

Mohit Bansal

Assistant Professor, Computer Science, UNC Chapel Hill

201 S. Columbia St.
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-3175

Email: mbansal@cs.unc.edu
Webpage: <http://www.cs.unc.edu/~mbansal>
Google Scholar Profile

Research Interests

Statistical Natural Language Processing (NLP), Machine Learning, Multimodal Artificial Intelligence.
Current focus: Multimodal, grounded, and embodied semantics (i.e., language with vision and speech, for robotics), human-like language generation and Q&A/dialogue, and interpretable and structured deep learning.

Education

University of California, Berkeley (2008-2013)

Ph.D. in Computer Science

Thesis: Surface Web Semantics for Structured Natural Language Processing

Advisor: Dan Klein. Committee members: Dan Klein, Marti Hearst, Line Mikkelsen, Nelson Morgan

University of California, Berkeley (2012)

Master of Science (M.S.) in Computer Science

Thesis: An All-Fragments Grammar for Simple and Accurate Parsing

Advisor: Dan Klein

Indian Institute of Technology, Kanpur (2004-2008)

Bachelor of Technology (B.Tech.) in Computer Science and Engineering

GPA: 3.96/4.00 (Institute and Department Rank 2)

Cornell University (Summer 2007)

CS490 (Independent Research and Reading)

GPA: 4.00/4.00

Advisors: Lillian Lee, Claire Cardie

Honors, Awards, and Funding

DARPA Young Faculty Award (2017)

Facebook ParlAI Faculty Research Award (2017)

Outstanding Paper Award, ACL (2017)

Amazon AWS Research Cloud Credits (2017)

UNC University Research Council (URC) Small Grant Program (2017)

Google Faculty Research Award (2016)

NVIDIA Hardware Grant (2016)

UNC Junior Faculty Development Award (2016)

Best Paper Award, ACL Representation Learning for NLP Workshop (2016)

Bloomberg Data Science Research Grant (2016)

NVIDIA Paper Award, NIPS Multimodal Machine Learning Workshop (2015)

NVIDIA Hardware Grant (2015)

Google Faculty Research Award (2014)
IBM Faculty Award (2014)
Best Paper Award Honorable Mention (top-5 paper), ACL (2014)
Best Reviewer Award, NAACL (2015), EMNLP (2012)
Outstanding Graduate Student Instructor Award, UC Berkeley (2011-2012)
Qualcomm Innovation Fellowship (2011)
Tong Leong Lim Pre-Doctoral Prize, EECS, UC Berkeley (2011)
Cornell Summer Research Fellowship, CS, Cornell University (2007)
INLAKS Fellowship – Award of Excellence at IITs (2005-2008)
OPJEMS Fellowship, IIT Kanpur (2007-2008)
Academic Excellence Award, IIT Kanpur (2004-2005 and 2005-2006)
Pandit Balajee G. H. Memorial Scholarship, IIT Kanpur (2004-2005)
IIT Joint Entrance Examination All-India-Rank 47 in 200,000 (2004)

Publications

(most conferences below have a ~25% acceptance rate)

Peer-reviewed Publications:

1. Source-Target Inference Models for Spatial Instruction Understanding
Hao Tan and **Mohit Bansal**
Proceedings of **AAAI 2018**. [pdf]
2. Interactive-Length Multi-Task Video Captioning with Cooperative Feedback
Han Guo, Ramakanth Pasunuru, and **Mohit Bansal**
Proceedings of **NIPS 2017** (demo papers).
3. Reinforced Video Captioning with Entailment Rewards
Ramakanth Pasunuru and **Mohit Bansal**
Proceedings of **EMNLP 2017** (short papers). [pdf]
4. Hierarchically-Attentive RNN for Album Summarization and Storytelling
Licheng Yu, **Mohit Bansal**, and Tamara Berg
Proceedings of **EMNLP 2017** (short papers). [pdf]
5. Video Highlight Prediction Using Audience Chat Reactions
Cheng-Yang Fu, Joon Lee, Mohit Bansal, and Alexander Berg
Proceedings of **EMNLP 2017** (short papers). [pdf]
6. Shortcut-Stacked Sentence Encoders for Multi-Domain Inference
Yixin Nie and Mohit Bansal
Proceedings of **RepEval Workshop, EMNLP 2017**. [pdf]
(*Top Single Model in Shared Task*)
7. Towards Improving Abstractive Summarization via Entailment Generation
Ramakanth Pasunuru, Han Guo, and Mohit Bansal
Proceedings of **Summarization Frontiers Workshop, EMNLP 2017**. [pdf]
8. Multi-Task Video Captioning with Video and Entailment Generation
Ramakanth Pasunuru and **Mohit Bansal**
Proceedings of **ACL 2017**. [pdf]
(*Outstanding Paper Award; 1.5% accep. rate*)

