Representing Uncertainty about Complex User Goals in Statistical Dialogue Systems

Paul A. Crook
Interaction Lab
Heriot-Watt University
Edinburgh, United Kingdom
p.a.crook@hw.ac.uk

Oliver Lemon
Interaction Lab
Heriot-Watt University
Edinburgh, United Kingdom
o.lemon@hw.ac.uk

Abstract

We point out several problems in scaling-up statistical approaches to spoken dialogue systems to enable them to deal with complex but natural user goals, such as disjunctive and negated goals and preferences. In particular, we explore restrictions imposed by current independence assumptions in POMDP dialogue models. This position paper proposes the use of Automatic Belief Compression methods to remedy these problems.

1 Introduction

One of the main problems for a spoken dialogue system is to determine the user's goal (*e.g.* plan suitable meeting times or find a good Indian restaurant nearby) under uncertainty, and thereby to compute the optimal next system dialogue action (*e.g.* offer a restaurant, ask for clarification). Recent research in statistical spoken dialogue systems (SSDS) has successfully addressed aspects of these problems through the application of Partially Observable Markov Decision Process (POMDP) approaches (Thomson and Young, 2010; Young et al., 2010). However POMDP SSDS are currently limited by an impoverished representation of user goals adopted to enable tractable learning.

Current POMDP SSDS state approximations make it impossible to represent some plausible user goals, *e.g.* someone who wants to know about nearby cheap restaurants *and* high-quality ones further away, or wants to schedule a meeting anytime this week except monday afternoon (also see Examples in Tables 1–3). This renders dialogue management sub-optimal and makes it impossible to deal adequately with the following types of user utterance: "I'm looking for French or Italian food", or "Not Italian, unless it's expensive". User utterances with negations and disjunctions of various sorts are very natural, and exploit the full power of natural language input. Moreover, work in dialogue system evaluation, *e.g.* (Walker et al., 2004; Lemon et al., 2006), shows that real user goals are generally *sets of items* with different features, rather than a single item. People like to explore possible trade offs between features of items.

A central challenge for the field of spoken dialogue systems is therefore to: develop realistic *large-scale* statistical approaches with an accurate, extended representation of user goals.

In this paper we propose that the independence assumptions that have guided POMDP SSDS design to date should be relaxed, user goal sets should be introduced and that the subsequent explosion in the size of the state space should be dealt with by employing *Automatic Belief Compression (ABC)* techniques.

2 POMDP SSDS

Partially Observable Markov Decision Processes (POMDPs) are Markov Decision Processes where the system's state is only *partially observable*, *i.e.* there is *uncertainty* as to what the true state is. The ability to account for uncertainty is crucial for spoken dialogue systems because their knowledge about the state is uncertain due to speech recognition errors and the fact that the user's goals are not directly observable. In POMDP models of spoken dialogue (Williams and Young, 2005; Thomson and Young, 2010; Young et al., 2010) the dialogue policy (what the system should say next) is based not on a single view of the current state of the conversation, but on a probability distribution over all possible states of the conversation. The optimal POMDP SSDS dialogue act thus automatically takes account of the uncertainty about the user's utterances and goals.

All work to date on POMDP SSDS has assumed *(i) that a user has a singular, fully constrained, fixed goal* and *(ii) that there is one source of un-*

Proceedings of SIGDIAL 2010: the 11th Annual Meeting of the Special Interest Group on Discourse and Dialogue, pages 209–212,
The University of Tokyo, September 24-25, 2010.

User goal: good Italian but not Thai food	User Goal: good Italian but not Thai food
POMDP SSDS (Example 1)	ABC-SSDS (Example 2)
U1: What restaurants are nearby? S1: There are 3 nearby. What type of food do you want? U2: *Not Italian, unless it's really good* S2: There are 2 Thai restaurants. ⟨*The compressed state (i.e. marginal likelihoods in Figure 1) indicates that Italian food is dis-preferred.*⟩ U3: No Thanks. Bye.	U'1: What restaurants are nearby? S'1: There are 3 nearby. What type of food do you want? U'2: *Not Italian, unless it's really good* S'2: Ok. Al Capone's is highly rated. There are also 2 Thai restaurants nearby. U'3: Great. I'll try Al Capone's.

