

injury] and [*iI, severity, major*]. This reification lets the system access different aspects of complex relations and events separately, including their predicates.

The architecture also incorporates meta-level predicates, at a higher level of abstraction than the domain level, for speech acts [1, 11]. Various taxonomies of speech acts have been developed, but we have found a set of six types suffices for our current purposes:

<i>inform</i> (<i>S,L,C</i>):	speaker <i>S</i> asks <i>L</i> to believe content <i>C</i> ;
<i>acknowledge</i> (<i>S,L,C</i>):	<i>S</i> tells <i>L</i> it has received and now believes content <i>C</i> ;
<i>question</i> (<i>S,L,C</i>):	<i>S</i> asks <i>L</i> a question <i>C</i> ;
<i>propose</i> (<i>S,L,C</i>):	<i>S</i> asks <i>L</i> to adopt goal <i>C</i> ;
<i>accept</i> (<i>S,L,C</i>):	<i>S</i> tells <i>L</i> it has adopted goal <i>C</i> ;
<i>reject</i> (<i>S,L,C</i>):	<i>S</i> tells <i>L</i> it has rejected goal <i>C</i> .

In working memory, all domain-level and meta-level concepts are embedded within elements that describe mental states: *belief*(*Ag, C*) or *goal*(*Ag, C*) for some agent *Ag* and content *C*, such as *belief*(*medic, [iI, type, injury]*). An element's content may be a triple, [*i, r, x*], a belief or goal (*nested* elements), an agent's belief that some attribute has a value, *belief_wh*(*Ag, [i, r]*), a belief about whether some propositional content is true, *belief_if*(*Ag, C'*), or a meta-level predicate, such as a speech act type.

Long-term memory contains generic knowledge in the form of rules. The body of each rule describes some situation or activity, in terms of a set of triples, and associates it with the predicate in the rule head. Predicates may be defined by decomposition into more basic structures, imposing an organization similar to that of hierarchical task networks [9]. Three types of structures appear in long-term memory. *Conceptual knowledge* describes classes of situations that are relevant to an agent's beliefs or goals. This content usually focuses on the domain level and involves relations that arise in states of the world. *Skills* describe agent activities, including conditions for application and either executable actions or a sequence of lower-level skills). Finally, *goal-generating knowledge* comprises domain-level rules that state the circumstances in which an agent should establish a new goal, such as stabilizing a person when he has been injured.

The architecture complements this domain content with meta-level knowledge. Some of this takes the form of skills, but important counterexamples are conceptual rules for *speech acts*. These refer to patterns of agents' beliefs and goals, *sans* domain-level predicates: they generalize over the *content* being communicated. For example, in the rule for *inform*,

$$\begin{aligned} \text{inform}(S,L,C) \leftarrow & \text{belief}(S,C), \\ & \text{goal}(S,\text{belief}(L,C)), \\ & \text{belief}(S,\text{belief}(L,C)). \end{aligned}$$

the variable *C* denotes the speech act's content, which may bind to any set of triples. Another important form of meta-level knowledge is the *dialog grammar*, which specifies valid patterns of speech acts. This includes rules that decompose a dialog into constituents, such as a speaker *S* proposing *P* to a listener *L*, followed by *L*'s acceptance *A* to *S*, followed by a continuation of the dialog. The grammar also includes rules for ensuring 'conceptual agreement' between the content of successive speech acts, such as that questions are followed by relevant answers.

3.2 Architectural Processing

As we have noted, the architecture deals with dynamic situations, explaining and reacting to inputs in terms of available knowledge. On each cycle, it incrementally extends the explanation of the dialog to date and applies new skills as appropriate.

Dialog Interpretation. The system's input is necessarily incomplete: many of the mental states relevant to understanding are unobserved, utterances may be misheard, and observations may be missed. The architecture must make plausible assumptions to extend its explanation of observed events; for this reason, abductive inference plays a key role in interpretation. The abduction mechanism prefers explanations that make fewer assumptions while providing greater explanatory coverage.

The process first attempts to construct a proof tree using a top-level rule (e.g., that a dialog is present which satisfies the grammar) without assumptions. If this does not succeed, it increases the tolerance to one default assumption, then two, and so forth until it finds a proof or exceeds a tolerance limit. Upon finding a proof, the module adds the assumed elements to working memory, making them available in future cycles. The abduction mechanism attempts to incorporate new utterances and other observations into working memory at the start of each cognitive cycle. Default assumptions can involve the beliefs and goals of participating agents, with the dialog grammar building on speech acts and other conceptual rules lower in the proof tree. Omitted speech acts, such as implicit acknowledgements, may also serve as default assumptions, appearing as terminal nodes in the explanation.

