

Grounding Natural Language References to Unvisited and Hypothetical Locations

Tom Williams¹, Rehj Cantrell², Gordon Briggs¹, Paul
Schermerhorn² and Matthias Scheutz¹

Human-Robot Interaction Laboratory

¹Tufts University

²Indiana University




Location-Based Spatial reference resolution

- What is it?
 - Determining the identities of spatial locations denoted by or referred to in natural language utterances.
- In natural language dialogues with robots, important for:
 - Navigating to locations
 - Discussing locations
 - Reasoning about locations
- Examples:
 - "Go to **the coffee shop**"
 - "Go to **the room at the end of the hall**"
 - "Go to **the classroom we just left**"



Scenario 1: one candidate referent


"Go to the room at the end of the hall"





Scenario 2: many candidate referents


"Go to the room at the end of the hall"





Scenario 3: no known candidate referents

"Go to the room at the end of the hall"





Previous approaches: known locations only

- Shimizu and Haas 2009
- Zender, Kruijff, and Kruijff-Korabayova 2009
- Kollar et al. 2010
- Matuszek, Fox, and Koscher 2010
- Chen and Mooney 2011
- Hemachandra et al. 2011

- Assume entire map is known after training period
- Cannot modify map once training period is over



Previous approaches: unknown locations

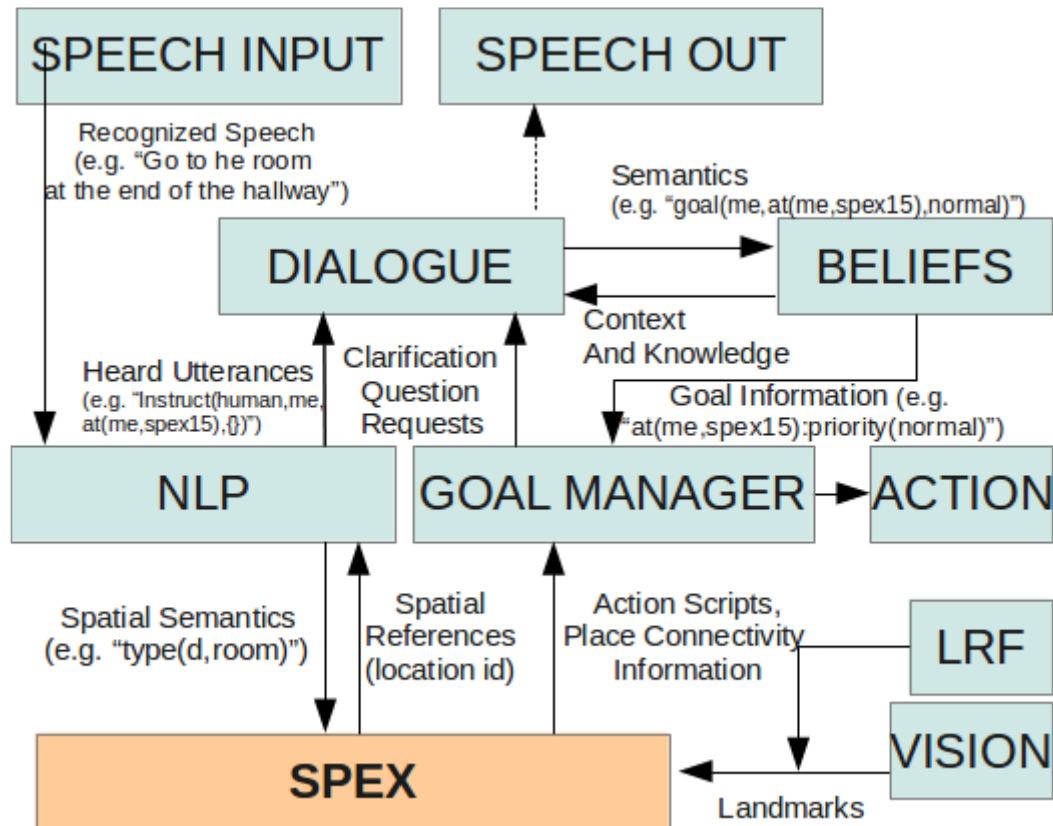
- Matuszek et al. 2012:
 - parse commands directly into action scripts
 - action scripts must be used immediately
 - assume map is known
 - cannot learn about places observed en route to unknown location.



Our approach


- Allow for reference resolution while environment is being explored
- Segment topological space using a set of percepts (e.g., observed doorways)
- When resolving a reference to an unknown location, posit its topological location without committing to its metric location
- Separate the processes of reference resolution and plan (action script) generation

SPEX: The SPatial EXpert





SPEX: reference resolution



"Go to **the room** at the end of **the hall** down to the right."


$\text{type}(R, \text{room}) \wedge \text{type}(H, \text{hall}) \wedge \text{at_end}(R, H) \wedge \text{to_right}(R, \text{here})$

H: 5

R: 9



SPEX: reference resolution



"Go to **the room** at the end of **the hall** down to the right."


$\text{type}(R, \text{room}) \wedge \text{type}(H, \text{hall}) \wedge \text{at_end}(R, H) \wedge \text{to_right}(R, \text{here})$

H: 5

R: 9



SPEX: positing unknown locations



"Go to **the room** at the end of **the hall** down to the right."


$\text{type}(R, \text{room}) \wedge \text{type}(H, \text{hall}) \wedge \text{at_end}(R, H) \wedge \text{to_right}(R, \text{here})$

H: 5

R: 9



SPEX: planning to visit unknown locations




Script 1: Go as far towards the destination as possible via landmarks with known positions

- 1) Exit room
- 2) Try again



SPEX: planning to visit unknown locations




Script 2: Travel the rest of the way using the destination's properties

- 1) Turn right
- 2) Travel to the end of the hallway
- 3) Try again



SPEX: consolidation and metric grounding



Script 2: Travel the rest of the way using the destination's properties

- 1) Turn right
- 2) Travel to the end of the hallway
- 3) Try again



Evaluation

- Gave SPEX prebuilt maps of a fully explored environment and of a partially explored environment.
- From each valid starting location within each environment, gave SPEX all valid descriptions of the other locations in the environment that fit a set of forms
 - e.g., “the room to your immediate left when exiting the breakroom”
 - “the room at the right end of the hallway”
 - “the third room on the right when facing left from your current position”
- Successfully resolved all references to known locations and generated appropriate action scripts to all unknown locations



Summary

- By allowing a robot to follow commands and engage in dialogue *while* exploring its environment, a robot is able to:
 - successfully resolve references to unknown locations
 - extend its world model based on dialogue
 - identify observed locations as those referenced in dialogue
 - travel to previously described (yet still unknown) locations
 - describe how two unknown locations are connected
 - pause an action sequence and resume it later from another location
 - return to a known location after visiting an unknown location



Future Work

- Handling uncertainty as to locations' properties during reference resolution
- Handling referential ambiguity resolution
- Deciding between exploration and inquiry
- Mapping and resolving references in more complex environments

This work was in part funded by NSF grant 111323 and in part funded by ONR grant #N00014-11-1-0493.