the investigated studies, we defined six criteria:

- **QA1** Is a well-defined methodology used?
- **QA2** Is the goal of the study clear?
- **QA3** Does the study discuss the state-of-the-art?
- **QA4** Is the approach applicable to different domains?
- **QA5** Does the study include a clear evaluation?
- **QA6** Does the study present an implementation or a clear illustrative problem of the approach?

Each paper was marked with one of three possible scores for each criteria (QA1-QA6): *Yes*, *No*, or *Partially*, which are weighted 1.0, 0.0 and 0.5, respectively. An exception was delineated for *QA5*, for which we defined the weights: 1.0 when an evaluation with real users has been conducted; 0.5

¹ <https://www.scopus.com/search/>

² <https://dl.acm.org/>

³ <https://dblp.uni-trier.de/>

⁴ <https://www.ieee.org/>

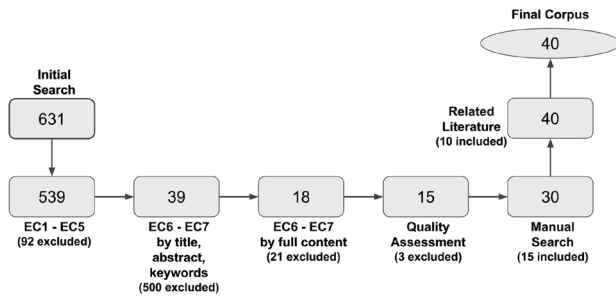


Fig. 1 Papers selection process

when the evaluation was conducted on synthetic scenarios (use cases); 0 when no evaluation was reported.

Data Extraction Strategy and Analysis

Aiming to extract only works that can answer our research questions, we extracted data by following three steps. First, all papers are evaluated against our EC by taking into account only their title, abstract and, keywords. Second, we carefully evaluate the works against our EC by taking into account their full content. Lastly, we apply the proposed quality assessment to the remaining papers, keeping only works that get a score equal to or above 3.0.

Results

In this section, we describe the process of extracting the works for answering our research questions. Our search has been conducted according to the methodology described in [Search Strategy](#). The data extraction has been conducted between December 2018 and November 2021. To focus on the most recent contributions while obtaining a satisfying amount of works to answer our research questions, we have limited our results to works that have been published since the year of 2014. Works retrieved in our manual search and from the related literature were also limited to papers starting from 2014.

Our search process is detailed in [Fig. 1](#). Initially, our search query retrieved a total of 631 papers.

Given our query definition ([Search Strategy](#)), which avoids the retrieval of approaches that are not plan-based, it is not surprising that we started with a small number of works. As discussed by Cohen [19] and by Petrick and Foster [18], plan-based dialogue systems still have not received due attention from the research community.

Next, we applied EC1 to EC5, which excluded 92 papers. EC6 and EC7 were first used to filter the papers' title, abstract, and keywords. Since *planning* is a rather common word, our query returned several works that do

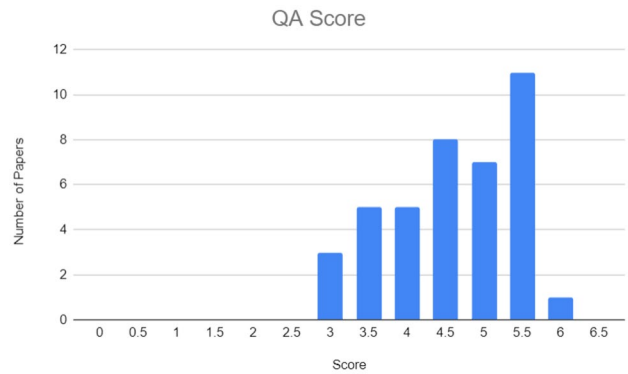


Fig. 2 Quality assessment: total of papers grouped by score

not address the AI planning community, but, instead, communities that use this term in different contexts (e.g., task planning or family planning) or communities that also may be able to conduct some kind of planning, such as machine learning. As those are not the focus of this survey, these papers were eliminated. This way, EC6 and EC7 excluded an extra 500. At this point, 39 papers remained. For the remaining papers, we then applied EC6 and EC7 by analyzing their entire content, which excluded another 21 works. In both steps considering EC6 and EC7, manual analysis was conducted.

Following our quality assessment, we further eliminated 3 papers that did not reach the minimum expected quality (3 out of 6.5), described in [Data Extraction Strategy and Analysis](#). Our manual search included another 15 works, which correspond to meaningful works from relevant authors. These works were obtained after analyzing the profile of the most cited authors within the initial corpus. By meaningful, we refer to works that, after full reading, were considered relevant to the topic of this SLR.

Finally, after carefully analyzing the resulting works, 10 papers have been included in our corpus as they were mentioned by the retrieved works and offered a meaningful contribution. Both works from the manual search and works retrieved from the related literature meet all our selection criteria and our quality assessment. [Figure 2](#) reports the total number of papers grouped by the score that they received during the quality assessment. Our final corpus is composed of 40 works and it is presented in [Table 2](#), which also describes the papers' publishers.

Answering the Research Questions

In this section, we discuss the answers to the research questions proposed in [Research Questions](#).

Table 2 Final corpus of surveyed papers

Id	Reference	Published at	Publisher
P01	Lison [32]	IWSDS	Springer
P02	Black et al. [33]	CLIMA	Springer
P03	Honold et al. [34]	IE	IEEE
P04	Black et al. [30]	AAMAS	ACM
P05	Panisson et al. [35]	ArgMAS	Springer
P06	Lee et al. [36]	PlatCon	IEEE
P07	Botea et al. [37]	DEEP-DIAL	AAAI
P08	Nothdurft et al. [38]	MMSYM	LiU Electronic Press
P09	Kominis and Geffner [39]	ICAPS	AAAI
P10	Nasihati et al. [40]	ICMI	ACM
P11	Biundo and Wendemuth [41]	KI-Künstliche Intelligenz	Springer
P12	Nothdurft et al. [42]	SIGDIAL	ACL
P13	Galescu et al. [43]	SIGDIAL	ACL
P14	Bercher et al. [44]	ICAPS	AAAI
P15	Geib et al. [45]	ICAPS	AAAI
P16	Petrick and Foster [18]	AAAI	AAAI
P17	Foster and Petrick [46]	IWSDS	Springer
P18	Nothdurft et al. [47]	IWSDS	Springer
P19	Behnke et al. [48]	ICCT	IEEE
P20	Petrick and Foster [49]	PlanSIG	University College
P21	Shams et al. [50]	IJCAI	AAAI
P22	Baskar and Lindgren [51]	PAAMS	Springer
P23	Behnke et al. [52]	Companion Technology	Springer
P24	Zhang and Stone [53]	AAAI	AAAI
P25	Lu et al. [54]	IROS	IEEE
P26	Morbini et al. [55]	IWSDS	Springer
P27	Marques and Rovastos [56]	HAIDM	HAIDM
P28	Franzoni et al. [57]	WI	ACM
P29	Cohen [19]	AAAI	AAAI
P30	Behnke et al. [58]	IJCAI	AAAI
P31	Garcia-Olaya et al. [59]	WAF	Springer
P32	Teixeira et al. [60]	ISWC	Springer
P33	Sreedharan et al. [61]	ICAPS	AAAI
P34	Muise et al. [62]	arXiv ⁵	-
P35	Petrick and Foster [63]	Knowledge Engineering Tools and Techniques for AI Planning	Springer
P36	Cohen [31]	SIGdial	ACL
P37	Pardo and Godo [64]	Journal of Logic and Computation	Oxford University Press
P38	Behnke et al. [65]	ICAPS	AAAI
P39	Teixeira et al. [66]	SAC	ACM
P40	Teixeira et al. [67]	CBMS	IEEE

⁵This work is currently under revision, but it was included in our corpus as it presents strong relevance for the domain

RQ1. What Types of Planning Have Been Used to Support Dialogue Management in Recent Plan-Based Dialogue Systems?

A range of different planning techniques is available for plan generation and the choice of one or another varies according to the problem approached. In dialogue systems,

examples of some aspects to be considered when choosing a planning approach are: (i) are actions deterministic, non-deterministic or probabilities are available for the outcomes? (ii) Is it possible to subdivide the dialogue into sub-dialogues or smaller tasks to be accomplished? (iii) Are there efficient planners available for the chosen planning approach?

Table 3 Planning techniques

Planning Technique	Paper ID
Classical	[18, 30, 31, 33, 39, 45, 46, 50, 55–57, 59, 63, 64]
Probabilistic	[19, 30, 32, 33, 45, 53–55]
HTN	[35, 36, 42, 47, 58, 65]
POMDP	[32, 48, 53, 54]
FOND	[37, 60–62, 66, 67]
Hybrid	[34, 44, 48, 52]

In answering RQ1, we obtain an overview of which planning techniques have been proven suitable to address dialogue management. It is important to note that we do not attempt to cover the whole variety of planning techniques; instead, we cover the ones identified in our corpus only. The planning techniques are listed in Table 3 and their description is summarized in the sequence. We would like to note that, for different reasons (e.g., planning was not the main focus of the approach or the paper presented a broad discussion on plan-based dialogue), some works have not specified a planning technique [38, 40, 41, 43, 49, 51].

