

Learning to Interpret Natural Language Commands through Human-Robot Dialog

Jesse Thomason

Shiqi Zhang, Raymond Mooney, Peter Stone

The University of Texas at Austin

Commanding Robots

- Autonomous robots in human environments



- Simplest to interact with via natural language

Our Task

- Command a robot operating in an office environment
- Robot autonomously wanders by default
- Robot can navigate to rooms and deliver items



System Goals

- Require little initial data
 - More domain independent
- Reason using composition
 - “Alice’s office”
- Robust to lexical variation
 - “bring”, “deliver”, “take”
- Execute the **right** action
 - Perform clarifications with user

Closest Previous Work

- Service robot that accepts commands (Kollar, 2013)
- Semantics match spans of words to known actions/people/locations
- Can learn new referring expressions through dialog

Human	Go to Alice's office
Robot	Where is "Alice's office"?
Human	Room 3

- This system would explicitly match "Alice's office" to room 3

Closest Previous Work

- When system sees “Bob’s office”, will have to ask where that is
- Want to take advantage of compositionality instead
 - Reason about possessive marker “’s” and what entities “office” picks out
- Need a more powerful formalism for representing sentence semantics
 - Want to keep initial training data light

Helpful Previous Work

- Augment a semantic parser through conversation logs (Artzi, 2011)

Human	I would like to fly out of boston arriving to new york and back from new york to boston
System	Leaving boston (CONFIRM:from(fl1,BOS)) on what date? (ASK: $\lambda x.departdate(fl1,x)$)

- Key idea for us: use known system semantic meanings to guess human utterance word meanings

Tag Token Sequence

token(s) ——— *syntax : semantics*

University of Washington Semantic Parsing Framework (SPF); (Artzi, 2011)

Known possibilities for each token stored in a lexicon

Use Combinatory Categorical Grammar (CCG)-driven parsing

bring ——— (S/PP)/NP : $\lambda x. \lambda P. (action(bring) \wedge patient(bring, x) \wedge P(bring))$

coffee ——— NP : *coffee*

to ——— PP/NP : $\lambda x. \lambda y. (recipient(y, x))$

Bob ——— NP : *bob*

Construct Meaning Hierarchically

$S : action(bring) \wedge patient(bring, coffee) \wedge recipient(bring, bob)$

$S/PP : \lambda P.(action(bring) \wedge patient(bring, coffee) \wedge P(bring))$

$PP : \lambda y.(recipient(y, bob))$

$(S/PP)/NP : \lambda x.\lambda P.(action(bring) \wedge patient(bring, x) \wedge P(bring))$

$NP : coffee$

$PP/NP : \lambda x.\lambda y.(recipient(y, x))$

$NP : bob$

bring

coffee

to

Bob

Tag Token Sequence – Missing Entry

bring — (S/PP)/NP : $\lambda x.\lambda P.(action(bring) \wedge patient(bring,x) \wedge P(bring))$

java — ?

to — PP/NP : $\lambda x.\lambda y.(recipient(y,x))$

Bob — NP : *bob*

Given semantic form, can guess about missing token syntax/semantics

Human	bring java to bob
Robot	what should I bring to bob?
Human	coffee

$S : action(bring) \wedge patient(bring,coffee) \wedge recipient(bring,bob)$

Tag Token Sequence – Missing Entry

bring — (S/PP)/NP : $\lambda x.\lambda P.(action(bring) \wedge patient(bring,x) \wedge P(bring))$
java — ?
to — PP/NP : $\lambda x.\lambda y.(recipient(y,x))$
Bob — NP : *bob*

Given form:

$action(bring) \wedge patient(bring,coffee) \wedge recipient(bring,bob)$

Lexicon entries
that produce
parts of this
form:

bring :- (S/PP)/NP : $\lambda x.\lambda P.(action(bring) \wedge patient(bring,x) \wedge P(bring))$
bring :- (S/NP)/NP : $\lambda x.\lambda y.(action(bring) \wedge recipient(bring,x) \wedge patient(bring,y))$
coffee :- NP : *coffee*
Bob :- NP : *bob*