Table 1: State-of-the-art POMDP SSDS vs ABC-SSDS: Example dialogues 1 and 2. U=User, S=System

certainty; errors in speech recognition. The traditional role of the Dialogue Manager (DM) is therefore one of reducing uncertainty until its belief is strongly focused on a single goal which strictly fits the user's stated constraints. Both the policies that DMs currently follow and the design decisions implemented to make SDSS tasks learnable use these assumptions (Thomson and Young, 2010; Young et al., 2010). The resulting systems force users into selecting a single domain object as their goal in order to progress the dialogue. The example POMDP SSDS dialogues, Examples 1 and 3 (Tables 1 and 2), show the problems inherent in this approach.

User goal: good Italian or good French food
S1: How can I help you? U1: What restaurants are nearby? S2: There are 6, what type of food do you want? U2: *Err... Italian or French* S3: Do you want Italian or French food? U3: Uh. either. S4: Do you want Italian or French food? U4: Ok , how about French then. S5: There are 2 good French restaurants nearby. What price range do you want? U5: What about Italian? S6: Did you say Italian? U6: QUIT

Table 2: POMDP SSDS: Example 3

3 Independence Assumptions

Even considering limited domains, POMDP state spaces grow very quickly. For example, consider finding a user's restaurant preference, which involves getting 4 pieces of information, *i.e.* food-type, city-area, price range, quality rating. Given 8 food types, 8 city areas, 3 price ranges and 3 quality ratings, coupled with 7 user actions and a 3^4 dialogue progress indicator[1] then the dialogue state space contains $8 \times 8 \times 3 \times 3 \times 7 \times 3^4 = 326,592$ states. A POMDP belief space is a probability distribution over all these dialogue states, *i.e.* a $326,592$ *dimensional* real valued ($\mathbb{R}$) space.

In order to render such large belief spaces tractable, the current state of the art in POMDP SSDS uses a variety of *handcrafted* compression techniques, such as making several types of independence assumption. For example, by assuming that users are only ever interested in one type of food or one location, and that their interests in food type, price range, quality, *etc.* are independent, the $326,592$ real valued state space can be reduced to a much smaller "summary space" (Williams and Young, 2005) consisting of, say, $4 \times \mathbb{R}$ values[2]. See Figure 1 for a graphical depiction of such assumptions[3].

As illustrated by Figure 1 the information lost due to the independence assumptions mean that these approaches are unable to support conversations such as that shown in Example 2 (Table 1).

4 Sets of User Goals

Getting rid of independence assumptions allows the DM to reason and ask questions about the user's requirements in a more rational way. It can, for example distinguish between the user wanting "excellent Italian" or "any Thai" versus only "excellent" restaurants – see Figure 1. However, the resulting high dimensional real valued state space can still only represent uncertainly over *singular* user goals (limited to *single* points in the feature space, *e.g.* an excellent Italian restaurant).

[1] Whether each piece of information is obtained, confirmed or unknown.

[2] By considering only the maximum marginal likelihood for each of the features.

[3] These apply after utterance U2/U'2 of Example 1.

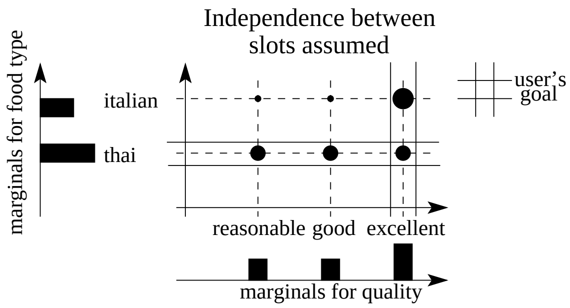

Figure 1: Assuming independence of features is equivalent to marginalising across features. Here, marginalisation incorrectly suppresses belief in Italian. Thai retains a uniform belief (which exists across all restaurant types not yet mentioned).

