

Generating Dialogue Agents via Automated Planning

Adi Botea
IBM Research

Christian Muise
IBM Research

Shubham Agarwal
IBM Research

Oznur Alkan
IBM Research

Ondrej Bajgar
IBM Watson

Elizabeth Daly
IBM Research

Akihiro Kishimoto
IBM Research

Luis Lastras
IBM Research

Radu Marinescu
IBM Research

Josef Ondrej
IBM Watson

Pablo Pedemonte
SilverGate – IBM Argentina

Miroslav Vodolán
IBM Watson

Abstract

Dialogue systems have many applications such as customer support or question answering. Typically they have been limited to shallow single turn interactions. However more advanced applications such as career coaching or planning a trip require a much more complex multi-turn dialogue. Current limitations of conversational systems have made it difficult to support applications that require personalization, customization and context dependent interactions. We tackle this challenging problem by using domain-independent AI planning to automatically create dialogue plans, customized to guide a dialogue towards achieving a given goal. The input includes a library of atomic dialogue actions, an initial state of the dialogue, and a goal. Dialogue plans are plugged into a dialogue system capable to orchestrate their execution. Use cases demonstrate the viability of the approach. Our work on dialogue planning has been integrated into a product, and it is in the process of being deployed into another.

1 Introduction

Dialogue agents are becoming increasingly pervasive across many industries. There is also a recognized demand for goal-oriented agents that are capable of bridging several services to assist human users in multi-turn dialogue (Ortiz 2018). It is natural, then, to view the generation of dialogue agents through the lens of a technology well-suited for multi-step and goal-oriented settings: automated planning.

State-of-the-art dialogue systems typically can fall into two categories, dialogue trees and conversation learners. ELIZA is a first dialogue system that used dialogue trees (Weizenbaum 1966). In ELIZA and more recent tree-based dialogue systems, the conversation is structured by a complex series of branches with options depending on the user responses, and the dialogue at each branching node is hard coded. Dialogue trees can rapidly become unwieldy as the choices become more complex. Additionally, a great deal of repetition may exist where similar dialogue interactions may be required at different points in the conversation.

Conversation learners use machine learning to compute the next response from examples or historical interactions (Ilievski et al. 2018; Gao et al. 2018). They allow the quick development of a complex dialogue without the laborious task of constructing a dialogue tree. Drawbacks include the

lack of control from the application owner and the potentially unpredictable results. In certain domains, uncertainty and the risk of erroneous responses may be acceptable, but in domains such as health, finance and human resources (HR) we require guarantees and predictability of results.

A modern dialogue agent should integrate important properties such as: multi-turn, goal-oriented dialogues; calls to external services whose output is needed in the dialogue; and handling contingencies. Informally, the latter covers the fact that, at certain points in a dialogue, the conversation could steer in more than one direction. For instance, the user could accept or reject a suggestion from the agent, and the follow-up dialogue could be very different in each case.

We introduce an approach to constructing dialogue agents with AI planning. Our dialogue plans feature the properties mentioned above. These properties are not necessarily new, but our novelty stems from computing such dialogue plans *automatically*, using domain-independent planning.

This allows to create plans tailored to specific scenarios (e.g., focused on achieving a given goal). The input includes a library of actions representing individual steps in a dialogue. These are the building blocks used to construct dialogue plans. As typical in AI planning, the input further contains a problem instance, which states the initial state and the goal of the dialogue plan to compute. The system computes a dialogue plan that starts from the initial state and continues until the goal is achieved.

A library of actions can be used to generate a range of dialogue plans, and a given subset of actions can be used across several domains. For instance, the user can ask about the weather in a range of domains, such as trip planning and family event planning. The ability to automatically compute dialogue plans allows us to easily maintain deployed dialogue systems. Fixing a bug, or slightly modifying the behaviour of an action further implies the need to recreate all dialogue plans impacted by the changes. Manually fixing a collection of plans could be tedious and error prone. Avoiding this is a key advantage to our approach.

Use cases show the feasibility and the scalability of the approach. In addition, our work has been integrated into a human resources product, and it is being added to another product, for career coaching.

2 Preliminaries

We introduce basic concepts needed in a dialogue system capable to integrate our dialogue plans. In the second part of this section we overview fully observable non-deterministic planning (FOND). This is the main planning approach we use, together with elements of contingent planning.

Dialogue In a dialogue, user utterances are classified into so-called *intents*. For example, a statement such as “*I wonder how the weather is in Paris*” could be classified into an intent such as #ASKED-ABOUT-WEATHER. We assume that intents are predefined (when the system is designed), and the system is pretrained for intent classification.