Classical Planning

The most simple planning technique is the *classical* approach. Traditional classical planners generate (offline) a sequence of deterministic actions (each action has only one outcome) that lead to the goal (Fig. 3).

A classical planning problem [68] is a 4-tuple $\langle A, O, I, G \rangle$ where A is a finite set of applicable actions, O is a finite set of domain objects, I is the initial state and G is the set of goal variables. Actions are in the form of $\langle Pre, Add, Del \rangle$, where *Pre* stands for precondition, and *Add* and *Del* stands for the predicates to be added and deleted as effects of the action execution. As can be viewed in Table 3, classical planning was the most common approach among the surveyed works.

Some works have supported classical planning with strategies that make it possible to adapt the plan according to the domain needs. Examples are: Franzoni et al. [57] that took into account emotional affordances in the planning model to prioritize actions, inform inconsistencies, and acquire some missing information; and Kominis et al. [39] that extended




Fig. 3 Rough representation of a classical plan. Nodes represent the states while edges represent the deterministic actions, being i the initial state and g the goal state




Fig. 4 Simple representation of an action with a probabilistic outcome, being i the initial state, and s_1 and s_2 the possible resulting states after the execution of action a

STRIPS with negation, conditional effects, and axioms to translate the dialogue problem with hidden initial state to a classical plan.

Some of the listed approaches, instead, do not focus on planning a dialogue itself but make an integration among planning and dialogue. For example, by evaluating goals and norms, the work of Shams et al. [50] reasons on the best plan and integrates this reasoning into a dialogue that provides explanations about its choice to the agent. In the collaborative approach of Geib et al. [45], after recognizing the plan (goal and subgoals) of the initiator agent, the supporter agent conducts a negotiation dialogue to define which actions it should execute. Only then, a plan is generated to organize the sequence among these actions. Also aiming to address a collaborative scenario, Pardo and Godo [64] propose an algorithm where multiple agents collaborate through dialogue to build a plan. In this work, the agents communicate relevant information to be taken into consideration when choosing an action for the plan.

Probabilistic Planning

The use of probabilistic reasoning was also frequent among the surveyed papers. Probabilistic planning is an extension of non-deterministic planning and it is used to address domains with some certain kind of uncertainty.

A probabilistic planning domain is a 4-tuple (S, A, γ, Pr) where S is a finite set of states, A is a finite set of actions, $\gamma : S \times A \rightarrow 2^S$ is the state-transition function, and Pr is the probability-transition function (in the form of $S \times A \times S \rightarrow [0, 1]$). Figure 4 shows a simple example of an action with a probabilistic outcome. That is, when action a is executed, the probability of reaching the state s_1 is 70% and state s_2 is 30%.

Besides being addressed in the effects of an action (e.g., a probability of occurrence is known to each of the different outcomes of the given action), probabilities can also represent the probability of success of a plan, guiding the search for a plan that maximizes the probability of reaching the goals. To some extent, an example of the former can be found in [55] that weights the information states with the

probability that they can be reached. The latter, on the other hand, is the case of [30], where the plan *guarantees* a certain probability of success of a persuasion strategy.

The works mentioned above have handled probabilistic problems by using classical planning techniques (**Classical Planning**). Instead of classical planning, some works have managed the probabilistic dialogue problem by using a partially observable Markov decision process (POMDP) [69]. A POMDP supports the state transition based on observations of the world state and their probabilities. Such observations form a belief of the current state (i.e., partial observability) and POMDP is able to generate a policy describing which action is optimal given a belief state. For example, Lison [32] relies on probabilistic rules (which can then be converted to POMDP) for updating the dialogue state, determining the possible actions available at that state and the reward that can be obtained with the execution of an action (each effect is assigned a probability). Instead, the architecture proposed in [48] integrates a POMDP module to monitor the dialogue and identify when a planned action does not correspond to the action that was determined by the POMDP.

Fully Observable Non-Deterministic (FOND) Planning

When the domain presents some level of uncertainty and probabilities are not available or are not convenient, non-deterministic methods can be applied. In fully observable non-deterministic (FOND) planning problems, the actions may have a set of different possible effects. Formally, a FOND planning problem is a 4-tuple (P, O, S_I, ϕ_{goal}) where P is a set of Boolean state propositions, O is an operator set, S_I is the initial state, and ϕ_{goal} is the goal state set.

FOND plans anticipate all possible states that can result from the execution of a non-deterministic action. Figure 5 shows a simple example of an action with non-deterministic outcome. Note that, different from the probabilistic action illustrated in Fig. 4, no probabilities are available for the outcomes. During plan execution, the agent is able to fully observe which effect has actually occurred in the real world.

FOND planning has been successfully adopted in [37, 62] and [61], which are related works. In these approaches, the solution to the FOND problem consists of a non-deterministic graph plan with actions to be executed in a dialogue system.

In [60, 66, 67], on the other hand, FOND planning is exploited with the aim of supporting dialogue management in health dialogues. To achieve this goal, domain-specific knowledge is translated into a planning problem. Each dialogue action is treated as a non-deterministic planning




Fig. 5 Simple representation of an action with a non-deterministic outcome, being i the initial state, and s_1 and s_2 the possible resulting states after the execution of action a

action. This way, FOND plans anticipate the different paths that the dialogue can take given the different answers that can be provided by the end-user.

Hierarchical Task Network (HTN)

Unlike classical planning that focuses on accomplishing goals, Hierarchical Task Networks (HTN) [70] aim to accomplish tasks. Tasks are high-level descriptions of some activity to be carried on and HTN provides methods for problem reduction. That is, *abstract tasks* are decomposed into sub-tasks until the achievement of a primary action that cannot be further decomposed, i.e., a *primitive task* (Fig. 6).

An HTN problem is a 4-tuple (S_0, T, O, M) where S_0 is the initial state, T is a set of initial tasks which defines the goal, O is a set of operators that define the achievable actions, and M is a set of decomposition methods.

Instead of a search through the state-space, HTN searches through the plan-space and its solutions are partial-order plans. In practice, HTN is one of the most used planning techniques and it was also frequent among the surveyed works. One reason for its substantial adoption within dialogue systems is that HTN assimilates




Fig. 6 Hierarchical Task Network planning: representation of an abstract task and the corresponding primitive tasks achieved through the decomposition process

to a human decision strategy, i.e., more abstract tasks are decomposed into smaller and easier actions to be accomplished.

An example of the use of HTN in a dialogue system is the work of Panisson et al. [35] that formalize HTN methods for an argumentative dialogue. By using an HTN planner, this approach can generate both: an immediate response or an optimal approach. Instead, the work of Nothdurft et al. [42] benefits from HTN planning to support the user to accomplish complex goals (tasks). In their approach, the HTN decomposition process of higher level tasks involves the end-user through a dialogue that asks the user to choose the sub-tasks.

Hybrid Planning

Hybrid planning integrates hierarchical actions that are partially adopted from HTN (**Hierarchical Task Network (HTN)**) and Partial-Order Causal-Link (POCL) [71].

A hybrid planning problem is a 6-tuple $(V, N_c, N_p, \delta, M, P^i)$ where V is a finite set of state variables, N_c is a finite set of compound task names, N_p is a finite set of primitive task names, δ is a function mapping the task names to their pre-conditions and effects, M is a finite set of decomposition methods, and P^i is the initial plan. In a hybrid planning problem, both abstract and primitive tasks can contain so-called *causal links*. This way, a hybrid plan presents a similar structure to an HTN plan (Fig. 6) with the difference that causal links will add some constraints on the ordering of the tasks (either abstract or primitive). For further details, we refer the reader to [71].

Similar to HTN planning, this type of planning has been reported as well suited for user-centered planning applications [72]. This happens because, to some extent, hybrid planning imitates the strategies that humans use for solving problems, i.e., in a hierarchical (top-down) manner [73]. However, a further advantage of hybrid planning is that it allows some flexibility to explain its decisions as a consequence of relying on causal reasoning.

An example of its use in a user-centered application is given in [44] for a scenario that supports a user to assemble a home-theater. Hybrid planning contributes to the execution of actions that are semantically connected, instead of actions that reach the goal but in an order that does not make sense to the user (domain-specific knowledge was required). In addition, when differences in the expected state and actual state are detected, the use of causal links helps to identify if replanning should be conducted.

RQ2. What are the most eminent Planning Characteristics that Motivate the Adoption of Planning for Dialogue Management?

The comprehension of AI planning characteristics is important to analyze whether this is a suitable strategy to be adopted for a given problem. Indeed, some planning characteristics, such as the possibility of explaining its action choices, are especially attractive to the dialogue community. This way, the aim of RQ2 is to provide a discussion on the planning characteristics that received significant attention among the surveyed works.

