Opportunistic Active Learning for Grounding Natural Language Descriptions

Jesse Thomason, Aishwarya Padmakumar,
Jivko Sinapov, Justin Hart, Peter Stone, and Raymond J. Mooney
University of Texas at Austin

CoRL 2017

Grounding Natural Language Descriptions

Grounded language learning bridges internal robot sensory information with natural language predicates (such as nouns and adjectives). It is expensive to obtain labels between objects and predicates.

Natural Language Setting

We use an object identification task to gather human language labels for objects. People described an object in natural language, then a robot asks about whether predicates in the description apply to surrounding objects, clarifying their meanings.

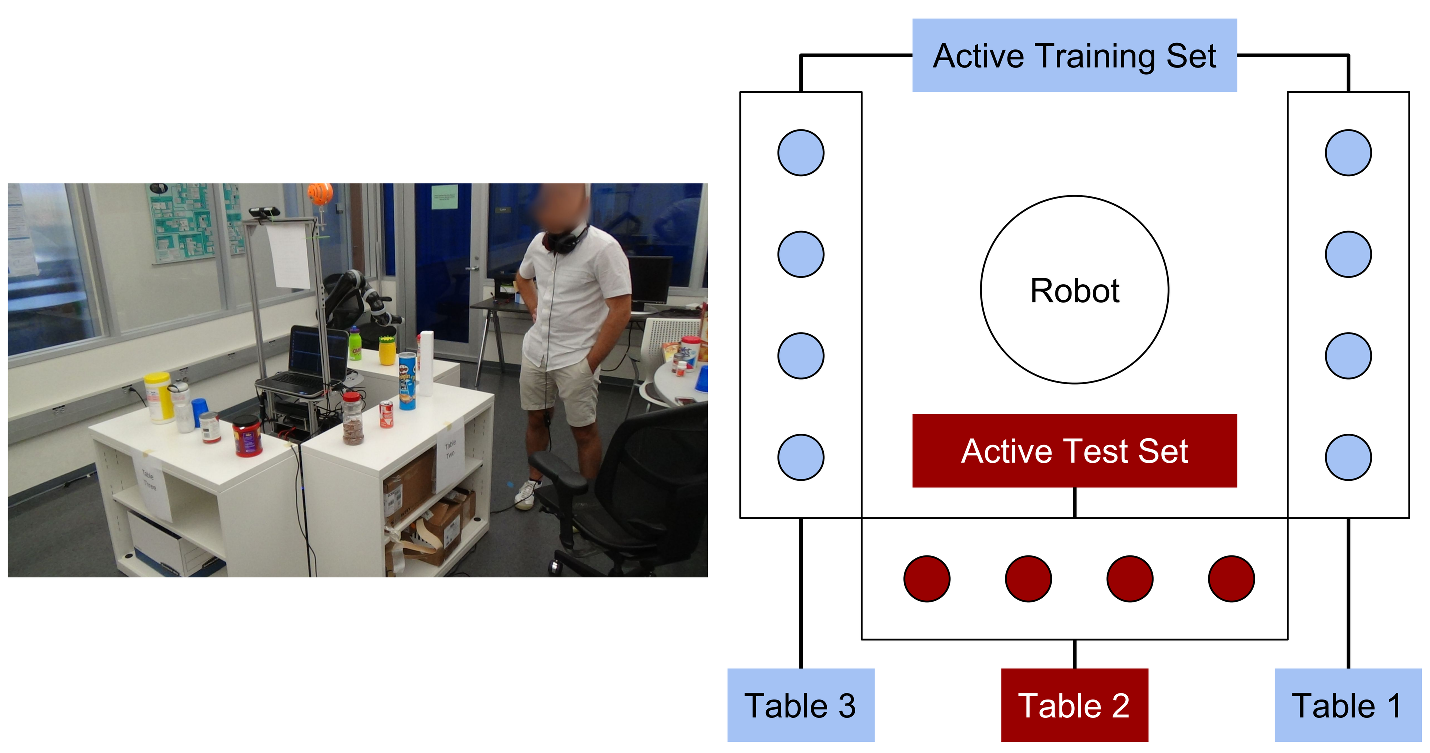

Participants described an object on Table 2 from the *active test set*, then answered the robot's questions about the objects in its *active training set* on Tables 1 and 3 before the robot guessed the target object.

