NAACL HLT 2019

The Combined Workshop on Spatial Language Understanding (SpLU) and Grounded Communication for Robotics (RoboNLP)

Proceedings of the Workshop

June 6, 2019
Minneapolis, MN

©2019 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 978-1-950737-09-3

Introduction

SpLU-RoboNLP 2019 is a combined workshop on spatial language understanding (SpLU) and grounded communication for robotics (RoboNLP) that focuses on spatial language, both linguistic and theoretical aspects and its application to various areas including and especially focusing on robotics. The combined workshop aims to bring together members of NLP, robotics, vision and related communities in order to initiate discussions across fields dealing with spatial language along with other modalities. The desired outcome is identification of both shared and unique challenges, problems and future directions across the fields and various application domains.

While language can encode highly complex, relational structures of objects, spatial relations between them, and patterns of motion through space, the community has only scratched the surface on how to encode and reason about spatial semantics. Despite this, spatial language is crucial to robotics, navigation, NLU, translation and more. Standardizing tasks is challenging as we lack formal domain independent meaning representations. Spatial semantics requires an interplay between language, perception and (often) interaction.

Following the exciting recent progress in visual language grounding, the embodied, task-oriented aspect of language grounding is an important and timely research direction. To realize the long-term goal of 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 in which spatial language understanding plays a great role. 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? We hope to learn about (and begin to answer) these questions as we delve deeper into spatial language understanding and grounding language for robotics.

We accepted 8 archival submissions and 12 cross-submissions.

Organizers:

James F. Allen, University of Rochester, IHMC
Jacob Andreas, Semantic Machines/MIT
Jason Baldridge, Google
Mohit Bansal, UNC Chapel Hill
Archna Bhatia, IHMC
Yonatan Bisk, University of Washington
Asli Celikyilmaz, Microsoft Research
Bonnie J. Dorr, IHMC
Parisa Kordjamshidi, Tulane University / IHMC
Matthew Marge, Army Research Lab
Jesse Thomason, University of Washington

Program Committee:

Malihe Alikhani, Rutgers University
Yoav Artzi, Cornell University
Jacob Arkin, University of Rochester
John A. Bateman, Universität Bremen
Mehul Bhatt, Örebro University
Jonathan Berant, Tel-Aviv University
Raffaella Bernardi, University of Trento
Steven Bethard, University of Arizona
Johan Bos, University of Groningen
Volkan Cirik, CMU
Guillem Collell, KU Leuven
Joyce Chai, Michigan State University
Angel Chang, Stanford University
Simon Dobnik, CLASP and FLOV, University of Gothenburg Sweden
Ekaterina Egorova, University of Zurich
Zoe Falomir, Universität Bremen
Daniel Fried, UCSF
Lucian Galescu, IHMC
Felix Gervits, Tufts
Hannaneh Hajishirzi, University of Washington
Casey Kennington, Boise State University
Jayant Krishnamurthy, Semantic Machines
Stephanie Lukin, Army Research Laboratory
Chris Mavrogiannis, Cornell
Dipendra Misra, Cornell University
Marie-Francine Moens, KU Leuven
Ray Mooney, University of Texas
Mari Broman Olsen, Microsoft
Martijn van Otterlo, Tilburg University, The Netherlands
Aishwarya Padmakumar, UT Austin
Natalie Parde, University of Illinois Chicago
Ian Perera, IHMC

James Pustejovsky, Brandeis University
Preeti Ramaraj, University of Michigan
Siva Reddy, Stanford
Kirk Roberts, The University of Texas
Anna Rohrbach, UC Berkeley
Marcus Rohrbach, FAIR
Manolis Savva, Princeton University
Jivko Sinapov, Tufts
Kristin Stock, Massey University of New Zealand
Alane Suhr, Cornell
Clare Voss, ARL
Xin Wang, University of California Santa Barbara
Shiqi Zhang, SUNY Binghamton
Victor Zhong, University of Washington

Invited Speakers:

Dhruv Batra, GaTech/FAIR
Joyce Chai, Michigan State University
Cynthia Matuszek, UMBC
Raymond J. Mooney, UT Austin
Martha Palmer, CU Boulder
Matthias Scheutz, Tufts
Stefanie Tellex, Brown
Dilek Hakkani-Tur, Amazon

Table of Contents

Corpus of Multimodal Interaction for Collaborative Planning
Miltiadis Marios Katsakioris, Helen Hastie, Ioannis Konstas and Atanas Laskov 1

¿Es un plátano? Exploring the Application of a Physically Grounded Language Acquisition System to Spanish
Caroline Kery, Francis Ferraro and Cynthia Matuszek 7

From Virtual to Real: A Framework for Verbal Interaction with Robots
Eugene Joseph 18

Learning from Implicit Information in Natural Language Instructions for Robotic Manipulations
Ozan Arkan Can, Pedro Zuidberg Dos Martires, Andreas Persson, Julian Gaal, Amy Loutfi, Luc De Raedt, Deniz Yuret and Alessandro Saffiotti 29

Multi-modal Discriminative Model for Vision-and-Language Navigation
Haoshuo Huang, Vihan Jain, Harsh Mehta, Jason Baldridge and Eugene Ie 40

Semantic Spatial Representation: a unique representation of an environment based on an ontology for robotic applications
Guillaume Sarthou, Aurélie Clodic and Rachid Alami 50

SpatialNet: A Declarative Resource for Spatial Relations
Morgan Ulinski, Bob Coyne and Julia Hirschberg 61

What a neural language model tells us about spatial relations
Mehdi Ghanimifard and Simon Dobnik 71

Workshop Program

Thursday, June 06, 2019

08:30–08:40 *Opening Remarks*
Workshop Chairs

08:40–12:30 ***Morning Session***

08:40–09:00 *Poster Spotlight (1 min madness)*

09:00–12:30 ***Morning Session***

09:00–09:45 *Invited Talk*
Joyce Chai

09:45–10:30 *Invited Talk*
Matthias Scheutz

11:00–11:45 *Invited Talk*
Martha Palmer

11:45–12:30 *Invited Talk*
Stefanie Tellex

12:30–14:00 ***Session Poster: Poster Session and Lunch***

Corpus of Multimodal Interaction for Collaborative Planning
Miltiadis Marios Katsakioris, Helen Hastie, Ioannis Konstas and Atanas Laskov

¿Es un plátano? Exploring the Application of a Physically Grounded Language Acquisition System to Spanish
Caroline Kery, Francis Ferraro and Cynthia Matuszek

From Virtual to Real: A Framework for Verbal Interaction with Robots
Eugene Joseph

Learning from Implicit Information in Natural Language Instructions for Robotic Manipulations
Ozan Arkan Can, Pedro Zuidberg Dos Martires, Andreas Persson, Julian Gaal, Amy Loutfi, Luc De Raedt, Deniz Yuret and Alessandro Saffiotti

Multi-modal Discriminative Model for Vision-and-Language Navigation
Haoshuo Huang, Vihan Jain, Harsh Mehta, Jason Baldridge and Eugene Ie

Thursday, June 06, 2019 (continued)

Semantic Spatial Representation: a unique representation of an environment based on an ontology for robotic applications
Guillaume Sarthou, Aurélie Clodic and Rachid Alami

SpatialNet: A Declarative Resource for Spatial Relations
Morgan Ulinski, Bob Coyne and Julia Hirschberg

What a neural language model tells us about spatial relations
Mehdi Ghanimifard and Simon Dobnik

14:00–17:30 ***Afternoon Session***

14:00–14:45 *Invited Talk*
Dhruv Batra

14:45–15:30 *Invited Talk*
Cynthia Matuszek

16:00–16:45 *Invited Talk*
Raymond Mooney

16:45–17:15 ***Best Paper Oral Presentations***

17:15–18:00 ***Continued Poster Session***