Candidates for
'java' lexical
entry:

:- (S/PP)/NP : $\lambda x.\lambda P.(action(bring) \wedge patient(bring,x) \wedge P(bring))$
:- (S/NP)/NP : $\lambda x.\lambda y.(action(bring) \wedge recipient(bring,x) \wedge patient(bring,y))$
:- NP : *coffee*
:- NP : *bob*

Tag Token Sequence – Missing Entry

bring — (S/PP)/NP : $\lambda x.\lambda P.(action(bring) \wedge patient(bring,x) \wedge P(bring))$

java — NP : *coffee*

to — PP/NP : $\lambda x.\lambda y.(recipient(y,x))$

Bob — NP : *bob*

With new lexicon entry, we can construct the correct semantic form

$S : action(bring) \wedge patient(bring,coffee) \wedge recipient(bring,bob)$

Meeting System Goals

- Require little initial data
 - Bootstrap parser with 5 expressions, 105 words
- Handle composition used by speakers
 - Use CCG-driven semantic parsing (Artzi, 2011)
- Robust to lexical variation
 - Incrementally train parser to obtain new words
- Execute the **right** action
 - Use dialog to clarify meanings with user (Kollar, 2013)

Mechanical Turk Experiment

- Users given one navigation and one delivery goal
 - Train/test goals chosen at random from possibilities
- Chat with robot's dialog agent until goal is understood

Mechanical Turk Interface





TASK TO COMPLETE

Dave Daniel wants the item in slot 5.

ROBOT	How can I help?
YOU	please bring the item in slot 5 to dave daniel
ROBOT	What should I bring to dave daniel?
YOU	a calander
ROBOT	I'm sorry, but I couldn't pinpoint what you meant by that.
ROBOT	What should I bring to dave daniel?
YOU	a day planner
ROBOT	I'm sorry, but I couldn't pinpoint what you meant by that.
ROBOT	What should I bring to dave daniel?
YOU	calendar
ROBOT	You want me to bring calendar to dave daniel?
YOU	yes
ROBOT	I thought so
ROBOT	Happy to help
YOU	<input type="text"/>