Entities are variables defined in a dialogue system. Besides general-purpose variables, such as places (e.g., *Paris*), people names, dates and numbers, designers can define entities specific to the domain at hand, with a range of values that an entity can take. A given value could possibly be defined in multiple ways, as a list of synonyms. Value assignments to entities can automatically be recognized in a user utterance, using Natural Language Understanding (NLU) (Manning and Schütze 1999; Jurafsky and Martin 2000; Florian et al. 2010). User utterances are annotated with recognized intents and entity instantiations. The sample utterance presented earlier would be annotated with both the intent mentioned, and with a variable assignment such as @PLACE = PARIS. As mentioned earlier, we allow calls to external services in the middle of a dialogue. Calls to external services are annotated with their inputs and outputs.

Such annotations will allow us to build more advanced *context* information. Informally, the context contains information that is relevant at a given point in a dialogue. See details in Section 4.

Fully Observable Non-deterministic Planning A *FOND problem* is defined by a set of *fluents* that can be true or false in the environment, an *initial state* represented as a set of fluents, a set of *non-deterministic actions* and a *goal* condition that must hold at the end of the plan.

The FOND representation is similar to classical planning, except that in FOND actions can have multiple effects (see also (Geffner and Bonet 2013) for additional details).

A *non-deterministic action* consists of a *precondition* that must be satisfied for the action to be applicable and a set of *outcomes*, one of which will be used during execution. Each outcome is a set of fluents or their negation, indicating if the fluent should be added or deleted from the state of the world.

Exactly one outcome occurs when an action is executed, and thus a solution to a FOND problem must adequately handle every possible outcome. Among several equivalent ways to define a solution, we will use one resembles more closely our dialogue plans.

Consider a new action **Done** whose precondition satisfies the goal. A *solution* to a FOND problem is a directed graph $(\mathcal{N}, \mathcal{E})$ where \mathcal{N} is a set of nodes and \mathcal{E} is a set of directed edges. Each node $n \in \mathcal{N}$ is associated with an action (including **Done**) and has one outgoing directed edge for each outcome of its corresponding action. \mathcal{E} is the union of all edges across all nodes. Nodes labeled by the **Done** action have no outgoing edges. We call these leaf nodes. A

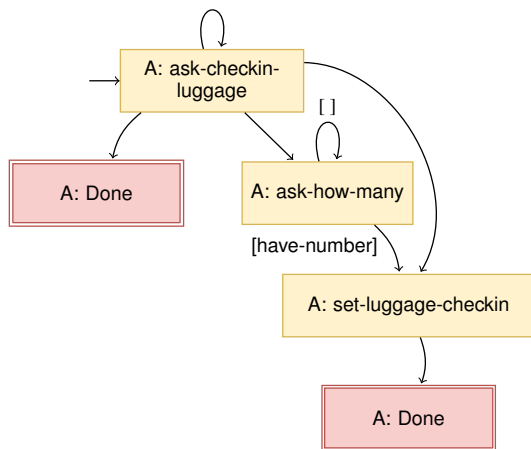


Figure 1: Toy dialogue plan. The plan has two goal states, shown as double-bordered boxes. Edge formulas shown for only one set of non-deterministic edges, to avoid clutter.

solution contains at least one leaf node and exactly one root node without any incoming edges.

The action associated with the root node must be applicable in the initial state. *Reachable states* are recursively defined by applying the action in a state reached at node n (starting with the initial state and root node, respectively). The outcome that occurs during execution dictates both the following state of the world, and the following node for the solution to proceed to. A solution must additionally have the following properties: for every reachable state and node pair (s, n) , the action corresponding to n must be applicable in s . For every reachable node n , some leaf node n' can be reached through some selection of outcomes.

In summary, a solution is a directed graph where the nodes correspond to actions the agent takes and edges correspond to how the uncertain environment responds. There must always be some path to arrive at a goal node.

3 Dialogue Plans

Definition 1. A *dialogue plan* is a structure $(\mathcal{V}, \mathcal{T}, v_0, G)$. $(\mathcal{V}, \mathcal{T})$ is a directed graph, with \mathcal{V} being the nodes and \mathcal{T} being the edges. One unique node, labeled v_0 , represents the *initial node* of the dialogue. $G \neq \emptyset$ is the set of nodes with no outgoing edges, called the goal nodes. Each node has an *action label* represented as a string. When a node has multiple outgoing edges, each such an edge has a Boolean formula where atoms are fluents from a set \mathcal{F} .