9. A Joint Speaker-Listener-Reinforcer Model for Referring Expressions
Licheng Yu, Hao Tan, **Mohit Bansal**, and Tamara L. Berg
Proceedings of **CVPR 2017**. [pdf]
(Spotlight; 8.0% accep. rate)
10. Navigational Instruction Generation as Inverse Reinforcement Learning with Neural Machine Translation
Andrea F. Daniele, **Mohit Bansal**, and Matthew R. Walter
Proceedings of **HRI 2017**. [pdf]
11. Contextual RNN-GANs for Abstract Reasoning Diagram Generation
Arnab Ghosh, Viveka Kulharia, Amitabha Mukerjee, Vinay Namboodiri, and **Mohit Bansal**
Proceedings of **AAAI 2017**. [pdf]
12. Coherent Dialogue with Attention-based Language Models
Hongyuan Mei, **Mohit Bansal**, and Matthew Walter
Proceedings of **AAAI 2017**. [pdf]
13. Interpreting Neural Networks to Improve Politeness Comprehension
Malika Aubakirova and **Mohit Bansal**
Proceedings of **EMNLP 2016** (short papers). [pdf]
14. Sort Story: Sorting Jumbled Images and Captions into Stories
Harsh Agrawal, Arjun Chandrasekaran, Dhruv Batra, Devi Parikh, and **Mohit Bansal**
Proceedings of **EMNLP 2016** (short papers). [pdf]
15. Question Relevance in VQA: Identifying Non-Visual And False-Premise Questions
Arijit Ray, Gordon Christie, **Mohit Bansal**, Dhruv Batra, and Devi Parikh
Proceedings of **EMNLP 2016** (short papers). [pdf]
16. Who did What: A Large-Scale Person-Centered Cloze Dataset
Takeshi Onishi, Hai Wang, **Mohit Bansal**, Kevin Gimpel, and David McAllester
Proceedings of **EMNLP 2016** (short papers). [pdf]
17. Charagram: Embedding Words and Sentences via Character n-grams
John Wieting, **Mohit Bansal**, Kevin Gimpel, and Karen Livescu
Proceedings of **EMNLP 2016**. [pdf]
18. End-to-end Relation Extraction using LSTMs on Sequences and Tree Structures
Makoto Miwa and **Mohit Bansal**
Proceedings of **ACL 2016**. [pdf]
19. Mapping Unseen Words to Task-Trained Embedding Spaces
Pranava Swaroop Madhyastha, **Mohit Bansal**, Kevin Gimpel, and Karen Livescu
Proceedings of **Workshop on Representation Learning for NLP, ACL 2016**. [pdf]
(Best Paper Award)
20. What to talk about and how? Selective Generation using LSTMs with Coarse-to-Fine Alignment
Hongyuan Mei, **Mohit Bansal**, and Matthew R. Walter
Proceedings of **NAACL 2016**. [pdf]
21. The Role of Context Types and Dimensionality in Learning Word Embeddings
Oren Melamud, David McClosky, Siddharth Patwardhan, and **Mohit Bansal**
Proceedings of **NAACL 2016**. [pdf]
22. We Are Humor Beings: Understanding and Predicting Visual Humor
Arjun Chandrasekaran, Ashwin Kalyan, Stanislaw Antol, **Mohit Bansal**, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh
Proceedings of **CVPR 2016**. [pdf]
(Spotlight; 9.7% accep. rate)