Opportunistic Active Learning

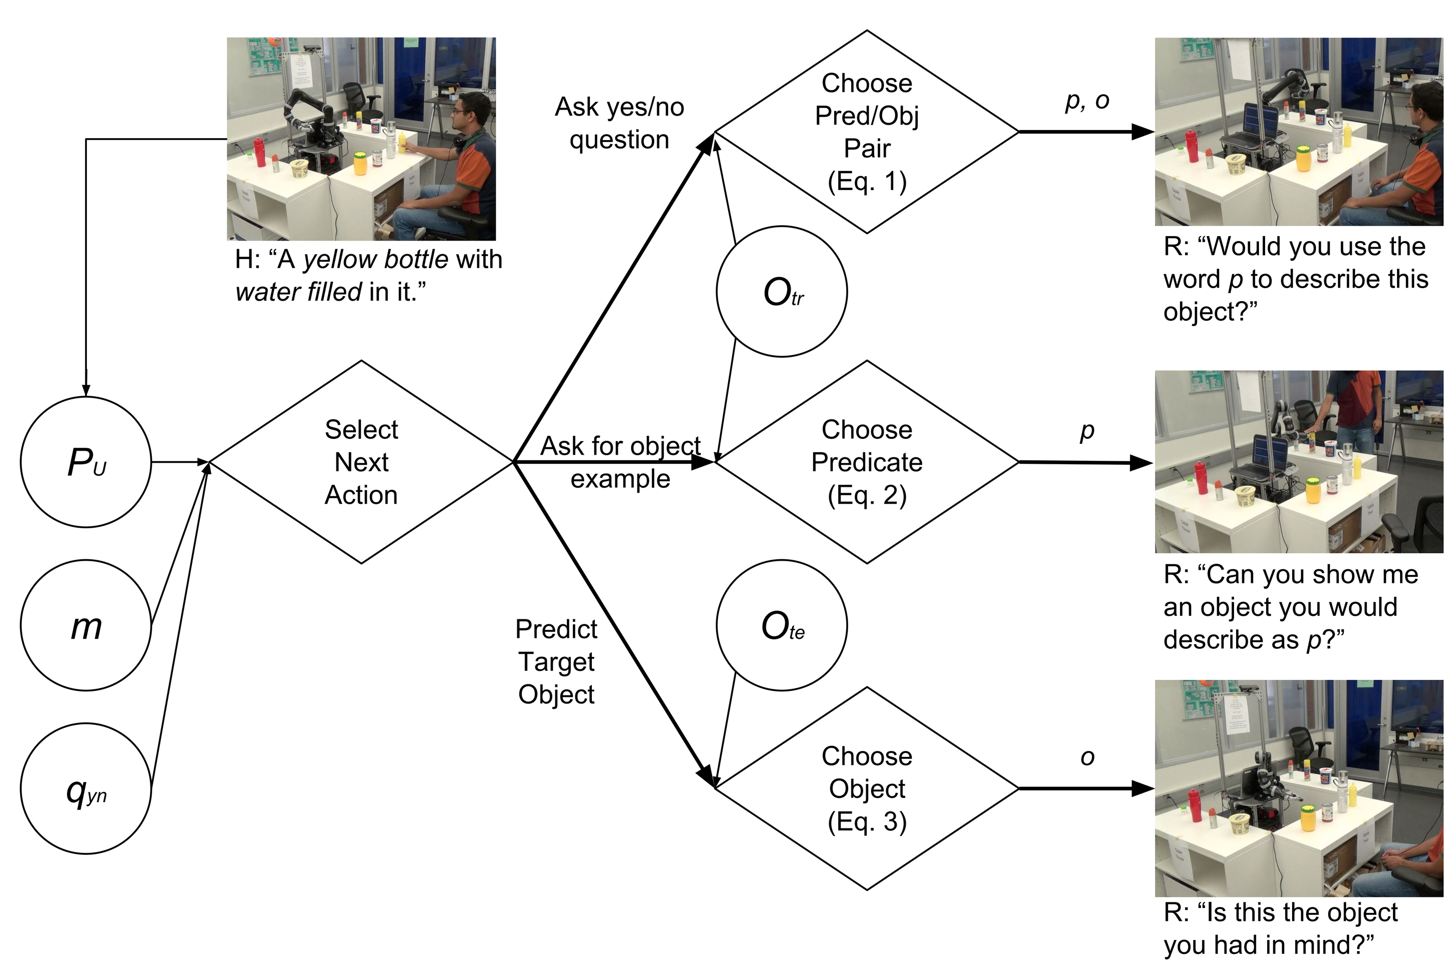

$$o_{\min}(p) = \text{argmin}_{o \in O_{tr}} \left(\kappa(p, o)\right)$$

Object with least confidence for *p*

Confidence in prediction of *p* for object *o*

$$prob(p) = \frac{1 - \kappa(p, o_{\min}(p))}{\sum_{q \in P \setminus \{p\}} 1 - \kappa(q, o_{\min}(q))}.$$

Confidence in *p* for the object with lowest confidence

Probability of choosing *p*

Experiments

We compare a *baseline* agent that can only ask about relevant predicates for the current user to an *inquisitive* one that can ask about any predicates it knows, hoping to improve interactions with future users. Learning is done over three rounds.

The objects used in each set, the active training sets left of each given active test set.

