Natural Language Understanding for Robots with Dialog Learning and Language Grounding

Harel Yedidsion, Aishwarya Padmakumar, Jesse Thomason,
Jivko Sinapov, Justin Hart, Peter Stone, and Raymond J. Mooney

The University of Texas at Austin

Commanding Robots in Natural Language

We develop a human-robot dialog system that enables a robot to learn from natural language interaction to better understand commands.
We integrate learning techniques from:

- Perceptual language grounding
- Semantic parsing
- Dialog management
- Opportunistic active learning

And implement them on a mobile robot.

The Challenge

Perceptual Language Grounding requires associating words and phrases in language to objects, properties and relations in the world as perceived by the robot's sensors.

"Get the red heavy mug!"

Exploring objects for all possible sensorimotor contexts takes time.

Exploration Policy - MOMDP

Amiri et al. (IJCAI 2018)
This work developed a method to accelerate the object exploration process, and to improve accuracy in predicate identification. The exploration process is modeled as a MOMDP and the model suggests the best exploratory action to take in order to update the belief.

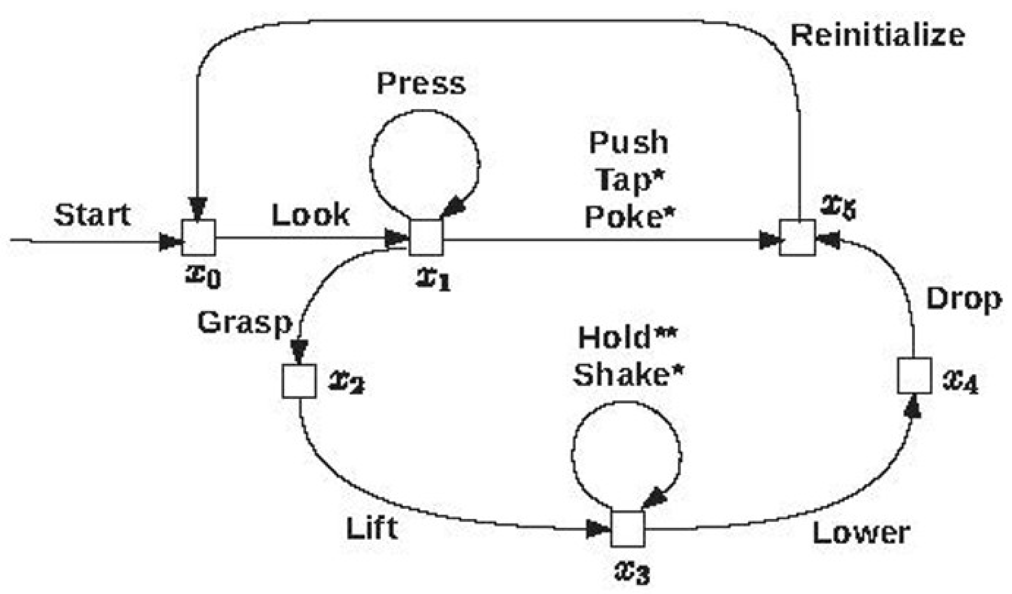

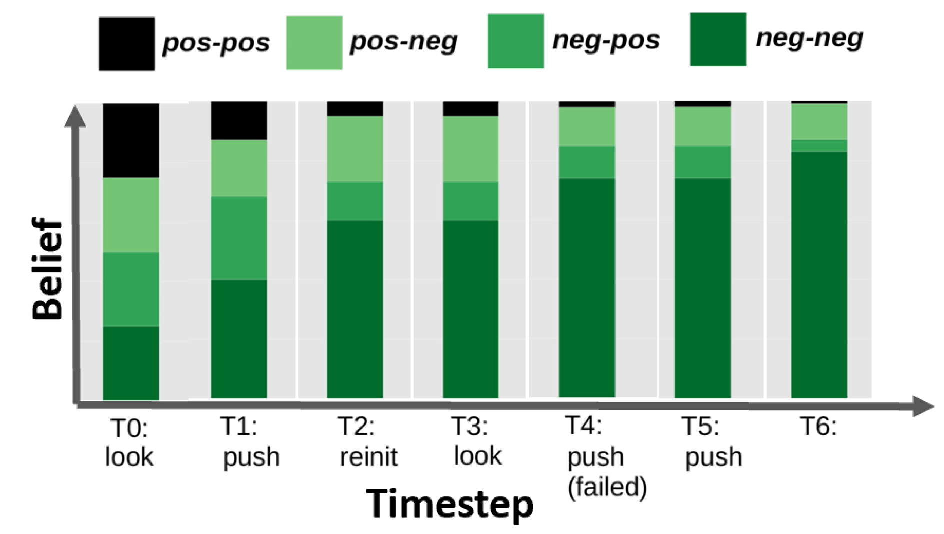

Exploration Policy - Linguistic Closeness

Thomason et al. (AAAI 2018)
Humans annotators suggested the best exploratory actions for some predicates, and then, based on the word embedding relations we can infer the right sensorimotor exploration process.