Goal-Driven Plans

Similar to a conversation, a plan aims to achieve a goal. In planning for human–machine dialogue, by anticipating the outcome of each dialogue action, a planner is capable of identifying the states that could be achieved and of choosing the best path that leads to the goal of the dialogue. In [RQ1. What Types of Planning Have Been Used to Support Dialogue Management in Recent Plan-Based Dialogue Systems?](#), we described the different types of planning approaches that have been used among the surveyed works to find a plan that reaches a goal in conversational scenarios.

Differently from traditional data-driven based dialogue systems, which frequently focus on achieving goals such as booking a restaurant or planning a trip [74], the plan-based surveyed works aimed at accomplishing a range of different goals that addressed tasks such as: supporting a user to assemble a home-theater [44]; building a workout plan [48]; a robot taking orders placed by human agents [53]; a bartender robot serving its customers [18]; car inspection [62]; supporting asthma patients [66]. Furthermore, to achieve the goal of the dialogue, we have observed the frequent use of some conversational strategies, such as collaboration [19, 39, 43, 45, 56, 64, 65], argumentation [30, 33, 35, 50, 64] or negotiation [35, 45, 57]. In fact, AI planning will always try to achieve a goal; the difference is in the strategy applied for it, which may vary according to the type of planning employed.

Classical planning presents a sequence of deterministic actions to reach the goal. However, instead of just selecting the shortest path to the goal, some of the surveyed works have exploited different strategies to define this sequence. For example, in the work of Franzoni et al. [57], action selection is affected by emotional affordances that might prioritize or exclude some actions. In the problem exploited by Shams et al. [50], on the other hand, selecting the shortest path might not be the most convenient solution as some

norms might influence action selection. Non-deterministic plans, instead, take into account the possible different outcomes that can result from an action execution. Probabilistic or POMDP planning can rely on probabilities to identify the best action towards the goal [32]. Meanwhile, as FOND planning treats all possible action outcomes equally and it will only know the result after action execution, it must anticipate all possible paths to reach the goal [62]. Finally, descriptions of HTN and hybrid planning strategies to reach the goal are given in [Hierarchical Task Network \(HTN\)](#) and [Hybrid Planning](#), respectively.

Among the surveyed works, only the work of Nasihati et al. [40] could not be categorized as goal-driven, but instead, as a chatbot [1]. Their system uses a dialogue manager that tries to engage infants to interact, maintain this engagement and promote responses from the baby with the aim of facilitating language (and also sign language) learning. The system does not have a final goal that must be achieved within a dialogue session. Finally, in some cases, the goal of the dialogue can change during a dialogue session. Lee et al. [36] exploited situation awareness to recognize possible changes in the goal of the dialogue and, if necessary, to build a new plan to accomplish the new recognized goal.

Addressing Large State Spaces

Traditional approaches for dialogue management, like finite state machines [1], require the dialogue author to handcraft the dialogue tree. Manually building extensive dialogue trees for complex dialogues is a time-consuming process. In addition, as the model size (i.e., number of variables, actions, states) increases, the complexity of the dialogue also increases, making this process error-prone and, in some cases, not feasible. Therefore, one of the main advantages of AI planning for dialogue is the automation of this process. As a consequence of this automated process, plan-based policies can implement much larger state spaces when compared to traditional approaches.

It is important to highlight that the planning type employed in the dialogue manager directly affects problem and solution sizes. Probabilistic approaches, for instance, suffer from a limitation on the number of domain variables that are allowed in the model. To handle this limitation, the POMDP approach in [54] has relied on commonsense reasoning to simplify and guide the dialogue manager. HTN planning, instead, is capable of reducing the problem into smaller tasks thanks to its decomposition feature. Therefore, instead of generating a single plan that includes large amounts of information, partial plans can be generated according to the progress of the dialogue [36]. Classical planning is also capable of addressing problem sizes that are very difficult to be solved by humans. An example is the

approach of Black et al. [33], which presents a solution for simple persuasion dialogues that present considerable complexity for hand-crafting. In further work [30], the authors showed a satisfactory plan generation time for argumentative dialogues with up to 13 arguments. While this does not seem a very huge problem size, the authors have verified that the approach covers realistically sized dialogues. FOND planning is another good candidate to handle huge state spaces. In [60] that specifies a health dialogue as a FOND problem, the authors showed how a very simple scenario can result in large state-spaces but that can be efficiently handled by a FOND planner. In a more recent work [67], the authors show the suitability of the approach to generate dialogue policies on real-time scenarios, addressing up to 160 slots. Botea et al. [37] that also exploited FOND planning demonstrated results of a synthetic scenario in which the solution size remained, in most cases, 4 times the model size. In some extreme cases, it grows up to 16 times the model size, without compromising the efficiency of the plan generation. In a further work [62], the authors showed satisfactory results for plans addressing up to 28 variables and resulting 482 nodes. Such problem size is intractable for human beings.

Although several data-driven approaches are also capable of generating dialogues with great complexity, these approaches require large amounts of training data, being limited to the capabilities of user simulators. In addition, updating the model to include extra variables in the dialogue means that more data must be provided for training.

It is important to note that the specification of the planning problem will affect the planning solution. With the aim of obtaining reasonable solutions for complex dialogues, aspects such as abstraction and the type of planning chosen for the problem must be carefully addressed for the dialogue domain. These aspects can be considered challenges in plan-based dialogue management and they are better discussed in [Dialogue Modeling](#). Finally, although planning is capable of implementing larger state spaces than several other techniques, a limitation with respect to scalability has been reported by some approaches and it will be discussed in [RQ3. What Challenges can be Identified in Recent Plan-Based Dialogue Modeling](#).

Explainability

Some problems require the implementation of agents that can explain their behavior [19]. In recent years, explainable AI planning (XAIP) has received substantial attention [75–78]. Explainability is a feature of AI planning that can be exploited for all types of planning and it can be provided from two perspectives: the developer and the end-user. From the developer's point of view, planning is explainable in the sense that it is possible to debug the resulting tree to understand each step that leads to the goal achievement.

Consequently, planning makes it possible to control the generated policy and to keep track of the behavior of the agent. Regarding the end-user, planning also makes it possible to provide explanations that justify its actions choice. This aspect is especially relevant if we take into account that an optimal solution provided by the planner can still differ from the user's mental model, that is, from an expected course of actions that is based on the user's domain knowledge. When no further explanation is provided in such situations, the plan solution might result in some confusion on why the system is taking a given action, endangering the user's trust in the system.

Among the surveyed works, explainability has been exploited in AI planning with aims like (i) explaining generated plans [38, 50], (ii) explaining the choice of an action [34, 35, 42, 44, 58], or, less frequently, (iii) explaining the unsolvability of a model [61]. A quick description of the use of explainability in these works is given below.

In Nothdurft et al. [38], the authors investigate appropriate interaction strategies for a proactive plan-based dialogue system. The authors conclude that a system should attempt to explain or justify its decisions to the user to align the perceived mental model to the actual system model. Later, in [42], the authors implemented proactive as well as requested explanations in a plan-based dialog system. The proposed approach builds an exercise plan together with the user, for which explanations can be given on each step, helping the user to understand why an action is being chosen and what should be changed if this is not the desired choice. In addition, as the plan is built dynamically with the user, when dead-ends are found and the planning system needs to roll back some previously made decision, this decision is also explained to the user. In Behnke et al. [58], a domain ontology is used both for plan generation and to enrich explanations on the plan steps.

The approach applied in Bercher et al. [44] and Honold et al. [34] deals with plan repair. As changing a plan during its execution might cause confusion to the user, the authors identified the need of providing explanations when requested by the user. By using formal proofs, which are translated with a template to a user-friendly language, their approach generates explanations on *why* a given step is present in the plan (e.g., it is a precondition for a further system action). Causal links that are used to build the plan also support the generation of the explanations.

To improve the agent's chance of success in finding the best plan to reach the desired outcome within a negotiation, the agent proposed in [35] justifies its actions (negotiation stance) with the arguments that were used to choose this action. Arguments have also been used for generating explanations in the work of Shams et al. [50]. In this normative approach, where norm compliance is weighted against goal

achievement, the agent explains why a given plan was chosen as the best one for execution.

Regarding explanations on model unsolvability, Sreedharan et al. [61] aimed at supporting dialogue authors during model acquisition with explanations that help them to understand why the designed model cannot be solved. The purpose of such explanations is to support authors in finding a fix to arrive at the solution.

Handling Uncertainty

Among the types of planning presented in RQ1. [What Types of Planning Have Been Used to Support Dialogue Management in Recent Plan-Based Dialogue Systems?](#), probabilistic ([Probabilistic Planning](#)) and FOND ([Fully Observable Non-Deterministic \(FOND\) Planning](#)) planning are the most appropriate types to address domains that present some level of uncertainty and for which it is not possible to define deterministic models. However, it is important to note that handling uncertainty in a dialogue problem is not limited to these approaches. In fact, the surveyed works exploited planning in different ways to address uncertainty, being it handled differently according to where it occurs in the dialogue.

