Mohit Bansal

John R. & Louise S. Parker Distinguished Professor
Computer Science, UNC Chapel Hill
Director, MURGe-Lab (UNC-NLP Group)

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

Natural language processing, computer vision and multimodal machine learning/AI, with a particular focus on multimodal generative models, reasoning and planning agents, faithful language generation, and interpretable, efficient, and generalizable 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)*
GPA: 4.00/4.00
Advisors: Lillian Lee, Claire Cardie

Research and Work Experience

UNC Chapel Hill, Computer Science Dept. *(2016 – present)*
John R. & Louise S. Parker Distinguished Professor
Director, MURGe Lab (UNC-NLP Group)

Toyota Technological Institute at Chicago *(2013 – 2016)*
Research Assistant Professor (3-year endowed position)

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

Honors, Awards, and Funding

AAAI Fellow *(2025)*

ACL Fellow *(2025)*

Presidential Early Career Award for Scientists and Engineers (PECASE) *(2025)*
(*nominated in 2019, and awarded ARO-ECASE funding in 2021 as a 'bridge program' due to White House PECASE announcement delays*)

TMLR Outstanding Paper Finalist *(2024)*

Provost's Kenan Senior Faculty Research and Scholarly Leave Award *(2023-2024)*

John R. & Louise S. Parker Permanent Distinguished Professorship *(2024-retirement)*

IIT Kanpur Young Alumni Award *(2023)*

Keynote Speaker, AACL-IJCNLP 2023, CoNLL 2023, INLG 2022, SouthNLP 2023, LxMLS 2024, IASNLP 2024

USC Distinguished Lecture *(2024)*

Provost's Kenan Senior Faculty Research and Scholarly Leave Award *(2023)*

Early Career Award for Scientists and Engineers (ECASE-Army) *(2021)*

CVPR Best Student Paper Honorable Mention *(2021)*

CoNLL Best Paper Runner-Up Award *(2021)*

EACL Best Long Paper Award Honorable Mention *(2021)*

UNC Phillip and Ruth Hettleman Prize for Artistic and Scholarly Achievement *(2020)*

John R. & Louise S. Parker Faculty Fellow/Scholar (Distinguished Professorship) *(2020)*

IJCAI Early Career Spotlight *(2020)* (previous years: 2016, 2017, 2018, 2019)

DARPA Director's Fellowship *(2019)*

Microsoft Investigator Fellowship *(2019)*

Amazon Machine Learning Research Award *(2019)*

NSF CAREER Award *(2019)*

Google Focused Research Award *(2019)*

ACL Best Short Paper Nomination *(2019)*

Salesforce Research Deep Learning Grant *(2018)*

Facebook Faculty Research Award *(2018)*

IBM Faculty Award *(2018)*

Army Research Office Young Investigator Award (ARO-YIP) *(2018)*

'Area Chair Favorites' Paper Award, COLING *(2018)*

Adobe Faculty Research Award *(2018)*

Verisk AI Faculty Research Award *(2018, 2019)*

DARPA Young Faculty Award (DARPA-YFA) *(2017)*

Facebook ParlAI Faculty Research Award *(2017)*

Outstanding Paper Award, ACL *(2017)*

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

Google Faculty Research Award *(2016)*

NVidia Hardware Grant *(2016, 2017, 2018)*

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)*

Google Faculty Research Award *(2014)*

IBM Faculty Award *(2014)*

Best Paper Award Honorable Mention (top-5 paper), ACL *(2014)*

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)*

Honors/Awards/Fellowships of Students:

Google PhD Fellowship, 2025 (Vaidehi Patil)

NDSEG Fellowship, 2025 (Zaid Khan)

Apple AI/ML PhD Fellowship, 2024 (Archiki Prasad)

Google PhD Fellowship, 2024 (David Wan)

Bloomberg PhD Fellowship, 2024 (David Wan) (declined)

CRA Outstanding Undergraduate Researcher Award Winner 2023 (Zineng Tang)

Bloomberg Research PhD Fellowship, 2023 (Jaemin Cho)

Google PhD Fellowship, 2022 (Swarnadeep Saha)

Apple AI/ML PhD Fellowship, 2022 (Yichen Jiang)

Google Research PhD Fellowship, 2021 (Peter Hase)

Bloomberg Research PhD Fellowship, 2021 (Shiyue Zhang)

Adobe Research Fellowship, 2021 (Jie Lei)

Microsoft Research PhD Fellowship, 2019 (Ramakanth Pasunuru)

Facebook PhD Fellowship Finalist, 2019 (Ramakanth Pasunuru)

Bloomberg Data Science PhD Fellowship, 2019 (Hao Tan)

NSF Graduate Research Fellowship, 2018 (Lisa Bauer)

NSF Graduate Research Fellowship, 2019 (Darryl Hannan)

Royster Society Kenan Fellowship, 2019 (Peter Hase)

CRA Outstanding Undergraduate Researcher Award Runner-Up, 2020 (Sweta Karlekar)

CRA Outstanding Undergraduate Researcher Award Finalist, 2020 (Han Guo)

CRA Outstanding Undergraduate Researcher Award Honorable Mention, 2021 (Zineng Tang, junior)

CRA Outstanding Undergraduate Researcher Award Honorable Mention, 2021 (Abhay Zala, junior)

CRA Outstanding Undergraduate Researcher Award Honorable Mention, 2019 (Han Guo, junior)

Other Funding/Grants:

NIH Advancing Health Research through Multimodal AI.(UNC Co-PI)

DARPA Environment-driven Conceptual Learning (ECOLE).(UNC PI)

ONR Science of Artificial Intelligence – Basic and Applied Research for the Naval Domain (Overall PI)

NSF-AI Institute on Engaged Learning (Core AI Lead)

NSF Future of Work at the Human-Technology Frontier (UNC Co-PI)

ONR Advancing Artificial Intelligence for the Naval Domain (UNC PI)

DARPA Machine Common Sense (MCS) (UNC PI)

DARPA Knowledge-directed Artificial Intelligence Reasoning Over Schemas (KAIROS) (UNC PI)

NSF-NIH SCH AURA Connecting Audio and Radio Sensing Systems to Improve Care at Home (UNC Co-PI)

Professional Service

Associate Editor-in-Chief: IEEE TPAMI Journal

Program Chair: EMNLP 2024

Member: ACL Executive Committee (2022-2024)

Member: ACM Doctoral Dissertation Award Committee (2021-2024)

Co-Organizer/Founder: ACL Doctoral Dissertation Award

Co-Organizer/Founder: ACL Mentoring Program

Action Editor: TACL Journal

Action Editor: Computation Linguistics (CL) Journal

Associate Editor: IEEE/ACM Transactions on Audio Speech and Language Processing (TASLP)

Editorial Board: Computer Speech and Language Journal

Americas Sponsorship Co-Chair for the ACL: 2020-2023

Senior Area Chair: ARR (ACL Rolling Review)

Senior Area Chair: EACL 2024

Senior Area Chair: EMNLP 2023

Senior Area Chair: IJCAI 2023

Senior Area Chair: AAAI 2023

Senior Area Chair: ACL 2022

Area Chair: ICLR 2021

Senior Area Chair: ACL 2021

Senior Area Chair: NAACL 2021

Senior Area Chair: AAAI 2021

Area Chair: IJCAI 2021

Senior Area Chair: EMNLP 2020

Senior Area Chair: ACL 2020

Program Co-Chair: CoNLL 2019

Senior Program Committee Member: AAAI 2020

Area Chair: NAACL 2019

Area Chair: EMNLP 2018

Tutorial Chair: NAACL 2018

Area Chair (Vision, Robotics, and Grounding): ACL 2017

Area Chair (Machine Learning): EMNLP 2017

Demonstration Chair: ACL 2017

Tutorial Chair: NAACL 2016

Area Chair: NAACL 2016

University Research Proposals Reviewer: NSF, ONR, ARO, ORAU

Organizer: EMNLP 2020 Workshop on Spatial Language Understanding (SpLU)

Organizer: NAACL 2019 Joint Workshop on Spatial Language Understanding & Language Grounding for Robotics (SpLU-RoboNLP)

Organizer: CVPR 2019 Workshop on Conceptual Captions

Organizer: NLP/ML Colloquium Series at UNC

Organizer: ACL 2017 Workshop on Language Grounding for Robotics (RoboNLP)

Organizer: Midwest Speech and Language Days 2015

Committee Member: Faculty Hiring Committee, CS, UNC Chapel HIll

Software and Datasets: Available for various papers on our lab's webpage: https://murgelab.cs.unc.edu/software.html

Publications

Citations = 43687; h-index = 100; i10-index=305

(All code/data available for various papers on our lab's webpage: https://murgelab.cs.unc.edu/software.html)

(Topic-based publication lists for Multimodal AI, NLG, Interpretability/Explainability, Efficiency, etc. available at: http://www.cs.unc.edu/ mbansal/#publications)

Peer-reviewed Publications:

340. Bifrost-1: Bridging Multimodal LLMs and Diffusion Models with Patch-level CLIP Latents
Han Lin, Jaemin Cho, Amir Zadeh, Chuan Li, Mohit Bansal
Proceedings of **NeurIPS 2025**. [pdf]

339. LASeR: Learning to Adaptively Select Reward Models with Multi-Armed Bandits
Duy Nguyen*, Archiki Prasad*, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **NeurIPS 2025**. [pdf]

338. 4D-LRM: Large Space-Time Reconstruction Model From and To Any View at Any Time
Ziqiao Ma, Xuweiyi Chen, Shoubin Yu, Sai Bi, Kai Zhang, Chen Ziwen, Sihan Xu, Jianing Yang, Zexiang Xu, Kalyan Sunkavalli, Mohit Bansal, Joyce Chai, Hao Tan
Proceedings of **NeurIPS 2025**. [pdf]

337. ReAgent-V: A Reward-Driven Multi-Agent Framework for Video Understanding
Yiyang Zhou, Yangfan He, Yaofeng Su, Siwei Han, Joel Jang, Gedas Bertasius, Mohit Bansal, Huaxiu Yao
Proceedings of **NeurIPS 2025**. [pdf]

336. MAgICoRe: Multi-Agent, Iterative, Coarse-to-Fine Refinement for Reasoning
Justin Chih-Yao Chen, Archiki Prasad, Swarnadeep Saha, Elias Stengel-Eskin, Mohit Bansal

Proceedings of **EMNLP 2025**. [pdf]

335. Video-RTS: Rethinking Reinforcement Learning and Test-Time Scaling for Efficient and Enhanced Video Reasoning
Ziyang Wang*, Jaehong Yoon*, Shoubin Yu, Md Mohaiminul Islam, Gedas Bertasius, Mohit Bansal
Proceedings of **EMNLP 2025**. [pdf]

334. Language Models Identify Ambiguities and Exploit Loopholes
Jio Choi, Mohit Bansal, Elias Stengel-Eskin
Proceedings of **EMNLP 2025**. [pdf]

333. RACCooN: Remove, Add, and Change Video Content with Auto-Generated Narratives
Jaehong Yoon*, Shoubin Yu*, Mohit Bansal
Proceedings of **EMNLP 2025**. [pdf]

332. Glider: Global and Local Instruction-Driven Expert Router
Pingzhi Li*, Prateek Yadav*, Jaehong Yoon, Jie Peng, Yi-Lin Sung, Mohit Bansal, Tianlong Chen
Proceedings of **EMNLP 2025**. [pdf]

331. MEXA: Towards General Multimodal Reasoning with Dynamic Multi-Expert Aggregation
Shoubin Yu*, Yue Zhang*, Ziyang Wang, Jaehong Yoon, Mohit Bansal
Findings of **EMNLP 2025**. [pdf]

