Dialogue Management for Task-Learning Human-Robot Dialogue

Sarah Fillwock, Changsong Liu, Joyce Y. Chai
Language and Interaction Research, Computer Science and Engineering Department
Michigan State University

Motivation

Recent research has explored how a robot can learn to perform a task from a human teacher, instead of from explicit programming

- An effective solution has been for the human to physically guide the robot through the steps of the task
- A more powerful solution is for the robot to learn the task by observing a human demonstration.

Dialogue provides the robot with a way to direct its learning and address any insufficiencies from vision

- It needs a dialogue management strategy to select actions that enable it to reach its goal

In this work, we identified sources of information that were relevant to the robot's task-learning goal and used them to direct the selection of the best response for the robot to take

Modeling Dialogue Behavior

Every utterance the robot hears is interpreted in terms of a Dialogue Act: labels that classify the intention of an utterance.

Dialogue Act	Utterance
social greet	hi
request teaching	can you teach me how to fold
describe action how	fold the right sleeve to the middle

The robot also has a set of dialogue rules that specify the sequence in which Dialogue Acts can occur.

Dialogue Act	Valid Responding Dialogue Acts
social greet	social greet, request teaching
request teaching	yes, no
describe action how	ok, ask object where, request repeat action

Task Structure

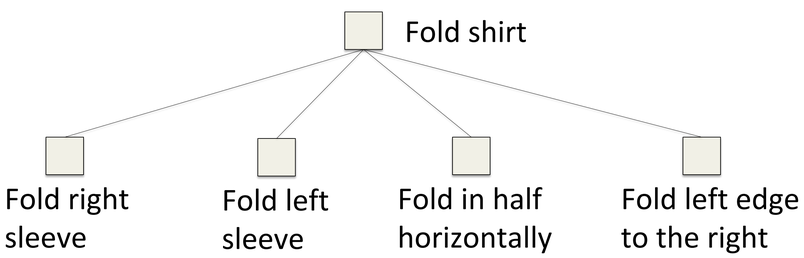

System Architecture

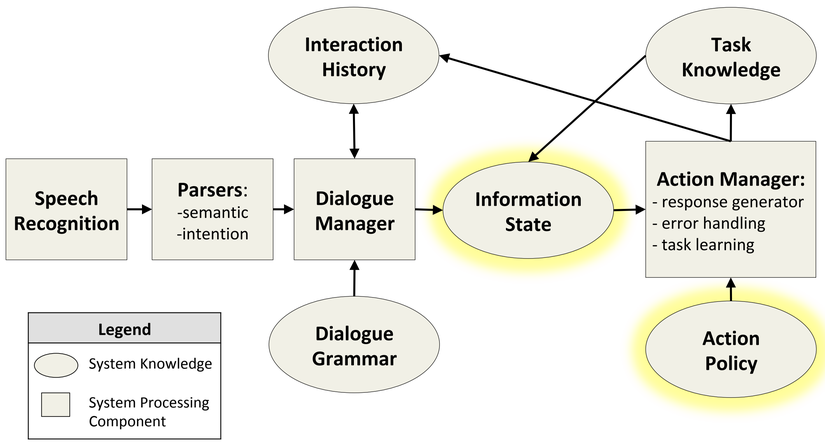

Information State

The robot needs to keep track of its understanding of the current dialogue situation and the completeness of its task knowledge.

The Information State contains 6 features for this information:

- Dialogue Act
- Valid Responding Dialogue Acts
- Speaker
- Consecutive Errors
- Dialogue parsing results
- the focus (current knowledge element - task, step, object - that has missing information)

Action Policy

The Action Policy is a set of hand-coded rules that match an Information State to the best action to take.

There are 3 categories of actions:

- Dialogue Action (D) – robot gives a spoken response
- Learning Action (L) – robot updates task structure
- Error-Handling Action (E) – robot handles a dialogue error from the human

Rule	Action
task is unknown	request teaching (D)
describe action how received AND unknown task is the current focus	add a step to the task (L)
unable to interpret human	handle non-understanding (E)

A Pilot Evaluation

25 dialogues between the robot and a human teacher

- 5 different sequences of steps for folding a shirt
- 5 dialogues per sequence

Results:

- Step Acquisition Accuracy – 97%
- Average Incorrect Information per Learned Step – 0.39
- Average Occurrences of Step Duplication – 0.08
- Average Dialogue Length – 26.04

Dialogue Sample	Robot's Selected Action(s)
R: hello	
H: hi	request-teaching (D)
R: can you teach me how to fold the shirt	
H: ok	request-first-step (D)
R: what is the first step	
H: fold in half vertically	add step (L), acknowledge-ok (D)
R: alright	
H: now fold the left sleeve to the middle	add step (L), ask-object-where (D)
R: where is the left sleeve	
H: here is the left sleeve	ground object (L), acknowledge-ok (D)
R: alright	

Future Work

Incorporate this dialogue management framework into a complete system with real-time visual processing and a hierarchical task structure

- The information state elements and action policy rules would have to be modified to capture the new task-learning dependencies
- The insufficiencies found during the evaluation would be helpful for guiding these modifications.

Furthermore, to enable more robust dialogue management, using reinforcement learning to obtain an optimal action policy will be explored.

This work is supported in part by an REU supplement to IIS-1208390 from the National Science Foundation