When the uncertainty is in the initial state, for example, exploring all possible alternatives for the truth initial state might be necessary. However, as reported by [39], this strategy increases the plan space and might result in scalability problems ([Scability](#)).

A few different strategies have been applied to address uncertainty on the current state of the dialogue. Bercher et al. [44], for example, rely on a *probability* distribution over the possible world states. A quasi-deterministic model that takes into account the probability that a variable last had is applied to any variable that has not received an updated explicit observation during the last interaction. Lison [32], instead, encodes the dialogue state as a Bayesian Network. The algorithm proposed by the author relies on high-level probabilistic rules that are based on prior domain knowledge. Instead of probabilities, *non-deterministic* approaches have also been adopted to represent uncertain states within the dialogue. In Botea et al. [37], non-deterministic planning was used to anticipate the different outcomes that can result from an action execution, predicting the possible resulting states and, therefore, the possible dialogue paths. Muise et al. [62] introduce a *determiner* that is in charge of identifying which non-deterministic outcome has actually occurred on execution time. Handling uncertainty on the current state has also been addressed with the use of *replanning* ([Dynamic Policy Through Replanning](#)). In Garcia et al. [59], a plan is built by assuming that each action execution will result in the most likely effect. When this is not the case, replanning is conducted.

Uncertainty may also come from the module that precedes the dialogue manager, i.e., the language understanding module. The uncertainty on the information received has been handled with confidence values [23, 48, 53]; properties like the *badASR* property in [23] that indicates low-confidence automatic speech recognition; or with fluents like the MAYBE-* fluent adopted in [37] to identify when a piece of information is uncertain. These strategies instigate the adoption of clarification questions, instead of the reproduction of a previously made question.

Additional techniques have been used to handle uncertainty in other parts of the plan-based dialogue. Common sense, for example, was exploited in the POMDP approaches discussed in [53, 54] to reduce the uncertainty in action outcomes. Aiming to handle the agent's uncertainty on its opponent's beliefs, Black et al. [30, 33] used conformant planning for compiling the uncertainty away and model the domain as a classical plan. Other works that have addressed uncertainty to some extent are [46, 56].

Mixed-Initiative

Mixed-initiative is a feature that is not a direct characteristic of AI planning, but that can be implemented and exploited by a plan-based dialogue manager. Mixed-initiative can be achieved with all types of planning and there are different possibilities to implement it; examples are: dialogue actions [61], predicates that work as flags to switch the initiative [55] or through replanning [63].

In the surveyed works, we could extract two types of mixed initiative: mixed-initiative planning (MIP) systems and mixed-initiative dialogues.

The purpose of MIP systems is to collaborate with the user for decision making with the aim of obtaining high-quality plans. To achieve this goal, the user itself is included in the planning process. In MIP, the system first interacts with the user to define a goal understandable by the planner. In the sequence, the interactive planner applies a strategy for refining the tasks that are required to accomplish this goal and it decides whether to include the user in each decision step. The MIP approach in [42] (also discussed in [47, 48, 52]), for example, includes the user in the process of generating a workout plan. Together with the user, the dialog system that relies on HTN planning decomposes the abstract actions into executable sub-actions. This approach reduces the cognitive load on the user to find a solution while building a plan that includes the user's preferences. However, although the generation of the plan is part of the MIP system, the dialogue that requests the choice of the next plan action, is system initiated.

Mixed-initiative dialogues, on the other hand, are able to support dialogue interactions that are initiated either by the system or by the user. Systems that implement

mixed-initiative dialogues are able to keep a more natural conversation when compared to system-initiated ones that present lower complexity. As previously mentioned, AI planning supports mixed-initiative and this feature has been exploited by some of the surveyed works in different ways.

The approach proposed by Galescu et al. [43], for example, supports mixed initiative among multi-agents that address collaborative problem solving. Since constant initiative switches can significantly increase the complexity of the dialogue, to facilitate and order the conversation among the agents, the initiative switch is restricted by an act that is called by the dialogue manager only when there are no further tasks pending or in progress.

The dialogue manager proposed by Morbini et al. [55], on the other hand, takes the initiative when it is necessary to obtain some information from the user, but it is also able to respond appropriately to user-initiated utterances. This behavior is achieved through the use of a precondition that informs on which initiative each operator can be activated. For example, operators with the precondition for system initiative mean that this operator can be activated at any moment in the dialogue, while, a precondition for user initiative means that this operator will be used to handle some user-initiated input.

Sreedharan et al. [61] addressed mixed initiative through the specification of a dialogue action to respond to user-initiated utterances. In addition, in their non-deterministic approach, every system action has an additional outcome to switch initiative based on the user input. In [62], mixed-initiative has been addressed similarly.

Finally, a common approach in classical planning to address user-initiated actions, i.e., actions that were not expected within the plan in execution, is replanning. This strategy has been adopted to address mixed-initiative in the bartending scenario developed by Petrick et al. [63]. Replanning is discussed in [Dynamic Policy Through Replanning](#).

Dynamic Policy Through Replanning

Handcrafted dialogue trees limit the flexibility of the dialogue in the sense that the dialogue must stick to a predefined path and, most times, ignore any information that is not expected at the current state. AI planning, on the other hand, presents the possibility of dynamically updating the dialogue policy through replanning.

Replanning can be conducted to all types of planning and upon different circumstances. Goal change, for example, requires a new plan to be built. In the work of [36], dynamic plans are built according to the current situation. When the situation changes, the goal also changes. Therefore, the existing plan must be abandoned and a new plan is generated to address the newly detected situation.

Execution failure also can be managed with replanning. An example is the approach proposed by [56], where each agent has a plan and tries to collaborate with the other agent by the means of a dialogue. When the agent's communicative plan fails (e.g., a belief about the other agent was wrong), it immediately updates its beliefs and replans. In the approach of [59], besides replanning if some error occurs during execution, replanning is also conducted if some unexpected event modifies the planned flow and invalidates the next planned action. In [67], instead, in case of a non-expected input, replanning can be conducted without affecting the ongoing dialogue session. Similarly, in [55], if a received event cannot be handled by the active operator, the dialogue manager simulates future dialogues that would result from the activation of each available operator and selects the most promising operator to handle this event.

In the bartender agent discussed in [63], the plan is also dynamically updated through replanning in a few situations: (i) on action execution failure or a not understood response (low confidence score); (ii) on over-answering, i.e., the user provides more information than what was asked at that state; and (iii) when the bartender agent notices a new customer that was not present during the generation of the current plan. In this last case, the new plan consists of an extension of the existing plan, where actions that were in progress are maintained and the new ones are added to the flow.

Replanning also enables the handling of large state spaces, which can be divided into smaller problems, feasible for runtime generation. This way, after reaching a milestone, a new plan can be built to address new variables. In the online approach proposed by Kominis and Geffner [39], for example, plans are built with a single action and, after executing and observing this action's outcome, replanning is conducted by taking into account the updated knowledge. This strategy is aimed to address multi-agents that collaborate to achieve a common goal.

Horizontal Ingredients

Besides the characteristics discussed above, two further ingredients received significant attention in the surveyed works. These ingredients are not direct features of AI planning, but they are worth some discussion. The first one is the *adaptability* of the dialogue policies. The second is the integration of *ontologies* as a knowledge source for the planning problems. Both aspects are discussed below.

Adaptability Examples of adaptability include: adapting plans to user's priorities or preferences [42, 44, 48, 51, 58, 65], to dynamic changes in the environment [32, 44, 57] or to the dialogue history [42, 48].

To address adaptability some of these works relied on ontologies (**Horizontal Ingredients**), replanning (**Dynamic Policy Through Replanning**), or online planning [32].

Concerning the latter, by considering that adapting a policy once it has been calculated can be quite difficult, online planning is a suitable strategy to address dynamic changes in the environment. Online approaches can more easily adapt to the actual state of the dialogue and plan the next action accordingly. The drawbacks rely on the limited time available for planning as these approaches plan at execution time and, therefore, must meet real-time constraints.

Ontology Plan-based dialogue systems are knowledge-based systems and the use of ontologies can provide additional expressiveness for representing the scenario as well as the possibility of decoupling planning variables from the relevant information gathered during the conversation, i.e., intents. Some of the surveyed works have exploited the integration of ontologies within their approaches with different purposes. Modeling the domain problem by relying on domain-specific ontologies, for example, was frequent among these works. In Lee et al. [36], a domain ontology is used by the planner to learn which is the goal state of a detected situation and how the components modeled in the ontology affect the environment so an appropriate plan can be generated. In Baskar and Lindgren [51], a domain ontology is exploited by the dialogue manager to provide information on the topic of the dialogue. Galescu et al. [43] also rely on ontology concepts to understand the context of the dialogue. However, in this approach, these concepts were mapped onto the domain ontology after being parsed from the user utterance with the support of a general-purpose ontology. In fact, this approach kept a greater focus on using the ontology to interpret the user input rather than on generating a plan.