330. Video-Skill-CoT: Skill-based Chain-of-Thoughts for Domain-Adaptive Video Reasoning Daeun Lee*, Jaehong Yoon*, Jaemin Cho, Mohit Bansal
Findings of **EMNLP 2025. (short)**. [pdf]

329. FLAMES: Improving LLM Math Reasoning via a Fine-Grained Analysis of the Data Synthesis Pipeline
Parker Seegmiller, Kartik Mehta, Soumya Saha, Chenyang Tao, Shereen Oraby, Arpit Gupta, Tagyoung Chung, Mohit Bansal, Nanyun Peng
Findings of **EMNLP 2025**. [pdf]

328. Localizing Factual Inconsistencies in Attributable Text Generation
Arie Cattan, Paul Roit, Shiyue Zhang, David Wan, Roee Aharoni, Idan Szpektor, Mohit Bansal, Ido Dagan
Proceedings of **TACL 2025**. [pdf]

327. Reliable and Responsible Foundation Models
Xinyu Yang, Junlin Han, Rishi Bommasani, ..., Pang Wei Koh, Yulia Tsvetkov, Andrew Gordon Wilson, Jiaheng Zhang, James Zou, Cihang Xie, Hao Wang, Philip Torr, Julian McAuley, David Alvarez-Melis, Florian Tramèr, Kaidi Xu, Suman Jana, Chris Callison-Burch, Rene Vidal, Filippos Kokkinos, Mohit Bansal, Beidi Chen, Huaxiu Yao
Proceedings of **TMLR 2025**. [pdf]

326. Learning to Generate Unit Tests for Automated Debugging
Archiki Prasad, Elias Stengel-Eskin, Justin Chih-Yao Chen, Zaid Khan, Mohit Bansal
Proceedings of **COLM 2025**. [pdf]

325. Retrieval-Augmented Generation with Conflicting Evidence
Han Wang, Archiki Prasad, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **COLM 2025**. [pdf]

324. Task-Circuit Quantization: Leveraging Knowledge Localization and Interpretability for Compression
Hanqi Xiao, Yi-Lin Sung, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **COLM 2025**. [pdf]

323. GenerationPrograms: Fine-grained Attribution with Executable Programs
David Wan, Eran Hirsch, Elias Stengel-Eskin, Ido Dagan, Mohit Bansal
Proceedings of **COLM 2025**. [pdf]

322. QAPyramid: Fine-grained Evaluation of Content Selection for Text Summarization
Shiyue Zhang*, David Wan*, Arie Cattan, Ayal Klein, Ido Dagan, Mohit Bansal
Proceedings of **COLM 2025**. [pdf]

321. A Multimodal Classroom Video Question-Answering Framework for Automated Understanding of Collaborative Learning
Nithin Sivakumaran*, Chia-Yu Yang*, Abhay Zala*, Shoubin Yu, Daeun Hong, Xiaotian Zou, Elias Stengel-Eskin, Dan Carpenter, Wookhee Min, Cindy Hmelo-Silver, Jonathan Rowe, James Lester, Mohit Bansal
Proceedings of **ICMI 2025**. [pdf]

320. CAPTURe: Evaluating Spatial Reasoning in Vision Language Models via Occluded Object Counting
Atin Pothiraj, Elias Stengel-Eskin, Jaemin Cho, Mohit Bansal
Proceedings of **ICCV 2025**. [pdf]

319. VEGGIE: Instructional Editing and Reasoning of Video Concepts with Grounded Generation
Shoubin Yu*, Difan Liu*, Ziqiao Ma*, Yicong Hong, Yang Zhou, Hao Tan, Joyce Chai, Mohit Bansal
Proceedings of **ICCV 2025**. [pdf]

318. SAME: Learning Generic Language-Guided Visual Navigation with State-Adaptive Mixture of Experts
Gengze Zhou, Yicong Hong, Zun Wang, Chongyang Zhao, Mohit Bansal, Qi Wu
Proceedings of **ICCV 2025**. [pdf]

317. M3DocRAG: Multi-modal Retrieval is What You Need for Multi-page Multi-document Understanding
Jaemin Cho, Debanjan Mahata, Ozan Irsoy, Yujie He, Mohit Bansal
Proceedings of **ICCV Findings 2025**. [pdf]

316. What Matters for Model Merging at Scale?
Prateek Yadav, Tu Vu, Jonathan Lai, Alexandra Chronopoulou, Manaal Faruqui, Mohit Bansal, Tsendsuren Munkhdalai
Proceedings of **TMLR 2025**. [pdf]

315. Multi-Attribute Steering of Language Models via Targeted Intervention
Duy Nguyen, Archiki Prasad, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **ACL 2025**. [pdf]

314. LAQuer: Localized Attribution Queries in Content-grounded Generation
Eran Hirsch, Aviv Slobodkin, David Wan, Elias Stengel-Eskin, Mohit Bansal, Ido Dagan
Proceedings of **ACL 2025**. [pdf]

313. Self-Consistency Preference Optimization
Archiki Prasad, Weizhe Yuan, Richard Yuanzhe Pang, Jing Xu, Maryam Fazel-Zarandi, Mohit Bansal, Sainbayar Sukhbaatar, Jason Weston, Jane Yu
Proceedings of **ICML 2025**. [pdf]

312. Generative AI Unlocks the Power of Interactive Storytelling for Science Teachers and Learners
Jeremy Roschelle, Mohit Bansal, Gautam Biswas, Cindy Hmelo-Silver, James Lester
Proceedings of **Social Innovations Journal, 2025**. [pdf]

311. VideoTree: Adaptive Tree-based Video Representation for LLM Reasoning on Long Videos
Ziyang Wang*, Shoubin Yu*, Elias Stengel-Eskin*, Jaehong Yoon, Feng Cheng, Gedas Bertasius, Mohit Bansal
Proceedings of **CVPR 2025**. [pdf]

310. Motion-Grounded Video Reasoning: Understanding and Perceiving Motion at Pixel Level
Andong Deng, Tongjia Chen, Shoubin Yu, Taojiannan Yang, Lincoln Spencer, Yapeng Tian, Ajmal Saeed Mian, Mohit Bansal, Chen Chen
Proceedings of **CVPR 2025**. [pdf]

309. REAL Sampling: Boosting Factuality and Diversity of Open-Ended Generation via Asymptotic Entropy
Haw-Shiuan Chang, Nanyun Peng, Mohit Bansal, Anil Ramakrishna, Tagyoung Chung
Proceedings of **TACL 2025**. [pdf]

308. A Survey on Model MoErging: Recycling and Routing Among Specialized Experts for Collaborative Learning
Prateek Yadav*, Colin Raffel*, Mohammed Muqeeth, Lucas Caccia, Haokun Liu, Tianlong Chen, Mohit Bansal, Leshem Choshen, Alessandro Sordoni
Proceedings of **TMLR 2025**. [pdf]

307. ComPEFT: Compression for Communicating Parameter Efficient Updates via Sparsification and Quantization
Prateek Yadav, Leshem Choshen, Colin Raffel, Mohit Bansal
Proceedings of **TMLR 2025**. [pdf]

306. DataEnvGym: Data Generation Agents in Teacher Environments with Student Feedback
Zaid Khan, Elias Stengel-Eskin, Jaemin Cho, Mohit Bansal
Proceedings of **ICLR 2025**. [pdf] ***(spotlight, top 5%)***

305. Ctrl-Adapter: An Efficient and Versatile Framework for Adapting Diverse Controls to Any Diffusion Model
Han Lin*, Jaemin Cho*, Abhay Zala, Mohit Bansal
Proceedings of **ICLR 2025**. [pdf] ***(oral, top 1.8%)***

304. System-1.x: Learning to Balance Fast and Slow Planning with Language Models
Swarnadeep Saha, Archiki Prasad, Justin Chih-Yao Chen, Peter Hase, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **ICLR 2025**. [pdf]

303. CREMA: Multimodal Compositional Video Reasoning via Efficient Modular Adaptation and Fusion
Shoubin Yu*, Jaehong Yoon*, Mohit Bansal
Proceedings of **ICLR 2025**. [pdf]

302. Adapt-inf: Scalable Lifelong Multimodal Instruction Tuning via Dynamic Data Selection
Adyasha Maharana*, Jaehong Yoon*, Tianlong Chen, Mohit Bansal
Proceedings of **ICLR 2025**. [pdf]

301. SAFREE: Training-Free and Adaptive Guard for Safe Text-to-Image And Video Generation
Jaehong Yoon*, Shoubin Yu*, Vaidehi Patil, Huaxiu Yao, Mohit Bansal
Proceedings of **ICLR 2025**. [pdf]

300. Unbounded: A Generative Infinite Game of Character Life Simulation
Jialu Li, Yuanzhen Li, Neal Wadhwa, Yael Pritch, David E. Jacobs, Michael Rubinstein, Mohit Bansal, Nataniel Ruiz
Proceedings of **ICLR 2025**. [pdf]

299. VEDIT: Latent Prediction Architecture For Procedural Video Representation Learning
Han Lin, Tushar Nagarajan, Nicolas Ballas, Mido Assran, Mojtaba Komeili, Mohit Bansal, Koustuv Sinha
Proceedings of **ICLR 2025**. [pdf]

298. Bootstrapping Language-Guided Navigation Learning with Self-Refining Data Flywheel
Zun Wang, Jialu Li, Yicong Hong, Songze Li, Kunchang Li, Shoubin Yu, Yi Wang, Yu Qiao, Yali Wang, Mohit Bansal, Limin Wang
Proceedings of **ICLR 2025**. [pdf]

297. AnyPrefer: An Automatic Framework for Preference Data Synthesis
Yiyang Zhou*, Zhaoyang Wang*, Tianle Wang*, Shangyu Xing, Peng Xia, Bo Li, Kaiyuan Zheng, Zijian Zhang, Zhaorun Chen, Wenhao Zheng, Xuchao Zhang, Chetan Bansal, Weitong Zhang, Ying Wei, Mohit Bansal, Huaxiu Yao
Proceedings of **ICLR 2025**. [pdf]

296. See It from My Perspective: Diagnosing the Western Cultural Bias of Large Vision-Language Models in Image Understanding
Amith Ananthram, Elias Stengel-Eskin, Mohit Bansal, Kathleen McKeown
Proceedings of **ICLR 2025**. [pdf]

295. Teaching Models to Balance Resisting and Accepting Persuasion Elias Stengel-Eskin, Peter Hase, Mohit Bansal
Proceedings of **NAACL 2025**. [pdf]

294. AdaCAD: Adaptively Decoding to Balance Conflicts between Contextual and Parametric Knowledge
Han Wang, Archiki Prasad, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **NAACL 2025**. [pdf]

293. Reverse Thinking Makes LLMs Stronger Reasoners
Justin Chih-Yao Chen, Zifeng Wang, Hamid Palangi, Rujun Han, Sayna Ebrahimi, Long Le, Vincent Perot, Swaroop Mishra, Mohit Bansal, Chen-Yu Lee, Tomas Pfister
Proceedings of **NAACL 2025**. [pdf]

292. On Positional Bias of Faithfulness for Long-form Summarization
David Wan, Jesse Vig, Mohit Bansal, Shafiq Joty
Proceedings of **NAACL 2025**. [pdf]

291. MAMM-Refine: A Recipe for Improving Faithfulness in Generation with Multi-Agent Collaboration
David Wan, Justin Chen, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **NAACL 2025**. [pdf]

290. DAM: Dynamic Adapter Merging for Continual Video QA Learning
Feng Cheng*, Ziyang Wang*, Yi-Lin Sung, Yan-Bo Lin, Mohit Bansal, Gedas Bertasius
Proceedings of **WACV 2025**. [pdf]

