

Jacob Andreas

Presentation by Rory Flukes and Ben Trevor

Current Status

- 5th-year PhD student in the Berkeley NLP group
- Advisor: Dan Klein
- Research interests:
 - “Language for interpretable AI”: how to turn the models developed through deep learning into something more easily understandable
 - “Language and reasoning”: modular deep learning architectures that can handle a variety of different linguistic tasks




Previous History

- B.S. from Columbia (2012)
 - Adviser: Michael Collins
- Master of Philosophy from University of Cambridge (2013)
 - Adviser: Stephen Clark
- Research positions and internships:
 - Semantic Machines
 - Information Sciences Institute
 - Google
 - Microsoft
 - Lawrence Berkeley National Labs




Influential Publications (as first author)

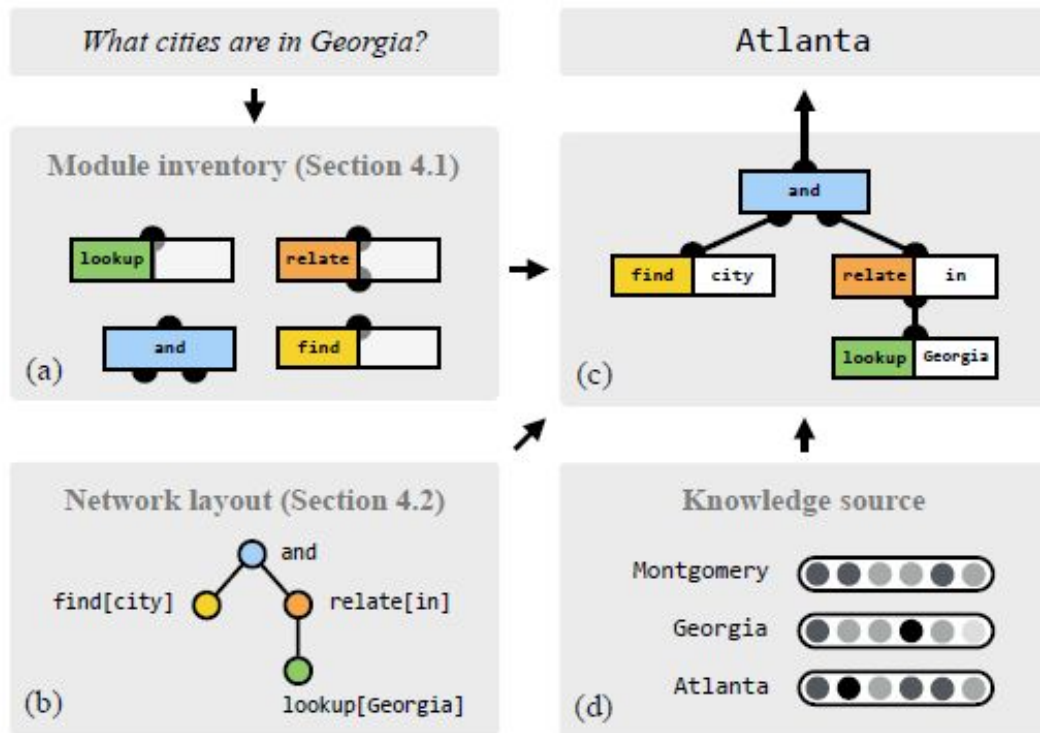
- “Learning to Compose Neural Networks for Question Answering” (NAACL 2016) (J Andreas, M Rohrbach, T Darrell, D Klein)
 - “Neural Module Networks” (IEEE Conference on Computer Vision and Pattern Recognition 2016) (J Andreas, M Rohrbach, T Darrell, D Klein)
 - “How Much do Word Embeddings Encode about Syntax?” (ACL 2014) (J Andreas, D Klein)
 - “Semantic Parsing as Machine Translation” (ACL 2013) (J Andreas, A Vlachos, S Clark)
- 