We use FOND planning to generate dialogue plans. Computing dialogue plans with AI planning provides a mechanism to construct the formulas associated to edges. As the multiple edges originating from a node are non-deterministic effects of a given action, we use the outcome of each effect as the formula of the corresponding edge.

Figure 1 shows a toy example of a dialogue plan inspired from a trip planning application. In this example, the agent asks the user whether the agent should check in any luggage

in an upcoming flight. The plan captures four possible options for the user’s answer: 1) no luggage should be checked in, in which case the dialogue can progress to the goal state at the left; 2) the user gives a positive reply, and provides the number of suitcases (e.g., “Yes, 2 pieces”); 3) the user gives a positive reply, without specifying the number; and 4) the user gives an irrelevant response, in which case the agent asks again. In case 2, the dialogue can progress to calling an external service that marks the corresponding field in a flight booking form, after which the dialogue progresses to a goal state. In case 3 the agent asks about the number. If a number is given, the dialogue progresses to case 2. Otherwise, the agent would ask again for a number. For simplicity, we skip details such as interrupting the dialogue after a finite number of iterations, in case that the user keeps providing meaningless answers.

4 Architecture Overview

Here, we present an architecture that integrates dialogue plans into an overall dialogue system, starting with the key definition required for maintaining the dialogue status.

Definition 2 (Context). Given a set of variables \mathcal{W} , a *context* is a partial instantiation of \mathcal{W} . In other words, a context contains instantiations to a subset of the variables in \mathcal{W} .

As mentioned, the context contains information available in the dialogue system at a given time. Variables in \mathcal{W} can include intents, entities, and variables to store the outputs of calls to external services (i.e., variables instantiated as described in Section 2). The context can contain additional variables, with rules about how and when to instantiate them. For instance, a trip planning domain can define variables such as LOCATION-DEST and LOCATION-ORIG. During the dialogue, the context can assign the automatically recognized value of the @PLACE entity to either LOCATION-DEST or LOCATION-ORIG, depending on the state of the dialogue.

Each goal that can be considered in some dialogue plan has a corresponding intent in the dialogue system, called a *top-level intent*. User utterances classified into a top-level intent trigger the execution of a dialogue plan with the corresponding goal.

Recall that a dialogue plan obtained from a planning system is a directed graph. A unique edge originating from a node represents a deterministic transition, and multiple edges from a given node represent the branches of a non-deterministic action. Each action in the plan (represented as a string) needs to be mapped into an actual code to execute. We call the code corresponding to an action a *transformer*. As such, each dialogue trace modelled in the dialogue plan has a corresponding sequence of transformers to call. As a dialogue progresses along a given trace, the context can change after each step. For instance, analysing a user utterance can lead to new entity instantiations. Likewise, calling an external service leads to new output. A transformer can consume context information (i.e., use context instantiations as input) and produce context information (i.e., populate the context with new instantiations). When the action at hand involves calling an external service (e.g., call a career-pathway recommender system), the corresponding transformer takes

as an argument the link to the API of the external service. The transformer calls the external service with the input at hand (e.g., user profile stored as a context variable) and places the results in dedicated context variables.

Consider a node n in the plan, with multiple outgoing edges, to a set of children nodes c_1, \dots, c_l . When the dialogue continues from a node with multiple outgoing edges (branches), such as n , the execution needs to decide what branch to choose. That is, we need a mechanism to observe part of the dialogue state (context) and make a decision based on that observation. At the end of applying the action corresponding to node n , the context allows to infer the current planning state.¹ We use the current planning state, the formulas defined for branches leading to the children nodes c_1, \dots, c_l , and the previous planning state (when the execution was at node n), to infer which branch should be followed. We assume that the effects of exactly one branch are consistent with the transition from the previous planning state to the current one.

Consider the example presented in Figure 1. The node corresponding to the action ASK-HOW-MANY has two outgoing branches, corresponding to two non-deterministic effects of whether the user provides a number or not. One branch is a self loop with no effects (i.e., no number provided) and the other progresses to a different state, with the number of luggage pieces given.

A deeper discussion on monitoring the execution of a plan is beyond the scope of this paper.

5 A Planning Model for Dialogue

AI planning problems are often expressed in a domain-independent language from the PDDL family. Our PDDL models for different dialogue domains share some commonalities in terms of the high-level design strategy. We present lessons learned when designing dialogue planning models.

Choosing the Level of Abstraction in PDDL The availability of the context, separately from a dialogue plan, eliminates the need to explicitly model in the PDDL problem description all objects (i.e., possible values of variables) that could occur in a dialogue. This abstraction avoids an artificial blow up in the problem size, and in the solving effort.