289. Improving Faithfulness of Text-to-Image Diffusion Models through Inference Intervention
Danfeng Guo, Sanchit Agarwal, Yu-Hsiang Lin, Jiun-Yu Kao, Tagyoung Chung, Nanyun Peng, Mohit Bansal
Proceedings of **WACV 2025**. [pdf]

288. Rethinking Machine Unlearning for Large Language Models
Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Xiaojun Xu, Yuguang Yao, Hang Li, Kush R. Varshney, Mohit Bansal, Sanmi Koyejo, Yang Liu
Proceedings of **Nature Machine Intelligence**. [pdf]

287. LACIE: Listener-Aware Finetuning for Confidence Calibration in Large Language Models
Elias Stengel-Eskin, Peter Hase, Mohit Bansal
Proceedings of **NeurIPS 2024**. [pdf]

286. SELMA: Learning and Merging Skill-Specific Text-to-Image Experts with Auto-Generated Data
Jialu Li*, Jaemin Cho*, Yi-Lin Sung, Jaehong Yoon, Mohit Bansal
Proceedings of **NeurIPS 2024**. [pdf]

285. GTBench: Uncovering the Strategic Reasoning Limitations of LLMs via Game-Theoretic Evaluations
Jinhao Duan, Renming Zhang, James Diffenderfer, Bhavya Kailkhura, Lichao Sun, Elias Stengel-Eskin, Mohit Bansal, Tianlong Chen, Kaidi Xu
Proceedings of **NeurIPS 2024**. [pdf]

284. Fundamental Problems With Model Editing: How Should Rational Belief Revision Work in LLMs?
Peter Hase, Thomas Hofweber, Xiang Zhou, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **TMLR 2024**. [pdf]

283. Unlearning Sensitive Information in Multimodal LLMs: Benchmark and Attack-Defense Evaluation
Vaidehi Patil, Yi-Lin Sung, Peter Hase, Jie Peng, Tianlong Chen, Mohit Bansal
Proceedings of **TMLR 2024**. [pdf]

282. Vision-and-Language Navigation Today and Tomorrow: A Survey in the Era of Foundation Models
Yue Zhang, Ziqiao Ma, Jialu Li, Yanyuan Qiao, Zun Wang, Joyce Chai, Qi Wu, Mohit Bansal, Parisa Kordjamshidi
Proceedings of **TMLR 2024**. [pdf]

281. FlexEControl: Flexible and Efficient Multimodal Control for Text-to-Image Generation
Xuehai He, Jian Zheng, Jacob Zhiyuan Fang, Robinson Piramuthu, Mohit Bansal, Vicente Ordonez, Gunnar A Sigurdsson, Nanyun Peng, Xin Eric Wang
Proceedings of **TMLR 2024**. [pdf]

280. A Simple LLM Framework for Long-Range Video Question-Answering
Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, Gedas Bertasius
Proceedings of **EMNLP 2024**. [pdf]

279. Explaining and Improving Contrastive Decoding by Extrapolating the Probabilities of a Huge and Hypothetical LM
Haw-Shiuan Chang, Nanyun Peng, Mohit Bansal, Anil Ramakrishna, Tagyoung Chung
Proceedings of **EMNLP 2024**. [pdf]

278. LLM Self-Correction with DeCRIM: Decompose, Critique, and Refine for Enhanced Following of Instructions with Multiple Constraints
Thomas Palmeira Ferraz, Kartik Mehta, Yu-Hsiang Lin, Haw-Shiuan Chang, Shereen Oraby, Sijia Liu, Vivek Subramanian, Tagyoung Chung, Mohit Bansal, Nanyun Peng
Findings of **EMNLP 2024**. [pdf]

277. Knowledge-Aware Reasoning over Multimodal Semi-structured Tables
Suyash Vardhan Mathur, Jainit Sushil Bafna, Kunal Kartik, Harshita Khandelwal, Manish Shrivastava, Vivek Gupta, Mohit Bansal, Dan Roth
Findings of **EMNLP 2024**. [pdf]

276. EnvGen: Generating and Adapting Environments via LLMs for Training Embodied Agents
Abhay Zala*, Jaemin Cho*, Han Lin, Jaehong Yoon, Mohit Bansal
Proceedings of **COLM 2024**. [pdf]

275. VideoDirectorGPT: Consistent Multi-scene Video Generation via LLM-Guided Planning
Han Lin, Abhay Zala, Jaemin Cho, Mohit Bansal
Proceedings of **COLM 2024**. [pdf]

274. DiagrammerGPT: Generating Open-Domain, Open-Platform Diagrams via LLM Planning
Abhay Zala, Han Lin, Jaemin Cho, Mohit Bansal
Proceedings of **COLM 2024**. [pdf]

273. Contrastive Region Guidance: Improving Grounding in Vision-Language Models without Training
David Wan, Jaemin Cho, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **ECCV 2024**. [pdf]

272. ReConcile: Round-Table Conference Improves Reasoning via Consensus among Diverse LLMs
Justin Chih-Yao Chen, Swarnadeep Saha, Mohit Bansal
Proceedings of **ACL 2024**. [pdf]

271. Inducing Systematicity in Transformers by Attending to Structurally Quantized Embeddings
Yichen Jiang, Xiang Zhou, Mohit Bansal
Proceedings of **ACL 2024**. [pdf]

270. The Unreasonable Effectiveness of Easy Training Data for Hard Tasks
Peter Hase, Mohit Bansal, Peter Clark, Sarah Wiegreffe
Proceedings of **ACL 2024**. [pdf]

269. Evaluating Very Long-Term Conversational Memory of LLM Agents
Adyasha Maharana, Dong-Ho Lee, Sergey Tulyakov, Mohit Bansal, Francesco Barbieri, Yuwei Fang
Proceedings of **ACL 2024**. [pdf]

268. Soft Self-Consistency Improves Language Model Agents
Han Wang*, Archiki Prasad*, Elias Stengel-Eskin*, Mohit Bansal
Proceedings of **ACL 2024 (short)**. [pdf]

267. RefineSumm: Self-Refining MLLM for Generating a Multimodal Summarization Dataset
Vaidehi Patil, Leonardo F. R. Ribeiro, Mengwen Liu, Mohit Bansal, Markus Dreyer
Proceedings of **ACL 2024**. [pdf]

266. Mementos: A Comprehensive Benchmark for Multimodal Large Language Model Reasoning over Image Sequences
Xiyao Wang, Yuhang Zhou, Xiaoyu Liu, Hongjin Lu, Yuancheng Xu, Feihong He, Jaehong Yoon, Taixi Lu, Gedas Bertasius, Mohit Bansal, Huaxiu Yao, Furong Huang
Proceedings of **ACL 2024**. [pdf]

265. ACUEval: Fine-grained Hallucination Evaluation and Correction for Abstractive Summarization
David Wan, Koustuv Sinha, Srini Iyer, Asli Celikyilmaz, Mohit Bansal, Ramakanth Pasunuru
Findings of **ACL 2024**. [pdf]

264. The Power of Summary-Source Alignments
Ori Ernst, Ori Shapira, Aviv Slobodkin, Sharon Adar, Mohit Bansal, Jacob Goldberger, Ran Levy, Ido Dagan
Findings of **ACL 2024**. [pdf]

263. MAGDi: Structured Distillation of Multi-Agent Interaction Graphs Improves Reasoning in Smaller Language Models
Justin Chih-Yao Chen*, Swarnadeep Saha*, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **ICML 2024**. [pdf]

262. ReGAL: Refactoring Programs to Discover Generalizable Abstractions
Elias Stengel-Eskin*, Archiki Prasad*, Mohit Bansal
Proceedings of **ICML 2024**. [pdf]

261. Position Paper: TrustLLM: Trustworthiness in Large Language Models
Lichao Sun, Yue Huang, Haoran Wang... Mohit Bansal, ...Yong Chen, Yue Zhao
Proceedings of **ICML 2024**. [pdf]

260. CoDi-2: In-Context, Interleaved, and Interactive Any-to-Any Generation
Zineng Tang, Ziyi Yang, Mahmoud Khademi, Yang Liu, Chenguang Zhu, Mohit Bansal
Proceedings of **CVPR 2024**. [pdf]
(selected spotlight)

259. Multimodal Representation Learning by Alternating Unimodal Adaptation
Xiaohui Zhang, Jaehong Yoon, Mohit Bansal, Huaxiu Yao
Proceedings of **CVPR 2024**. [pdf]

258. Rethinking Interactive Image Segmentation with Low Latency, High Quality, and Diverse Prompts
Qin Liu, Jaemin Cho, Mohit Bansal, Marc Niethammer
Proceedings of **CVPR 2024**. [pdf]

257. Branch-Solve-Merge Improves Large Language Model Evaluation and Generation
Swarnadeep Saha, Omer Levy, Asli Celikyilmaz, Mohit Bansal, Jason Weston, Xian Li
Proceedings of **NAACL 2024**. [pdf]

256. ADaPT: As-Needed Decomposition and Planning with Language Models
Archiki Prasad, Alexander Koller, Mareike Hartmann, Peter Clark, Ashish Sabharwal, Mohit Bansal, Tushar Khot
Findings of **NAACL 2024**. [pdf]

255. Prompting Vision-Language Models For Aspect-Controlled Generation of Referring Expressions
Danfeng Guo, Sanchit Agarwal, Arpit Gupta, Jiun-Yu Kao, Emre Barut, Tagyoung Chung, Jing Huang, Mohit Bansal
Findings of **NAACL 2024**. [pdf]

254. Merging by Matching Models in Task Subspaces
Derek Tam, Mohit Bansal, Colin Raffel
Proceedings of **TMLR 2024**. [pdf]

253. Low-Cost Algorithmic Recourse for Users With Uncertain Cost Functions
Prateek Yadav, Peter Hase, Mohit Bansal
Proceedings of **TMLR 2024**. [pdf]

252. Unified Embeddings for Multimodal Retrieval via Frozen LLMs
Ziyang Wang, Heba Elfardy, Markus Dreyer, Kevin Small, Mohit Bansal
Findings of **EACL 2024**. [pdf]

251. Hierarchical and Dynamic Prompt Compression for Efficient Zero-shot API Usage
Yichen Jiang, Marco Vecchio, Mohit Bansal, Anders Johannsen
Findings of **EACL 2024**. [pdf]

250. Exposing and Addressing Cross-Task Inconsistency in Unified Vision-Language Models
Adyasha Maharana, Amita Kamath, Christopher Clark, Mohit Bansal, Aniruddha Kembhavi
Proceedings of **TMLR 2024**. [pdf]

249. VLN-Video: Utilizing Driving Videos for Outdoor Vision-and-Language Navigation
Jialu Li, Aishwarya Padmakumar, Gaurav Sukhatme, Mohit Bansal
Proceedings of **AAAI 2024**. [pdf]

248. Can Sensitive Information Be Deleted From LLMs? Objectives for Defending Against Extraction Attacks
Vaidehi Patil*, Peter Hase*, Mohit Bansal
Proceedings of **ICLR 2024**. [pdf]
(selected spotlight)

247. Davidsonian Scene Graph: Improving Reliability in Fine-grained Evaluation for Text-to-Image Generation
Jaemin Cho, Yushi Hu, Roopal Garg, Peter Anderson, Ranjay Krishna, Jason Baldridge, Mohit Bansal, Jordi Pont-Tuset, Su Wang
Proceedings of **ICLR 2024**. [pdf]

246. Merge, Then Compress: Demystify Efficient SMoE with Hints from Its Routing Policy
Pingzhi Li, Zhenyu Zhang, Prateek Yadav, Yi-Lin Sung, Yu Cheng, Mohit Bansal, Tianlong Chen
Proceedings of **ICLR 2024**. [pdf]
(selected spotlight)