To achieve a substantial gain in the flexibility of SSDS we need to allow user's goals that are *sets* of points. Maintaining beliefs over "sets of goals" allows a POMDP DM to refine its belief in the user's requirements (managing speech recognition errors) without forcing the user to specify a singular tightly constrained goal. The disadvantage of this approach is a further expansion of the state space.

5 Automatic Belief Compression

To allow for expansion of the state space, whilst keeping its size tractable for policy learning, we suggest replacing handcraft approaches with *Automatic Belief Compression (ABC)* techniques.

We propose to use proven, principled statistical learning methods for automatically reducing the dimensionality of belief spaces, but which preserve the useful distributions within the full space.

Two complementary methods that we are currently investigating are **VDC** (Poupart, 2005) and **E-PCA** (Roy and Gordon, 2002; Roy et al., 2005). These methods have been applied successfully in a real-time daily living assistant with over 10^6 states (St-Aubin et al., 2000; Hoey and Poupart, 2005; Poupart et al., 2006) and to robotic navigation by (Roy and Gordon, 2002; Roy et al., 2005). They:

- reduce the dimensionality of state spaces that were previously intractable for POMDP solution methods, and
- automatically compress the representation of belief space distributions to take advantage of sparsity between likely distributions.

The tight coupling between some dialogue states and actions (*e.g.* a user's goal state `travel-from-London` and system act `confirm-from-London`) has led some researchers to conclude that compression techniques, such as state aggregation, are not useful in the dialogue domain (Williams and Young, 2007). However, such tight coupling may not exist for all states, indeed VDC has already been applied to a small spoken dialogue system problem (Poupart, 2005) where it was shown that compressions could be found *without* losing any information[4]. Further, for POMDP approaches the *state* is not the dialogue state but the belief distribution over dialogue states. Incompressibility at the dialogue state level does not rule out compressibility of belief distributions. Finally, our introduction of sets for user goals should provide additional possibilities for compression.

Our aim in applying ABC methods is to allow POMDP SSDS to handle the much larger state spaces that are required to achieve the expressiveness which we believe will be a real benefit to users. We plan to do this for real world tasks, *e.g.* a city search over 1000s of entities with an uncompressed belief space of the order of $10^8 \times \mathbb{R}$.

6 Target Dialogues

In general, when a user starts a dialogue they rarely have a singular goal in mind (Walker et al., 2004; Lemon et al., 2006). Their goal is not a fixed point in the domain but instead can be thought of as a (possibly disconnected) set of points, for example either a nearby cheap restaurant or high-quality one further away. The set represents trade offs that the particular user is interested in. People rarely communicate their goals in terms of such distributions or trade offs, preferring to provide information in a piecemeal manner and thus incrementally explore the domain.

In Examples 1–4 (Tables 1–3) we contrast the operation of a current state-of-the-art POMDP SSDS with our proposed *ABC-SSDS* system. The user's goal in Examples 3 and 4 (Tables 2 and 3) is to explore what restaurants are nearby, with a preference for French or Italian. Current POMDP SSDS approaches assume that any spread of probability mass in the belief space represents uncertainty which needs to be resolved. This generates problems for the POMDP SSDS in Example 3 since the user is forced into specifying one food type at a time, resulting in an unwieldy confirma-

[4]Compressing a test problem of 433 states to 31 basis functions, *i.e.* a summary space of 31 states.

User goal: good Italian or good French food
S'1: How can I help you?
U'1: What restaurants are nearby?
S'2: There are 6, what type of food do you want?
U2': *Err... Italian or French*
S'3: Ok, there are 2 good French restaurants nearby, and one good Italian.
U'4: OK. Which is best quality?
S'3: Mamma Mia's has the best rating.
U'5: Great. I'll go there!