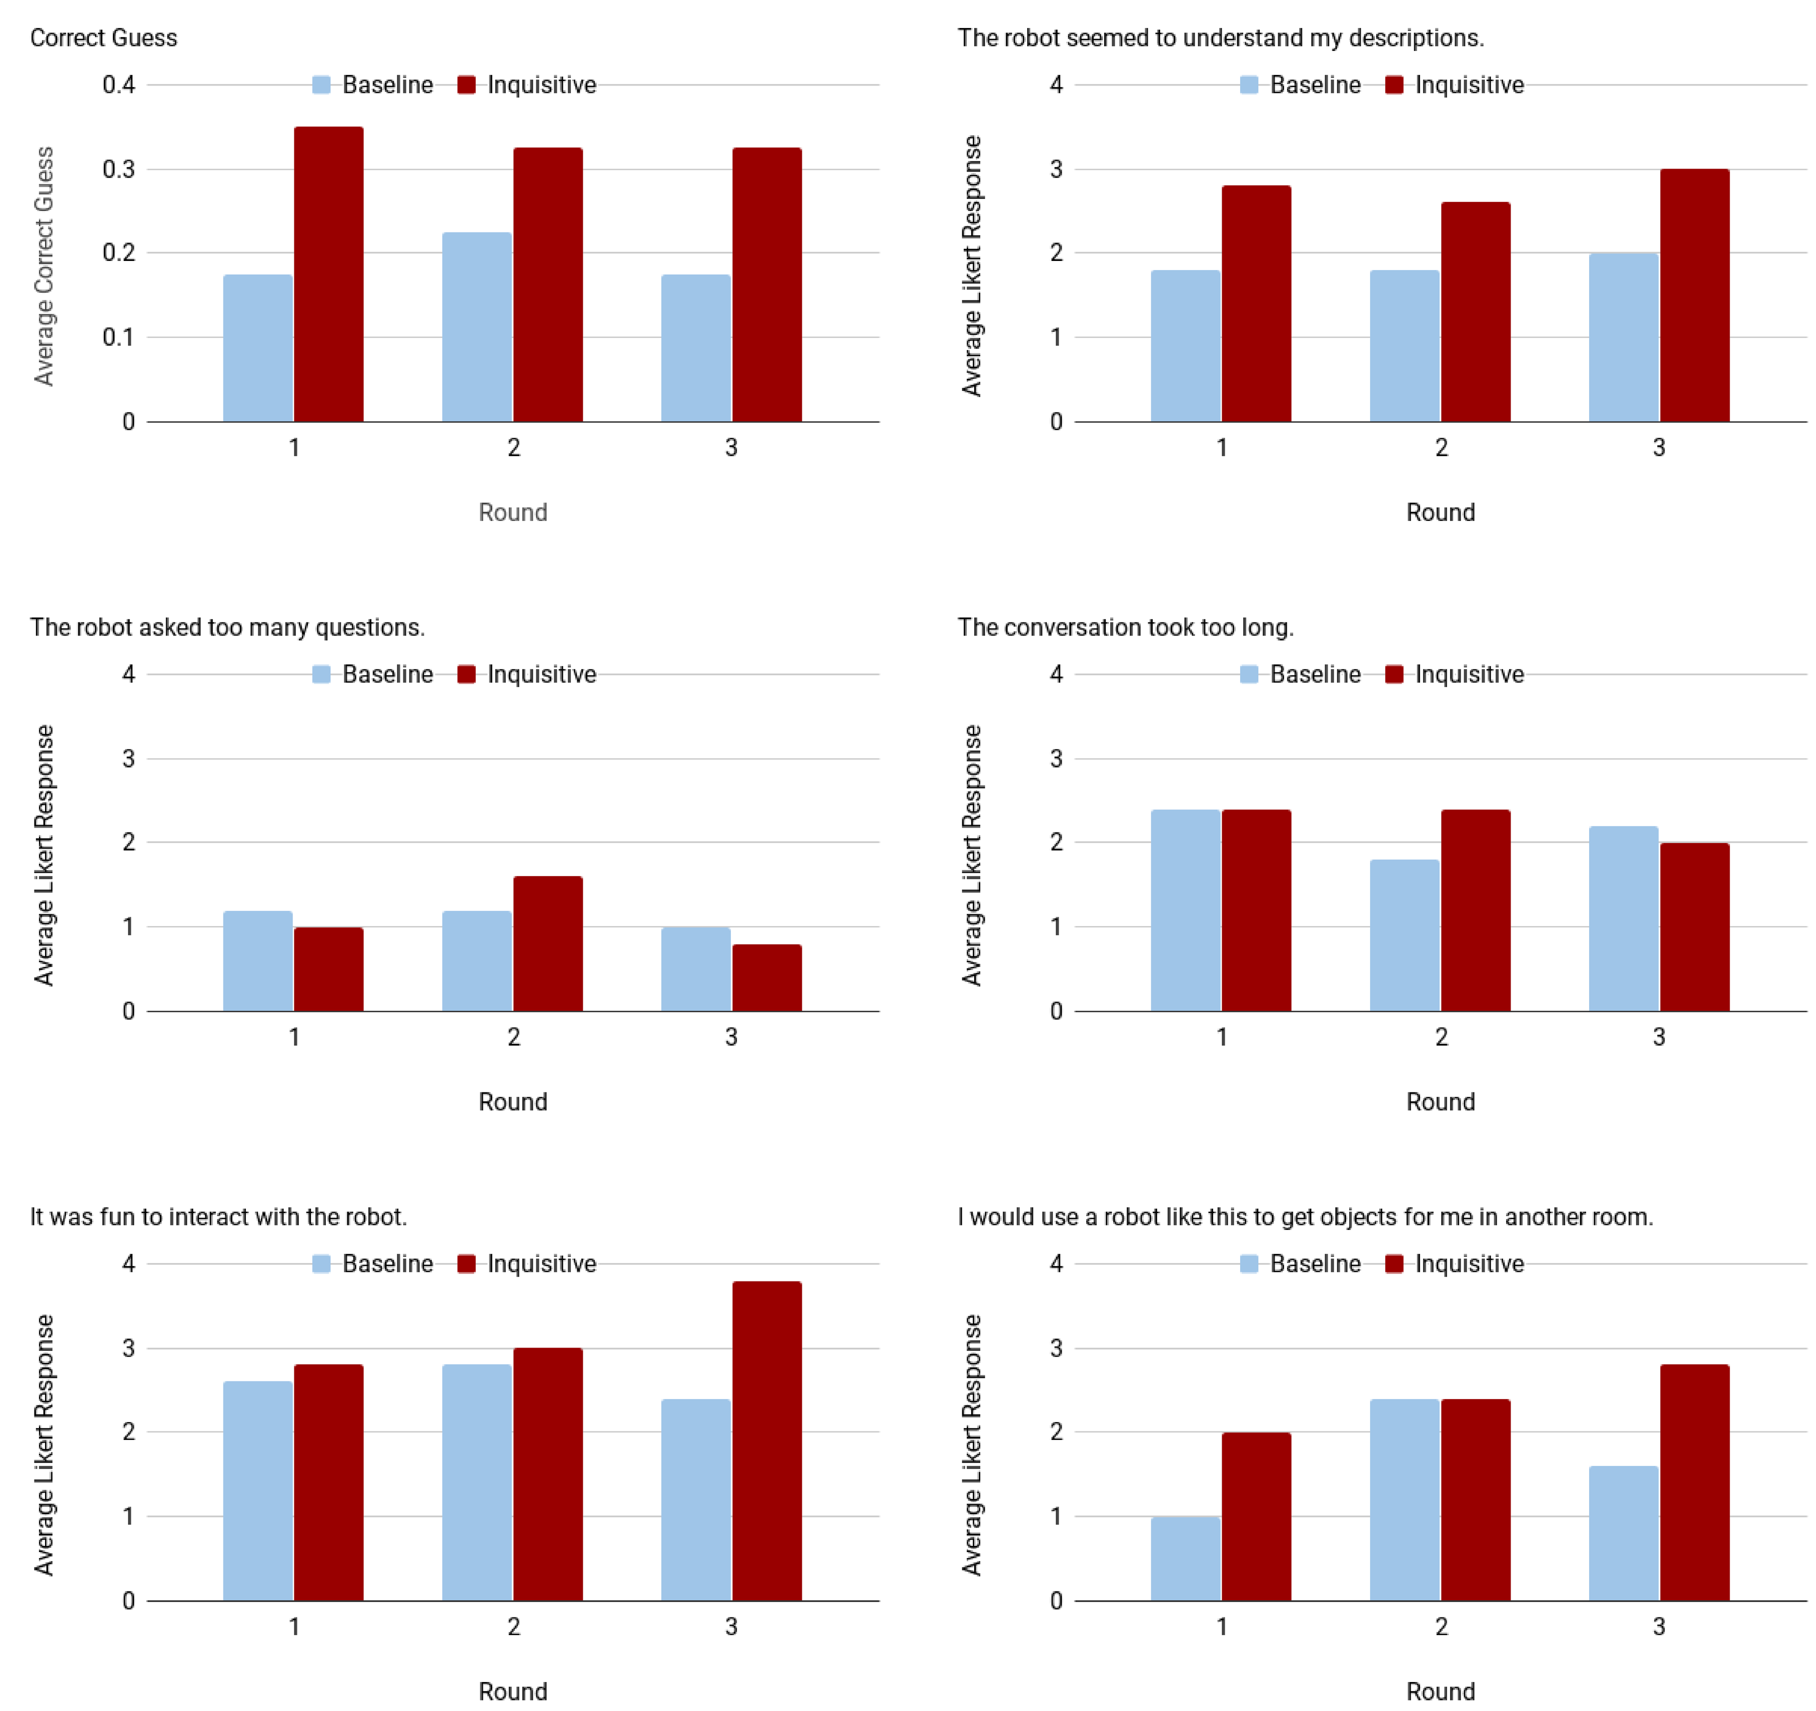

In round 1, the agents differed only in the number of questions they could ask. In round 2, the *inquisitive* agent could both ask more questions and ask about off-topic predicates. In round 3, the agents differed only in their training so far, and both had their number of questions fixed to $m = 3$.

Conclusions

Opportunistic active learning can be used both to accomplish a current user's object identification task more reliably and to improve downstream performance by asking off-topic (*inquisitive*) questions.

jesse@cs.utexas.edu