245. Rephrase, Augment, Reason: Visual Grounding of Questions for Vision-Language Models
Archiki Prasad, Elias Stengel-Eskin, Mohit Bansal
Proceedings of **ICLR 2024**. [pdf]

244. D2 Pruning: Message Passing for Balancing Diversity and Difficulty in Data Pruning
Adyasha Maharana, Prateek Yadav, Mohit Bansal
Proceedings of **ICLR 2024**. [pdf]

243. ECoFLaP: Efficient Coarse-to-Fine Layer-Wise Pruning for Vision-Language Models
Yi-Lin Sung, Jaehong Yoon, Mohit Bansal
Proceedings of **ICLR 2024**. [pdf]

242. Analyzing and Mitigating Object Hallucination in Large Vision-Language Models
Yiyang Zhou, Chenhang Cui, Jaehong Yoon, Linjun Zhang, Zhun Deng, Chelsea Finn, Mohit Bansal, Huaxiu Yao
Proceedings of **ICLR 2024**. [pdf]

241. ReCEval: Evaluating Reasoning Chains via Correctness and Informativeness
Archiki Prasad, Swarnadeep Saha, Xiang Zhou, Mohit Bansal
Proceedings of **EMNLP 2023**. [pdf]

240. HistAlign: Improving Context Dependency in Language Generation by Aligning with History
David Wan, Shiyue Zhang, Mohit Bansal
Proceedings of **EMNLP 2023**. [pdf]

239. Data Factors for Better Compositional Generalization
Xiang Zhou, Yichen Jiang, Mohit Bansal
Proceedings of **EMNLP 2023**. [pdf]

238. Generating Summaries with Controllable Readability Levels
Leonardo F. R. Ribeiro, Mohit Bansal, Markus Dreyer
Proceedings of **EMNLP 2023**. [pdf]

237. An Empirical Study of Multimodal Model Merging
Yi-Lin Sung, Linjie Li, Kevin Lin, Zhe Gan, Mohit Bansal, Lijuan Wang
Findings of **EMNLP 2023**. [pdf]

236. Debiasing Multimodal Models via Causal Information Minimization
Vaidehi Patil, Adyasha Maharana, Mohit Bansal
Findings of **EMNLP 2023**. [pdf]

235. Visual Programming for Text-to-Image Generation and Evaluation
Jaemin Cho, Abhay Zala, Mohit Bansal
Proceedings of **NeurIPS 2023**. [pdf]

234. CoDi: Any-to-Any Generation via Composable Diffusion
Zineng Tang, Ziyi Yang, Chenguang Zhu, Michael Zeng, Mohit Bansal
Proceedings of **NeurIPS 2023**. [pdf]

233. Can Language Models Teach Weaker Agents? Teacher Explanations Improve Students via Theory of Mind
Swarnadeep Saha, Peter Hase, Mohit Bansal
Proceedings of **NeurIPS 2023**. [pdf]

232. Resolving Interference When Merging Models
Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, Mohit Bansal
Proceedings of **NeurIPS 2023**. [pdf]

231. PanoGen: Text-Conditioned Panoramic Environment Generation for Vision-and-Language Navigation Jialu Li, Mohit Bansal. Proceedings of **NeurIPS 2023**. [pdf]

230. Does Localization Inform Editing? Surprising Differences in Causality-Based Localization vs. Knowledge Editing in Language Models
Peter Hase, Mohit Bansal, Been Kim, Asma Ghandeharioun
Proceedings of **NeurIPS 2023**. [pdf]
(selected spotlight)

229. Self-Chained Image-Language Model for Video Localization and Question Answering
Shoubin Yu, Jaemin Cho, Prateek Yadav, Mohit Bansal
Proceedings of **NeurIPS 2023**. [pdf]

228. Paxion: Patching Action Knowledge in Video-Language Foundation Models
Zhenhailong Wang, Ansel Blume, Sha Li, Genglin Liu, Jaemin Cho, Zineng Tang, Mohit Bansal, Heng Ji
Proceedings of **NeurIPS 2023**. [pdf]
(selected spotlight)

227. Adaptive Contextual Perception: How to Generalize to New Backgrounds and Ambiguous Objects
Zhuofan Ying, Peter Hase, Mohit Bansal
Proceedings of **NeurIPS 2023**. [pdf]

226. DALL-Eval: Probing the Reasoning Skills and Social Biases of Text-to-Image Generative Transformers
Jaemin Cho, Abhay Zala, Mohit Bansal
Proceedings of **ICCV 2023**. [pdf]

225. Unified Coarse-to-Fine Alignment for Video-Text Retrieval
Ziyang Wang, Yi-Lin Sung, Feng Cheng, Gedas Bertasius, Mohit Bansal
Proceedings of **ICCV 2023**. [pdf]

224. Scaling Data Generation in Vision-and-Language Navigation
Zun Wang, Jialu Li, Yicong Hong, Yi Wang, Qi Wu, Mohit Bansal, Stephen Gould, Hao Tan, Yu Qiao
Proceedings of **ICCV 2023**. [pdf]
(selected oral)

223. Extractive is not Faithful: An Investigation of Broad Unfaithfulness Problems in Extractive Summarization
Shiyue Zhang*, David Wan*, Mohit Bansal
Proceedings of **ACL 2023**. [pdf]

222. MeetingQA: Extractive Question-Answering on Meeting Transcripts
Archiki Prasad, Trung Bui, Seunghyun Yoon, Hanieh Deilamsalehy, Franck Dernoncourt, Mohit Bansal
Proceedings of **ACL 2023**. [pdf]

221. Revealing Single Frame Bias for Video-and-Language Learning
Jie Lei, Tamara Berg, Mohit Bansal
Proceedings of **ACL 2023**. [pdf]

220. MixCE: Training Autoregressive Language Models by Mixing Forward and Reverse Cross-Entropies
Shiyue Zhang, Shijie Wu, Ozan Irsoy, Steven Lu, Mohit Bansal, Mark Dredze, David Rosenberg
Proceedings of **ACL 2023**. [pdf]

219. Non-Sequential Graph Script Induction via Multimedia Grounding
Yu Zhou, Sha Li, Manling Li, Xudong Lin, Shih-Fu Chang, Mohit Bansal, Heng Ji
Proceedings of **ACL 2023**. [pdf]

218. Exploring Continual Learning for Code Generation Models
Prateek Yadav, Qing Sun, Hantian Ding, Xiaopeng Li, Dejiao Zhang, Ming Tan, Parminder Bhatia, Xiaofei Ma, Ramesh Nallapati, Murali Krishna Ramanathan, Mohit Bansal, Bing Xiang
Proceedings of **ACL 2023** (short). [pdf]

217. Exclusive Supermask Subnetwork Training for Continual Learning
Prateek Yadav, Mohit Bansal
Findings of **ACL 2023**. [pdf]

216. MURMUR: Modular Multi-Step Reasoning for Semi-Structured Data-to-Text Generation
Swarnadeep Saha, Xinyan Velocity Yu, Mohit Bansal, Ramakanth Pasunuru, Asli Celikyilmaz
Findings of **ACL 2023**. [pdf]

215. Evaluating the Factual Consistency of Large Language Models Through Summarization
Derek Tam, Anisha Mascarenhas, Shiyue Zhang, Sarah Kwan, Mohit Bansal, Colin Raffel
Findings of **ACL 2023**. [pdf]

214. Can Sequence-to-Sequence Transformers Naturally Understand Sequential Instructions?
Xiang Zhou, Aditya Gupta, Shyam Upadhyay, Mohit Bansal, Manaal Faruqui
Proceedings of ***SEM 2023**. [pdf]

213. On Conditional and Compositional Language Model Differentiable Prompting
Jonathan Pilault, Can Liu, Mohit Bansal, Markus Dreyer
Proceedings of **IJCAI 2023**. [pdf]

212. UDOP: Unifying Vision, Text, and Layout for Universal Document Processing
Zineng Tang, Ziyi Yang, Guoxin Wang, Yuwei Fang, Yang Liu, Chenguang Zhu, Michael Zeng, Cha Zhang, Mohit Bansal
Proceedings of **CVPR 2023**. [pdf]

211. Improving Vision-and-Language Navigation by Generating Future-View Image Semantics
Jialu Li, Mohit Bansal
Proceedings of **CVPR 2023**. [pdf]

210. Hierarchical Video-Moment Retrieval and Step-Captioning
Abhay Zala*, Jaemin Cho*, Satwik Kottur, Xilun Chen, Barlas Oguz, Yashar Mehdad, Mohit Bansal
Proceedings of **CVPR 2023**. [pdf]

209. VindLU: A Recipe for Effective Video-and-Language Pretraining
Feng Cheng, Xizi Wang, Jie Lei, David Crandall, Mohit Bansal, Gedas Bertasius
Proceedings of **CVPR 2023**. [pdf]

208. Vision Transformers are Parameter-Efficient Audio-Visual Learners
Yan-Bo Lin, Yi-Lin Sung, Jie Lei, Mohit Bansal, Gedas Bertasius
Proceedings of **CVPR 2023**. [pdf]

207. Summarization Programs: Interpretable Abstractive Summarization with Neural Modular Trees
Swarnadeep Saha, Shiyue Zhang, Peter Hase, Mohit Bansal
Proceedings of **ICLR 2023**. [pdf]

206. Faithfulness-Aware Decoding Strategies for Abstractive Summarization
David Wan, Mengwen Liu, Kathleen McKeown, Markus Dreyer and Mohit Bansal
Proceedings of **EACL 2023**. [pdf]

205. GrIPS: Gradient-free, Edit-based Instruction Search for Prompting Large Language Models
Archiki Prasad, Peter Hase, Xiang Zhou, Mohit Bansal
Proceedings of **EACL 2023**. [pdf]

204. Do Language Models Have Beliefs? Methods for Detecting, Updating, and Visualizing Model Beliefs
Peter Hase, Mona Diab, Asli Celikyilmaz, Xian Li, Zornitsa Kozareva, Veselin Stoyanov, Mohit Bansal, Srinivasan Iyer
Proceedings of **EACL 2023**. [pdf]

203. Social Commonsense for Explanation and Cultural Bias Discovery
Lisa Bauer, Hanna Leth Tischer and Mohit Bansal
Proceedings of **EACL 2023**. [pdf]

202. Enhancing Multi-Document Summarization with Cross-Document Graph-based Information Extraction
Zixuan Zhang, Heba Elfardy, Markus Dreyer, Kevin Small, Heng Ji and Mohit Bansal
Proceedings of **EACL 2023**. [pdf]

201. DeepMaven: Deep Question Answering on Long-Distance Movie/TV Show Videos with Multimedia Knowledge Extraction and Synthesis
Yi Fung, Han Wang, Tong Wang, Ali Kebarighotbi, Prem Natarajan, Mohit Bansal, Heng Ji
Proceedings of **EACL 2023**. [pdf]

200. Perceiver-VL: Efficient Vision-and-Language Modeling with Iterative Latent Attention
Zineng Tang*, Jaemin Cho*, Jie Lei, Mohit Bansal
Proceedings of **WACV 2023**. [pdf]

199. Evaluating and Improving Factuality in Multimodal Abstractive Summarization
David Wan, Mohit Bansal
Proceedings of **EMNLP 2022**. [pdf]

198. Mutual Exclusivity Training and Primitive Augmentation to Induce Compositionality
Yichen Jiang*, Xiang Zhou*, Mohit Bansal
Proceedings of **EMNLP 2022**. [pdf]

197. Are Hard Examples also Harder to Explain? A Study with Human and Model-Generated Explanations
Swarnadeep Saha, Peter Hase, Nazneen Rajani, Mohit Bansal
Proceedings of **EMNLP 2022 (short)**. [pdf]