Learning to Compose Neural Networks for Question Answering

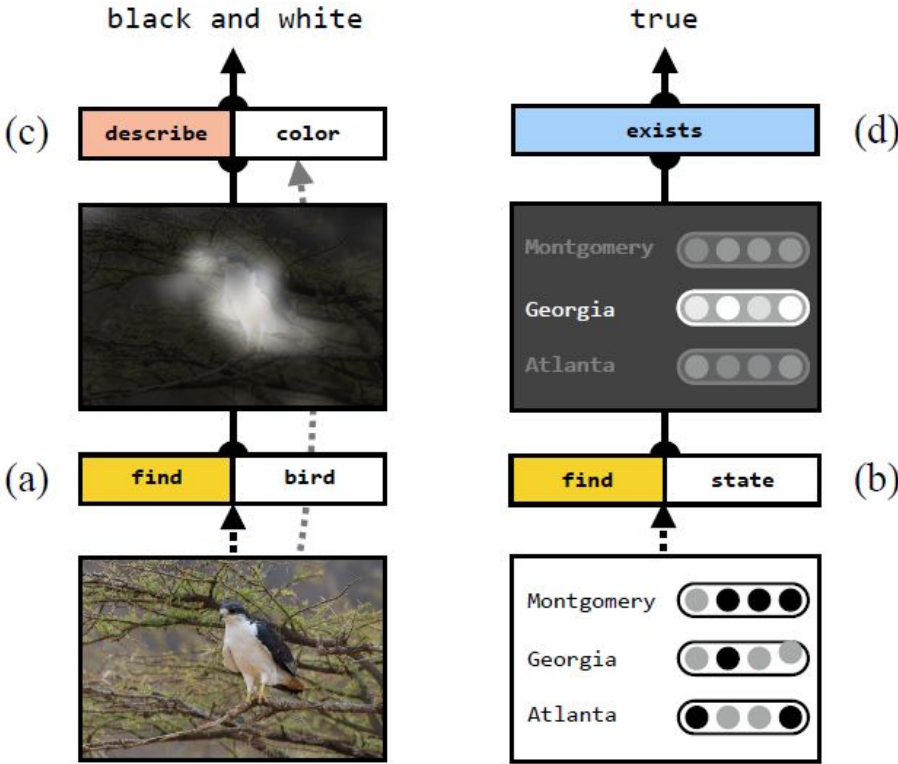
Question: Can a question-answering system learn to dynamically generate the structure of a neural module network based on the syntactic structure of a natural language question, and then apply that network to the knowledge base to answer the question?



Overall Concept



Neural Net Modules as Attention Mechanism



Two-Step Process

Input: set of (world, question, answer) triples

Variables: w = world representation, x = question, y = answer, z = network layout, θ = collection of model parameters (weights, etc.)

Step 1: Learn the Layout Model $\rightarrow p(z|x;\theta_L)$

Step 2: Learn the Execution Model $\rightarrow p(y|w;\theta_E)$

Step 1 is new in this research; Step 2 described in “Neural Module Networks” (Andreas et al., 2016)



Experiments

Two tasks:

- Visual Question Answering challenge (VQA) (Antol et al., 2015)
 - Questions about images; “What color is the bird?”

- GeoQA (Krishnamurthy and Kollar, 2013)
 - Questions about geography; “What cities are in Georgia?”



Visual Question Answering challenge results

	test-dev				test-std
	Yes/No	Number	Other	All	All
Zhou (2015)	76.6	35.0	42.6	55.7	55.9
Noh (2015)	80.7	37.2	41.7	57.2	57.4
Yang (2015)	79.3	36.6	46.1	58.7	58.9
NMN	81.2	38.0	44.0	58.6	58.7
D-NMN	81.1	38.6	45.5	59.4	59.4

Visual Question Answering challenge results



What is in the sheep's ear?

```
(describe[what]
 (and find[sheep]
      find[ear]))
```

tag



What color is she wearing?

```
(describe[color]
 find[wear])
```

white



What is the man dragging?

```
(describe[what]
 find[man])
```

boat (board)

GeoQA results

Model	Accuracy	
	GeoQA	GeoQA+Q
LSP-F	48	–
LSP-W	51	–
NMN	51.7	35.7
D-NMN	54.3	42.9

GeoQA results

Is Key Largo an island?

```
(exists (and lookup[key-largo] find[island]))
```

yes: correct

What national parks are in Florida?

```
(and find[park] (relate[in] lookup[florida]))
```

everglades: correct

What are some beaches in Florida?


```
(exists (and lookup[beach]
              (relate[in] lookup[florida])))
```

yes (daytona-beach): wrong parse

What beach city is there in Florida?


```
(and lookup[beach] lookup[city]
      (relate[in] lookup[florida]))
```

