

Reasoning About Decisions / Reasoning About Language

Jacob Andreas

University of California, Berkeley

jda@cs.berkeley.edu

Sequence-to-sequence models for language interpretation and generation have become ubiquitous. The supervised sequence-to-sequence paradigm is powerful: flexible enough to handle many kinds of perceptual and discourse context, and expressive enough to model (some) long-range linguistic structure. But it also has significant limitations: it empirically favors generic utterances over informative ones, and is fundamentally limited to imitating the communicative strategies employed by annotators. This talk will explore first steps towards overcoming these limitations by reasoning explicitly about communicative context.

We'll begin with a family of "neuralized" rational speech acts models that combine learned semantics with inference-driven pragmatics, and see how to apply these models to tasks as diverse as image captioning, instruction generation, and visual navigation. Next, we'll turn to a family of less traditional NLP problems, and look at ways of using the same modeling tools to use language as a scaffold for model interpretability and few-shot concept learning.