196. Investigating the Role of Language for Action Learning in Interactive Visual Environments
Arjun Akula, Spandana Gella, Aishwarya Padmakumar, Mahdi Namazifar, Mohit Bansal, Jesse Thomason, Dilek Hakkani-Tur
Proceedings of **EMNLP 2022 (short)**. [pdf]

195. Analyzing the Limits of Self-Supervision in Handling Bias in Language
Lisa Bauer, Karthik Gopalakrishnan, Spandana Gella, Yang Liu, Mohit Bansal, Dilek Hakkani-Tur
Findings of **EMNLP 2022**. [pdf]

194. An Empirical Survey of Data Augmentation for Limited Data Learning in NLP
Jiaao Chen*, Derek Tam*, Colin Raffel, Mohit Bansal, Diyi Yang
Proceedings of **TACL**. [pdf]

193. TVLT: Textless Vision-Language Transformer
Zineng Tang*, Jaemin Cho*, Yixin Nie*, Mohit Bansal
Proceedings of **NeurIPS 2022**. [pdf]
(selected oral)

192. VisFIS: Visual Feature Importance Supervision with Right-for-the-Right-Reason Objectives
Zhuofan Ying*, Peter Hase*, Mohit Bansal
Proceedings of **NeurIPS 2022**. [pdf]

191. LST: Ladder Side-Tuning for Parameter and Memory Efficient Transfer Learning
Yi-Lin Sung, Jaemin Cho, Mohit Bansal
Proceedings of **NeurIPS 2022**. [pdf]

190. VidIL: Language Models with Image Descriptors are Strong Few-Shot Video-Language Learners
Zhenhailong Wang, Manling Li, Ruochen Xu, Luowei Zhou, Jie Lei, Xudong Lin, Shuohang Wang, Ziyi Yang, Chenguang Zhu, Derek Hoiem, Shih-Fu Chang, Mohit Bansal, Heng Ji
Proceedings of **NeurIPS 2022**. [pdf]

189. T-few: Few-Shot Parameter-Efficient Fine-Tuning is Better and Cheaper than In-Context Learning
Haokun Liu*, Derek Tam*, Mohammed Muqeeth*, Jay Mohta, Tenghao Huang, Mohit Bansal, Colin Raffel
Proceedings of **NeurIPS 2022**. [pdf]

188. WinoGAViL: Gamified Association Benchmark to Challenge Vision-and-Language Models
Yonatan Bitton, Nitzan Bitton Guetta, Ron Yosef, Yuval Elovici, Mohit Bansal, Gabriel Stanovsky, Roy Schwartz
Proceedings of **NeurIPS 2022 (datasets/benchmarks track)**. [pdf]
(selected oral)

187. StoryDALL-E: Adapting Pretrained Text-to-Image Transformers for Story Continuation
Adyasha Maharana, Darryl Hannan, Mohit Bansal
Proceedings of **ECCV 2022**. [pdf]

186. ECLIPSE: Efficient Long-range Video Retrieval using Sight and Sound
Yan-Bo Lin, Jie Lei, Mohit Bansal, Gedas Bertasius
Proceedings of **ECCV 2022**. [pdf]
(selected oral)

185. GraDA: Graph Generative Data Augmentation for Commonsense Reasoning
Adyasha Maharana, Mohit Bansal
Proceedings of **COLING 2022**. [pdf]

184. GRAVL-BERT: Graphical Visual-Linguistic Representations for Multimodal Coreference Resolution
Danfeng Guo, Arpit Gupta, Sanchit Agarwal, Jiun-Yu Kao, Shuyang Gao, Arijit Biswas, Chien-Wei Lin, Tagyoung Chung, Mohit Bansal
Proceedings of **COLING 2022**. [pdf]

183. How Robust is Neural Machine Translation to Language Imbalance in Multilingual Tokenizer Training?
Shiyue Zhang, Vishrav Chaudhary, Naman Goyal, James Cross, Guillaume Wenzek, Mohit Bansal, Francisco Guzman
Proceedings of **AMTA 2022**. [pdf]

182. FactPEGASUS: Factuality-Aware Pre-training and Fine-tuning for Abstractive Summarization
David Wan, Mohit Bansal
Proceedings of **NAACL 2022**. [pdf]

181. CoSIm: Commonsense Reasoning for Counterfactual Scene Imagination
Hyounghun Kim*, Abhay Zala*, Mohit Bansal
Proceedings of **NAACL 2022**. [pdf]

180. Masked Part-Of-Speech Model: Does Modeling Long Context Help Unsupervised POS-tagging?
Xiang Zhou, Shiyue Zhang, Mohit Bansal
Proceedings of **NAACL 2022**. [pdf]

179. FactGraph: Evaluating Factuality in Summarization with Semantic Graph Representations
Leonardo Ribeiro, Mengwen Liu, Iryna Gurevych, Markus Dreyer, Mohit Bansal
Proceedings of **NAACL 2022**. [pdf]

178. On Curriculum Learning for Commonsense Reasoning
Adyasha Maharana, Mohit Bansal
Proceedings of **NAACL 2022 (short)**. [pdf]

177. Proposition-Level Clustering for Multi-Document Summarization
Ori Ernst, Avi Caciularu, Ori Shapira, Ramakanth Pasunuru, Mohit Bansal, Jacob Goldberger, Ido Dagan
Proceedings of **NAACL 2022**. [pdf]

176. Interactive Query-Assisted Summarization via Deep Reinforcement Learning
Ori Shapira, Ramakanth Pasunuru, Mohit Bansal, Ido Dagan, Yael Amsterdamer
Proceedings of **NAACL 2022**. [pdf]

175. Enhanced Knowledge Selection for Grounded Dialogues via Document Semantic Graphs
Sha Li, Mahdi Namazifar, Di Jin, Mohit Bansal, Heng Ji, Yang Liu, Dilek Hakkani-Tur
Proceedings of **NAACL 2022**. [pdf]

174. CLEAR: Improving Vision-Language Navigation with Cross-Lingual, Environment-Agnostic Representations
Jialu Li, Hao Tan, Mohit Bansal
Findings of **NAACL 2022**. [pdf]

173. Fine-grained Image Captioning with CLIP Reward
Jaemin Cho, Seunghyun Yoon, Ajinkya Kale, Franck Dernoncourt, Trung Bui, Mohit Bansal
Findings of **NAACL 2022 (short)**. [pdf]

172. Multimodal Intent Discovery from Livestream Videos
Adyasha Maharana, Quan Hung Tran, Franck Dernoncourt, Seunghyun Yoon, Trung Bui, Walter W Chang, Mohit Bansal
Findings of **NAACL 2022**. [pdf]

171. Efficient Few-Shot Fine-Tuning for Opinion Summarization
Arthur Bra?inskas, Ramesh Nallapati, Mohit Bansal, Markus Dreyer
Findings of **NAACL 2022**. [pdf]

170. SETSum: Summarization and Visualization of Student Evaluations of Teaching
Yinuo Hu*, Shiyue Zhang*, Viji Sathy, A. T. Panter, Mohit Bansal
Proceedings of **NAACL 2022 (demo)**. [pdf]

169. RESIN-11: Schema-guided Event Prediction for 11 Newsworthy Scenarios
Xinya Du, Zixuan Zhang, Sha Li, Pengfei Yu, Hongwei Wang, Tuan Lai, Xudong Lin, Ziqi Wang, Iris Liu, Ben Zhou, Haoyang Wen, Manling Li, Darryl Hannan, Jie Lei, Hyounghun Kim, Rotem Dror, Haoyu Wang, Michael Regan, Qi Zeng, Qing Lyu, Charles Yu, Carl Edwards, Xiaomeng Jin, Yizhu Jiao, Ghazaleh Kazeminejad, Zhenhailong Wang, Chris Callison-Burch, Mohit Bansal, Carl Vondrick, Jiawei Han, Dan Roth, Shih-Fu Chang, Martha Palmer, Heng Ji
Proceedings of **NAACL 2022 (demo)**. [pdf]

168. EnvEdit: Environment Editing for Vision-and-Language Navigation
Jialu Li, Hao Tan, Mohit Bansal
Proceedings of **CVPR 2022**. [pdf]

167. VL-Adapter: Parameter-Efficient Transfer Learning for Vision-and-Language Tasks
Yi-Lin Sung, Jaemin Cho, Mohit Bansal
Proceedings of **CVPR 2022**. [pdf]

166. VIMPAC: Video Pre-Training via Masked Token Prediction and Contrastive Learning
Hao Tan, Jie Lei, Thomas Wolf, Mohit Bansal
Proceedings of **T4V Workshop, CVPR 2022**. [pdf]

165. How can NLP Help Revitalize Endangered Languages? A Case Study and Roadmap for the Cherokee Language
Shiyue Zhang, Ben Frey, Mohit Bansal
Proceedings of **ACL 2022**. [pdf]

164. Explanation Graph Generation via Pre-trained Language Models: An Empirical Study with Contrastive Learning
Swarnadeep Saha, Prateek Yadav, and Mohit Bansal
Proceedings of **ACL 2022**. [pdf]

163. Distributed NLI: Learning to Predict Human Opinion Distributions for Language Reasoning
Xiang Zhou*, Yixin Nie*, Mohit Bansal
Findings of **ACL 2022**. [pdf]

162. When Can Models Learn From Explanations? A Formal Framework for Understanding the Roles of Explanation Data
Peter Hase, Mohit Bansal
Proceedings of **LNLS Workshop, ACL 2022**. [pdf]

161. CAISE: Conversational Agent for Image Search and Editing
Hyounghun Kim, Doo Soon Kim, Seunghyun Yoon, Franck Dernoncourt, Trung Bui, Mohit Bansal
Proceedings of **AAAI 2022**. [pdf]

160. MuMuQA: Multimedia Multi-Hop News Question Answering via Cross-Media Knowledge Extraction and Grounding
Revanth Gangi Reddy, Xilin Rui, Manling Li, Xudong Lin, Haoyang Wen, Jaemin Cho, Lifu Huang, Mohit Bansal, Avirup Sil, Shih-Fu Chang, Alexander Schwing, Heng Ji
Proceedings of **AAAI 2022**. [pdf]

159. Scientific Chart Summarization: Datasets and Improved Text Modeling
Hao Tan, Chen-Tse Tsai, Yujie He, Mohit Bansal
Proceedings of **SDU Workshop, AAAI 2022**. [pdf]

158. CLIP-ViL: How Much Can CLIP Benefit Vision-and-Language Tasks?
Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, Kurt Keutzer
Proceedings of **ICLR 2022**. [pdf]

157. VidLanKD: Improving Language Understanding via Video-Distilled Knowledge Transfer
Zineng Tang, Jaemin Cho, Hao Tan, Mohit Bansal
Proceedings of **NeurIPS 2021**. [pdf]

156. QVHighlights: Detecting Moments and Highlights in Videos via Natural Language Queries
Jie Lei, Tamara L. Berg, Mohit Bansal
Proceedings of **NeurIPS 2021**. [pdf]

155. Search Methods for Sufficient, Socially-Aligned Feature Importance Explanations with In-Distribution Counterfactuals
Peter Hase, Harry Xie, Mohit Bansal
Proceedings of **NeurIPS 2021**. [pdf]

154. VALUE: A Multi-Task Benchmark for Video-and-Language Understanding Evaluation
Linjie Li*, Jie Lei*, Zhe Gan, Licheng Yu, Yen-Chun Chen, Rohit Pillai, Yu Cheng, Luowei Zhou, Xin Eric Wang, William Yang Wang, Tamara Lee Berg, Mohit Bansal, Jingjing Liu, Lijuan Wang, Zicheng Liu
Proceedings of **NeurIPS 2021 (benchmark track)**. [pdf]