Assume, for instance, that at some point in a trip planning dialogue, the destination has been set. In the PDDL modeling, it is sufficient to encode that the destination is known, with no need to explicitly name the destination. That is, we use a predicate such as HAVE-LOCATION-DEST, as opposed to HAVE-LOCATION-DEST ?LOC. The latter would have to be instantiated into many grounded fluents, one for each possible destination. In contrast, the former is grounded into exactly one instantiated fluent, with corresponding savings in the problem size and difficulty.

A PDDL model abstracts away some, but not necessarily all information included in the context. As mentioned in Section 4, part of the predicates used in the PDDL model are

¹To achieve this, each fluent from the planning problem is also defined as a context variable, with a rule about how to instantiate (evaluate) it to TRUE or FALSE.

mirrored with corresponding variables in the context, to be able to decide on what branches to continue with the execution of a dialogue plan.

Basic Fluents Following the previous discussion on using the right level of abstraction, we introduce the following types of fluents for the PDDL model:

- OK-* To indicate if a Boolean flag holds true.
- HAVE-* To indicate we have a context value.
- MAYBE-* To indicate uncertainty of a context value.
- GOAL Specially designated fluent for the goal.

For a context variable, such as LOCATION-DEST, the fluents HAVE-LOCATION-DEST and MAYBE-LOCATION-DEST make a 3-valued logic (at most one can be true). We found the latter mode to be essential for tailored dialogue that responds appropriately to uncertain data (e.g., asking “*You’ll be traveling to Berlin, right?*” instead of “*Where will you be traveling to?*” when MAYBE-LOCATION-DEST holds and we have some idea what the location should be).

The GOAL fluent captures the fact that we typically achieve the dialogue goal by means of executing a particular action (e.g., booking a trip or making a successful recommendation). We elaborate on this further in Section 6.

Basic Actions We have identified two key action types that are shared across the dialogue domains: dialogue actions and service actions. The model could optionally include other actions, such as an auxiliary action at the end of every plan, to indicate the termination of the dialogue, but the two pervasive types are what we discuss.

Dialogue actions correspond to sending messages to the end-user in a conversation. We assume that the executor of a plan has a way to map a given dialogue action (along with the current context) to an utterance that should be sent to the end user. If the dialogue action has more than one outcome, it is presumed to be a message that warrants a response from the user, and the user’s response will correspond to the various action outcomes.

In the deployed dialogue plans we have created, the execution of a dialogue action sends the message to the end-user using a common messaging protocol, and the response in situations with more than one outcome is assessed using off-the-shelf NLU technology (e.g., services for natural language disambiguation and entity extraction). The effects of an outcome for a dialogue action can encode *whether* various types of information are available. Separately from the dialogue plan, the context will store the actual values of those types of information.

Eliciting information from the user is an important feature in multi-turn, goal-oriented dialogues. Dedicated fluents encode whether a given type of information has successfully been elicited. Fluents such as HAVE-EMPLOYEE-NAME (in a HR dialogue where a manager can ask about the performance of various team members) and HAVE-LOCATION-DEST, mentioned earlier, are prime examples of this. Once again, during execution, the context is updated to reflect the actual values that have been acquired or modified.

The preconditions of a dialogue action dictate *when* such an utterance or question would be posed to the user. For example, querying a user’s destination location only makes

sense if we do not already have it. More subtly, an action such as ASK-USER-DEST would be predicated on *not* having MAYBE-LOCATION-DEST hold in the state, as we would instead prefer the action CONFIRM-USER-DEST.

Service actions refer to the actions in the model that do not directly correspond to messages that are sent to the end-user. These include system checks that have multiple outcomes associated (essentially embedding key components of logic into the process of compiling the dialogue agent) or even web API calls that may be required as part of the conversation. An example would be looking up the weather using an online RESTful service. The outcomes of a service action correspond to the possible responses we might expect and wish to handle as part of the conversation.

The specific implementation details are beyond the scope of this paper, but essential to a service action being used as part of a plan, we assume that the executor is capable of making the RESTful API calls (or similar such service actions), and resolving the outcome. Part of this resolution process is to update the context with new information as appropriate, and maintain the corresponding state of the world from the view of the planner’s abstraction.

It is worth emphasizing the role of the outcomes from the dialogue designer’s perspective. There may be countless ways that a user could respond to a question, and similarly countless error codes that a RESTful endpoint might return. However, the task of the dialogue designer is to *only specify the outcomes that contribute to changes in state and/or conversation*. This means that a large variety of possible outcomes are categorized together.