Aiming to keep coherent models within a cognitive system, the work of Behnke et al. [58] relies on an ontology as a common knowledge source, i.e., the *mutual knowledge model*. This approach does not plan a dialogue, but instead, integrates the user through a dialogue on a plan generation. By relying on the mutual knowledge model, separated models are automatically generated (and also extended) for the plan and the dialogue. In this process, the ontology is first extended from the translation of an existing planning domain that is modeled as an HTN. Next, a reasoner makes inferences on concepts subsumption. The resulting inferences are then described as additional decomposition methods in a step that expands the planning domain. For the dialogue, the topmost elements of the ontology are used as entry points for a topic. Besides keeping a shared vocabulary and model for both, this model can reproduce possible domain updates and reduce costs on maintenance. This approach has been further discussed in [42, 47, 48, 52].

Another interesting integration of planning and ontology can also be found in the work of Franzoni et al. [57], which models emotional affordances in an ontology. This ontology supports the plan by describing relations among some

emotional components that make it possible to reason and to identify the agent's emotional state to answer accordingly. In addition, a relation of time with the emotions supports the scheduling of the planned actions that should respect the actual time that the emotion is identified. Finally, in Teixeira et al. [60], the authors discuss the integration of a reasoner to support the tracking of the dialogue state to be matched with the state that is expected in the plan. In more recent contributions [66, 67], the authors have integrated an ontology that models the goal-oriented paradigm. This ontology is then translated to a planning problem.

RQ3. What Challenges can be Identified in Recent Plan-Based Dialogue Systems?

There are several open problems in planning research [79] and some of them are also extended to plan-based dialogue management. Motivated by RQ3, this section provides a discussion on the main challenges that could be identified in the surveyed works.

Dialogue Modeling

Modeling the behavior of the dialogue as a planning problem is a complex task and it involves expertise in both dialogue modeling and AI planning. The automated generation of efficient policies depends on the specification provided; a poor specification will result in poor policies. Indeed, some authors [18, 19] recognize that research on plan-based dialogue is still in its early stages and some works [59, 62] highlight that plan-based dialogue requires careful analysis on domain modeling to overcome challenges on the representation of complex aspects that concern the dialogue domain in the planning problem (e.g., initiative-switch). By analyzing the surveyed works, we were able to identify that modeling a dialogue as a planning problem can be challenging from two perspectives: (i) modeling decisions, i.e., which strategies are the most appropriate to generate the expected policy; and, (ii) model acquisition, i.e., how a new planning problem can be generated for a new dialogue domain.

Modeling decisions Considering the different range of dialogue applications that can be instantiated, it is not possible to define a single strategy that would work for all of them. Therefore, decisions on how to model the information as a planning problem must be carefully analyzed according to the purpose of the dialogue application and to what is expected from the resulting agent. In fact, the complexity of the policy might increase according to the modeling strategy adopted. Nesting non-deterministic effects within a FOND planning problem, for example, can result in a longer search time [62]. Meanwhile, several levels of decomposition tasks in an HTN planning problem will give origin to

longer policies, which might or might not be desired in the given domain [35].

The *type of planning* employed is, indeed, one of the aspects that must be taken into account when modeling a plan-based dialogue. Planning types employed in the surveyed works have been discussed in [RQ1. What Types of Planning Have Been Used to Support Dialogue Management in Recent Plan-Based Dialogue Systems?](#). Nonetheless, we would like to note that the size of the model might be affected by the planning type selected for the dialogue problem. The model size that can be measured by elements such as the number of actions and variables addressed is a concern of several authors [36, 58, 59, 62] since it may affect the quality of the generated plan and impact the plan generation time, constraining real-time scenarios. HTN approaches, for instance, are able to address relatively large state spaces. POMDP approaches, on the other hand, must restrict the number of variables to remain efficient. Meanwhile, with a declarative specification of the problem, the FOND approach in [62] showed efficiency for problems with a significant number of both actions and variables. In general, it is possible to identify that significant improvement on the model size can be obtained by abstracting the problem.

The *abstraction level* implemented in the planning problem presents significant relevance. Abstraction can reduce scalability problems, as discussed in [Scability](#), and it may also influence the quality of the generated plan. This aspect has been confirmed by some of the surveyed works. An example is given by Garcia et al. [59] that compared a unified and more abstract domain with simplified (less abstract) specific domains. The authors identified that the former was able to find plans much faster than the specific ones.

With the aim of abstracting the domain and reducing the search space for complex and especially huge domains, an efficient strategy might be the adoption of HTN planning. As discussed in [Hierarchical Task Network \(HTN\)](#), HTN generates an abstract plan with abstract actions and relies on subsequent decomposition methods for generating an executable plan, i.e., a plan without abstract actions [36, 44, 52]. However, hierarchical decomposition is not suitable for every domain and, like other aspects, it must be analyzed for each case scenario.

Interesting abstraction strategies can be found in argumentative approaches like [33, 35]. In these scenarios, actions are modeled with a high level of abstraction as the content of the arguments is not relevant for ordering the actions. A great advantage of these models is their reusability in quite different argumentative scenarios.

Another possibility is to abstract the values given to the variables since, as discussed in [37, 62], anticipating all possible values for a variable may become intractable and, in several cases, knowing whether the variable has a value or

not is enough to proceed with the next action selection. Of course, such strategy does not apply to domains where these values influence the next action.

Another aspect to be considered while modeling a plan-based dialogue is the decision on how to *handle errors* and *non-expected states*. Some discussion on this aspect was given in [Dynamic Policy Through Replanning](#).

In general, most of the surveyed works presented a satisfactory level of *generality*, meaning that the modeling decisions adopted in their definition are capable of addressing different dialogue domains without requiring a complete redefinition.

Model acquisition By analyzing the surveyed works, it is possible to note that model acquisition has not received as much attention as modeling decisions have received. Although automated policy generation can result in richer policies when compared to handcrafted dialogue trees, handcoding the planning problem (or parts of it) for every new dialogue domain may still imply high costs for building a new dialogue system and, considering that not many dialogue authors have knowledge on automated planning, this is likely a reason for the low adoption of such a powerful approach.

In fact, to guarantee reliable policies, some approaches rely on human specification (handcoding) of the whole planning domain or of complex parts of the dialogue. Examples of the former can be found in [\[63\]](#), where the actions were crafted for the bartending domain; and in [\[40\]](#) that uses specific actions to be executed by an avatar/robot for multimodal interactions with infants. Morbini et al. [\[55\]](#), on the other hand, have focused on supporting efficient policy authoring. Aspects such as mixed initiative and topic selection (through a reward function) are automated. However, this approach requires the dialogue author to craft the complex parts of the dialogue (e.g., specific moves within a subtopic) and the authors still report that significant human effort was required for modeling new dialogues.

In the work of Behnke et al. [\[58\]](#), ontology reasoning is used to automatically derive further decomposition methods for hierarchical tasks. However, the approach requires as input an initial planning domain already containing a description of the domain problem. In [\[66, 67\]](#), an ontology is translated into a planning problem. The ontology models the domain-independent knowledge that is common to any goal-oriented dialogue system. Therefore, when generating a new dialogue system, the domain-specific knowledge must be populated into the assertional box of the ontology. The approach of Leet et al. [\[36\]](#), instead, relies on a domain ontology as the knowledge source for decomposition tasks. However, too few details are given on how this decomposition is accomplished.

Some works like [\[30, 39, 43, 56\]](#) specify sets of communicative acts (including their preconditions and effects) that

address some specific type of dialogue (e.g., negotiation, argumentation). The modeled actions can be applied to different domains that implement these types of dialogue by loading the variables specific to the conversation topic (e.g., arguments, resources).

Finally, a work that has paid significant attention to model acquisition is the work of Muise et al. [\[62\]](#). By using an interface for declarative dialogue design, dialogue authors are in control of the behavior of the agent, being able to edit it as desired. Although the proposed approach is capable of automatically generating complex FOND problems for dialogue, dialogue authors are not required to understand AI planning to build a new agent. With the aim of improving explainability during the model acquisition process, concepts from XAIP have been integrated into this approach in [\[61\]](#).

Scalability

Although automatically generated plans can address significantly larger state spaces if compared to handcrafted techniques, scalability restrictions for large problems are still a topic that gets some of the attention of the AI planning community [\[79\]](#). As a dialogue can take several paths to achieve its goal, such limitation might also constrain plan-based dialogue managers when dealing with nontrivial dialogues. Indeed, early approaches proposed for plan-based dialogue suffered from the limited performance of the planners available by then. Currently, the amount of information that can be encoded in a plan-based dialogue may vary according to the coding scheme of the dialogue acts and to the planning approach employed.

Limitations on offline approaches are common since these models must compute a whole policy offline. For some planning approaches, this process might become impracticable for large state spaces. Given their uncertain nature, probabilistic planning approaches, for instance, allow the modeling of only a small number of variables in the state space to ensure high-quality policies. An alternative to deal with this limitation and improve the plan performance on probabilistic spoken scenarios was presented in [\[53\]](#) and in [\[54\]](#). The method employed in these works exploited commonsense with the aim of reducing the number of possible worlds, so a POMDP solver is able to calculate accurate action policies with less uncertainty and at a reasonable time.