Final Task

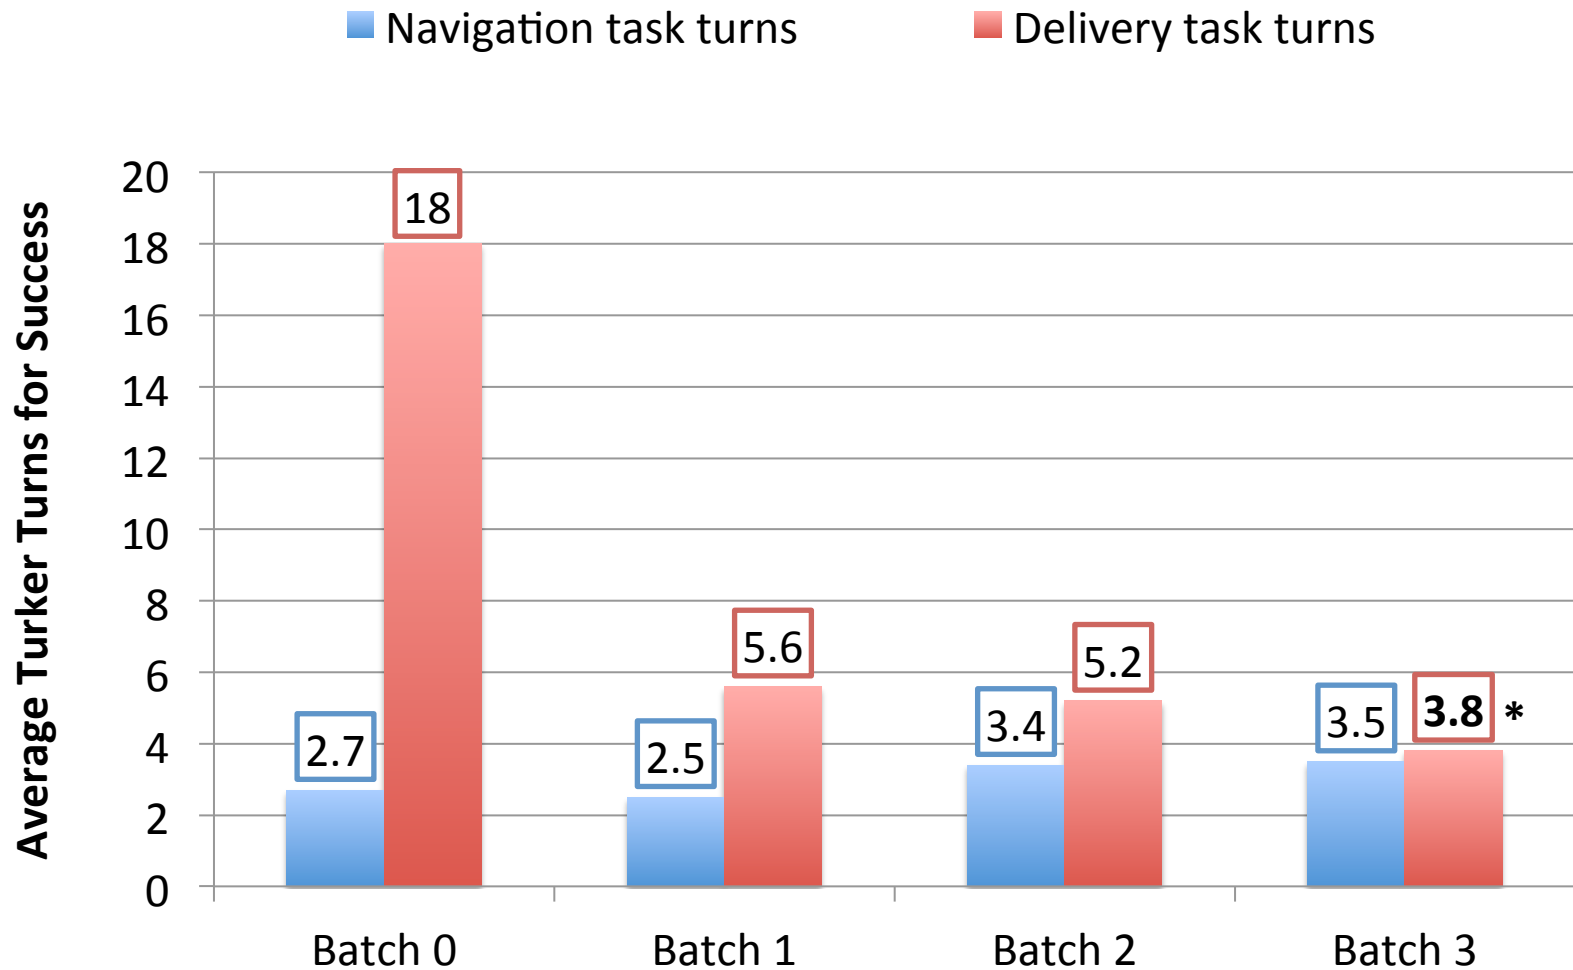
Items available to robot:

 1	 2
 3	 4
 5	

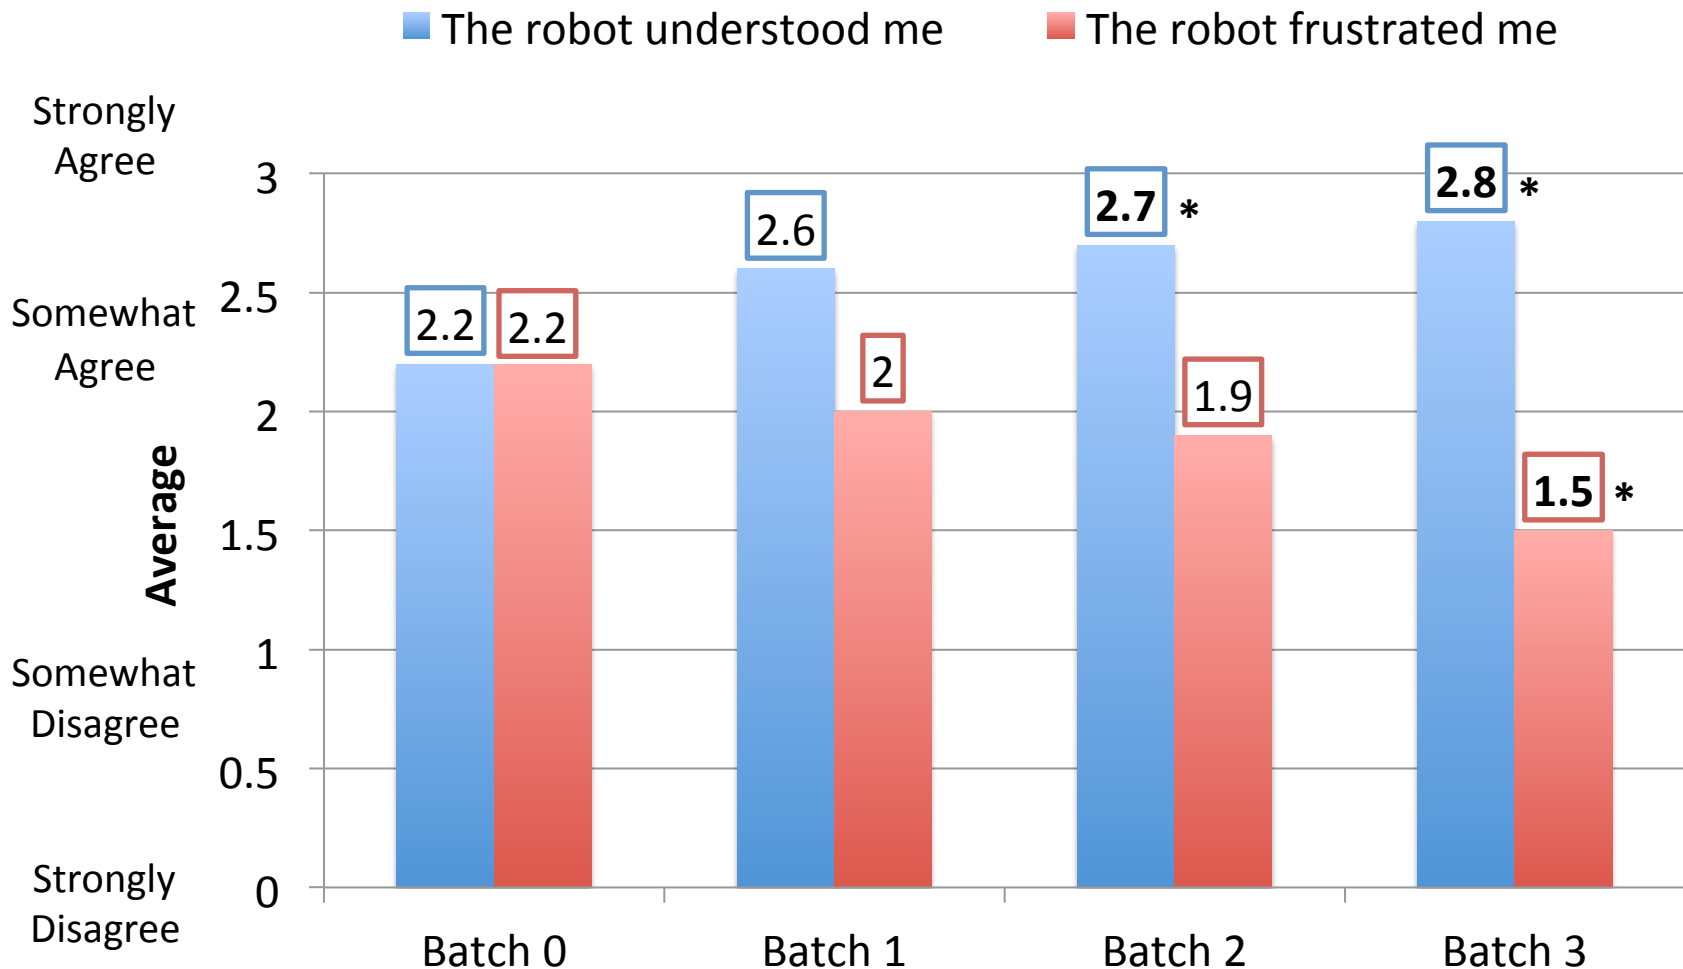
Large-Scale Experiment

- Tested in 4 phases
- ~50 users received test goals, ~50 train goals
 - Unique users in each phase
- System incrementally trained via train goal conversations only