An example for the service action might be mapping all error codes of the weather service into one NO-WEATHER-SERVICE outcome. An example for a dialogue question might be all of the ways the user could respond in the affirmative. It was a prevailing design philosophy of the dialogue agent modelling that we should only consider the outcomes that are *required for conversation*, and including a catch-all outcome as needed when the response is unclear (e.g., when the NLU cannot understand the end-user response).

6 Advanced PDDL Features

Having the base encoding in hand, we now describe some of the advanced encoding features that we have identified and deployed for the dialogue agents we have created. These stem from common patterns observed in addressing the pain points of dialogue designers.

Forced Followup Generally speaking, the declarative nature of planning can offer massive savings to the process of dialogue design (and we demonstrate as such later in Section 7). That said, there are some limited forms of imperative-style specification that we found to be common in the domains we have investigated. Almost exclusively, these took the form of immediate followup functionality: examples include responding quickly with an affirmation, running a complex service action with many outcomes after a particular response was received, etc. Here, we detail the modelling strategy used for such situations.

Forced followup is a modelling feature that introduces a new set of fluents that are incorporated into the actions in a particular way. We use two new fluents: (1) FORCED-FOLLOWUP- t indicates if there is a forced followup that must occur of type t ; and (2) FORCE-REASON- r indicates what the reason is for the forced followup.

Examples we have considered for type t include DIALOGUE (immediately respond with a message), CHECK (run a system check), and ABORT (to abort the conversation and hand off to a human operator). The type of forced followup allows us to predicate some subset of the actions with the ability to handle the forced followup. If only one action can handle a particular type, then the model should ensure that it is the *only* applicable action. We achieve this by having the negation of all types that an action a cannot handle as a precondition for a . For many of the actions, there will be no type of forced followup that they can handle, which means they have a negated precondition for every type (easily specified using quantified preconditions in PDDL).

We assume that actions which handle a particular type of forced followup always remove the appropriate fluent as part of their effects (i.e., the FORCED-FOLLOWUP- t is deleted in every outcome of the action). This ensures that the remaining actions in the domain are subsequently re-enabled.

The FORCE-REASON- r fluents in some sense mirror the FORCED-FOLLOWUP- t fluents, as they are both added and deleted at the same time, but they additionally provide a higher fidelity to the followup mechanism. As a grounded example, one of the domains (discussed later in Section 7) uses DIALOGUE as a type for forced followup with reasons spanning a range of errors (such as BAD-WEATHER, BAD-DATES, etc), warnings (e.g., NO-WEATHER-SERVICE), and affirmations (e.g., AFFIRM-OK). The action description in the model makes use of the lifted representation for PDDL, and thus only one action is needed to handle the range of forced responses corresponding to each reason. The example action schema for a forced followup of type *dialogue* would be (HANDLE-FORCED-DIALOGUE ?R - REASON)

We found that the task of declaratively specifying a dialogue agent was greatly simplified by allowing for this single-step imperative pattern to be used directly. It essentially empowers the dialogue designer to specify the immediate followup for key outcomes on certain actions, and additionally had the benefit of simplifying the execution of the dialogue plans (as utterances from the agent to the end-user need not be placed on outcomes).

Handling Multiple Intents If desired, one can build a dialogue plan with a disjunctive goal. Such a plan would satisfy any one of a collection of goals (top-level intents) considered in the disjunctive goal. In such a case, we introduce additional auxiliary fluents and actions to address them: for each intent i , we have a fluent INTENT- i that indicates if the user has that intent, and a corresponding action ASSERT-INTENT- i is introduced with the following properties:

1. Only applicable when both INTENT- i and the necessary condition for intent i to be satisfied holds.
2. Adds the goal fluent GOAL as its only effect.

In the domains we have experienced, it is enough to assume that only a single intent needs to be confirmed.

7 Use Cases

We evaluate our generic approach to computing dialogue plans with AI planning in four domains: human resources, career coaching, trip planning, and a synthetic domain. The first two correspond to two products. The synthetic domain allows to perform a more focused scalability evaluation.

Human Resources We present a HR application with dialogues related to employee professional performance. Regular employees could have a dialogue about their own performance. In addition, a manager could have a chat about the performance of any team member, and the team as a whole.

The performance is defined along a number of criteria called capabilities, such as technical skills and knowledge.

In this domain, dialogue snippets can be partitioned into three main categories: questions with static answers related to the domain (e.g., “*How can I use this tool?*”); general-purpose chitchat (e.g., “*Hello*”); and goal-oriented, multi-turn conversations. The first two categories are simple snippets of one question and one answer. The third one is handled with the approach presented in this paper.