[none] (daytona-beach): wrong module behavior

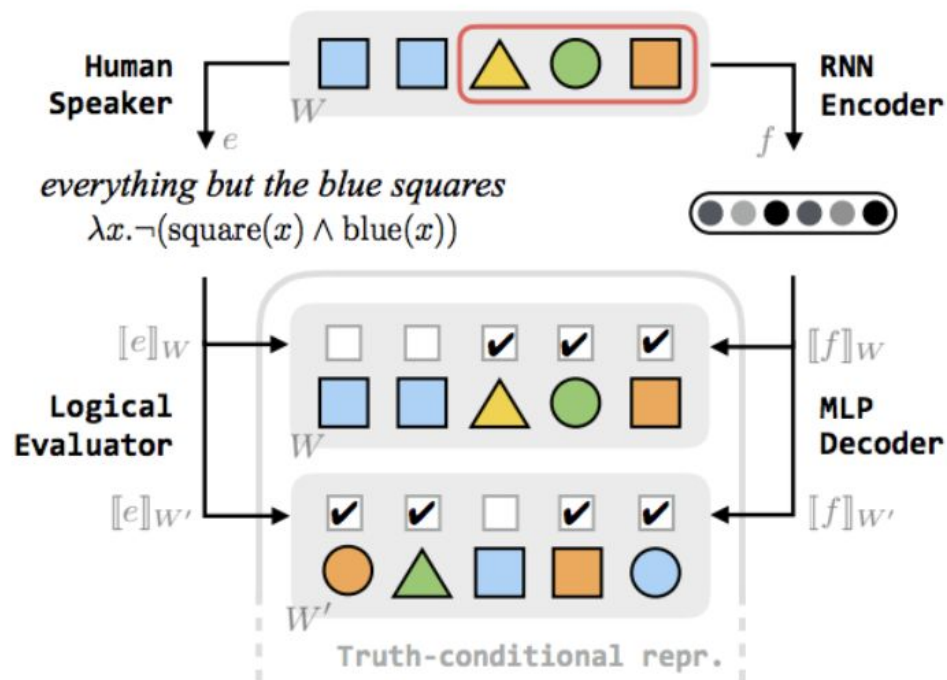


Analogs of Linguistic Structure in Deep Representations (EMNLP 2017)

- Examine message structure of an RNN trained on communication game and compare to human-produced referring expressions
- Question: are neural representations capable of spontaneously developing a “syntax” with functional analogues to properties of natural language?
- Don't care about RNN performance on task (it is nearly perfect) - but rather care about what kinds of messages the model computes in order to reach level of accuracy




Task Overview



Analogs of linguistic structure in deep representations. Jacob Andreas and Dan Klein. Figure 1.

Two- Part Question

1. High Level - Does the RNN employ a human-like communicative strategy?
2. Low Level - Do RNN representations have interpretable compositional structure?

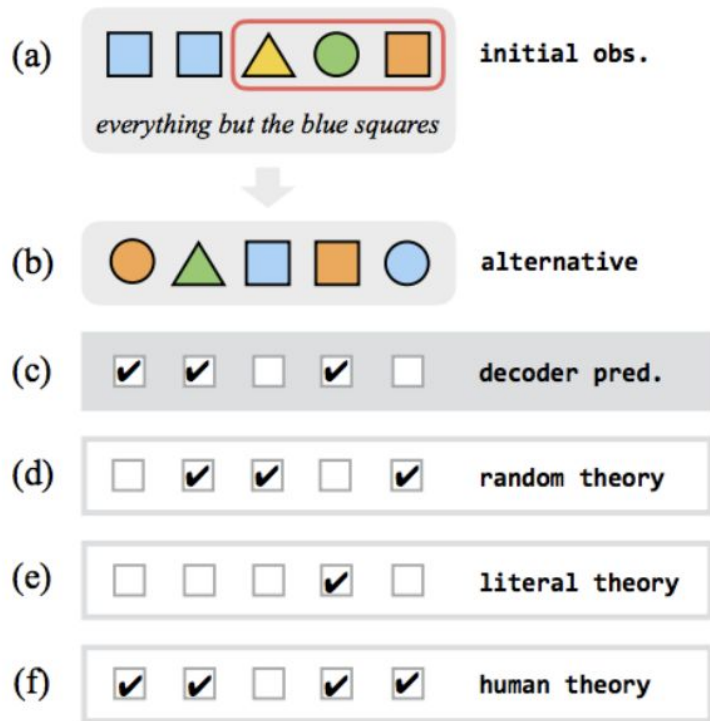


Approach

- Problem: how can you judge semantic equivalence between natural language and vector representations?
- Represent meaning via truth-condition (same if they designate same subset of objects for given game scenario)
- sample collection of scenes from dataset, and record collection of representations by each technique
 - measure agreement between techniques on individual scenes or objects



Testing Model Agreement



Results for High Level Semantics

- how often do messages generated by encoder have same meaning as messages generated by human for same context?

