

ACL 2017

**The First Workshop on Language
Grounding for Robotics**

Proceedings of the Workshop

August 3, 2017
Vancouver, Canada

©2017 The Association for Computational Linguistics

Order copies of this and other ACL proceedings from:

Association for Computational Linguistics (ACL)
209 N. Eighth Street
Stroudsburg, PA 18360
USA
Tel: +1-570-476-8006
Fax: +1-570-476-0860
acl@aclweb.org

ISBN 978-1-945626-64-7

Introduction

After the remarkable successes of recent work visually grounded models of language, the embodied and task-oriented aspects of language learning stand as a natural next challenge. As autonomous robotic agents become increasingly capable and are deployed to progressively more complex environments, expressive, accessible interfaces are becoming essential to realizing the potential of such technologies. Natural language is immediately available to non-expert users and expressive enough to represent complex actions and plans. Can we give instructions to robotic agents to assist with navigation and manipulation tasks in remote settings? Can we talk to robots about the surrounding visual world, and help them interactively learn the language needed to finish a task? To build robots that we can converse with in our homes, offices, hospitals, and warehouses, it is essential that we develop new techniques for linking language to action in the real world.

While the opportunity is clear, enabling effective interaction between users and autonomous agents requires addressing some of the core open challenges in NLP while studying new domains and tasks. This workshop aims to explore these challenges, bringing together members of the NLP, robotics, and vision communities to focus on language grounding in robots and other interactive goal-driven systems. The program features twelve new articles and seven cross-submissions from related areas, to be presented as both posters and talks. We are also excited to host remarkable invited speakers, including Regina Barzilay, Joyce Chai, Karl Moritz Hermann, Hadas Kress-Gazit, Terence Langendoen, Percy Liang, Ray Mooney, Nicholas Roy, Stefanie Tellex and Jason Weston.

We thank the program committee, the ACL workshop chairs Wei Xu and Jonathan Berant, the invited speakers, and our sponsors DeepMind and Facebook.

—Mohit Bansal, Cynthia Matuszek, Jacob Andreas, Yoav Artzi and Yonatan Bisk, organizers

Organizers:

Mohit Bansal, UNC Chapel Hill
Cynthia Matuszek, UMBC
Jacob Andreas, UC Berkeley
Yoav Artzi, Cornell
Yonatan Bisk, ISI, USC

Program Committee:

Antoine Bordes, Facebook AI Research
Ankur Parikh, Google Research
Devi Parikh, Georgia Tech
Dhruv Batra, Georgia Tech
Dieter Fox, Univ of Washington
Dipendra Misra, Cornell
Edward Grefenstette, Google DeepMind
Eunsol Choi, Univ of Washington
Hadas Kress-Gazit, Cornell
Hannaneh Hajishirzi, Univ of Washington
Hongyuan Mei, JHU
Jason Weston, Facebook AI Research
Jayant Krishnamurthy, AI2
Jesse Dodge, CMU
Jesse Thomason, UT Austin
Jivko Sinapov, UT Austin
Jonathan Berant, Tel-Aviv
Joyce Chai, MSU
Julia Hockenmaier, UIUC
Karthik Narasimhan, MIT
Kenton Lee, Univ of Washington
Licheng Yu, UNC Chapel Hill
Lisa A Hendricks, UC Berkeley
Lucy Vanderwende, MSR
Marcus Rohrbach, UC Berkeley
Mark Yatskar, Univ of Washington
Matthew Walter, TTI-Chicago
Raia Hadsell, Google DeepMind
Ray Mooney, UT Austin
Siva Reddy, Univ of Edinburgh
Subhashini Venugopalan, UT Austin
Thomas Kollar, Amazon
Tom Kwiatkowski, Google Research
Tom Williams, Tufts
Yejin Choi, Univ of Washington