Goals considered in multi-turn dialogues include: (1) presenting the performance rating of a given individual for a given capability; (2) giving an explanation about the value of a performance rating for an individual and a capability; (3) recommending learning resources to improve the performance of a given individual in a given capability; (4) presenting the performance rating of a team for a given capability; (5) giving an explanation about the value of a performance rating for a team and a capability; and (6) identifying the strongest and the weakest performer in a team, for a given capability.

Such goals often require a multi-turn conversation, for person disambiguation and capability disambiguation. For instance, a person name given in the user utterance is used to identify a person record in the team member database. A first name provided in a user utterance might correspond to zero, one or several person records in the database. In the first and the third case, the dialogue needs to continue with disambiguating the person. Likewise, when no capability is specified, or the capability is ambiguous, the dialogue needs to disambiguate it.

An early implementation of the goal-oriented conversations used one Java program for goals 1 and 4; one Java program for goals 2 and 5; one Java program for goal 3; and one Java program for goal 6 (using a total of 3,702 lines of code). These Java programs are essentially implementing hard-coded dialogue plans. This is hard to maintain and to extend to new goals, and the portability to a different domain is very limited.

We have replaced such multiple individual Java apps, based on hardcoded specific dialogue plans, with one single application (using only 1,166 lines of code). The application analyzes the goal at hand, identified from the intent of the user utterance, and constructs a planning instance accordingly. The instance can be fed into an off-the-shelf planning




Figure 2: Dialogue plan for career goals and pathways.

system to obtain a dialogue plan on demand. Alternatively, dialogue plans can be precomputed and stored into a library, indexed on the goals they address. The dialogue plan is executed and monitored in the system.

Career Coaching We consider two dialogue goals in career coaching: reaching a point where the user has eventually chosen a long-term career goal; and eventually choosing a career pathway towards that goal. Thus, the system implements calls to two APIs: one for a career goal recommender, and one for a pathway recommender. Each API implements calls that provide: a list of recommendations; an explanation associated with a given item (goal or pathway) included in a recommendation; additional details about a given item included in a recommendation.

For long-term career goals, a service call to a recommender system provides a number of recommendations. These are ordered, and the top three are presented to the user. The user can choose a career goal, or request additional information (e.g., an explanation of why a given career goal is included), or reject the career goals currently provided. In the first case, the goal is achieved. In the second case, additional information is provided and the dialogue continues recursively, from the state where the user is required to choose between the three options again. In the third case, additional information is elicited from the user, regarding the reason of not liking any career goal. Based on the newly elicited information, the short list of three recommendations is re-computed, and the dialogue continues recursively.

Another dialogue plan focuses on helping the end user choose a career pathway. Career pathways are computed with a call to an external service. Pathways are presented to the user, which can accept or reject the career pathway at hand. In the former case, the goal is reached and the dialogue concludes. In the latter case, information about the reason of rejecting the pathway is elicited from the user. Potential reasons can be related to a job role along the pathway, or constraints associated with roles (e.g., the user might dislike management roles). The context is updated, a career pathway is recomputed, and the dialogue continues recursively.

Computing a career pathways towards a long-term career goal requires the career goal as an input. As such, the two dialogues could be chained in a sequence. They can also




Figure 3: Generated dialogue for the Trip Planning use case.

be generated as independent dialogue plans. Such variations can easily be obtained with very small modifications in the PDDL problem instance definition. For instance, if the goal is to choose a career pathway, and a career goal is already available in the initial state, there is no need to have the dialogue focused on choosing a career goal. Otherwise, the two plans will automatically be chained in one larger dialogue plan. Figure 2 illustrates this combined dialogue plan.

Trip Planning The goal of our trip planning system is to provide booking services while considering weather situation relevant for the trip. The system needs to collect departure and arrival destinations and date range of the trip. We provide in the supplementary material the actual PDDL domain and problem instance used to drive the dialogue. For the previous two use cases, the PDDL specification is closed.

The system is requesting the trip parameters through the natural dialogue and verifies their correctness (e.g. validity of provided locations). After enough information is collected, the system uses an external call to check the weather situation for the specified destination and dates. If the weather is evaluated as inferior, the system informs the user and suggests changing the trip parameters.

From the automated planning perspective, the interesting aspect of this use case is handling the uncertainty of collected system parameters. The uncertainty comes from two sources: (1) the location recognition can be ambiguous due to NLU errors; and (2) the system can hypothesize about arrival and destination locations based on historical data and the actual user location. This leads to the planner to confirming information with the user that it has some certainty about, and soliciting information from scratch when it does not have a sense as to what the true value is.

