

Cognitive Models of Failure and Recovery in Natural Language Interactions - A Joint Actions Approach

Arthi Murugesan

NRC/NRL
Postdoctoral fellow,
4555 Overlook Ave,
Washington D.C.

Arthi.Murugesan.ctr
@ nrl.navy.mil

Derek Brock

Naval Research Lab,
4555 Overlook Ave,
Washington D.C.,
20375, USA.

Derek.Brock@
nrl.navy.mil

Wende K. Frost

Naval Research Lab,
4555 Overlook Ave,
Washington D.C.,
20375, USA.

Wende.Frost@
nrl.navy.mil

Dennis Perzanowski

Naval Research Lab,
4555 Overlook Ave,
Washington D.C.,
20375, USA.

Dennis.Perzanowski@
nrl.navy.mil

Abstract

Natural language interaction, like any other joint action, is a coordination problem involving agents who work together to convey and thus coordinate their interaction goals. Joint actions frequently fail, as agents act on their best guesses of what is intended by the other person. The ability of agents to correct each other, and recover from failures, makes it possible for joint actions to succeed even in highly error prone situations. In the modeling work presented here, a sequence of interrelated modules, originally developed in the Polyscheme cognitive architecture to understand simple commands to video application, is modified to implement error discovery and accommodate possible user-initiated repairs.

1 Introduction

Natural language can be viewed as a collaborative means for expressing and understanding intentions using a body of widely shared conventions. The challenge of conveying an intention from one agent to another, for example, from a speaker to an addressee, can be characterized as a coordination problem that participants must work together to solve. People rely on a procedural convention for collaborating with each other (Clark 1996) that can be summarized as follows: 1) make the focus of the coordination problem explicit or salient; 2) pose a problem one expects the addressee will be able to solve; and 3) frame the problem in a manner that makes it easy for the addressee to solve it.

Previous modeling work by Murugesan et al. (2011) demonstrates how a sequence of interrelated cognitive models can simulate the stages of reasoning involved in understanding simple commands issued to a video monitoring system. This paper builds on the previous work and describes how agents can initiate repairs to recover from failures in each of these stages of reasoning.

2 Natural Language Interactions as Joint Actions

All agents that perform joint actions must rely on certain heuristic presumptions regarding the set of actions they expect to carry out together. In the case of conversation, this includes posing and understanding the problem, working out the intention and acting upon the expected intention. The heuristic presumptions of *salience* and *solvability* are modeled in the Polyscheme cognitive architecture developed by Cassimatis (2006). New modeling work related to initiating repairs is discussed in the following two sections.

3 Repairs in Salience

Clark's principle of joint salience suggests, roughly, that the ideal solution to a coordination problem is one that is most prominent between the agents with respect to their common ground. Thus, for example, when the model's user enters "...the red car..." it is expected that these words are intended to make objects that correspond to this phrase more prominent than other objects in the knowledge and experiences the user shares with the interactive system that is being addressed, which in our case is an interactive video monitoring application.

However, when the same user enters a word the application does not know, for e.g. "... the ted car ..." due to a typo 't' instead of 'r', the model recognizes that it is unable to identify the user's intention because the word "ted" is not in the common ground shared by the user and the application (see figure 1). The model responds by showing the user a message saying "I do not recognize the word ted."

```
<constraint>
  IsA(?word, WordUtteranceEvent, E, ?w) ^
  Orthography(?word, ?orth, E, ?w) ^
  -IsA(?orth, LexicalEntry, E, ?w)
==>
  EncounteredUnknownWord(?word, E, ?w) ^
  -InSharedLexiconWithUser(?orth, E, ?w)
</constraint>
```

Figure 1. A sample constraint from the model that identifies an unknown word.

The user now has the option of recovering by either rephrasing the utterance with words known to the system, or in the case of advanced users, adding the specific unknown word and its syntactic, semantic and common sense implications to the common ground.

4 Repairs in Solvability

The first stage in solving the coordination problem posed by a natural language utterance involves parsing it, forming its semantic interpretation and combining the semantic knowledge with relevant world knowledge in the common ground. In the second stage, the listener reasons further to identify the intention or goal behind the speaker's actions, the actions in this case being the speaker's words.

4.1 Repairs in Stage 1 – Natural Language Understanding

Sentence processing may terminate abruptly due to any of several causes for failure, the most common being an inability to form a valid parse of the sentence. On failure, the model reports the problem in parsing to the user, and initiates a repair by asking the user to enter a simpler or more grammatically correct sentence.

The process of understanding the semantics or intended meaning of a sentence within the context of domain knowledge may also result in inconsistencies. For example, a contradiction arises when "...the stalled car passed the truck..." (i.e., a car previously referentially identified in this way) is combined with simple common sense knowledge

that stalled objects do not move. The model again initiates a repair by identifying the contradiction, reporting that a stalled car cannot be motion. The user can then alter the input (e.g., "the silver car passed the truck") or, more elaborately, make changes to the domain rules associated with this input (e.g., sometimes stalled cars are towed and can thus be in motion).

4.2 Repairs in Stage 2 – Task Recognition

When one agent's intentions must be understood and acted upon by another, addressees ordinarily presume the speaker has a practical outcome or task in mind that they can recognize and help achieve. For example, when a user says, "Show me the red car passing the black car," the monitoring application's model recognizes that the user expects it to find and display a corresponding scene. But coordinating tasks specified in this way can fail in at least two ways: 1) the intended task may not be correctly recognized — when the user says "Show me the next stop of the bus", the literal meaning of the bus at a signal light is not intended (conversational implicatures) or 2) the application may not be able to perform the identified task—for example, the application currently set up to display only one scene is incapable of responding to "Show me everywhere the red turns left." The model is able to identify when it is incapable of performing the task and allows the user to revise or repair the command.

Conclusion

This paper presents various stages in which a natural language interaction can fail and introduces the notion that cognitive models can be created to accommodate error recovery initiated by an agent participating in the conversation.

References

- Herbert H. Clark. 1996. *Using Language*. Cambridge University Press, Cambridge, UK.
- Nicholas L. Cassimatis. 2006. A Cognitive Substrate for Achieving Human-Level Intelligence. *AI Magazin* 27(2): 45-56.
- Arthi Murugesan, Derek P. Brock, Wende K. Frost and Dennis Perzanowski. 2011. *Accessing Previously Shared Interaction States through Natural Language*. Springer-Verlag, HCII 2011, CCIS 173: 590-594.