In some offline approaches like [\[33\]](#) and [\[30\]](#) that addressed classical planning, the size of the problem and the search space grew exponentially with the number of domain variables, resulting in a concern on *memory usage*. Instead, the offline approach in [\[37\]](#) was able to minimize the scalability issue by modeling the domain variables with a high level of abstraction. The declarative representation employed

in this approach benefited the model scalability, as reported in a scalability analysis that was conducted for a synthetic domain. In a more recent work [62], the authors exhibit a further scalability analysis that shows an exponential scale-up on the size of the generated dialogue graph with respect to the declarative specification. The FOND planner applied has proven to be very efficient in computing the solutions, taking satisfactory generation time. The same planner was used in [67], also delivering feasible results.

Alternatively, some online planning approaches, i.e., approaches that plan only the next action on runtime based on the current state, have been proposed to avoid scalability problems. These approaches are very convenient to dialogues in open-ended domains that change constantly (e.g., robotics) and that require a fast policy update. For instance, by using high-level probabilistic rules based on prior domain knowledge, the online algorithm proposed by Lison [32] aims at speeding up the action selection process. Unfortunately, when dealing with complex problems, the author reports that the proposed approach is not able to scale to real-time requirements (*runtime performance* limitation). Another online approach was proposed in [39], where an online algorithm aims at finding an action sequence for multi-agent problems with a hidden initial state. Although this algorithm shows efficiency for problems with tens of possible initial states, a scalability limitation is still reported for larger problems.

Dynamic planning through replanning ([Dynamic Policy Through Replanning](#)) is also an alternative that can minimize scalability issues for dialogue approaches. With replanning, it is possible to avoid the anticipation of everything that can go wrong during a plan execution and, consequently, keep a smaller search space. As an example, for better scalability (among other motivations), Garcia et al. [59] have opted for planning strategies that apply corrective actions instead of relying on probabilities for every action outcome. By *repairing* a plan in case of execution failures, the approach in [44] is also likely to avoid or minimize scalability problems.

To sum up, offline approaches can address a greater number of domain variables when compared to online approaches. However, they suffer from some limitations concerning policy updates and adaptation; it is harder to change a policy already built. Meanwhile, online approaches restrict the number of domain variables, aiming to address runtime processing. However, as reported in the works mentioned in this section, the plan generation time still requires improvement to address real-world problems.

In general, recent works have shown improvement with respect to scalability problems, but their applicability to domains different from the proposed ones still has to be further explored. This way, how to better exploit current planning techniques to achieve better scalability in conversational scenarios opens the interest to new research perspectives.

Table 4 Metrics used for evaluation

Metric	Paper reference
average search time	[30, 33, 58, 59, 62, 67]
success rate/probability of success	[30, 54]
number of actions in the executed plan	[59]
number of times it is necessary to replan	[59]
question-asking cost	[54]

Learning Over Time

Learning from experience is an active topic being exploited over different research fields. While for some risky domains like healthcare, to keep a stable and well-defined problem definition might be preferred, some systems, like the ones implemented in companion robots, might become more interesting to the end-user if new actions and behaviors are learned over time.

As previously discussed in [Dynamic Policy Through Replanning](#), plan-based dialogue systems are capable of implementing dynamic policies through online planning [32, 39] or replanning. However, in both cases, the updated policy is limited to a previously specified planning problem and, unless an external learning module is integrated into the system architecture, planning models do not address learning. As a consequence, learning new actions or states from experience remains an open challenge for plan-based dialogue. Possible research directions to overcome this challenge include the integration of planning with reinforcement learning techniques [80].

Evaluation

The evaluation of dialogue systems is well-known as a challenging and subjective task [81]. Similarly, evaluating plan-based dialogue managers is a challenging task as no benchmarks nor official protocols are available. Making direct comparisons over different dialogue management techniques can be infeasible as their different nature and purposes would make the comparison unfair. Indeed, just a little more than half (52%) of the surveyed works have conducted some type of evaluation.

Among the works that evaluated the proposed strategies, it was possible to identify two types of analysis: with real users or with user simulations. For the latter, interaction errors or uncertain responses were also simulated with the aim of better reproducing real scenarios [54, 59]. Some works have focused on evaluating *technical planning aspects* [30, 33, 58, 59, 61, 62, 67]. Among them, we identified the aspects listed in Table 4 as the most relevant measured ones. Some derivations of these metrics were given by the comparison of the *model size* with respect to the *plan*

generation time [30, 33, 59, 61, 62, 67], and of the *solution size* with respect to the *model size* [37, 62, 67]. Although the other works might have discussed one or more of these aspects, no evaluation was provided. This brings to attention the need of designing evaluation protocols for validating both the effectiveness and efficiency of plan-based dialogue systems.

A few works, instead, conducted *subjective* evaluations that concerned *dialogue quality* and/or *user satisfaction* [42, 44, 47, 55, 65, 66]. In these works, questionnaires have been submitted to the users after a certain period of using the system. The results were either compared to different system settings (e.g., [44]) or a different version of the same system (e.g., [55, 65]).

Finally, rather than evaluating resulting plans or dialogues, Nasihati et al. [40] have focused on evaluating the effectiveness of their approach with respect to the designated application purpose, i.e., behavior change.

Discussion

This study has revealed some interesting aspects regarding the research on plan-based dialogue systems. First of all, the relatively low number of works that were retrieved reveals that this topic is still under-investigated. Although our search was limited to works dated since 2014, a greater number of works was expected. In fact, it has already been brought to light by other authors that research on plan-based dialogue systems is still in its early stages [18, 19]. Plan-based dialogue requires further research to overcome challenges such as the representation of complex aspects that concern the dialogue domain in the planning problem. Cohen [19] also highlights the fact that current approaches have yet to be improved to handle dialogues able to keep the context of the conversation and unexpected initiative switches. Such capabilities can be implemented on plan-based approaches. However, as could be observed in most of the surveyed works, they have not yet reached such a maturity level. Furthermore, trending topics in dialogue systems such as empathy [82], emotional information or sentiment analysis [83, 84], and continual learning [85] have yet to be exploited by the community.

Regarding the application domains, we observed that plan-based approaches follow a different line when compared to traditional reinforcement learning approaches, which frequently focus on addressing problems like booking trips or restaurants [24, 80]. Among the surveyed works, most present domain-independent approaches and just a few addressed a strategy that was very specific to the proposed domain [18, 40, 52]. However, some lines of research have gained greater emphasis. They are: *robotics* [18, 32, 40, 45, 46, 53, 54], *companion* technology [34, 38, 41, 42, 44, 47, 48, 52], *cognitive* systems [51, 57, 58], and *healthcare* [51,

55, 59, 66]. A possible reason for that can be associated with the fact that such systems, especially companion, cognitive and health systems, are expected to be predictable to gain the user's trust. Explainability, as discussed in [Explainability](#), is also a factor that motivates the adoption of planning in such systems. Meanwhile, robotics have long exploited the adoption of AI planning in different contexts [17, 86].

Another important aspect that came to our attention is the lack of a common strategy to evaluate and compare the proposed approaches. As discussed in [Evaluation](#), several works lack evaluation at all and no pattern could be identified among the ones that have conducted some kind of evaluation. This way, we highlight the evaluation of plan-based dialogue systems as an open topic for future research. It is necessary to find ways to evaluate aspects such as the quality of the resulting policies, costs of building plan-based dialogue managers, and costs of re-planning a policy. In addition, some comparisons to different techniques for dialogue management, such as finite-state or data-driven approaches, would provide a more clear picture of the actual contribution of plan-based approaches. Another aspect that would reinforce the quality of plan-based dialogue systems is the user's satisfaction with resulting dialogues. As discussed in [42], dialogue systems are user-centered applications and, building systems that do achieve the goal, but in ways that might confuse the user, may reduce the user's trust in the system. This might have as a consequence the user abandoning the system, especially for domains in which the user has a choice on whether to use the system or not.

Finally, this study presents some limitations. First, the time constraint defined for the retrieved works might have left out some relevant literature and authors. However, as stated in the goals of this SLR ([Motivation and Study Goals](#)), our focus was on discussing the latest contributions only. The search query employed might also have limited this SLR to some extent. The inclusion of additional keywords, as previously discussed, would result in a greater number of works. To handle this threat and to try to not leave out some interesting contributions, we manually searched for further works on the topic. However, we are aware that some relevant papers might still be left out. As a last limitation, we highlight possible biases in the selection of the papers and possible imprecisions on the data extraction. These aspects might have emerged as a consequence of the subjectivity of the analysis carried out and due to the authors' previous knowledge on the topic.

Conclusion

This work provided a systematic literature review of the recent advances on plan-based dialogue managers. This SLR addressed research questions concerning (i) the types of planning recently exploited for dialogue management; (ii)

the planning features that present significant relevance for a dialogue manager; and, (iii) the challenges on the development of plan-based dialogue managers.

