

Jacob Andreas

jda@cs.berkeley.edu

people.eecs.berkeley.edu/~jda

Education

- 2013– *Doctor of Philosophy, University of California, Berkeley*
Adviser: Dan Klein
Thesis: *Learning from language*
- 2012–2013 *Master of Philosophy, University of Cambridge*
with distinction
Adviser: Stephen Clark
Thesis: *Compositional models of natural language semantics*
- 2008–2012 *Bachelor of Science, Columbia University in the City of New York*
summa cum laude
Thesis adviser: Michael Collins
Thesis: *Toward semantic machine translation*

Recent Employment

- 2014– *Research Scientist, Semantic Machines*
- 2012 *Visiting Research Assistant, Information Sciences Institute*
- 2011 *Intern, Google (Search Quality)*
- 2010 *Intern, Microsoft (Live Labs)*
- 2009 *Intern, Lawrence Berkeley National Labs (Advanced Computing for Science)*

Fellowships, Awards & Honors

- 2016 Best paper, NAACL 2016:
Learning to compose neural networks for question answering [10]
- 2017 Best paper honorable mention, ICML 2017:
Modular multitask reinforcement learning with policy sketches [5]
- 2016–2018 Facebook Graduate Fellowship
- 2017 Huawei–Berkeley Artificial Intelligence Research Fellowship
- 2013–2016 National Science Foundation Graduate Fellowship
- 2012–2013 Winston Churchill Scholarship
- 2008–2012 C. Prescott Davis Scholarship, Columbia
- 2013 M.Phil. Dissertation Prize, Computer Laboratory, Cambridge
- 2012 Theodore R. Bashkow Prize (for computer science research), Columbia
- 2012 Russell C. Mills Prize (for computer science coursework), Columbia

Conference papers

- [1] *Learning with latent language.*
Jacob Andreas, Dan Klein and Sergey Levine. *NAACL*, 2018.
- [2] *Unified pragmatic models for generating and following instructions.*
Daniel Fried, Jacob Andreas and Dan Klein. *NAACL*, 2018.
- [3] *Learning to reason: End to end module networks for visual question answering.*
Ronghang Hu, Jacob Andreas, Marcus Rohrbach, Trevor Darrell and Kate Saenko. *ICCV*, 2017. (Spotlight presentation.)
- [4] *Analogs of linguistic structure in deep representations.*
Jacob Andreas and Dan Klein. *EMNLP*, 2017.
- [5] *Modular multitask reinforcement learning with policy sketches.*
Jacob Andreas, Dan Klein and Sergey Levine. *ICML*, 2017.
(Best paper honorable mention.)
- [6] *Translating neuralese.*
Jacob Andreas, Anca Dragan and Dan Klein. *ACL*, 2017.
- [7] *A minimal span-based constituency parser.*
Mitchell Stern, Jacob Andreas and Dan Klein. *ACL*, 2017.
- [8] *Modeling relationships in referential expressions with compositional modular networks.*
Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell and Kate Saenko. *CVPR*, 2017. (Spotlight presentation.)
- [9] *Reasoning about pragmatics with neural listeners and speakers.*
Jacob Andreas and Dan Klein. *EMNLP*, 2016.
- [10] *Learning to compose neural networks for question answering.*
Jacob Andreas, Marcus Rohrbach, Trevor Darrell and Dan Klein. *NAACL*, 2016.
(Best paper.)
- [11] *Neural module networks.*
Jacob Andreas, Marcus Rohrbach, Trevor Darrell and Dan Klein. *CVPR*, 2016.
(Oral presentation.)
- [12] *On the accuracy of self-normalized log-linear models.*
Jacob Andreas,* Maxim Rabinovich,* Dan Klein and Michael I. Jordan. *NIPS*, 2015.
- [13] *Alignment-based compositional semantics for instruction following.*
Jacob Andreas and Dan Klein. *EMNLP*, 2015.
- [14] *When and why are log-linear models self-normalizing?*
Jacob Andreas and Dan Klein. *NAACL*, 2015.

