

Bhuwan Dhingra

Assistant Professor

Computer Science Department, Duke University

Research Scientist

Google

Email: bdhingra@google.com

Home: <http://www.cs.cmu.edu/~bdhingra/>

Github: <https://github.com/bdhingra>

Education

- 2015–2020 Ph.D. in Language & Information Technologies,
Carnegie Mellon University
Supervisors: William W. Cohen and Ruslan Salakhutdinov
Thesis: End-to-end learning with text and knowledge bases
- 2012–2013 M.Tech., Electrical Engineering, Indian Institute of Technology, Kanpur
- 2008–2012 B.Tech., Electrical Engineering, Indian Institute of Technology, Kanpur

Professional Experience

- 2021-Present Assistant Professor, Computer Science, Duke University, Durham USA
- 2020-Present Research Scientist, Google AI, Seattle USA
- 2019-2020 Student Researcher, Google Research, Pittsburgh USA
- 2018 Summer Research Intern, Google AI, New York USA
- 2017 Summer Research Intern, Google Brain, Mountain View USA
- 2016 Summer Research Intern, Microsoft Research, Redmond USA
- 2013–2015 Systems Engineer, Qualcomm Research, San Diego USA
- 2011, 2012 Summer Research Intern, Qualcomm Research, San Diego USA

Selected Honours and Awards

- 2019 Siemens FutureMakers PhD Fellowship
- 2019 CMU 3-Minute Thesis Championship Winner
- 2013 Israel Planning & Budgeting Committee Summer Scholarship
- 2012 Best Student Paper Award, IEEE SCES, Allahabad, India
- 2011 Todai IIT Undergraduate Student Scholarship, IIT Kanpur
- 2010 Summer Undergraduate Research Grant for Excellence, IIT Kanpur
- 2009–2012 All India Engineering Entrance Examination Merit Scholarship

Teaching

Fall 2021 CS590, Introduction to NLP (*graduate*)

Duke University

Publications

Refereed Conferences

- [1] Reasoning over virtual knowledge bases with open predicate relations.
Haitian Sun, Pat Verga, Bhuwan Dhingra, Ruslan Salakhutdinov, William W. Cohen.
The Thirty-eighth International Conference on Machine Learning (ICML, 2021).
- [2] Fool Me Twice: Entailment from Wikipedia Gamification.
Julian Eisenschlos, Bhuwan Dhingra, Jannis Bulian, Benjamin Börschinger, Jordan Boyd-Graber.
19th Annual Conference of the North American Chapter of the ACL (NAACL, 2021).
- [3] Differentiable open-ended commonsense reasoning.
Bill Yuchen Lin, Haitian Sun, Bhuwan Dhingra, Manzil Zaheer, Xiang Ren, William W. Cohen
19th Annual Conference of the North American Chapter of the ACL (NAACL, 2021).
- [4] ToTTo: A Controlled Table-To-Text Generation Dataset.
Ankur P. Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui,
Bhuwan Dhingra, Diyi Yang, Dipanjan Das.
Conference on Empirical Methods in Natural Language Processing (EMNLP, 2020).
- [5] Weakly- and Semi-supervised Evidence Extraction.
Danish Pruthi, Bhuwan Dhingra, Graham Neubig, Zachary C. Lipton.
Findings of Empirical Methods in Natural Language Processing (Findings of EMNLP, 2020).
- [6] Differentiable Reasoning over a Virtual Knowledge Base.
Bhuwan Dhingra, Manzil Zaheer, Vidhisha Balachandran, Graham Neubig,
Ruslan Salakhutdinov, William W. Cohen.
International Conference on Learning Representations (ICLR, 2020).
- [7] Learning to Deceive with Attention-Based Explanations.
Danish Pruthi, Mansi Gupta, Bhuwan Dhingra, Graham Neubig, Zachary C. Lipton.
58th annual meeting of Association of Computational Linguistics (ACL, 2020).
- [8] PubMedQA: A Dataset for Biomedical Research Question Answering
Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W. Cohen, Xinghua Lu.
Conference on Empirical Methods in Natural Language Processing (EMNLP, 2019).
- [9] Handling Divergent Reference Texts when Evaluating Table-to-Text Generation.
Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang,
Dipanjan Das, William W. Cohen.
57th annual meeting of Association of Computational Linguistics (ACL, 2019).
- [10] Combating Adversarial Misspellings with Robust Word Recognition.
Danish Pruthi, Bhuwan Dhingra, Zachary C. Lipton.
57th annual meeting of Association of Computational Linguistics (ACL, 2019).
- [11] Text Generation with Exemplar-based Adaptive Decoding.
Hao Peng, Ankur Parikh, Manaal Faruqui, Bhuwan Dhingra, Dipanjan Das.
17th Annual Conference of the North American Chapter of the ACL (NAACL, 2019).