The analysis conducted in this SLR reveals that the amount of works that have addressed plan-based dialogue management in recent years is considerably small; but these works have addressed relatively complex problems, not being limited to simplistic dialogues (e.g., single-turn interactions). Moreover, by analyzing which planning aspects have been exploited in the retrieved works, it was possible to identify the contributions of AI planning on the generation of robust policies. Some examples are: automated action selection towards a goal, dynamic update of the dialogue policy, and explanatory dialogues.

On the other hand, relevant challenges like scalability limitations on complex dialogues, dialogue modeling, and the lack of a shared evaluation strategy were identified. This opens up an opportunity for further research efforts to overcome these aspects.

In this SLR we have also discussed some of the application domains within the analyzed works. Our analysis revealed that plan-based approaches tend to address goal-oriented domains, but that follow a different line with respect to traditional domains commonly exploited in data-driven approaches (e.g., hotel booking). Robotics, companion, cognitive and health systems were among the most frequent domains in the analyzed corpus. Finally, we have discussed the limitations of our approach, which include the time constraint that was set for the retrieved works (works starting from 2014), limitations in the search query employed, and possible biases in the papers selection or data extraction. These limitations leave the opportunity for an extension of this work, which could consider further aspects that are relevant to the community.

Declarations

Conflict of Interest The authors declare that they have no conflict of interest.

Ethical Approval This article does not contain any studies with human participants or animals performed by any of the authors.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Jurafsky D, Martin JH. Dialog systems and chatbots. *Speech Lang Process.* 2017;3.
- Papangelis A, Galatas G, Tsiakas K, Lioulemes A, Zikos D, Makedon F. A dialogue system for ensuring safe rehabilitation. In: *International Conference on Universal Access in Human-Computer Interaction*, Springer. 2014;349–358.
- Sonntag D, Schulz C, Reuschling C, Galarraga L. Radspeech's mobile dialogue system for radiologists. In: *Proceedings of the 2012 ACM international conference on Intelligent User Interfaces*, ACM. 2012;317–318.
- Ma Y, Nguyen KL, Xing FZ, Cambria E. A survey on empathetic dialogue systems. *Information Fusion.* 2020;64:50–70.
- Howard N, Cambria E. Intention awareness: improving upon situation awareness in human-centric environments. *Hum Centric Comput Inf Sci.* 2013;3:9.
- Li W, Shao W, Ji S, Cambria E. Bieru: Bidirectional emotional recurrent unit for conversational sentiment analysis. *Neurocomputing.* 2022;467:73–82.
- Rheu M, Shin JY, Peng W, Huh-Yoo J. Systematic review: trust-building factors and implications for conversational agent design. *Int J Hum Comput Interact.* 2021;37(1):81–96.
- LuperFoy S, Loehr D, Duff D, Miller K, Reeder F, Harper L. An architecture for dialogue management, context tracking, and pragmatic adaptation in spoken dialogue systems. In: *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics-Volume 2*, Association for Computational Linguistics. 1998;794–801.
- Stefanelli M, Barahona LMR. Health care dialogue systems: Practical and theoretical approaches to dialogue management.
- Laranjo L, Dunn AG, Tong HL, Kocballi AB, Chen J, Bashir R, Surian D, Gallego B, Magrabi F, Lau AY, et al. Conversational agents in healthcare: a systematic review. *J Am Med Inform Assoc.* 2018;25(9):1248–58.
- Rojas-Barahona LM, Giorgino T. Adaptable dialog architecture and runtime engine (adarte): a framework for rapid prototyping of health dialog systems. *Int J Med Inform.* 2009;78:S56–S68.
- Papangelis A, Gatchel R, Metsis V, Makedon F. An adaptive dialogue system for assessing post traumatic stress disorder. In: *Proceedings of the 6th International Conference on Pervasive Technologies Related to Assistive Environments*, ACM. 2013;49.
- Xu H, Peng H, Xie H, Cambria E, Zhou L, Zheng W. End-to-end latent-variable task-oriented dialogue system with exact log-likelihood optimization. *World Wide Web.* 2020;23(3):1989–2002.
- Wei Z, Liu Q, Peng B, Tou H, Chen T, Huang X, Wong KF, Dai X. Task-oriented dialogue system for automatic diagnosis. In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. 2018;201–207.
- Young S, Gašić M, Thomson B, Williams JD. Pomdp-based statistical spoken dialog systems: A review. *Proc IEEE.* 2013;101(5):1160–79.
- Price WN. Big data and black-box medical algorithms. *Sci Transl Med.* 2018;10(471):eaao5333.
- Ghallab M, Nau D, Traverso P. *Automated Planning: theory and practice*. Elsevier; 2004.
- Petrick RP, Foster ME. Using general-purpose planning for action selection in human-robot interaction. In: *2016 AAAI Fall Symposium Series*. 2016.
- Cohen PR. Back to the future for dialogue research: A position paper. *arXiv preprint arXiv:1812.01144*. 2018.
- Lee K, Lee YS, Nam Y. A model of fsm-based planner and dialogue supporting system for emergency call services. *J Supercomput.* 2018;74(9):4603–12.

21. Muise C, McIlraith SA, Belle V. Non-deterministic planning with conditional effects. In: Twenty-Fourth International Conference on Automated Planning and Scheduling. 2014.
22. Bertoli P, Cimatti A, Pistore M, Roveri M, Traverso P. Mbp: a model based planner. In: Proc. of the IJCAI'01 Workshop on Planning under Uncertainty and Incomplete Information. 2001.
23. Petrick RP, Bacchus F. A knowledge-based approach to planning with incomplete information and sensing. *AIPS*. 2002;2:212–22.
24. Young SJ. Probabilistic methods in spoken–dialogue systems. *Philos Trans R Soc London, Ser A*. 2000;358(1769):1389–1402.
25. Weld H, Huang X, Long S, Poon J, Han SC. A survey of joint intent detection and slot-filling models in natural language understanding. *CoRR abs/2101.08091*. 2021.
26. Gatt A, Krahmer E. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *J Artif Intell Res*. 2018;61:65–170.
27. Austin JL. How to do things with words. Oxford University Press; 1975.
28. Kitchenham B. Procedures for performing systematic reviews. Keele, UK, Keele University. 2004;33(2004):1–26.
29. Petticrew M, Roberts H. Systematic reviews in the social sciences: A practical guide. John Wiley & Sons; 2008.
30. Black E, Coles AJ, Hampson C. Planning for persuasion. In: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, International Foundation for Autonomous Agents and Multiagent Systems. 2017;933–942.
31. Cohen PR. Foundations of collaborative task-oriented dialogue: What's in a slot? In: Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue. 2019;198–209.
32. Lison P. Towards online planning for dialogue management with rich domain knowledge. In: *Natural Interaction with Robots, Knowbots and Smartphones*. Springer 2014;111–123.
33. Black E, Coles A, Bernardini S. Automated planning of simple persuasion dialogues. In: *International Workshop on Computational Logic and Multi-Agent Systems*, Springer 2014;87–104.
34. Honold F, Bercher P, Richter F, Nothdurft F, Geier T, Barth R, Hörnle T, Schüssel F, Reuter S, Rau M, et al. Companion-technology: towards user- and situation-adaptive functionality of technical systems. In: 2014 International Conference on Intelligent Environments, *IEEE* 2014;378–381.
35. Panisson AR, Farias G, Freitas A, Meneguzzi F, Vieira R, Bordini RH. Planning interactions for agents in argumentation-based negotiation. In: Proc. of 11th Int. Workshop on Argumentation in Multi-Agent Systems. 2014.
36. Lee K, Kim JK, Park MW, Kim L. Situation based dynamic planning for dialog support with hierarchical knowledge. In: 2016 International Conference on Platform Technology and Service (PlatCon), *IEEE* 2016;1–4.
37. Botea A, Muise C, Agarwal S, Alkan O, Bajgar O, Daly E, Kishimoto A, Lastras L, Marinescu R, Ondrej J, et al. Generating dialogue agents via automated planning. *arXiv preprint arXiv:1902.00771*. 2019.
38. Nothdurft F, Ultes S, Minker W. Finding appropriate interaction strategies for proactive dialogue systems—an open quest. In: Proceedings of the 2nd European and the 5th Nordic Symposium on Multimodal Communication, August 6-8, 2014, Tartu, Estonia. Number 110, Linköping University Electronic Press. 2015;73–80.
39. Kominis F, Geffner H. Multiagent online planning with nested beliefs and dialogue. In: Twenty-Seventh International Conference on Automated Planning and Scheduling. 2017.
40. Nasihati Gilani S, Traum D, Merla A, Hee E, Walker Z, Manini B, Gallagher G, Petitto L.A. Multimodal dialogue management for multiparty interaction with infants. In: Proceedings of the 2018 on International Conference on Multimodal Interaction, *ACM* 2018;5–13.
41. Biundo S, Wendemuth A. Companion-technology for cognitive technical systems. *KI-Künstliche Intelligenz*. 2016;30(1):71–5.
42. Nothdurft F, Behnke G, Bercher P, Biundo S, Minker W. The interplay of user-centered dialog systems and ai planning. In: Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue. 2015;344–353.
43. Galescu L, Teng CM, Allen J, Perera I. Cogent: A generic dialogue system shell based on a collaborative problem solving model. In: Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue. 2018;400–409.
44. Bercher P, Biundo S, Geier T, Hoernle T, Nothdurft F, Richter F, Schattenberg B. Plan, repair, execute, explain—how planning helps to assemble your home theater. In: Twenty-Fourth International Conference on Automated Planning and Scheduling. 2014.
45. Geib C, Craenen B, Petrick RP. Generating collaborative behaviour through plan recognition and planning. In: Proceedings of the ICAPS 2016 Workshop on Distributed and Multi-Agent Planning (DMAP). 2016;98–105.
46. Foster ME, Petrick RP. Separating representation, reasoning, and implementation for interaction management: Lessons from automated planning. In: *Dialogues with Social Robots*. Springer 2017;93–107.
47. Nothdurft F, Bercher P, Behnke G, Minker W. User involvement in collaborative decision-making dialog systems. In: *Dialogues with Social Robots*. Springer 2017;129–141.
48. Behnke G, Nielsen F, Schiller M, Bercher P, Kraus M, Minker W, Glimm B, Biundo S. Sloth—the interactive workout planner. In: 2017 International Conference on Companion Technology (ICCT), *IEEE* 2017;1–6.
49. Petrick R, Foster ME. Action selection for interaction management: Opportunities and lessons for automated planning. 2016.
50. Shams Z, De Vos M, Oren N, Padget J. Normative practical reasoning via argumentation and dialogue. In: Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI-16), *AAAI Press*. 2016.
51. Baskar J, Lindgren H. Cognitive architecture of an agent for human-agent dialogues. In: *International Conference on Practical Applications of Agents and Multi-Agent Systems*, Springer 2014;89–100.
52. Behnke G, Nielsen F, Schiller M, Ponomaryov D, Bercher P, Glimm B, Minker W, Biundo S. To plan for the user is to plan with the user: Integrating user interaction into the planning process. In: *Companion Technology*. Springer 2017;123–144.
53. Zhang S, Stone, P. Corpp: Commonsense reasoning and probabilistic planning, as applied to dialog with a mobile robot. In: Twenty-Ninth AAAI Conference on Artificial Intelligence. 2015.
54. Lu D, Zhang S, Stone P, Chen X. Leveraging commonsense reasoning and multimodal perception for robot spoken dialog systems. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), *IEEE* 2017;6582–6588.
55. Morbini F, DeVault D, Sagae K, Gerten J, Nazarian A, Traum D. Flores: a forward looking, reward seeking, dialogue manager. In: *Natural interaction with robots, knowbots and smartphones*. Springer 2014;313–325.
56. Marques T, Rovatsos M. Toward domain-independent dialogue planning. 2015.
57. Franzoni V, Milani A, Vallverdú J. Emotional affordances in human-machine interactive planning and negotiation. In: Proceedings of the International Conference on Web Intelligence, *ACM* 2017;924–930.
58. Behnke G, Ponomaryov D, Schiller M, Bercher P, Nothdurft F, Glimm B, Biundo S. Coherence across components in cognitive systems—one ontology to rule them all. In: Twenty-Fourth International Joint Conference on Artificial Intelligence. 2015.
59. García-Olaya A, Fuentetaja R, García-Polo J, González JC, Fernández F. Challenges on the application of automated planning for comprehensive geriatric assessment using an autonomous social robot. In: *Workshop of Physical Agents*, Springer. 2018;179–194.