23. Towards Universal Paraphrastic Sentence Embeddings
John Wieting, **Mohit Bansal**, Kevin Gimpel, and Karen Livescu
Proceedings of **ICLR 2016**. [pdf]
(Oral; 5.7% accep. rate)
24. Listen, Attend, and Walk: Neural Mapping of Navigational Instructions to Action Sequences
Hongyuan Mei, **Mohit Bansal**, and Matthew R. Walter
Proceedings of **AAAI 2016**. [pdf]
(Nvidia Paper Award in NIPS 2015 Multimodal Machine Learning workshop)
25. Machine Comprehension with Syntax, Frames, and Semantics
Hai Wang, **Mohit Bansal**, Kevin Gimpel, and David McAllester
Proceedings of **ACL 2015** (short papers). [pdf]
26. From Paraphrase Database to Compositional Paraphrase Model and Back
John Wieting, **Mohit Bansal**, Kevin Gimpel, and Karen Livescu
Proceedings of **TACL** (presented at **EMNLP 2015**). [pdf]
27. Dependency Link Embeddings: Continuous Representations of Syntactic Substructures
Mohit Bansal
Proceedings of **Workshop on Vector Space Modeling for NLP, NAACL 2015**. [pdf]
(selected oral)
28. Deep Multilingual Correlation for Improved Word Embeddings
Ang Lu, Weiran Wang, **Mohit Bansal**, Kevin Gimpel, and Karen Livescu
Proceedings of **NAACL 2015** (short papers). [pdf]
29. A Sense-Topic Model for Word Sense Induction with Unsupervised Data Enrichment
Jing Wang, **Mohit Bansal**, Kevin Gimpel, Brian Ziebart, and Clement Yu
Proceedings of **TACL** (presented at **NAACL 2015**). [pdf]
30. Weakly-Supervised Learning with Cost-Augmented Contrastive Estimation
Kevin Gimpel and **Mohit Bansal**
Proceedings of **EMNLP 2014**. [pdf]
31. Tailoring Continuous Word Representations for Dependency Parsing
Mohit Bansal, Kevin Gimpel, and Karen Livescu
Proceedings of **ACL 2014** (short papers). [pdf]
32. Structured Learning for Taxonomy Induction with Belief Propagation
Mohit Bansal, David Burkett, Gerard de Melo, and Dan Klein
Proceedings of **ACL 2014**. [pdf]
(Best Paper Award Honorable Mention – top-5 paper)
33. What are you talking about? Text-to-Image Coreference
Chen Kong, Dahua Lin, **Mohit Bansal**, Raquel Urtasun, and Sanja Fidler
Proceedings of **CVPR 2014**. [pdf]
34. Good, Great, Excellent: Global Inference of Semantic Intensities
Gerard de Melo and **Mohit Bansal**
Proceedings of **TACL** (presented at **ACL 2013**). [pdf]
35. Coreference Semantics from Web Features
Mohit Bansal and Dan Klein
Proceedings of **ACL 2012**. [pdf]
36. Unsupervised Translation Sense Clustering
Mohit Bansal, John DeNero, and Dekang Lin
Proceedings of **NAACL 2012**. [pdf]