Dialog Generation. Following the interpretation stage, the architecture matches the conditions of goal-generating rules against the current working memory elements, instantiating their arguments and adding new top-level goals to memory.¹ After this, it selects a top-level goal and finds a skill with this goal in its head and with conditions that match working memory. The generation module repeats this step recursively, generating a path down through the skill hierarchy that should achieve the top-level goal. Upon reaching a primitive skill, the architecture instantiates its variables and carries out its associated actions. On the next cycle, the system may select the same top-level goal, repeating this process, but typically the conditions of some skills along the previous path will no longer be satisfied, so the generation module follows a slightly different route. This causes the agent to carry out subskills and actions in a particular sequence. The architecture appears to be reactive, but the influence of top-level goals provides continuity. The result is hierarchical behavior in which the system traverses different branches of an AND-tree on each cycle.

The generation mechanism produces different behavior depending on whether the selected goal was produced by abductive inference or domain-level goal generation. In the former case, the system may have adopted a goal that leads it to access a meta-level skill to invoke the *inform* speech act, such as *goal(medic, belief(advisor, [inj], injury_type, bleeding))*. Obtaining this answer might in turn require domain-level actions for information gathering. In contrast, a domain-level goal might, in the process of executing domain-specific skills, lead to the agent requesting assistance.

¹ The abductive inference process may also introduce top-level goals as default assumptions.

4 Experience with the Architecture

We can now discuss our experience with the meeting support system introduced in Section 2.1. Consider again the interaction in Table 1. After the architecture processes utterance “John, do leg lifts” by the therapist (*th*), its working memory contains an explanation that includes:²

```
belief(th, propose(th, john, [[e1, exercise_type, leg_lifting], [e1, agent, john]]))
belief(john, propose(th, john, [[e1, exercise_type, leg_lifting], [e1, agent, john]]))
goal(th, [[e1, exercise_type, leg_lifting], [e1, agent, john]])
goal(th, goal(john, [[e1, exercise_type, leg_lifting], [e1, agent, john]]))
belief(john, goal(th, [[e1, exercise_type, leg_lifting], [e1, agent, john]]))
belief(john, goal(th, goal(john, [[e1, exercise_type, leg_lifting], [e1, agent, john]])))
```

In other words, after the utterance, the system believes that both the therapist and John believe a speech act occurred in which the speaker (therapist) proposes that the listener (John) do a leg-lifting exercise, that the therapist has the goal that John do leg lifting, that the therapist has the goal for John to adopt the goal of leg lifting, and that John also believes the therapist has those two goals. The system then considers possible new goals, but does not generate any. Nor does dialog generation produce anything at this point, completing the cognitive cycle.

The next utterance is John’s response, “Okay. I will do leg lifts,” in which he accepting the therapist’s proposal. The system starts a new cycle, first expanding its explanation by adding the elements:

```
belief(john, accept(john, th, [[e1, exercise_type, leg_lifting], [e1, agent, john]]))
belief(th, accept(john, th, [[e1, exercise_type, leg_lifting], [e1, agent, john]]))
goal(john, [[e1, exercise_type, leg_lifting], [e1, agent, john]])
goal(john, belief(th, goal(john, [[e1, exercise_type, leg_lifting], [e1, agent, john]])))
belief(th, goal(john, [[e1, exercise_type, leg_lifting], [e1, agent, john]]))
goal(system, [[e1, exercise_type, leg_lifting], [e1, agent, john]])
```

At this point, the system believes that both the therapist and John believe an accept speech act has occurred, that John has adopted the leg-lifting goal and wants the therapist to believe this, that the therapist believes John has adopted the goal, and, since the system is aiding two parties with a shared goal, it adopts the goal as well. This interpretation step is followed by goal generation, which produces a goal to command the TV to play a leg-lifting tutorial for the patient:

```
goal(system, system_message(tv, tutorial, nil))
```

During the generation stage, the system acts on this goal and sends the command to the TV, which plays the corresponding video tutorial.

On the next cognitive cycle, the system receives no utterance from the human users. However, it receives a signal from the motion detector, leading it to add to working memory the literal:

```
observation(motion, [[ep2, type, leg_lift], [ep2, agent, john]])
```

where *ep2* is a new constant that denotes a leg lift event whose *agent* is John. In

² For readability, we omit the top level predicate *belief(system, Content)* and only show *Content*.

response, the system expands its explanation by adding a number of new beliefs:

```
belief(system, [ep2, type, leg_lift])
belief(system, [ep2, agent, john])
belief(system, [e1, current_state, active])
belief(system, [e1, reps_done, 1])
belief(system, [e1, last_rep_time, 1382124783.0])
```

The system now thinks the leg lifting event is ongoing and the first leg lift has occurred, so it adds a time stamp for the last repetition (leg lift) of the activity, as its knowledge states that leg lifting involves ten repetitions. Goal generation then produces a goal for the system to utter encouragement to the patient:

```
goal(system, support(system, john, activity_start))
```

During the dialog generation stage, the system acts on this goal by producing the utterance “Great start, John!”

We lack the space to completely analyze the remaining interaction, but it is important to note how the system reacts to divergences from the above sequence of events. For instance, consider a case in which the therapist starts instead by proposing “John lie down”, followed by John’s counter proposal “No, I will do leg lifts.” In this case, lacking a shared goal, the system would not play the video tutorial, but instead remind John of the therapist’s goal by uttering “John, the therapist wants you to lie down.” Alternatively, consider a variation in which the interaction starts with the same two utterances by the therapist and John about doing leg lifts, the system playing the tutorial, followed by no signal from the motion sensor (say, because John is not wearing it). In this case, after failing to receiving the expected signal for some time, the system would generate a goal to utter “John, you need to strap on the motion detector.”