60. Teixeira MS, Dragoni M, Eccher C. A planning strategy for dialogue management in healthcare. In: SWH 2019 - Second International Workshop on Semantic Web Meets Health Data Management. 2019.
61. Sreedharan S, Chakraborti T, Muise C, Khazaeni Y, Kambhampati S. d3wa+—a case study of xaip in a model acquisition task for dialogue planning. *Proceedings of the International Conference on Automated Planning and Scheduling*. 2020;30:488–97.
62. Muise C, Chakraborti T, Agarwal S, Bajgar O, Chaudhary A, Lastras-Montano LA, Ondrej J, Vodolan M, Wiecha C. Planning for goal-oriented dialogue systems. *arXiv preprint arXiv:1910.08137*. 2019.
63. Petrick RP, Foster ME. Knowledge engineering and planning for social human–robot interaction: A case study. In: *Knowledge Engineering Tools and Techniques for AI Planning*. Springer. 2020;261–277.
64. Pardo P, Godo L. A temporal argumentation approach to cooperative planning using dialogues. *J Log Comput*. 2018;28(3):551–80.
65. Behnke G, Bercher P, Kraus M, Schiller M, Mickleit K, Häge T, Dorna M, Dambier M, Manstetten D, Minker W, et al. New developments for robot-assisting novice users even better in diy projects. *Proceedings of the International Conference on Automated Planning and Scheduling*. 2020;30:343–7.
66. Teixeira MS, Maran V, Dragoni M. The interplay of a conversational ontology and ai planning for health dialogue management. In: *Proceedings of the 36th Annual ACM Symposium on Applied Computing*. 2021;611–619.
67. Teixeira MS, Maran V, Dragoni M. Towards semantic-awareness for information management and planning in health dialogues. In *IEEE 34th International Symposium on Computer-Based Medical Systems (CBMS)*. IEEE. 2021:372–7.
68. Fisher MD, Gabbay DM, Vila L. *Handbook of temporal reasoning in artificial intelligence*, vol. 1. Elsevier; 2005.
69. Cassandra AR. *Exact and approximate algorithms for partially observable Markov decision processes*. Brown University. 1998.
70. Erol K, Hendler J, Nau DS. Htn planning: Complexity and expressivity. *AAAI*. 1994;94:1123–8.
71. Bercher P, Höller D, Behnke G, Biundo S. More than a name? on implications of preconditions and effects of compound htn planning tasks. In: *Proceedings of the Twenty-second European Conference on Artificial Intelligence*, IOS Press. 2016;225–233.
72. Behnke G, Höller D, Bercher P, Biundo S. Change the plan—how hard can that be? In: *Twenty-Sixth International Conference on Automated Planning and Scheduling*. 2016.
73. Byrne R. Planning meals: Problem-solving on a real data-base. *Cognition*. 1977;5(4):287–332.
74. Williams J, Raux A, Henderson M. The dialog state tracking challenge series: A review. *Dialogue & Discourse*. 2016;7(3):4–33.
75. Cashmore M, Collins A, Krarup B, Krivic S, Magazzeni D, Smith D. Towards explainable ai planning as a service. In: *ICAPS Workshop on Explainable AI Planning (XAIP)*. 2019.
76. Allen J, André E, Cohen PR, Hakkani-Tür D, Kaplan R, Lemon O, Traum D. Challenge discussion: advancing multimodal dialogue. *The Handbook of Multimodal-Multisensor Interfaces: Language Processing, Software, Commercialization, and Emerging Directions—Volume*. 2019;3:191–217.
77. Sreedharan S, Srivastava S, Smith D, Kambhampati S. Why can't you do that hal? explaining unsolvability of planning tasks. In: *Proc. IJCAI*. 2019.
78. Krarup B, Krivic S, Magazzeni D, Long D, Cashmore M, Smith DE. Contrastive explanations of plans through model restrictions. *J Artif Int Res*. 2021;72:533–612.
79. Geffner H, Bonet B. A concise introduction to models and methods for automated planning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*. 2013;8(1):1–141.
80. Peng B, Li X, Gao J, Liu J, Wong KF, Su SY. Deep dyna-q: Integrating planning for task-completion dialogue policy learning. *arXiv preprint arXiv:1801.06176*. 2018.
81. Deriu J, Rodrigo A, Otegi A, Echegoyen G, Rosset S, Agirre E, Cieliebak M. Survey on evaluation methods for dialogue systems. *Artif Intell Rev*. 2020;1–56.
82. Lin Z, Xu P, Winata GI, Siddique FB, Liu Z, Shin J, Fung P. Caire: An end-to-end empathetic chatbot. *Proceedings of the AAAI Conference on Artificial Intelligence*. 2020;34:13622–3.
83. Takatsu H, Ando R, Matsuyama Y, Kobayashi T. Sentiment analysis for emotional speech synthesis in a news dialogue system. In: *Proceedings of the 28th International Conference on Computational Linguistics*. 2020;5013–5025.
84. Young T, Pandelea V, Poria S, Cambria E. Dialogue systems with audio context. *Neurocomputing*. 2020;388:102–9.
85. Liu B, Mazumder S. Lifelong and continual learning dialogue systems: learning during conversation. *Proceedings of AAAI-2021*. 2021.
86. Leidner D, Borst C, Dietrich A, Beetz M, Albu-Schäffer A. Classifying compliant manipulation tasks for automated planning in robotics. In: *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE. 2015;1769–1776.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.