Figure 3 shows a high-level view of the generated dialogue plan that comes from a model with only nine actions. The node symbols indicate the type of action that corresponds to that part of the plan: either dialogue, API call, or system action (the latter two being specific examples of non-dialogue actions discussed in Section 5). Even in this limited setting, we can observe how complex behaviour can be captured in the generated dialogue agent from a simple declarative specification.

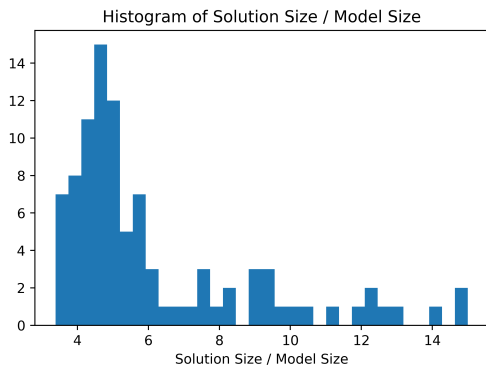


Figure 4: Plot of ratio of Solution size and Model size

Scalability Analysis Scalability is a major advantage of using a declarative representation for goal-oriented dialogue agents. To demonstrate this empirically, we created synthetic domains and problems mirroring the properties we observed in the existing dialogue encodings of our three use cases above. We measure the model size of the generated problems and solution size of the computed dialogue agents as the number of unique actions used in the solution and the total number of actions in the dialogue plan respectively.

We setup our experiments by populating a domain with random non-deterministic actions and a problem with random initial and goal states. Action precondition and effects are generated through random sampling from a set of fluents. We mirror the characteristics of actions inherent to a real dialogue system by: (1) keeping size of action preconditions within the range of 1-5 inclusive; (2) randomly sampling the effect type as either *select* (exactly one of 2-5 fluents will become true) or *assign* (1-4 fluents are randomly flipped); (3) randomly sampling 1-5 fluents for the initial state; and (4) randomly sampling 1-2 fluents for the goal state. As parameters to the random problem generator, we provide the number of actions and fluents.

The two action types correspond to typical dialogue actions that we have encountered. The *select* action mirrors the determination of one type of response the user could provide (typically with a range of 2 to 5 possibilities). The *assign* action mirrors dialogue actions where many aspects of the context can be assessed simultaneously (and thus multiple non-deterministic aspects are considered simultaneously). We have additionally confirmed qualitatively that the generated plans appear to contain a similar structure as the known dialogue plans (e.g., with the same expected plan length).

In total we generated 100 instances. Figure 4 shows a histogram of the ratio of solution size divided by model size. In most of the instances, solution size is at least 4 times the size of the model, and in extreme cases it can grow to 16 times the model size. This confirms our assertion that complex dialogue systems can be efficiently designed with very compact declarative representations.

8 Related Work

In the past, Kuijpers and Dockx (1998) presented an approach where a pre-existing library of plans can be used by

an agent in a dialogue. Steedman and Petrick (2007) advocated the use of AI planning to facilitate mixed-initiative collaborative discourse. More recently, (Nothdurft et al. 2015) described a system capable of explaining the decisions of the planning system to the user. Ortiz (2018) presented a holistic approach to building conversational assistants which includes a planning component to assist the interaction with human users. However, none of these previous approaches employ AI planning to generate customized dialog plans.

Garoufi and Koller (2010) showed recently how to generate whole sentences word-by-word with AI planning techniques and described an approach to generate multi-turn navigational commands to be followed by a human user in a conversational interface to achieve a specified goal.

Black, Coles, and Bernardini (2014) developed a formalism that describes persuasion dialogues as state transitions in a planning domain. In this setup, an automated planner was able to find optimal persuasion strategies. However, the execution monitoring is greatly simplified due to the abstract nature of simulated dialogues.

Petrick and Foster (2013) considered execution monitoring issues that arise from non-determinism of real-world dialogue systems. Their proposed execution monitor expects deterministic behavior of actions in the world and rebuilds the plan if inconsistencies between measurements and the expectations are observed (via various sensors). In contrast, our execution monitor accounts for non-determinism explicitly within the generated plan. Therefore, our plans can be static and easier to debug.

Another direction of research is focused on using (deep) reinforcement learning for end-to-end dialogue generation (Dhingra et al. 2017; Peng et al. 2018). In contrast to our model-based approach, training these data centric systems is costly and requires many interactions with human users.

The work that perhaps is the closest in spirit with ours is that by (Williams 2007; Thomson et al. 2007; Bui et al. 2010) who focus on using (factored) POMDPs to manage spoken dialogue systems in various domains. However, these systems, which are quite sensitive to the uncertainty inherent in the spoken utterances, require considerable training to learn good policies and thus drive the dialogue.