153. Finding a Balanced Degree of Automation for Summary Evaluation
Shiyue Zhang, Mohit Bansal
Proceedings of **EMNLP 2021**. [pdf]

152. ExplaGraphs: An Explanation Graph Generation Task for Structured Commonsense Reasoning
Swarnadeep Saha, Prateek Yadav, Lisa Bauer, Mohit Bansal
Proceedings of **EMNLP 2021**. [pdf]

151. Integrating Visuospatial, Linguistic, and Commonsense Structure into Story Visualization
Adyasha Maharana, Mohit Bansal
Proceedings of **EMNLP 2021**. [pdf]

150. Inducing Transformer's Compositional Generalization Ability via Auxiliary Sequence Prediction Tasks
Yichen Jiang, Mohit Bansal
Proceedings of **EMNLP 2021**. [pdf]

149. Continual Few-Shot Learning for Text Classification
Ramakanth Pasunuru, Veselin Stoyanov, Mohit Bansal
Proceedings of **EMNLP 2021**. [pdf]

148. FastIF: Scalable Influence Functions for Efficient Model Interpretation and Debugging
Han Guo, Nazneen Fatema Rajani, Peter Hase, Mohit Bansal, Caiming Xiong
Proceedings of **EMNLP 2021**. [pdf]

147. NDH-Full: Learning and Evaluating Navigational Agents on Full-Length Dialogue
Hyounghun Kim, Jialu Li, Mohit Bansal
Proceedings of **EMNLP 2021**. [pdf]

146. Learning and Analyzing Generation Order for Undirected Sequence Models
Yichen Jiang, Mohit Bansal
Findings of **EMNLP 2021**. [pdf]

145. Improving and Simplifying Pattern Exploiting Training
Derek Tam*, Rakesh R Menon*, Mohit Bansal, Shashank Srivastava and Colin Raffel
Proceedings of **EMNLP 2021 (short)**. [pdf]

144. iFacetSum: Coreference-based Interactive Faceted Summarization for Multi-Document Exploration
Eran Hirsch, Alon Eirew, Ori Shapira, Avi Caciularu, Arie Cattan, Ori Ernst, Ramakanth Pasunuru, Hadar Ronen, Mohit Bansal, Ido Dagan
Proceedings of **EMNLP 2021 (demo)**. [pdf]

143. Summary-Source Proposition-level Alignment: Task, Datasets and Supervised Baseline
Ori Ernst, Ori Shapira, Ramakanth Pasunuru, Michael Lepioshkin, Jacob Goldberger, Mohit Bansal, Ido Dagan
Proceedings of **CoNLL 2021**. [pdf]
(CoNLL Best Paper Runner-Up)

142. To what extent do human explanations of model behavior align with actual model behavior?
Grusha Prasad, Yixin Nie, Mohit Bansal, Robin Jia, Douwe Kiela, Adina Williams
Proceedings of **BlackboxNLP Workshop, EMNLP 2021**. [pdf]

141. An Overview of Uncertainty Calibration for Text Classification and the Role of Distillation
Han Guo, Ramakanth Pasunuru, Mohit Bansal
Proceedings of **RepL4NLP Workshop, ACL 2021**. [pdf]

140. Unifying Vision-and-Language Tasks via Text Generation
Jaemin Cho, Jie Lei, Hao Tan, Mohit Bansal
Proceedings of **ICML 2021**. [pdf]

139. EmailSum: Abstractive Email Thread Summarization
Shiyue Zhang, Asli Celikyilmaz, Jianfeng Gao, Mohit Bansal
Proceedings of **ACL 2021**. [pdf]

138. Continuous Language Generative Flow
Zineng Tang, Shiyue Zhang, Hyounghun Kim, Mohit Bansal
Proceedings of **ACL 2021**. [pdf]

137. mTVR: Multilingual Moment Retrieval in Videos
Jie Lei, Tamara Berg, Mohit Bansal
Proceedings of **ACL 2021** (short papers). [pdf]

136. I like fish, especially dolphins: Addressing Contradictions in Dialogue Modeling
Yixin Nie, Mary Williamson, Mohit Bansal, Douwe Kiela, Jason Weston
Proceedings of **ACL 2021**. [pdf]

135. InfoSurgeon: Cross-Media Fine-grained Information Consistency Checking for Fake News Detection
Yi Fung, Christopher Thomas, Revanth Gangi Reddy, Sandeep Polisetty, Heng Ji, Shih-Fu Chang, Kathleen McKeown, Mohit Bansal, Avi Sil
Proceedings of **ACL 2021**. [pdf]

134. Analysis of Tree-Structured Architectures for Code Generation
Samip Dahal, Adyasha Maharana, Mohit Bansal
Findings of **ACL 2021** (short papers). [pdf]

133. ChrEnTranslate: Cherokee-English Machine Translation Demo with Quality Estimation and Corrective Feedback
Shiyue Zhang, Benjamin Frey and Mohit Bansal
Proceedings of **ACL 2021** (demo papers). [pdf]

132. Disentangling Online Chats with DAG-structured LSTMs
Duccio Pappadopulo*, Lisa Bauer*, Marco Farina, Ozan ?rsoy, Mohit Bansal
Proceedings of ***SEM 2021**. [pdf]

131. multiPRover: Generating Multiple Proofs for Improved Interpretability in Rule Reasoning
Swarnadeep Saha, Prateek Yadav and Mohit Bansal
Proceedings of **NAACL 2021**. [pdf]

130. Improving Generation and Evaluation of Visual Stories via Semantic Consistency
Adyasha Maharana, Darryl Hannan and Mohit Bansal
Proceedings of **NAACL 2021**. [pdf]

129. DeCEMBERT: Learning from Noisy Instructional Videos via Dense Captions and Entropy Minimization
Zineng Tang*, Jie Lei* and Mohit Bansal
Proceedings of **NAACL 2021**. [pdf]

128. Improving Cross-Modal Alignment in Vision Language Navigation via Syntactic Information
Jialu Li, Hao Tan and Mohit Bansal
Proceedings of **NAACL 2021** (short papers). [pdf]

127. Dynabench: Rethinking Benchmarking in NLP
Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts and Adina Williams
Proceedings of **NAACL 2021**. [pdf]

126. Enriching Transformers with Structured Tensor-Product Representations for Abstractive Summarization
Yichen Jiang, Asli Celikyilmaz, Paul Smolensky, Paul Soulos, Sudha Rao, Hamid Palangi, Roland Fernandez, Caitlin Smith, Mohit Bansal, and Jianfeng Gao
Proceedings of **NAACL 2021**. [pdf]

125. Efficiently Summarizing Text and Graph Encodings of Multi-Document Clusters
Ramakanth Pasunuru, Mengwen Liu, Mohit Bansal, Sujith Ravi and Markus Dreyer
Proceedings of **NAACL 2021**. [pdf]

124. Extending Multi-Document Summarization Evaluation to the Interactive Setting
Ori Shapira, Ramakanth Pasunuru, Hadar Ronen, Mohit Bansal, Yael Amsterdamer and Ido Dagan
Proceedings of **NAACL 2021**. [pdf]

123. Robustness Gym: Unifying the NLP Evaluation Landscape
Karan Goel, Nazneen Fatema Rajani, Jesse Vig, Zachary Taschdjian, Mohit Bansal and Christopher Re
Proceedings of **NAACL 2021** (demo papers). [pdf]

122. ERNIE-NLI: Analyzing the Impact of Domain-Specific External Knowledge on Enhanced Representations for NLI
Lisa Bauer, Lingjia Deng, Mohit Bansal
Proceedings of **DeeLIO Workshop, NAACL 2021**. [pdf]

121. GENE: Global Event Network Embedding
Qi Zeng, Manling Li, Tuan Lai, Heng Ji, Mohit Bansal, Hanghang Tong
Proceedings of **TextGraphs Workshop, NAACL 2021**. [pdf]

120. The Effect of Pretraining on Extractive Summarization for Scientific Documents
Yash Gupta, Pawan Sasanka, Shikha Bordia, Arjun Manoharan, Deepak Mittal, Ramakanth Pasunuru, Manish Shrivastava, Maneesh Singh, Mohit Bansal, Preethi Jyothi
Proceedings of **Scholarly Document Processing Workshop, NAACL 2021**. [pdf]

119. Less is More: ClipBERT for Video-and-Language Learning via Sparse Sampling
Jie Lei*, Linjie Li*, Luowei Zhou, Zhe Gan, Tamara L. Berg, Mohit Bansal, Jingjing Liu
Proceedings of **CVPR 2021**. [pdf]
(CVPR Best Student Paper Honorable Mention)

118. Identify, Align, and Integrate: Matching Knowledge Graphs to Commonsense Reasoning Tasks
Lisa Bauer and Mohit Bansal
Proceedings of **EACL 2021**. [pdf]

117. Hidden Biases in Unreliable News Detection Datasets
Xiang Zhou, Heba Elfardy, Christos Christodoulopoulos, Thomas Butler and Mohit Bansal
Proceedings of **EACL 2021**. [pdf]
(EACL Best Long Paper Honorable Mention)

116. FixMyPose: Pose Correctional Captioning and Retrieval
Hyounghun Kim*, Abhaysinh Zala*, Graham Burri, and Mohit Bansal
Proceedings of **AAAI 2021**. [pdf]

115. Data Augmentation for Abstractive Query-Focused Multi-Document Summarization
Ramakanth Pasunuru, Asli Celikyilmaz, Michel Galley, Chenyan Xiong, Yizhe Zhang, Mohit Bansal, and Jianfeng Gao
Proceedings of **AAAI 2021**. [pdf]

114. Dual Reinforcement-Based Specification Generation for Image De-Rendering
Ramakanth Pasunuru, David Rosenberg, Gideon Mann, and Mohit Bansal
Proceedings of **Scientific Document Understanding Workshop, AAAI 2021**. [pdf]

113. ChrEn: Cherokee-English Machine Translation for Endangered Language Revitalization
Shiyue Zhang, Benjamin Frey, and Mohit Bansal
Proceedings of **EMNLP 2020**. [pdf]

112. Vokenization: Improving Language Understanding via Contextualized, Visually-Grounded Supervision
Hao Tan and Mohit Bansal
Proceedings of **EMNLP 2020**. [pdf]

111. What Can We Learn from Collective Human Opinions on Natural Language Inference Data?
Yixin Nie, Xiang Zhou, and Mohit Bansal
Proceedings of **EMNLP 2020**. [pdf]

110. What Is More Likely To Happen Next? Video-and-Language Future Event Prediction
Jie Lei, Licheng Yu, Tamara Berg, and Mohit Bansal
Proceedings of **EMNLP 2020**.

109. ConjNLI: Natural Language Inference Over Conjunctive Sentences
Swarnadeep Saha, Yixin Nie, and Mohit Bansal
Proceedings of **EMNLP 2020**.

108. PRover: Proof Generation for Interpretable Reasoning over Rules
Swarnadeep Saha, Sayan Ghosh, Shashank Srivastava and Mohit Bansal
Proceedings of **EMNLP 2020**. [pdf]

107. DORB: Dynamically Optimizing Multiple Rewards with Bandits
Ramakanth Pasunuru, Han Guo, and Mohit Bansal
Proceedings of **EMNLP 2020**.

106. The Curse of Performance Instability in Analysis Datasets: Consequences, Source, and Suggestions
Xiang Zhou, Yixin Nie, Hao Tan, and Mohit Bansal
Proceedings of **EMNLP 2020**. [pdf]

105. Leakage-Adjusted Simulatability: Can Models Generate Non-Trivial Explanations of Their Behavior in Natural Language?
Peter Hase, Shiyue Zhang, Harry Xie, and Mohit Bansal
Findings of **EMNLP 2020**. [pdf]

104. ArraMon: A Joint Navigation-Assembly Instruction Interpretation Task in Dynamic Environments
Hyounghun Kim, Abhaysinh Zala, Graham Burri, Hao Tan, and Mohit Bansal
Findings of **EMNLP 2020**.