- [12] GLoMo: Unsupervisedly Learned Relational Graphs as Transferable Representations.
Zhilin Yang*, Jake Zhao*, Bhuwan Dhingra, Kaiming He,
William W. Cohen, Ruslan Salakhutdinov, Yann LeCun.
32nd Conference on Neural Information Processing Systems (NeurIPS, 2018).
- [13] Open Domain Question Answering Using Early Fusion of Knowledge Bases and Text.
Haitian Sun*, Bhuwan Dhingra*, Manzil Zaheer, Kathryn Mazaitis,
Ruslan Salakhutdinov, William W. Cohen.
Conference on Empirical Methods in Natural Language Processing (EMNLP, 2018).
- [14] Neural Models for Reasoning over Multiple Mentions using Coreference.
Bhuwan Dhingra, Qiao Jin, Zhilin Yang, William Cohen, Ruslan Salakhutdinov.
16th Annual Conference of the North American Chapter of the ACL (NAACL, 2018).
- [15] Simple and Effective Semi-Supervised Question Answering.
Bhuwan Dhingra*, Danish Pruthi*, Dheeraj Rajagopal*.
16th Annual Conference of the North American Chapter of the ACL (NAACL, 2018).
- [16] Towards End-to-end reinforcement learning of dialogue agents for information access.
Bhuwan Dhingra, Lihong Li, Xiujun Li, Jianfeng Gao, Yun-Nung Chen, Faisal Ahmed, Li Deng.
55th annual meeting of Association of Computational Linguistics (ACL, 2017).
- [17] Gated Attention Readers for Text Comprehension.
Bhuwan Dhingra*, Hanxiao Liu*, Zhilin Yang, William W. Cohen, Ruslan Salakhutdinov.
55th annual meeting of Association of Computational Linguistics (ACL, 2017).
- [18] Words or Characters? Fine-grained Gating for Reading Comprehension.
Zhilin Yang, Bhuwan Dhingra, Ye Yuan, Junjie Hu, William W. Cohen, Ruslan Salakhutdinov.
International Conference on Learning Representations (ICLR, 2017).
- [19] Bootstrapping Distantly Supervised IE using Joint Learning and Small Well-structured Corpora.
Lidong Bing, Bhuwan Dhingra, Kathryn Mazaitis, Jong Hyuk Park, William W. Cohen.
Thirty-First AAAI Conference on Artificial Intelligence (AAAI, 2017).
- [20] Tweet2Vec: Character-Based Distributed Representations for Social Media.
Bhuwan Dhingra, Zhong Zhou, Dylan Fitzpatrick, Michael Muehl, William W. Cohen.
54th annual meeting of Association of Computational Linguistics (ACL, 2016).
- [21] Stock market prediction using hidden markov model.
Aditya Gupta, Bhuwan Dhingra.
IEEE Student's Conference on Engineering & Systems (SCES, 2012).

Master's Dissertation

- [1] Local Quadrature Reconstruction on Smooth Manifolds.
Bhuwan Dhingra.
M.Tech. Dissertation. Indian Institute of Technology Kanpur (2013).
<https://www.cse.iitk.ac.in/users/manifolds/theses/bhuwand/>

Refereed Workshops

- [1] Investigating the Effect of Background Knowledge on Natural Questions.
Vidhisha Balachandran, Bhuwan Dhingra, Haitian Sun, Michael Collins, William W. Cohen.
The 2nd Workshop on Knowledge Extraction and Integration
for Deep Learning Architectures, (DeeLIO, 2021).