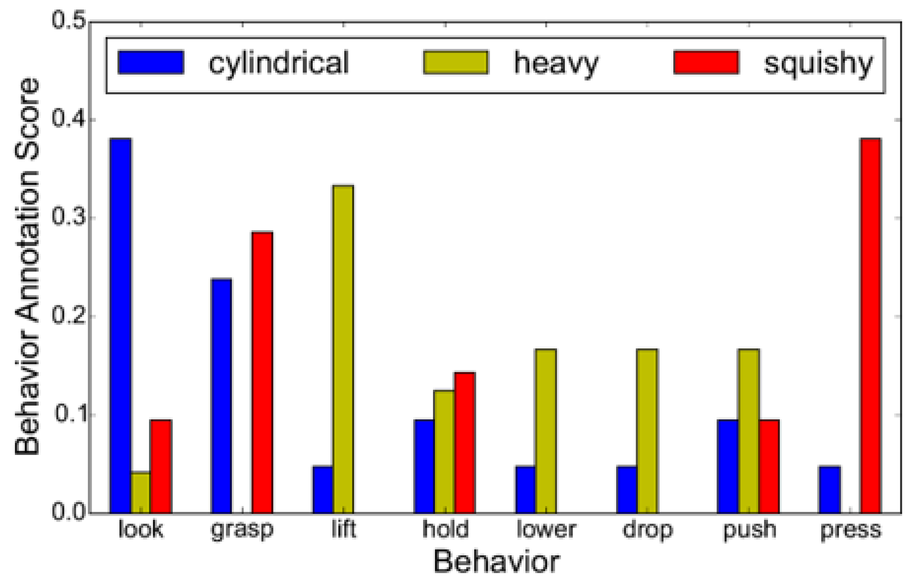

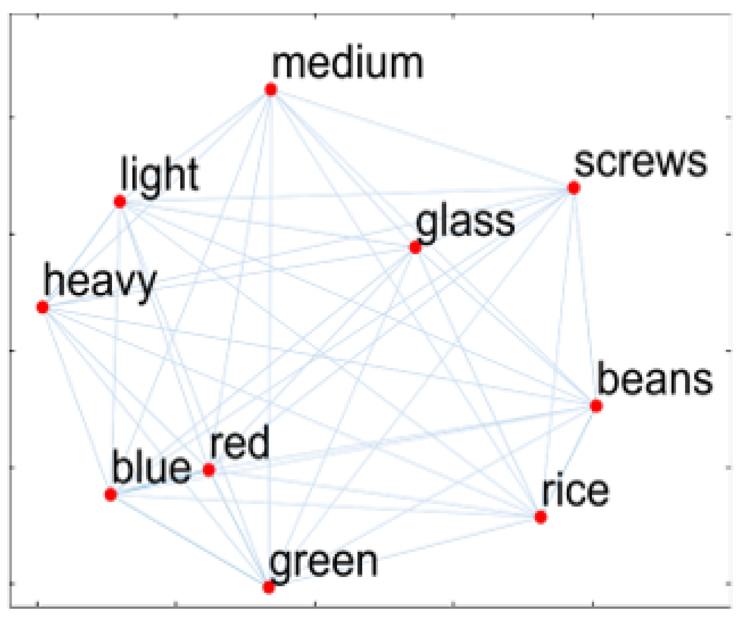

Word Embedding

Opportunistic Active Learning

Thomason et al. (CoRL 2017)
Asking off topic questions to improve future interactions in object retrieval tasks.

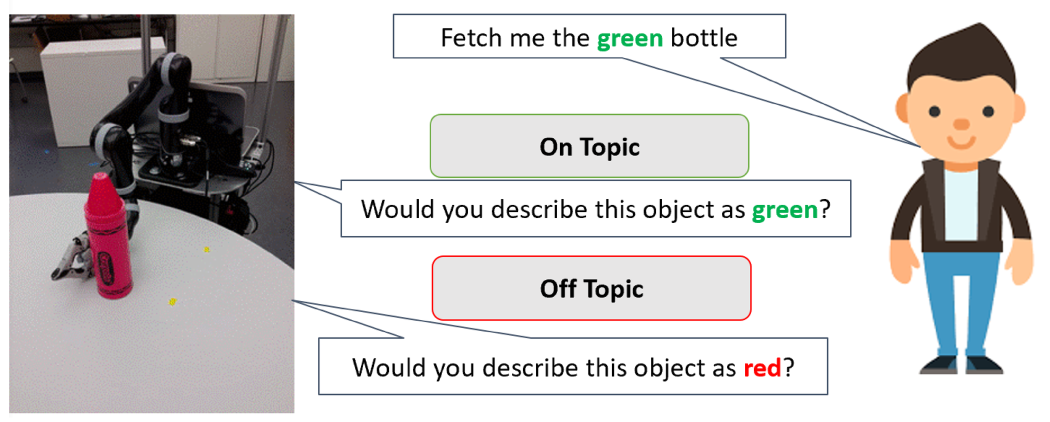

The inquisitive agent achieved better accuracy and rated as more fun to interact with.

Opportunistic Active Learning

Padmakumar et al. (EMNLP 2018)
Using reinforcement learning to learn a policy for opportunistic active learning using the Visual Genome dataset.