103. HoVer: A Dataset for Many-Hop Fact Extraction And Claim Verification
Yichen Jiang*, Shikha Bordia*, Zheng Zhong, Charles Dognin, Maneesh Singh, and Mohit Bansal
Findings of **EMNLP 2020**.

102. Adversarial Augmentation Policy Search for Domain and Cross-Lingual Generalization in Reading Comprehension
Adyasha Maharana and Mohit Bansal
Findings of **EMNLP 2020**. [pdf]

101. FENAS: Flexible and Expressive Neural Architecture Search
Ramakanth Pasunuru and Mohit Bansal
Findings of **EMNLP 2020** (short papers).

100. TVR: A Large-Scale Dataset for Video-Subtitle Moment Retrieval
Jie Lei, Licheng Yu, Tamara L. Berg, Mohit Bansal
Proceedings of **ECCV 2020**. [pdf]

99. Diagnosing the Environment Bias in Vision-and-Language Navigation
Yubo Zhang*, Hao Tan*, and Mohit Bansal
Proceedings of **IJCAI 2020**. [pdf]

98. Evaluating Explainable AI: Which Algorithmic Explanations Help Users Predict Model Behavior?
Peter Hase and Mohit Bansal
Proceedings of **ACL 2020**. [pdf]

97. Towards Robustifying NLI Models Against Lexical Dataset Biases
Xiang Zhou and Mohit Bansal
Proceedings of **ACL 2020**.

96. Adversarial NLI: A New Benchmark for Natural Language Understanding
Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela
Proceedings of **ACL 2020**. [pdf]

95. Dense-Caption Matching and Frame-Selection Gating for Temporal Localization in VideoQA
Hyounghun Kim, Zineng Tang, and Mohit Bansal
Proceedings of **ACL 2020**.

94. MART: Memory-Augmented Recurrent Transformer for Coherent Video Paragraph Captioning
Jie Lei, Liwei Wang, Yelong Shen, Dong Yu, Tamara Berg, and Mohit Bansal
Proceedings of **ACL 2020**.

93. TVQA+: Spatio-Temporal Grounding for Video Question Answering
Jie Lei, Licheng Yu, Tamara L. Berg, and Mohit Bansal
Proceedings of **ACL 2020**. [pdf]

92. Simple Compounded-Label Training for Fact Extraction and Verification
Yixin Nie*, Lisa Bauer*, and Mohit Bansal
Proceedings of **Fact Extraction and VERification (FEVER) workshop, ACL 2020**. [pdf]

91. Multi-Source Domain Adaptation for Text Classification via DistanceNet-Bandits
Han Guo, Ramakanth Pasunuru, and Mohit Bansal
Proceedings of **AAAI 2020**. [pdf]

90. ManyModalQA: Modality Disambiguation and QA over Diverse Inputs
Darryl Hannan, Akshay Jain, and Mohit Bansal
Proceedings of **AAAI 2020**. [pdf]

89. AvgOut: A Simple Output-Probability Measure to Eliminate Dull Responses
Tong Niu and Mohit Bansal
Proceedings of **AAAI 2020**. [pdf]

88. Modality-Balanced Models for Visual Dialogue
Hyounghun Kim, Hao Tan, and Mohit Bansal
Proceedings of **AAAI 2020**. [pdf]

87. Enabling Robots to Understand Incomplete Natural Language Instructions Using Commonsense Reasoning
Haonan Chen, Hao Tan, Alan Kuntz, Mohit Bansal, Ron Alterovitz
Proceedings of **ICRA 2020**. [pdf]

86. LXMERT: Learning Cross-Modality Encoder Representations from Transformers
Hao Tan and Mohit Bansal
Proceedings of **EMNLP 2019**. [pdf]

85. Self-Assembling Modular Networks for Interpretable Multi-Hop Reasoning
Yichen Jiang and Mohit Bansal
Proceedings of **EMNLP 2019**. [pdf]

84. Addressing Semantic Drift in Question Generation for Semi-Supervised Question Answering
Shiyue Zhang and Mohit Bansal
Proceedings of **EMNLP 2019**. [pdf]

83. Revealing the Importance of Semantic Retrieval for Machine Reading at Scale
Yixin Nie, Songhe Wang, and Mohit Bansal
Proceedings of **EMNLP 2019**. [pdf]

82. Automatically Learning Data Augmentation Policies for Dialogue Tasks
Tong Niu and Mohit Bansal
Proceedings of **EMNLP 2019** (short papers). [pdf]

81. Continual and Multi-Task Architecture Search
Ramakanth Pasunuru and Mohit Bansal
Proceedings of **ACL 2019**. [pdf]

80. Avoiding Reasoning Shortcuts: Adversarial Evaluation, Training, and Model Development for Multi-Hop QA
Yichen Jiang and Mohit Bansal
Proceedings of **ACL 2019**. [pdf]

79. Explore, Propose, and Assemble: An Interpretable Model for Multi-Hop Reading Comprehension
Yichen Jiang, Nitish Joshi, Yen-Chun Chen, and Mohit Bansal
Proceedings of **ACL 2019**. [pdf]

78. Expressing Visual Relationships via Language
Hao Tan, Franck Dernoncourt, Zhe Lin, Trung Bui, and Mohit Bansal
Proceedings of **ACL 2019**. [pdf]

77. Improving Visual Question Answering by Referring to Generated Paragraph Captions
Hyounghun Kim and Mohit Bansal
Proceedings of **ACL 2019** (short papers). [pdf]
(ACL Best Short Paper Nominee)

76. PaperRobot: Incremental Draft Generation of Scientific Ideas
Qingyun Wang, Lifu Huang, Zhiying Jiang, Kevin Knight, Heng Ji, Mohit Bansal, and Yi Luan
Proceedings of **ACL 2019**. [pdf]

75. Learning to Navigate Unseen Environments: Back Translation with Environmental Dropout
Hao Tan, Licheng Yu, and Mohit Bansal
Proceedings of **NAACL 2019**. [pdf]
(1st Rank Model in Room-to-Room Vision-Language-Navigation Leaderboard)

74. AutoSeM: Automatic Task Selection and Mixing in Multi-Task Learning
Han Guo, Ramakanth Pasunuru, and Mohit Bansal
Proceedings of **NAACL 2019**. [pdf]

73. Crowdsourcing Lightweight Pyramids for Manual Summary Evaluation
Ori Shapira, David Gabay, Yang Gao, Hadar Ronen, Ramakanth Pasunuru, Mohit Bansal, Yael Amsterdamer, and Ido Dagan
Proceedings of **NAACL 2019** (short papers). [pdf]

72. Multi-Target Embodied Question Answering
Licheng Yu, Xinlei Chen, Georgia Gkioxari, Mohit Bansal, Tamara L. Berg, and Dhruv Batra
Proceedings of **CVPR 2019**. [pdf]

71. Efficient Generation of Motion Plans from Attribute-Based Natural Language Instructions Using Dynamic Constraint Mapping
Jae Sung Park, Biao Jia, Mohit Bansal, Dinesh Manocha
Proceedings of **ICRA 2019**. [pdf]

70. Combining Fact Extraction and Verification with Neural Semantic Matching Networks
Yixin Nie, Haonan Chen, and Mohit Bansal
Proceedings of **AAAI 2019**. [pdf]

69. Analyzing Compositionality-Sensitivity of NLI Models
Yixin Nie, Yicheng Wang, and Mohit Bansal
Proceedings of **AAAI 2019**. [pdf]

68. DSTC7-AVSD: Scene-Aware Video-Dialogue Systems with Dual Attention
Ramakanth Pasunuru, Mohit Bansal
Proceedings of **Dialog System Technology Challenges Workshop, AAAI 2019**. [pdf]
(selected oral, rank-3)

67. Closed-Book Training to Improve Summarization Encoder Memory
Yichen Jiang and Mohit Bansal
Proceedings of **EMNLP 2018**. [pdf]

66. SafeCity: Understanding Diverse Forms of Sexual Harassment Personal Stories
Sweta Karlekar and Mohit Bansal
Proceedings of **EMNLP 2018** (short papers). [pdf]

65. Commonsense for Generative Multi-Hop Question Answering Tasks
Lisa Bauer, Yicheng Wang, and Mohit Bansal
Proceedings of **EMNLP 2018**. [pdf]

64. Game-Based Video-Context Dialogue
Ramakanth Pasunuru and Mohit Bansal
Proceedings of **EMNLP 2018**. [pdf]

63. TVQA: Localized, Compositional Video Question Answering
Jie Lei, Licheng Yu, Mohit Bansal, and Tamara Berg
Proceedings of **EMNLP 2018**. [pdf]

62. Incorporating Background Knowledge into Video Description Generation
Spencer Whitehead, Heng Ji, Mohit Bansal, Shih-Fu Chang, and Clare Voss
Proceedings of **EMNLP 2018**. [pdf]

61. Adversarial Over-Sensitivity and Over-Stability Strategies for Dialogue Models
Tong Niu and Mohit Bansal
Proceedings of **CoNLL 2018**. [pdf]

60. Combining Fact Extraction and Claim Verification in an NLI Model
Yixin Nie, Haonan Chen, and Mohit Bansal
In **Fact Extraction and Verification (FEVER) Workshop, EMNLP 2018**. [pdf]
(1st Rank Model in Shared Task)

59. Dynamic Multi-Level, Multi-Task Learning for Sentence Simplification
Han Guo, Ramakanth Pasunuru and Mohit Bansal
Proceedings of **COLING 2018**. [pdf]
('Area Chair Favorites' Paper Award)

58. Polite Dialogue Generation Without Parallel Data
Tong Niu and Mohit Bansal
Proceedings of **TACL 2018**. [pdf]

57. Fast Abstractive Summarization with Reinforce-Selected Sentence Rewriting
Yen-Chun Chen and Mohit Bansal
Proceedings of **ACL 2018**. [pdf]

56. Soft, Layer-Specific Multi-Task Summarization with Entailment and Question Generation
Han Guo, Ramakanth Pasunuru, and Mohit Bansal
Proceedings of **ACL 2018**. [pdf]

55. #MeToo: Neural Detection and Explanation of Language in Personal Abuse Stories
Sweta Karlekar and Mohit Bansal
Proceedings of **WiNLP 2018 (Widening NLP Workshop), NAACL 2018**. [pdf]

54. Object Ordering with Bidirectional Matchings for Visual Reasoning
Hao Tan and Mohit Bansal
Proceedings of **NAACL 2018** (short papers). [pdf]
(Top Image Leaderboard Position)

53. Multi-Reward Reinforced Summarization with Saliency and Entailment
Ramakanth Pasunuru and Mohit Bansal
Proceedings of **NAACL 2018** (short papers). [pdf]

52. Detecting Linguistic Characteristics of Alzheimer's Dementia by Interpreting Neural Models
Sweta Karlekar, Tong Niu, and Mohit Bansal
Proceedings of **NAACL 2018** (short papers). [pdf]

51. Robust Machine Comprehension Models via Adversarial Training
Yicheng Wang and Mohit Bansal
Proceedings of **NAACL 2018** (short papers). [pdf]

50. Punny Captions: Witty Wordplay in Image Descriptions
Arjun Chandrasekaran, Devi Parikh, and Mohit Bansal
Proceedings of **NAACL 2018** (short papers). [pdf]

49. Joint Modeling of Text and Acoustic-Prosodic Cues for Neural Parsing
Trang Tran, Shubham Toshniwal, Mohit Bansal, Kevin Gimpel, Karen Livescu, and Mari Ostendorf
Proceedings of **NAACL 2018**. [pdf]

48. MAttNet: Modular Attention Network for Referring Expression Comprehension
Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara Berg
Proceedings of **CVPR 2018**. [pdf]

47. Source-Target Inference Models for Spatial Instruction Understanding
Hao Tan and Mohit Bansal
Proceedings of **AAAI 2018**. [pdf]

46. Retweet Wars: Tweet Popularity Prediction via Multimodal Regression
Ke Wang, Mohit Bansal, and Jan-Michael Frahm
Proceedings of **WACV 2018**. [pdf]

45. Interactive-Length Multi-Task Video Captioning with Cooperative Feedback
Han Guo, Ramakanth Pasunuru, and Mohit Bansal
Proceedings of **NIPS 2017** (demo papers).