- [15] *Unsupervised transcription of piano music.*
Taylor Berg-Kirkpatrick, Jacob Andreas and Dan Klein. *NIPS*, 2014.
(Spotlight presentation.)
- [16] *Grounding language with points and paths in continuous spaces.*
Jacob Andreas and Dan Klein. *CoNLL*, 2014.
- [17] *How much do word embeddings encode about syntax?*
Jacob Andreas and Dan Klein. *ACL*, 2014.
- [18] *Semantic parsing as machine translation.*
Jacob Andreas, Andreas Vlachos and Stephen Clark. *ACL*, 2013.
- [19] *Parsing graphs with hyperedge replacement grammars.*
David Chiang, Jacob Andreas, Daniel Bauer, Karl Moritz Hermann, Bevan Jones and Kevin Knight. *ACL*, 2013.
- [20] *Semantics-based machine translation with hyperedge replacement grammars.*
Bevan Jones,* Jacob Andreas,* Daniel Bauer,* Karl Moritz Hermann,* and Kevin Knight. *COLING*, 2012.
- [21] *Annotating agreement and disagreement in threaded discussion.*
Jacob Andreas, Sara Rosenthal and Kathleen McKeown. *LREC*, 2012.
- [22] *Semi-automated annotation for prepositional phrase attachment.*
Sara Rosenthal, William Lipovsky, Kathleen McKeown, Kapil Thadani and Jacob Andreas. *LREC*, 2010.

Refereed workshop papers

- [23] *Learning to plan without a planner.*
Jacob Andreas, Mitchell Stern and Dan Klein. *NIPS—Neural Abstract Machines and Program Induction*, 2016.
- [24] *A generative model of vector space semantics.*
Jacob Andreas and Zoubin Ghahramani. *ACL—Continuous Vector Space Models and their Compositionality*, 2013.
- [25] *Detecting influencers in written online conversations.*
Or Biran, Sara Rosenthal, Jacob Andreas, Kathleen McKeown and Owen Rambow. *NAACL—Language and Social Media*, 2012.
- [26] *Fuzzy syntactic reordering for phrase-based statistical machine translation.*
Jacob Andreas, Nizar Habash and Owen Rambow. *EMNLP—Machine Translation*, 2011.
- [27] *Corpus creation for new genres: a crowdsourced approach to PP attachment.*
Mukund Jha, Jacob Andreas, Kapil Thadani, Sara Rosenthal and Kathleen McKeown. *NAACL—Creating Speech and Language Data with Mechanical Turk*, 2010.

Teaching & Advising

Teaching

- 2016 *Artificial Intelligence* (cs188), Berkeley
- 2011 *Emerging Scholars Program* (COMS 3998), Columbia

Assistant teaching

- 2016 *Artificial Intelligence* (cs188), Berkeley
- 2012 *Computer Science Theory* (COMS 3261), Columbia

Guest lectures

- 2017 *Formal semantics* in Natural Language Processing (info159), Berkeley
- 2014, 2017 *Computational semantics* in Syntax & Semantics (ling121), Berkeley
- 2016 *Language and behavior* in Alg. Human–Robot Interaction (cs294), Berkeley
- 2016 *Natural language processing* in Artificial Intelligence (cs188), Berkeley
- 2016 *Language and vision* in Computer Vision (cs280), Berkeley
- 2015 *Recurrent neural networks* in Applied NLP (info256), Berkeley
- 2014 *Grounded semantics* in Natural Language Processing (cs288), Berkeley
- 2012 *Computability, the halting problem, reductions* in CS Theory (COMS 3261), Columbia

Invited Talks & Panels

- 2018 CLASP Workshop on Dialogue and Perception
- 2018 Workshop on Visual Question Answering and Visual Dialog (CVPR)
- 2018 *Formal semantics for informal worlds*
Symposium: Perceptrons and Syntactic Structures at 60
Society for Computation in Linguistics (LSA)
- 2017 *Panelist*, Workshop on Emergent Communication (NIPS)
- 2017 *Translating neuralese*
University of Amsterdam, AI2 “NLP Highlights” podcast, Facebook Research
- 2017 *Structure and interpretation of neural codes*
Stanford
- 2016 *Modular neural architectures for perception and communication*
MIT, Harvard, Allen Institute for AI, University of Washington,
Microsoft Research, TTI Chicago, Google Research
- 2015 *Language understanding as guided planning*
Workshop on Algorithms for Human–Robot Interaction
- 2015 *Unsupervised transcription of piano music* (with T. Berg-Kirkpatrick)
Berkeley Center for New Music and Audio Technology

Professional Activities & Service

- 2014– *Reviewing: ACL, NAACL, EMNLP, NIPS, ICML, ICLR, EACL, PAMI, SCiL*
- 2017 *Organizing Committee, ACL Workshop on Language Grounding for Robotics*
- 2016 *Organizing Committee, NAACL Student Research Workshop*
- 2017– *Mentor, Berkeley AI Undergraduate Mentoring Program*
- 2014–2015 *President, Berkeley CS Graduate Student Association*
- 2010–2012 *Programming coach, 2Train Robotics—FIRST 395*