37. Web-Scale Features for Full-Scale Parsing
Mohit Bansal and Dan Klein
Proceedings of **ACL 2011**. [pdf]
38. Gappy Phrasal Alignment by Agreement
Mohit Bansal, Chris Quirk, and Robert C. Moore
Proceedings of **ACL 2011**. [pdf]
39. The Surprising Variance in Shortest-Derivation Parsing
Mohit Bansal and Dan Klein
Proceedings of **ACL 2011** (short papers). [pdf]
40. Mention Detection: Heuristics for the OntoNotes annotations
Jonathan K. Kummerfeld, **Mohit Bansal**, David Burkett, and Dan Klein
Proceedings of **CoNLL 2011** (shared task). [pdf]
41. Simple, Accurate Parsing with an All-Fragments Grammar
Mohit Bansal and Dan Klein
Proceedings of **ACL 2010**. [pdf]
42. Efficient Parsing for Transducer Grammars
John DeNero, **Mohit Bansal**, Adam Pauls, and Dan Klein
Proceedings of **NAACL 2009**. [pdf]
43. The power of negative thinking: Exploiting label disagreement in the min-cut classification framework
Mohit Bansal, Claire Cardie, and Lillian Lee
Proceedings of **COLING 2008** (short papers). [pdf]
44. Estimating hybrid frequency moments of data streams
Sumit Ganguly, **Mohit Bansal**, and Shruti Dube
Proceedings of **FAW 2008**, LNCS 5059, pp. 55-66.
Also in the Journal of Combinatorial Optimization (**JOCO**). [pdf]
45. Text Processing for Text to Speech Systems in Indian Languages
Anand Raj, Tanuja Sarkar, Satish Pammi, Santhosh Yuvaraj, **Mohit Bansal**, SP Kishore, and Alan W Black
Proceedings of **ISCA SSW6 2007**. [pdf]

Theses:

1. Surface Web Semantics for Structured Natural Language Processing
Mohit Bansal
Ph.D. Thesis. EECS, UC Berkeley. Committee: Dan Klein (advisor), Marti Hearst, Line Mikkelsen, Nelson Morgan. [pdf]
2. An All-Fragments Grammar for Simple and Accurate Parsing
Mohit Bansal
M.S. Thesis. EECS, UC Berkeley. Advisor: Dan Klein. [pdf]

Patents:

1. Techniques for Generating Translation Clusters
John DeNero and **Mohit Bansal** (Google Research)
Publication number: US20130275118 A1 (Oct 17, 2013).

Teaching

Instructor, Advanced Topics in Natural Language Processing: Grounded Language for Robotics (COMP 790.139), UNC Chapel Hill, Spring 2017.

Instructor, Natural Language Processing (COMP 790.139), UNC Chapel Hill, Fall 2016.

Guest Lecturer, Computational Linguistics (CMSC 35050, *Instructor*: John Goldsmith), University of Chicago, Spring 2015 – ‘Automatic Taxonomy Induction’.

Guest Lecturer, Robotics and Artificial Intelligence (TTIC 31170, *Instructor*: Matthew Walter), TTI-Chicago, University of Chicago, Spring 2015 – ‘Automatic Taxonomy Induction’.

Guest Lecturer, Visual Recognition with Text (CSC 2523, *Instructor*: Sanja Fidler), University of Toronto, Winter 2015 – short course on ‘Topics, Trends, and Resources in NLP’ [slides].

GSI, Introduction to Artificial Intelligence (CS188, *Instructor*: Dan Klein), UC Berkeley, Fall 2011. Received an Outstanding Graduate Student Instructor Award by UC Berkeley for excellence in teaching.

GSI, Advanced Topics in Artificial Intelligence (CS194-10, *Instructors*: Pieter Abbeel, Dan Klein, Jitendra Malik), UC Berkeley, Spring 2009. Sole TA for new course with 30 advanced students.

Students/Interns

UNC Advisees:

Robbie Allen (UNC, PhD)
Lisa Bauer (UNC, PhD)
Yen-Chun Chen (UNC, MS)
Han Guo (UNC, BS)
Sweta Karlekar (UNC, BS)
Jie Lei (UNC, PhD; co-advised with Tamara Berg)
Yixin Nie (UNC, PhD)
Ramakanth Pasunuru (UNC, PhD)
Hao Tan (UNC, PhD)
Licheng Yu (UNC, PhD; co-advised with Tamara Berg)

Other Advisees:

Arjun Chandrasekaran (Georgia Tech, PhD; advisor = Devi Parikh) (PhD Thesis Committee Member)
Tong Niu (Duke, MS)

Past Advisees:

Malika Aubakirova (UChicago, BS)
Dhivya Eswaran (IIT-Madras, BTech → CMU, PhD)
Rasool Fakoor (UT-Arlington, PhD → MSR)
Arnab Ghosh (IIT Kanpur, BTech → Oxford, PhD)
Yuchen He (UIUC, PhD)
Myungin Kim (UChicago, MS)
Zuyao Li (USC, MS → Google/Nest)
Ang Lu (Tsinghua, BS → CMU, MS)
Pranava S. Madhyastha (UPC Barcelona, PhD)
Hongyuan Mei (UChicago/TTIC, MS → JHU, PhD) (MS Thesis Co-Advisor)
Aravind L Srinivas (IIT Madras, BTech → UC Berkeley, PhD)
Ryan Stout (UIUC, MS)
Trang Tran (UWash, PhD)
Jing Wang (UIC, PhD → Conversant, Scientist)
John Wieting (UIUC/TTIC, MS → CMU, PhD)
Zhengyang Wu (GeorgiaTech, BS)

Professional Service

Tutorial Co-chair: NAACL 2018

Area Co-chair (Vision, Robotics, and Grounding): ACL 2017

Area Co-chair (Machine Learning): EMNLP 2017

Demonstration Co-chair: ACL 2017

Tutorial Co-chair: NAACL 2016

Area Co-chair: NAACL 2016

Program Committee Member/Reviewer:

Conferences: EMNLP (best reviewer award in 2012; top 9%), NAACL (best reviewer award in 2015; top 25%), ACL, NIPS, ICLR, IJCAI, EACL, COLING, *SEM, IJCNLP, ICON

Journals: TACL, TPAMI, TALIP

Recent Workshops: ACL Workshop of Women in Natural Language Processing (2017), ACL Workshop on Representation Learning for NLP (2017), EACL Workshop on Ethics in Natural Language Processing (2017), NAACL Multilingual and Crosslingual Methods in NLP (2016), NAACL Human-Computer Question Answering (2016), ACL Evaluating Vector-Space Representations for NLP (2016), NAACL Vector Space Modeling for NLP (2015).

University Research Proposals: ORAU

Organizer: NLP/ML Colloquium Series at UNC

Organizer: Workshop on Language Grounding for Robotics at ACL 2017

Organizer: Midwest Speech and Language Days 2015

Committee Member: Graduate (PhD) Admissions Committee, CS, UNC Chapel Hill

Committee Member: Graduate (PhD) Admissions Committee, EECS, UC Berkeley

Panel Member: National Science Foundation (NSF) Review Panels

Software and Datasets: Available for various papers on my webpage: <http://ttic.uchicago.edu/~mbansal>

Recent Invited Talks

‘Multi-Task and Reinforcement Learning for Entailment-Based Natural Language Generation’, *JHU, Bloomberg, Google, Facebook, UIUC* (May-Oct, 2017)

‘Structured Learning of World Knowledge for Natural Language Semantics’, *CMU, MSR, Rutgers, UC Davis, UC Irvine, UNC Chapel Hill, UT Austin, Virginia Tech* (Feb-Apr, 2016)

‘Neural Attention Models for Natural Language Grounding and Generation’, *IIT-Delhi, IIT-Kanpur* (Sep-Oct, 2015)

‘Improving Neural Embeddings via Paraphrase, Translational, and Syntactic Knowledge’, *Columbia University, Google Research, NYU* (Apr, 2015)

‘Semantic World Knowledge for NLP’, *UToronto, UMichigan, Virginia Tech*, (Nov-Jan, 2015)

Research and Work Experience

UNC Chapel Hill, Computer Science Dept. (2016 – present)

Assistant Professor

Toyota Technological Institute at Chicago (2013 – 2016)

Research Assistant Professor (3-year endowed position)

IARPA Babel Project, Swordfish team (*Feb 2014 – June 2014*)
Member/Consultant

University of California, Berkeley (*2008 – 2013*)
Graduate Student Researcher (Advisor: Dan Klein)

Google Research, Mountain View (*Summer 2011*)
Research Intern (with John DeNero and Dekang Lin)

Microsoft Research, Redmond (*Summer 2010*)
Research Intern (with Chris Quirk and Bob Moore)

Cornell University, CS division (*Summer 2007*)
Research Intern (with Lillian Lee and Claire Cardie)