Table 3: Proposed *ABC-SSDS*: Example 4

tion step (S6 of Example 3) where the user is assumed to have *changed their mind*. In contrast, the proposed *ABC-SSDS* system can believe that the user has requested information on the combined set of French and Italian restaurants.

In Examples 1 and 2 (both shown in Table 1) the user's goal is to explore restaurants nearby, including only well-rated Italians. Here the standard POMDP SSDS is forced by its "summary space" (see marginals in Figure 1) to incorrectly represent the user's goal after U2 "Not Italian, unless it's really good" by ruling out all Italian restaurants[5]. The *ABC-SSDS* user is able to find the restaurant of their choice, whereas the POMDP SSDS user's choice is artificially restricted, and they quit having failed to find a suitable item.

The *ABC-SSDS* style of dialogue is clearly more efficient than that of current POMDP SSDS. It seems likely that users of such a system may also find the style of the conversation more natural, and may be more confident that their eventual choices really meet their goals (Walker et al., 2004).

All of these hypotheses remain to be explored in our future empirical work.

7 Conclusion

We present several problems for current POMDP approaches to spoken dialogue systems, concerning the representation of complex, but natural, user goals. We propose the development of principled *automatic* methods for dimensionality reduction, in place of the ad-hoc assumptions and *hand-crafted* compressions currently used.

In parallel we are also exploring: (i) what approaches are required for *updating* beliefs over *sets* in real time – in principle a method similar to *user goal state partitioning* (Young et al., 2010) would appear to be sufficient, (ii) what exploitable bounds exist on the sets of goals that are communicable and (iii) to what extent the complexity of user goal sets can be traded off against the overall user experience.

[5]There are several ways to try to remedy this, but all have problems.

Acknowledgments

Thanks to Dr. Jesse Hoey, the SIGdial reviewers and the Engineering and Physical Sciences Research Council (EPSRC) project EP/G069840/1.

References

J. Hoey and P. Poupart. 2005. Solving POMDPs with Continuous or Large Discrete Observation Spaces. In *IJCAI*.

O. Lemon, K. Georgila, and J. Henderson. 2006. Evaluating Effectiveness and Portability of Reinforcement Learned Dialogue Strategies with real users: the TALK TownInfo Evaluation. In *IEEE/ACL Spoken Language Technology*.

P. Poupart, N. Vlassis, and J. Hoey. 2006. An Analytic Solution to Discrete Bayesian Reinforcement Learning. In *ICML*.

P. Poupart. 2005. *Exploiting Structure to Efficiently Solve Large Scale Partially Observable Markov Decision Processes*. Ph.D. thesis, Dept. Computer Science, University of Toronto.

N. Roy and G. Gordon. 2002. Exponential Family PCA for Belief Compression in POMDPs. In *NIPS*.

N. Roy, G. Gordon, and S. Thrun. 2005. Finding Approximate POMDP Solutions Through Belief Compression. *Artificial Intelligence Research*, 22(1-40).

R. St-Aubin, J. Hoey, and C. Boutilier. 2000. Approximate policy construction using decision diagrams. In *NIPS*.

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

M. Walker, S. Whittaker, A. Stent, P. Maloor, J. Moore, M. Johnston, and G. Vasireddy. 2004. User tailored generation in the match multimodal dialogue system. *Cognitive Science*, 28:811–840.

J. Williams and S. Young. 2005. Scaling Up POMDPs for Dialog Management: The "Summary POMDP" Method. In *Proc. ASRU*.

J. Williams and S. Young. 2007. Scaling POMDPs for spoken dialog management. *IEEE Transactions on Audio, Speech, and Language Processing*, 15(7):2116 –2129, Sept.

S. Young, M. Gašić, S. Keizer, F. Mairesse, B. Thomson, and K. Yu. 2010. The Hidden Information State model: a practical framework for POMDP based spoken dialogue management. *Computer Speech and Language*, 24(2):150–174.