9 Summary

Dialogue agents capable of handling multi-turn, goal-oriented conversations are becoming increasingly important across a range of domains, including human resources, career coaching, and personal assistants. We have presented an approach to constructing dialogue plans automatically. Given a library of individual actions available, our system produces dialogue plans customized to achieving a given goal. Dialogue plans are further plugged into a dialogue system that can orchestrate their execution during a conversation with a user. We have shown that our approach is viable and scalable. Our work has been applied in one product, and is in the process of being integrated into a second one.

Future work includes building dialogue agents in additional domains. Dynamically interleaving dialogue agents, to allow the user to temporarily change the topic in the middle of a dialogue, is another important topic for future work.

References

- Black, E.; Coles, A.; and Bernardini, S. 2014. Automated planning of simple persuasion dialogues. In *International Workshop on Computational Logic and Multi-Agent Systems*, 87–104. Springer.
- Bui, T.; Zwiers, J.; Poel, M.; and Nijholt, A. 2010. Affective dialogue management using factored pomdps. *Interactive Collaborative Information Systems. Studies in Computational Intelligence* 281(1):207–236.
- Dhingra, B.; Li, L.; Li, X.; Gao, J.; Chen, Y.-N.; Ahmed, F.; and Deng, L. 2017. Towards end-to-end reinforcement learning of dialogue agents for information access. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, volume 1, 484–495.
- Florian, R.; Pitrelli, J. F.; Roukos, S.; and Zitouni, I. 2010. Improving mention detection robustness to noisy input. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing, EMNLP '10*, 335–345. Stroudsburg, PA, USA: Association for Computational Linguistics.
- Gao, J.; Wong, K.; Peng, B.; Liu, J.; and Li, X. 2018. Deep dyna-q: Integrating planning for task-completion dialogue policy learning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, 2182–2192.
- Garoufi, K., and Koller, A. 2010. Automated planning for situated natural language generation. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, 1573–1582. Association for Computational Linguistics.
- Geffner, H., and Bonet, B. 2013. *A Concise Introduction to Models and Methods for Automated Planning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.
- Ilievski, V.; Musat, C.; Hossmann, A.; and Baeriswyl, M. 2018. Goal-oriented chatbot dialog management bootstrapping with transfer learning. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, 4115–4121.
- Jurafsky, D., and Martin, J. H. 2000. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 1st edition.
- Kuijpers, B., and Dockx, K. 1998. An intelligent man-machine dialogue system based on ai planning. *Applied Intelligence* 8(3):235–245.
- Manning, C. D., and Schütze, H. 1999. *Foundations of Statistical Natural Language Processing*. Cambridge, MA, USA: MIT Press.
- Nothdurft, F.; Behnke, G.; Bercher, P.; Biundo, S.; and Minker, W. 2015. The interplay of user-centered dialog systems and AI planning. In *Proceedings of the SIGDIAL 2015 Conference, The 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, 344–353.
- Ortiz, C. L. 2018. Holistic conversational assistants. *AI Magazine* 39(1):88–90.
- Peng, B.; Li, X.; Gao, J.; Liu, J.; and Wong, K.-F. 2018. Deep dyna-q: Integrating planning for task-completion dialogue policy learning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, 2182–2192. Association for Computational Linguistics.
- Petrick, R. P., and Foster, M. E. 2013. Planning for social interaction in a robot bartender domain. In *ICAPS*.
- Steedman, M., and Petrick, R. P. A. 2007. Planning dialogue actions. In *Proceedings of the 8th Sigdial Workshop on Discourse and Dialogue (SIGDIAL)*, 265–272.
- Thomson, B.; Schatzmann, J.; Weilhammer, K.; Ye, H.; and Young, S. 2007. Training a real-world pomdp-based dialogue system. In *Proceedings of the Workshop on Bridging the Gap: Academic and Industrial Research in Dialog Technologies, NAACL-HLT-Dialog '07*, 9–16. Stroudsburg, PA, USA: Association for Computational Linguistics.
- Weizenbaum, J. 1966. Eliza – a computer program for the study of natural language communication between man and machine. *Commun. ACM* 9(1):36–45.
- Williams, J. D. 2007. Applying pomdps to dialog systems in the troubleshooting domain. In *Proceedings of the Workshop on Bridging the Gap: Academic and Industrial Research in Dialog Technologies, NAACL-HLT-Dialog '07*, 1–8. Association for Computational Linguistics.