Mechanical Turk Dialog Turns



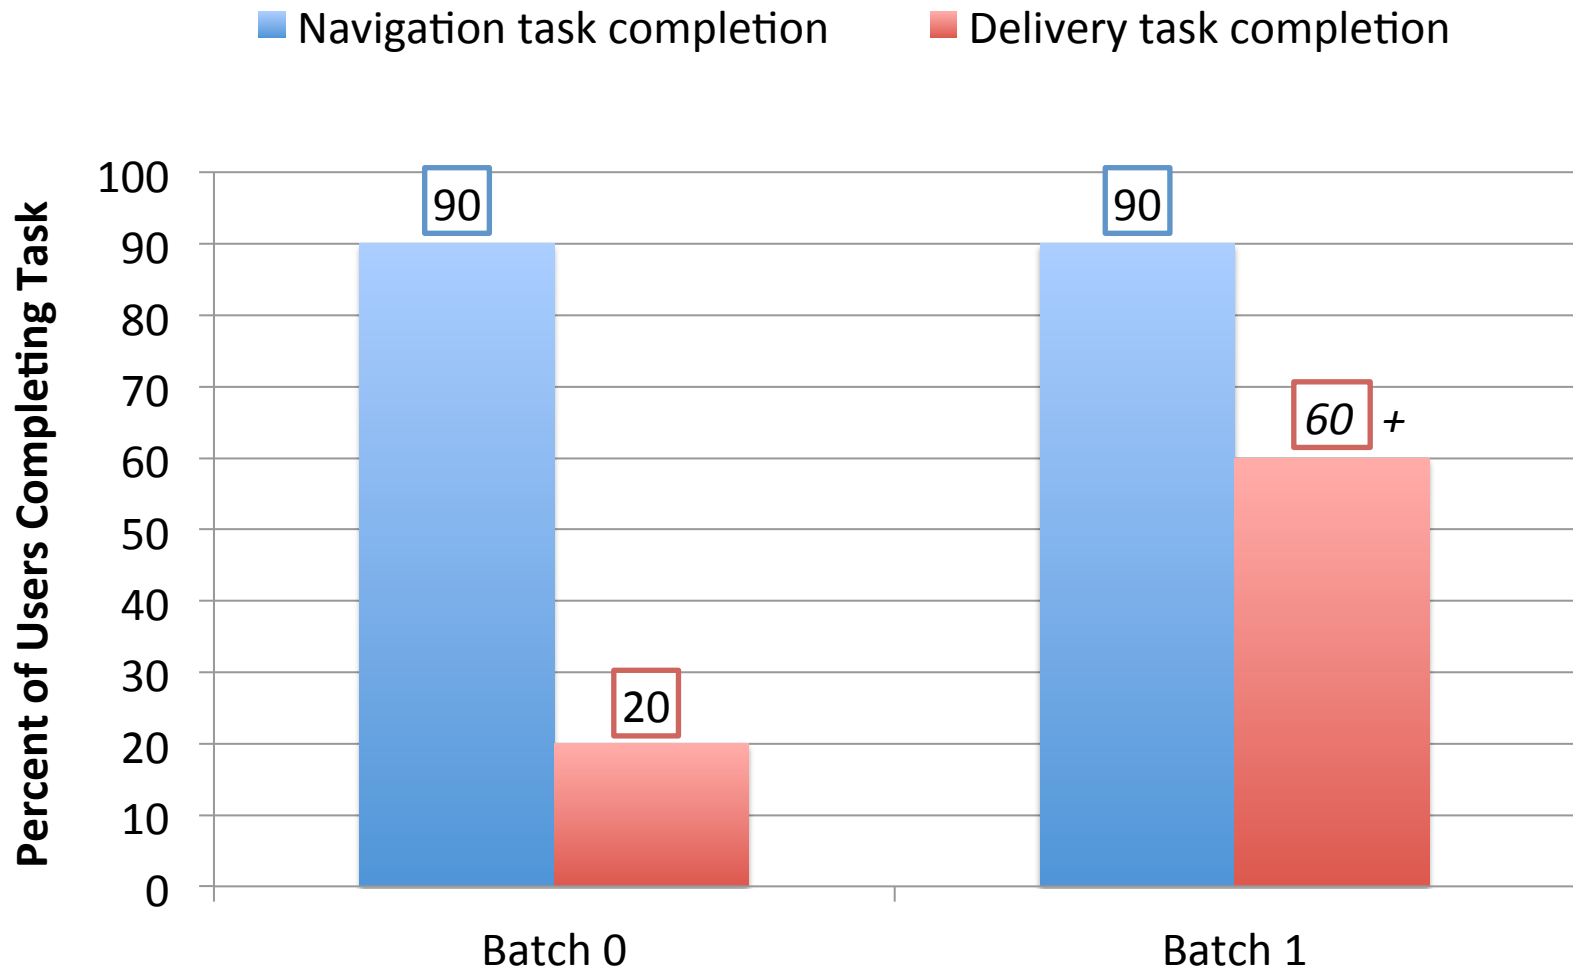
Mechanical Turk Survey Responses



Robot Experiment

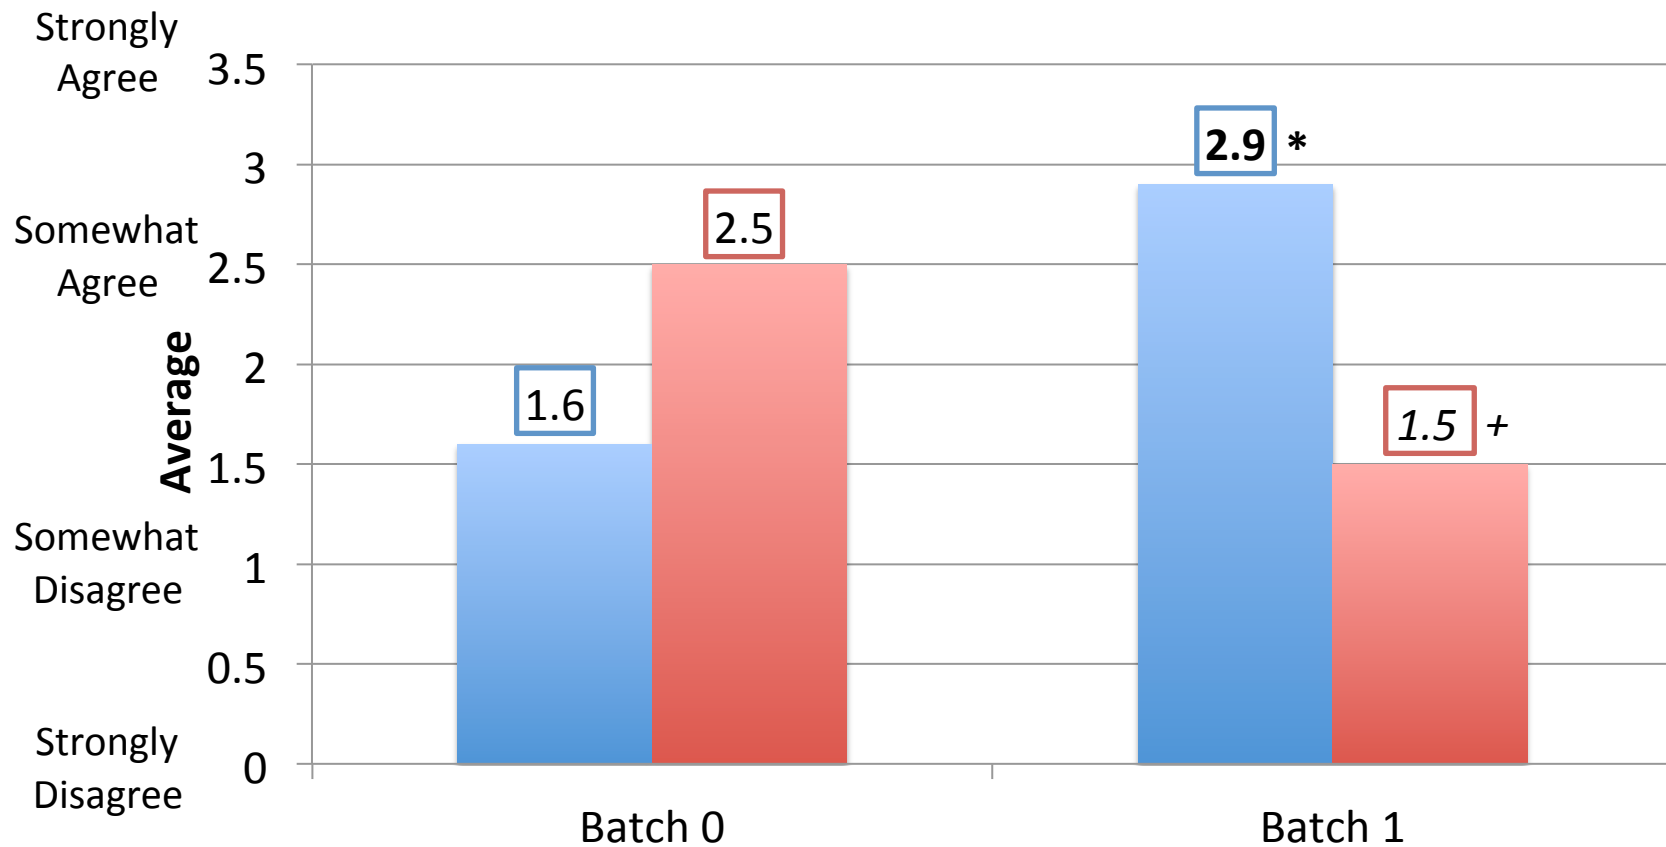
- Same setup, but real robot and fewer users
 - Users type to robot to mimic Mechanical Turk setup
- 10 users in initial test batch
- System interacted freely with people on the floor for four days as training (34 conversations in total)
- 10 users in the second test batch, after retraining

Office Robot Dialog Completion



Office Robot Survey Responses

■ The robot understood me ■ The robot frustrated me



Conclusions

- Lexical acquisition reduces dialog lengths for multi-argument predicates like delivery
- Causes users to perceive the system as more understanding
- Leads to less user frustration
- Allows improving language understanding without large, annotated corpora

Learning to Interpret Natural Language Commands through Human-Robot Dialog

Jesse Thomason

Shiqi Zhang, Raymond Mooney, Peter Stone

The University of Texas at Austin

Related Work

- Command processing has taken many forms
- Specify tasks step-by-step (Meriçli, 2014)
 - Assumes particular words in particular order
- Specify low-level action sequences (Misra, 2014; Tellex, 2011)
 - Uses a parser trained on a huge corpus
- Map language to action specifications (Matuszek, 2013)
 - Cannot learn new words/expressions

Future Work

- Perceptual grounding (`blue', `left of')
- Predicate invention (`ruddy')
- Learning a multi-objective dialog policy that trades off learning and user satisfaction