Learning to Interpret Natural Language Commands through Human-Robot Dialog

Jesse Thomason, Shiqi Zhang, Raymond Mooney, Peter Stone

University of Texas at Austin

15 BUENOS AIRES

Commanding Robots in Human Environments

We present a system for commanding robots in human environments that learns new words and referring expressions on-the-fly from dialog with human interlocutors. To improve domain independence, our approach requires little initial training data. Like other systems that accept language commands (Kollar, 2013), we use dialog to clarify user intentions. These clarifications are fed back to our language understanding component as training data to improve semantic parsing by lexical acquisition, as done post-hoc from existing conversations in previous work (Artzi, 2011).

Robot Dialog Agent

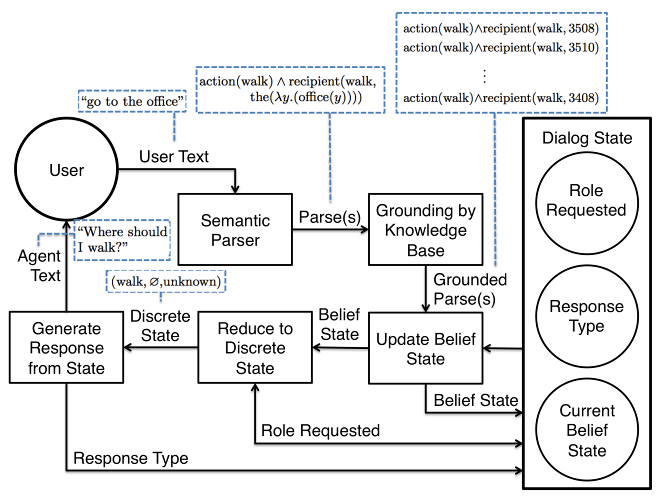

Dialog agent workflow (left) and Segbot platform (right). Dashed boxes show processing of user command "go to the office".

Large-Scale Experiment

Mechanical Turk users were given one navigation and one delivery goal in train/test (80/20%) conditions, and chatted with the robot's dialog agent until the goal was understood. Our approach was tested in 4 phases, with about 50 users each per train/test set. The system was batch trained on data from train goal conversations.

The Mechanical Turk interface for the delivery task.

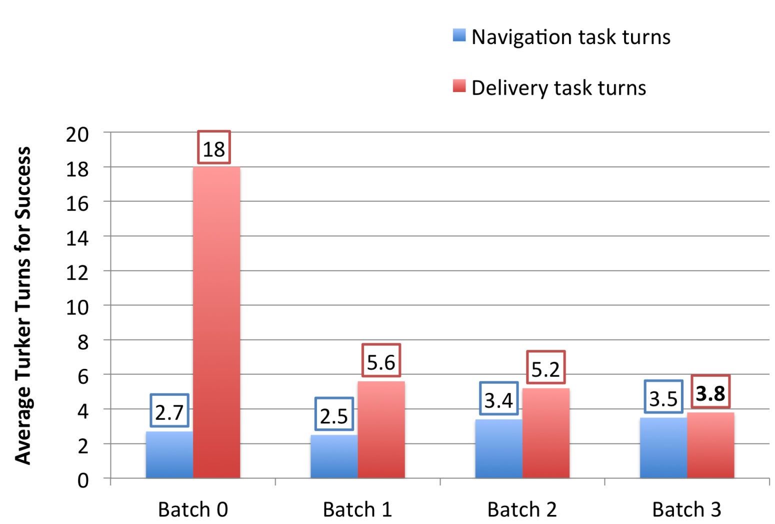

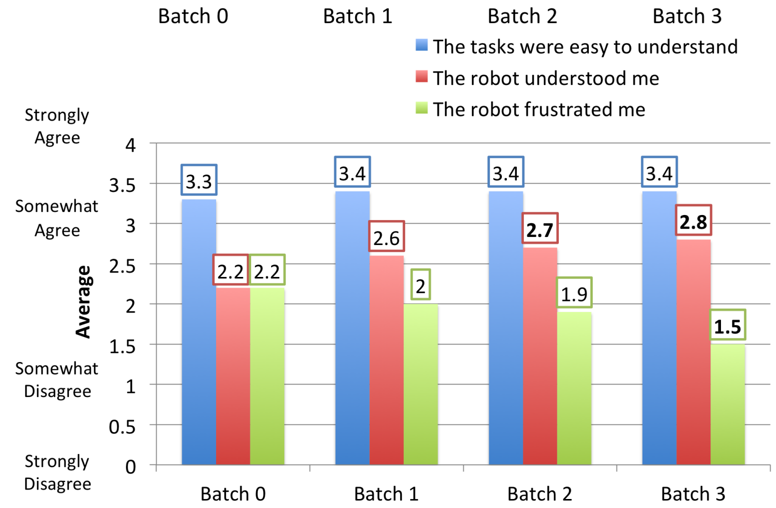

Bold averages differ statistically significantly ($p < 0.05$) from batch 0.