44. Reinforced Video Captioning with Entailment Rewards
Ramakanth Pasunuru and Mohit Bansal
Proceedings of **EMNLP 2017** (short papers). [pdf]

43. Hierarchically-Attentive RNN for Album Summarization and Storytelling
Licheng Yu, Mohit Bansal, and Tamara Berg
Proceedings of **EMNLP 2017** (short papers). [pdf]

42. Video Highlight Prediction Using Audience Chat Reactions
Cheng-Yang Fu, Joon Lee, Mohit Bansal, and Alexander Berg
Proceedings of **EMNLP 2017** (short papers). [pdf]

41. 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)

40. Towards Improving Abstractive Summarization via Entailment Generation
Ramakanth Pasunuru, Han Guo, and Mohit Bansal
Proceedings of **Summarization Frontiers Workshop, EMNLP 2017**. [pdf]
(Contributed Talk)

39. 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)

38. 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)

37. 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]

36. Contextual RNN-GANs for Abstract Reasoning Diagram Generation
Arnab Ghosh, Viveka Kulharia, Amitabha Mukerjee, Vinay Namboodiri, and Mohit Bansal
Proceedings of **AAAI 2017**. [pdf]

35. Coherent Dialogue with Attention-based Language Models
Hongyuan Mei, Mohit Bansal, and Matthew Walter
Proceedings of **AAAI 2017**. [pdf]

34. Interpreting Neural Networks to Improve Politeness Comprehension
Malika Aubakirova and Mohit Bansal
Proceedings of **EMNLP 2016** (short papers). [pdf]

33. 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]

32. 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]

31. 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]

30. Charagram: Embedding Words and Sentences via Character n-grams
John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu
Proceedings of **EMNLP 2016**. [pdf]

29. End-to-end Relation Extraction using LSTMs on Sequences and Tree Structures
Makoto Miwa and Mohit Bansal
Proceedings of **ACL 2016**. [pdf]

28. 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)

27. 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]

26. 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]

25. 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)

24. Towards Universal Paraphrastic Sentence Embeddings
John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu
Proceedings of **ICLR 2016**. [pdf]
(Oral; 5.7% accep. rate)

23. 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)

22. Machine Comprehension with Syntax, Frames, and Semantics
Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester
Proceedings of **ACL 2015** (short papers). [pdf]

21. From Paraphrase Database to Compositional Paraphrase Model and Back
John Wieting, Mohit Bansal, Kevin Gimpel, Karen Livescu, and Dan Roth
Proceedings of **TACL** (presented at **EMNLP 2015**). [pdf]

20. Dependency Link Embeddings: Continuous Representations of Syntactic Substructures
Mohit Bansal
Proceedings of **Workshop on Vector Space Modeling for NLP, NAACL 2015**. [pdf]
(Selected oral)

19. 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]

18. 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]

17. Accurate Vision-based Vehicle Localization using Satellite Imagery
Hang Chu, Hongyuan Mei, Mohit Bansal, and Matthew R. Walter
Proceedings of **NIPS 2015 Workshop on Transfer and Multi-Task Learning**. [pdf]

16. Weakly-Supervised Learning with Cost-Augmented Contrastive Estimation
Kevin Gimpel and Mohit Bansal
Proceedings of **EMNLP 2014**. [pdf]

15. Tailoring Continuous Word Representations for Dependency Parsing
Mohit Bansal, Kevin Gimpel, and Karen Livescu
Proceedings of **ACL 2014** (short papers). [pdf]

14. 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)

13. 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]

12. Good, Great, Excellent: Global Inference of Semantic Intensities
Gerard de Melo and Mohit Bansal
Proceedings of **TACL** (presented at **ACL 2013**). [pdf]

11. Coreference Semantics from Web Features
Mohit Bansal and Dan Klein
Proceedings of **ACL 2012**. [pdf]

10. Unsupervised Translation Sense Clustering
Mohit Bansal, John DeNero, and Dekang Lin
Proceedings of **NAACL 2012**. [pdf]

9. Web-Scale Features for Full-Scale Parsing
Mohit Bansal and Dan Klein
Proceedings of **ACL 2011**. [pdf]

8. Gappy Phrasal Alignment by Agreement
Mohit Bansal, Chris Quirk, and Robert C. Moore
Proceedings of **ACL 2011**. [pdf]

7. The Surprising Variance in Shortest-Derivation Parsing
Mohit Bansal and Dan Klein
Proceedings of **ACL 2011** (short papers). [pdf]

6. Mention Detection: Heuristics for the OntoNotes annotations
Jonathan K. Kummerfeld, Mohit Bansal, David Burkett, and Dan Klein
Proceedings of **CoNLL 2011** (shared task). [pdf]

5. Simple, Accurate Parsing with an All-Fragments Grammar
Mohit Bansal and Dan Klein
Proceedings of **ACL 2010**. [pdf]

4. Efficient Parsing for Transducer Grammars
   John DeNero, Mohit Bansal, Adam Pauls, and Dan Klein
   Proceedings of **NAACL 2009**. [pdf]

3. 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]

2. 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]

1. 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]

Patents:

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

Recent Invited Talks/Keynotes

Keynote slides also available at link.

Distinguished Speaker Series, University of Virginia

Distinguished Lecture (university-wide), Georgetown University

India Keynotes: ACM CODS 2025; CVIP 2025

Keynote, ICDM 2025 Workshop on Reasoning, Agents, Retrieval, and Attribution [New: 'Multimodal Retrieval for Understanding and Generation across Diverse Sources']

Keynote, EMNLP 2025 5th New Frontiers in Summarization Workshop [New: 'Attributable, Conflict-Robust, Multimodal Summarization with Multi-Source Retrieval']

Keynote, ICCV 2025 4th Workshop on What is Next in Multimodal Foundation Models ['Multimodal Generative Models: Unification and Composable Generalization']

Keynote, 28th European Conference on Artificial Intelligence (ECAI), 2025

Distinguished Lecture Series, MSU

Distinguished Lecture Series, GMU

Keynote, Midwest Speech and Language Days (MSLD) 2025

Keynote, NLP@Michigan Day, 2025

Distinguished Lecture Series, StonyBrook

Distinguished Speaker Series, USF

Keynote, Amazon Research Days

Keynote, TTIC Multimodal AI Workshop

Keynote Speaker, Lisbon Machine Learning Summer School (LxMLS) 2024

Keynote Speaker, 12th International Advanced Summer School on NLP (IASNLP) 2024

Keynote Speaker, Pinterest Labs ML Symposium 2024

Keynote Speaker, TTIC Summer Workshop on Multimodal AI

Invited Speaker, Korea NLP/LM Workshop

USC Distinguished Lecture Series

Keynote Speaker, SouthNLP Symposium 2024

Invited Speaker, IndoML 2023

IISc Bangalore (Kotak AI-ML Centre)

Keynote Speaker, 27th Conference on Computational Natural Language Learning (CoNLL) 2023

Keynote Speaker, AACL-IJCNLP 2023

Invited Speaker, TTI-Chicago 20th-Anniversary Symposium

Distinguished Speaker Series, Auburn University

Penn State NLP Colloquium Series

TWIML AI Podcast: 'Unifying Vision and Language Models'

ACL 2023 Narrative Understanding Workshop

CVPR 2023 Explainable AI for Computer Vision (XAI4CV) Workshop

IBM Neuro-Symbolic AI Workshop 2023

MBZUAI AI Quorum's Inaugural NLP Symposium

Invited Talks, Stanford, UT Austin, ODSC, ICML Workshop, COLING Workshop, 2022

Keynote Speaker, 15th International Natural Language Generation Conference (INLG), 2022

Open-Domain Retrieval Under Multimodal Settings Workshop (O-DRUM), CVPR 2022

Robustness in Sequential Data Workshop (ROSE), CVPR 2022

Automatic Summarization for Creative Writing Workshop, COLING 2022

Indian Symposium on Machine Learning (IndoML), 2021

Fact Extraction and VERification (FEVER) Workshop, EMNLP 2021

Closing the Loop Between Vision and Language (CLVL) Workshop, ICCV 2021

Human Interaction for Robotic Navigation Workshop, ICCV 2021

CVIT Summer School on Artificial Intelligence, 2021

Advances in Language and Vision Research (ALVR) Workshop, NAACL 2021

Person in Context Workshop, CVPR 2021

VQA Workshop, CVPR 2021

IJCAI 2020 Early Career Spotlight Talk

Singapore Symposium on Natural Language Processing (SSNLP 2020) ['Towards Knowledge-Robust and Multimodally-Grounded NLP': video]

3rd Workshop on Neural Generation and Translation (WNGT @ EMNLP 2019) ['Knowledgeable and Multimodal Language Generation': slides]

1st Workshop on Beyond Vision and Language: Integrating Knowledge from Real-World (LANTERN @ EMNLP 2019) ['Knowledgeable and Dynamic Spatio-Temporal Language+Vision+Robotics': slides]

Workshop on Machine Reading for Question Answering (MRQA @ EMNLP 2019) ['Interpretability and

Robustness for Multi-Hop QA': slides]

4th Workshop on Representation Learning for NLP (RepL4NLP @ ACL 2019) ['Knowledgeable and Adversarially-Robust Representation Learning': slides]

RSS-2018 Natural Human-Robot Communication Workshop ['Spatially-Grounded, Personable, and Sensible Human-Robot Communication']

'Grounded, Personable, and Adversarial Dialog Models', *GeorgiaTech, Google Assistant and Dialog Workshop, UNC-SAS-NVidia Deep Learning Symposium, Kenan Rethinc Machine Learning Symposium, IBM NCTEC Conference* (June-Dec, 2018)

'Multi-Task and Reinforcement Learning for Entailment-Based Natural Language Generation', *JHU, Bloomberg, Google, Facebook, UIUC, UWash, Amazon, Triangle ML Day, RTI* (May 2017-June 2018)

'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)

Teaching

Instructor, Multimodal AI: Connecting Language to Vision and Robotics (COMP 690), UNC Chapel Hill, Spring 2026.

Instructor, Multimodal AI: Connecting Language to Vision and Robotics (COMP 590+790), UNC Chapel Hill, Spring 2025.

Instructor, Connecting Language to Vision and Robotics (COMP 590+790), UNC Chapel Hill, Spring 2023.

Instructor, Connecting Language to Vision and Robotics (COMP 590+790), UNC Chapel Hill, Fall 2021.

Instructor, Natural Language Processing (COMP 786), UNC Chapel Hill, Fall 2020.

Instructor, First-Year Honors Undergraduate: Human and Artificial Intelligence Through the Prism of Language, Fall 2019.

Instructor, Advanced Topics in Natural Language Processing: Recent Progress in Different Learning Paradigms (COMP 790.139), UNC Chapel Hill, Spring 2019.

Instructor, Advanced Topics in Natural Language Processing: Conversational Models (COMP 790.139), UNC Chapel Hill, Spring 2018.

Instructor, Natural Language Processing (COMP 790.139), UNC Chapel Hill, Fall 2017.

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.