	Theory	Objects	Worlds	Tables
All	Random	0.50	0.00	0.00
	Literal	0.74	0.27	0.05
	Human	0.92	0.63	0.35

Analogs of linguistic structure in deep representations. Jacob Andreas and Dan Klein. Table 1.

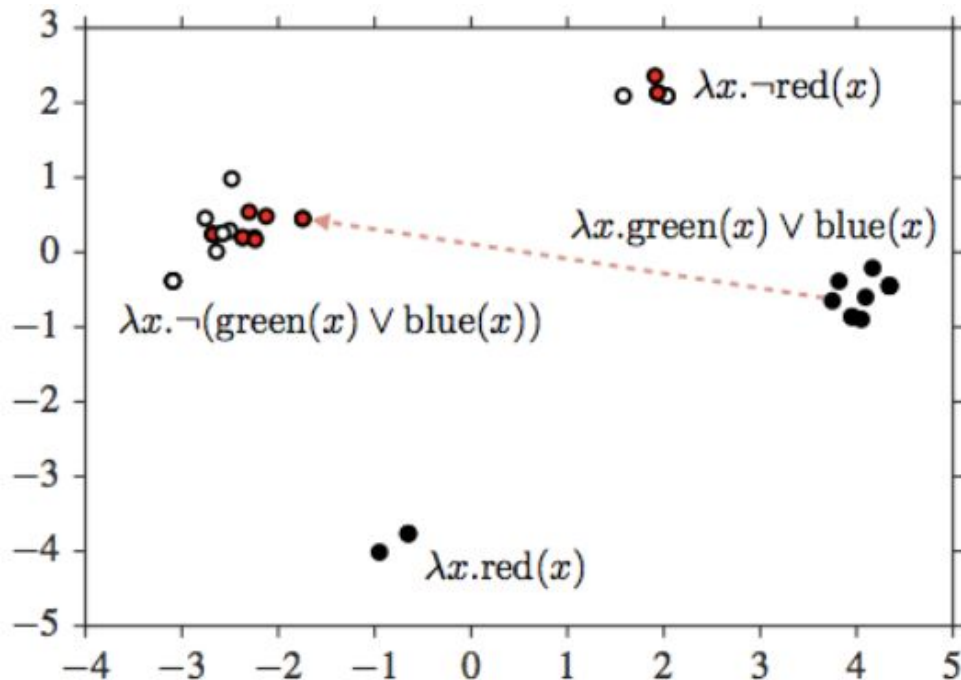
Identifying Natural Language Constructions in Vector Space

- Want to find structural regularities in vector space that can be tied to logic constructions: particularly negation, conjunction, and disjunctions
- Negation: find examples (e, f, e', f') where:
 - e' is the negation of e .
 - $\text{rep}(e) = \text{rep}(f)$ (i.e they mean the same thing)
 - $\text{rep}(e') = \text{rep}(f')$
 - Thus, can determine f' behaves like the negation of f
 - Try to find transformation N s.t $\text{rep}(Nf) = \text{rep}(e')$
 - Determine how often linear operator N actually corresponds to logical negation(f')
- Perform similar processes for logical conjunction and disjunction

Results for Identifying Low Level Structures

	Theory	Objects	Worlds	Tables
Neg.	Random	0.50	0.00	0.00
	Literal	0.50	0.12	0.03
	Negation	0.97	0.81	0.45
Disj.	Random	0.50	0.00	0.00
	Literal	0.58	0.09	0.01
	Disjunction	0.92	0.54	0.19
Conj.	Random	0.50	0.00	0.00
	Literal	0.81	0.19	0.01
	Conjunction	0.90	0.56	0.37

Analogs of linguistic structure in deep representations. Jacob Andreas and Dan Klein. Table 2.



Analogs of linguistic structure in deep representations. Jacob Andreas and Dan Klein. Figure 3.

Conclusion

- Future Work: What happens when multiple transformations are applied hierarchically?
 - extending the techniques in this paper to explore recursive structure
 - abstract from communication game: apply techniques on more general application
- Under the right conditions, RNN representations exhibit compositional structure of natural language
- highlight the usefulness of a denotational perspective from formal semantics when interpreting the behavior of deep models