This interaction illustrates our dialog system’s ability to respond appropriately by drawing on its beliefs about the mental state of the users, such as whether they adopted the same goal. We have also seen that it responds differently depending on the environmental conditions, such as when it fails to receive a signal about detected motion. Together, these clarify that the system is reactive to both the external world and to internal models of the participating agents.

5 Discussion

Dialogue has received increased attention in recent years, but much of the research has focused on the level of speech. However, the literature reports a number of higher-level dialog managers, one of the most advanced being RavenClaw [2]. This system separates the domain level from some domain-independent aspects of dialogue management, such as turn taking, timing, and error handling, but it appears to encode these strategies in procedural terms. In contrast, our aim is to extract domain-independent principles and encode them as knowledge that interacts directly with domain content. Another difference is that, although both systems integrate dialog interpretation with generation, RavenClaw focuses more heavily on the latter, while our architecture devotes more of its resources to interpretation.

Our approach has more in common with some older research on dialog. Colla-gen [10] also operated over hierarchical plan structures and constructed models of agents' beliefs for use during interpretation and generation. One key difference is that the earlier system did not separate meta-level from domain knowledge. Also, despite sharing high-level assumptions, our abductive inference mechanism leads to very different operational details. We should also mention TRIPS [5], an integrated system that carried out dialogs with users to help them generate plans. TRIPS used knowledge to interpret user input and generate appropriate responses, but it was designed for the task of plan creation, while our architecture should be able to support any collaborative task given suitable domain-specific knowledge.

Our architecture and prototype systems take some promising steps towards robust, task-oriented dialog systems, but considerable work remains. We have remarked on plans to replace our text-based interfaces with ones that use spoken language, introducing uncertainty about the meaning of utterances, which would arrive as a set of alternatives with associated probabilities. We believe our abductive approach to interpretation can be adapted to handle such inputs, but whether it work as envisioned is an empirical question. We should also extend the architecture to correct faulty assumptions and misunderstandings [8], which will involve adding methods for belief revision to the interpretation module. More generally, we hope to embed our dialog knowledge and mechanisms within broader cognitive systems that interact with other agents. We believe that our architecture's representations and processes are relevant to other tasks that involve social cognition, such as providing assistance in the absence of verbal communication and in dealing with self-interested agents who utilize ignorance and deception [3] to achieve their goals.

6 Concluding Remarks

We have presented an architecture for task-oriented dialog that integrates interpretation and generation, and two implemented systems based on it. We discussed results obtained from trial runs with the meeting support system and demonstrated how the system interprets the current situation, and together with background meta and domain-level knowledge, supports the users by participating in the dialog and issuing commands to actuators. In addition to integrating the dialog interpretation and generation processes, it also provides a clear separation of meta-level knowledge from domain-level knowledge, an advantageous feature in a cognitive architecture.

Acknowledgements

This research was supported in part by Grant N00014-09-1-1029 from the Office of Naval Research and by a gift from Ericsson. We thank Chitta Baral, Paul Bello, Will Bridewell, Herb Clark, Tolga Könik, Nimish Radia, Ted Selker, David Stracuzzi, Chihiro Suga, and Richard Weyrauch for discussions that influenced the approach we have reported here.

References

- [1] Austin JL (1962) How to do things with words. Harvard University Press, Cambridge, MA
- [2] Bohus D, Rudnicky A (2009) The RavenClaw dialog management framework: Architecture and systems. *Computer Speech & Language* 23(3):332–361
- [3] Bridewell W, Isaac A (2011) Recognizing deception: A model of dynamic belief attribution. In: *Advances in Cognitive Systems: Papers from the 2011 AAAI Fall Symposium*, pp 50–57
- [4] Clark HH (1996) Using language. Cambridge University Press, Cambridge, UK
- [5] Ferguson G, Allen JF (1998) TRIPS: An integrated intelligent problem-solving assistant. In: *Proceedings of the 15th National Conference on Artificial Intelligence*, pp 567–572
- [6] Gabaldon A, Langley P, Meadows B (2013) Integrating meta-level and domain-level knowledge for interpretation and generation of task-oriented dialogue. In: Klenk M, Laird J (eds) *Second Annual Conference on Advances in Cognitive Systems*
- [7] Langley P (2012) The cognitive systems paradigm. *Advances in Cognitive Systems* 1:3–13
- [8] McRoy S, Hirst G (1995) The repair of speech act misunderstandings by abductive inference. *Computational linguistics* 21(4):435–478
- [9] Nau DS, Cao Y, Lotem A, Munoz-Avila A (2001) The SHOP planning system. *AI Magazine* 22:91–94
- [10] Rich C, Sidner CL, Lesh N (2001) Collagen: Applying collaborative discourse theory to human-computer interaction. *AI Magazine* (22):15–25
- [11] Searle J (1969) *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, New York, NY