Table of Contents

<i>Grounding Language for Interactive Task Learning</i> Peter Lindes, Aaron Mininger, James R. Kirk and John E. Laird	1
<i>Learning how to Learn: An Adaptive Dialogue Agent for Incrementally Learning Visually Grounded Word Meanings</i> Yanchao Yu, Arash Eshghi and Oliver Lemon	10
<i>Guiding Interaction Behaviors for Multi-modal Grounded Language Learning</i> Jesse Thomason, Jivko Sinapov and Raymond Mooney	20
<i>Structured Learning for Context-aware Spoken Language Understanding of Robotic Commands</i> Andrea Vanzo, Danilo Croce, Roberto Basili and Daniele Nardi	25
<i>Natural Language Grounding and Grammar Induction for Robotic Manipulation Commands</i> Muhannad Alomari, Paul Duckworth, Majd Hawasly, David C. Hogg and Anthony G. Cohn ...	35
<i>Communication with Robots using Multilayer Recurrent Networks</i> Bedřich Pišl and David Mareček	44
<i>Grounding Symbols in Multi-Modal Instructions</i> Yordan Hristov, Svetlin Penkov, Alex Lascarides and Subramanian Ramamoorthy	49
<i>Exploring Variation of Natural Human Commands to a Robot in a Collaborative Navigation Task</i> Matthew Marge, Claire Bonial, Ashley Foots, Cory Hayes, Cassidy Henry, Kimberly Pollard, Ron Artstein, Clare Voss and David Traum	58
<i>A Tale of Two DRAGGNs: A Hybrid Approach for Interpreting Action-Oriented and Goal-Oriented Instructions</i> Siddharth Karamcheti, Edward Clem Williams, Dilip Arumugam, Mina Rhee, Nakul Gopalan, Lawson L.S. Wong and Stefanie Tellex	67
<i>Are Distributional Representations Ready for the Real World? Evaluating Word Vectors for Grounded Perceptual Meaning</i> Li Lucy and Jon Gauthier	76
<i>Sympathy Begins with a Smile, Intelligence Begins with a Word: Use of Multimodal Features in Spoken Human-Robot Interaction</i> Jekaterina Novikova, Christian Dondrup, Ioannis Papaioannou and Oliver Lemon	86
<i>Towards Problem Solving Agents that Communicate and Learn</i> Anjali Narayan-Chen, Colin Graber, Mayukh Das, Md Rakibul Islam, Soham Dan, Sriraam Natara- jan, Janardhan Rao Doppa, Julia Hockenmaier, Martha Palmer and Dan Roth	95

Workshop Program

Grounding Language for Interactive Task Learning

Peter Lindes, Aaron Mininger, James R. Kirk and John E. Laird

Learning how to Learn: An Adaptive Dialogue Agent for Incrementally Learning Visually Grounded Word Meanings

Yanchao Yu, Arash Eshghi and Oliver Lemon

Guiding Interaction Behaviors for Multi-modal Grounded Language Learning

Jesse Thomason, Jivko Sinapov and Raymond Mooney

Structured Learning for Context-aware Spoken Language Understanding of Robotic Commands

Andrea Vanzo, Danilo Croce, Roberto Basili and Daniele Nardi

Natural Language Grounding and Grammar Induction for Robotic Manipulation Commands

Muhannad Alomari, Paul Duckworth, Majd Hawasly, David C. Hogg and Anthony G. Cohn

Communication with Robots using Multilayer Recurrent Networks

Bedřich Pišl and David Mareček

Grounding Symbols in Multi-Modal Instructions

Yordan Hristov, Svetlin Penkov, Alex Lascarides and Subramanian Ramamoorthy

Exploring Variation of Natural Human Commands to a Robot in a Collaborative Navigation Task

Matthew Marge, Claire Bonial, Ashley Fouts, Cory Hayes, Cassidy Henry, Kimberly Pollard, Ron Artstein, Clare Voss and David Traum

A Tale of Two DRAGGNs: A Hybrid Approach for Interpreting Action-Oriented and Goal-Oriented Instructions

Siddharth Karamcheti, Edward Clem Williams, Dilip Arumugam, Mina Rhee, Nakul Gopalan, Lawson L.S. Wong and Stefanie Tellex

Are Distributional Representations Ready for the Real World? Evaluating Word Vectors for Grounded Perceptual Meaning

Li Lucy and Jon Gauthier

Sympathy Begins with a Smile, Intelligence Begins with a Word: Use of Multimodal Features in Spoken Human-Robot Interaction

Jekaterina Novikova, Christian Dondrup, Ioannis Papaioannou and Oliver Lemon

Towards Problem Solving Agents that Communicate and Learn

Anjali Narayan-Chen, Colin Graber, Mayukh Das, Md Rakibul Islam, Soham Dan, Sriraam Natarajan, Janardhan Rao Doppa, Julia Hockenmaier, Martha Palmer and Dan Roth