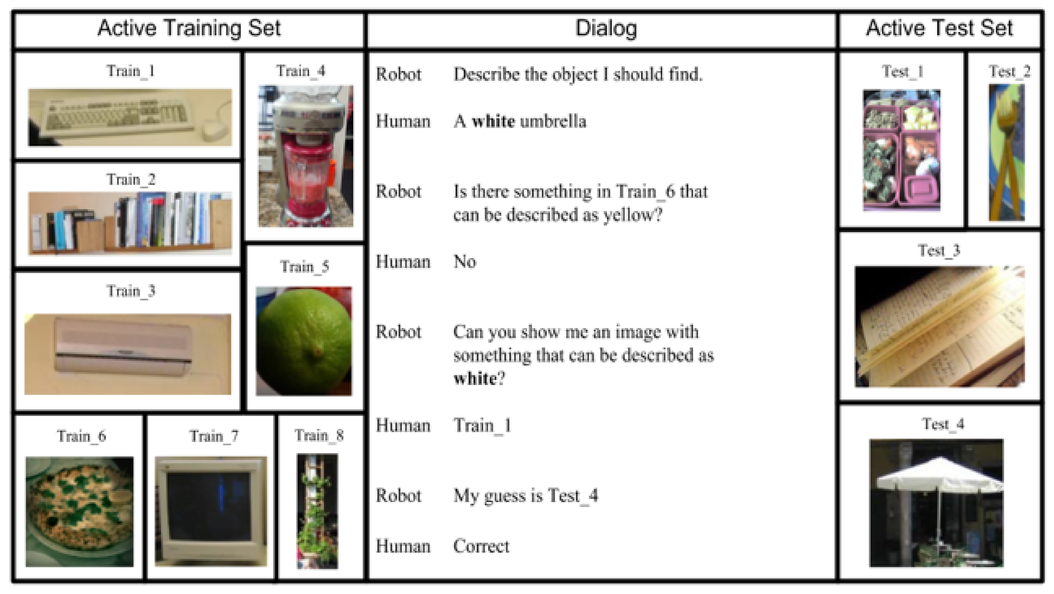

The agent needs to identify which object is best described by a given phrase

Policy	Success rate	Average Dialog Length
Learned	**0.44**	**12.95**
Static	0.29	16

The learned policy proved more accurate than the static policy with a shorter dialog length.

Robotic Implementation