Discovering New Word Meanings

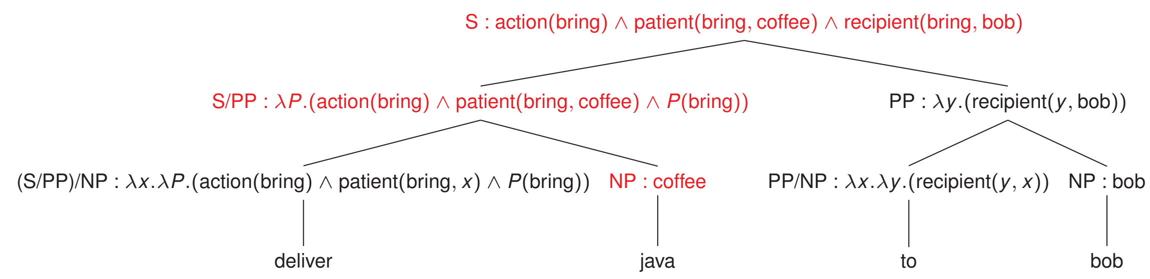

A CCG-driven λ-calculus parse of the command "deliver java to bob" fails (black) because 'java' is not known in the lexicon, but the action and recipient are recognized. After lexical acquisition (below), the parse succeeds (red).

Action	Patient	Recipient
bring	?	bob

→

Human	Deliver java to Bob
System	What should I bring to Bob?
Human	Coffee

↙

Action	Patient	Recipient
bring	coffee	bob

→ action(bring) ∧ patient(bring, coffee) ∧ recipient(bring, bob)

The system uses dialog to clarify the missed patient argument of the bring command. Given the action, patient, and recipient, it constructs the logical form the original command "Deliver java to Bob" may have had.

triggering semantics	lexical entry
action(bring)	bring :- (S/PP)/NP : $\lambda x.\lambda P.$(action(bring) ∧ patient(bring, x) ∧ P(bring))
action(bring)	bring :- (S/NP)/NP : $\lambda x.\lambda y.$(action(bring) ∧ patient(bring, y) ∧ recipient(bring, x))
coffee	coffee :- NP : coffee
bob	coffee :- NP : bob
⋮	⋮

Candidates for a new lexical entry for 'java' are based on the implied semantic form from user dialog. Only the entry "NP : coffee" leads to a parse matching that form, so a lexicon entry is added for "java :- NP : coffee".

Embodied Experiment

A real robot carried out tasks in an identical experiment with fewer users. An initial test batch of 10 users completed tasks, then the robot interacted freely with people on the floor for four days as training (34 conversations in total). It was retrained from these conversations, and then a second test batch of 10 users completed the same test tasks as the initial users.

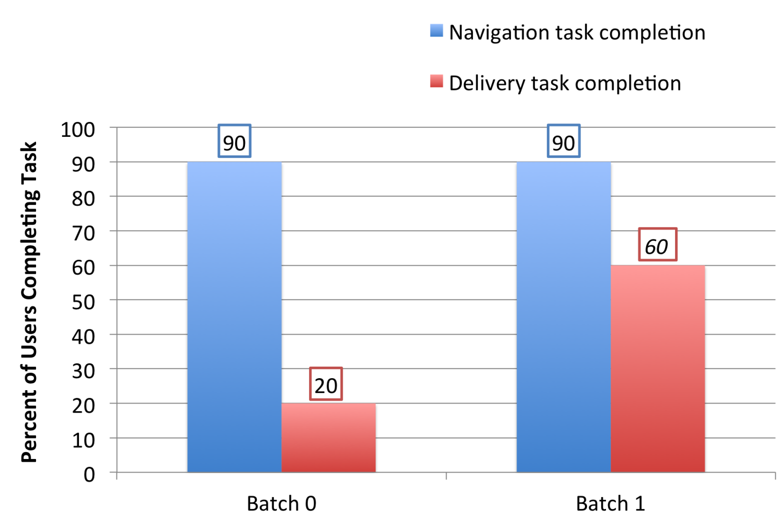

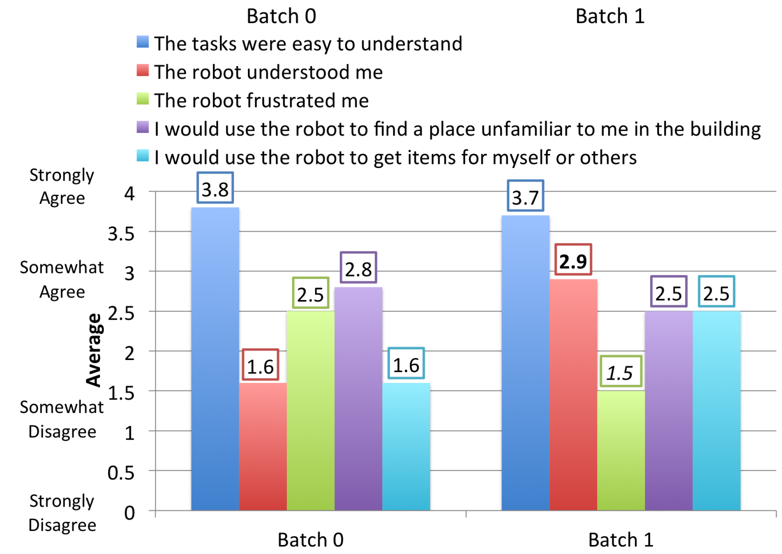

Bold averages differ statistically significantly ($p < 0.05$) from batch 0; *italics* trend different ($p < 0.1$).

Conclusions and Future Work

Lexical acquisition done by retraining a parser from user dialogs reduces dialog lengths for multi-argument predicates, causes users to perceive the system as more understanding, and leads to less user frustration. This method enables improving language understanding without large, annotated corpora. Future work includes perceptual grounding, predicate invention, and learning a multi-object dialog policy that trades off learning and user satisfaction.

jesse@cs.utexas.edu