Thomason et al. (RSS 2018)
Finally we integrate these components on a mobile robot (BWIbot).

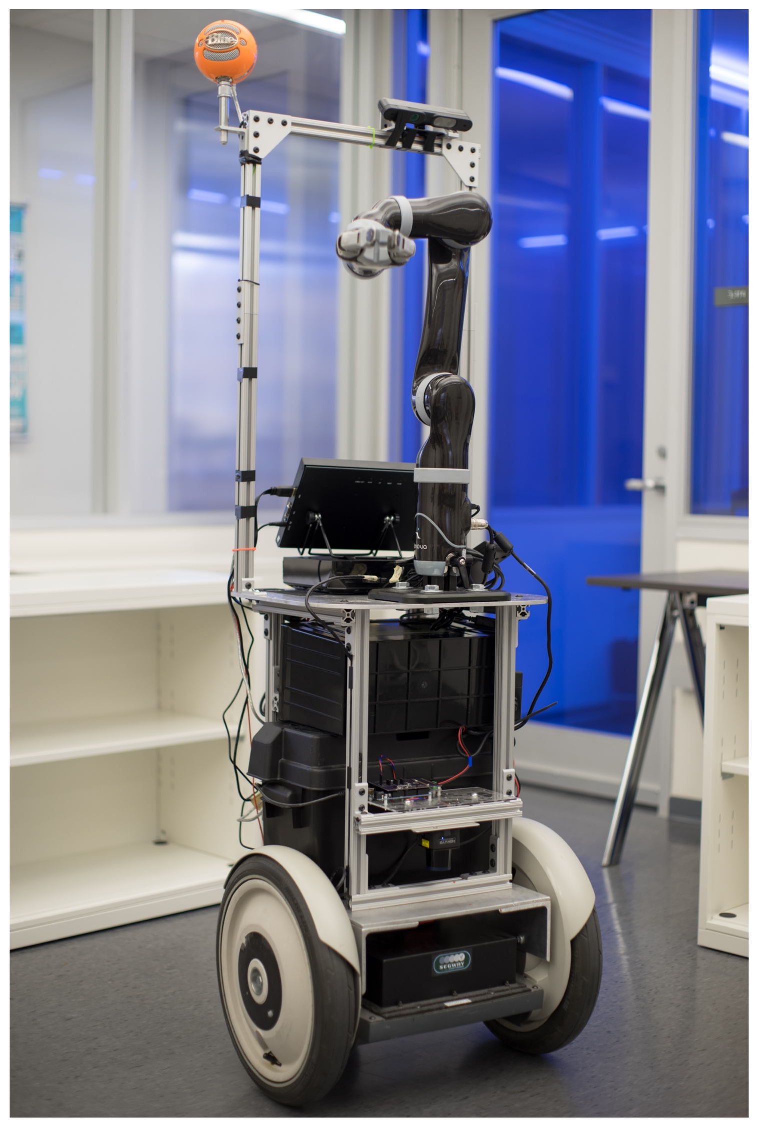

The robot can parse natural language commands for object delivery tasks, ground new predicates by asking clarification questions, navigate autonomously, grasp and deliver objects.