- [2] Probing Biomedical Embeddings from Language Models.
Qiao Jin, Bhuwan Dhingra, William W. Cohen, Xinghua Lu.
3rd Workshop for Evaluating Vector Space Representations for NLP, NAACL (2019).
- [3] AttentionMeSH: Simple, Effective and Interpretable Automatic MeSH Indexer.
Qiao Jin*, Bhuwan Dhingra*, William W. Cohen, Xinghua Lu.
6th BioASQ Workshop, EMNLP (2018).
- [4] Embedding Text in Hyperbolic Spaces.
Bhuwan Dhingra, Christopher Shallue, Mohammad Norouzi, Andrew Dai, George Dahl.
12th Workshop on Graph-Based Natural Language Processing, NAACL (2018).
- [5] Answering Cloze-style Software Questions Using Stack Overflow.
Ezra Winston, Bhuwan Dhingra, Kathryn Mazaitis, Graham Neubig, William Cohen.
Workshop on Machine Learning for Programming, FLoC (2018).
- [6] Using Graphs of Classifiers to Impose Constraints on Semi-supervised Relation Extraction.
Lidong Bing, William W. Cohen, Bhuwan Dhingra, Richard C. Wang.
Workshop on Automatic Knowledge Base Completion, NAACL (2016).

Technical reports

- [1] Time-Aware Language Models as Temporal Knowledge Bases.
Bhuwan Dhingra, Jeremy R. Cole, Julian Eisenschlos
Daniel Gillick, Jacob Eisenstein, William W. Cohen.
<https://arxiv.org/abs/2106.15110> (2021).
- [2] Evaluating Explanations: How much do explanations from the teacher aid students?
Danish Pruthi, Bhuwan Dhingra, Livio Baldini Soares, Michael Collins
Zachary C Lipton, Graham Neubig, William W. Cohen.
<https://arxiv.org/abs/2012.00893> (2020).
- [3] Quasar: Datasets for Question Answering by Search and Reading.
Bhuwan Dhingra, Kathryn Mazaitis, William W. Cohen.
<https://arxiv.org/abs/1707.03904> (2017).
- [4] Linguistic Knowledge as Memory for Recurrent Neural Networks.
Bhuwan Dhingra, Zhilin Yang, William W. Cohen, Ruslan Salakhutdinov.
<https://arxiv.org/abs/1703.02620> (2017).
- [5] A Comparative Study of Word Embeddings for Reading Comprehension.
Bhuwan Dhingra, Hanxiao Liu, Ruslan Salakhutdinov, William W. Cohen.
<https://arxiv.org/abs/1703.00993> (2017).
- [6] Question Answering from Unstructured Text by Retrieval and Comprehension.
Yusuke Watanabe, Bhuwan Dhingra, Ruslan Salakhutdinov.
<https://arxiv.org/abs/1703.08885> (2017).
- [7] A User Simulator for Task-Completion Dialogues.
Xiujun Li, Zachary C. Lipton, Bhuwan Dhingra, Lihong Li, Jianfeng Gao, Yun-Nung Chen.
<https://arxiv.org/abs/1612.05688> (2016).

Patents

- [1] End-to-end learning of dialogue agents for information access.
Lihong Li, Bhuwan Dhingra, Jianfeng Gao, Xiujun Li, Yun-Nung Chen, Li Deng, Faisal Ahmed.
U.S. Patent Application No. 15/406,425, 2018.
- [2] Context aware system with multiple power consumption modes.
Tadeusz Jarosinski, Shankar Sadasivam, Ryan Carey, Jinwon Lee, Bhuwan Dhingra,
Abhijeet Bisain, Vitor Carvalho, Rajeev Jain, Muralidhar Akula, and Ashwin Swaminathan.
U.S. Patent 9,622,177, issued April 11, 2017.
- [3] Smartphone motion classifier.
Vitor Carvalho, Bhuwan Dhingra, Edward Harrison Teague,
Siddika Parlak Polatkan, Shankar Sadasivam, and Carlos Manuel Puig.
U.S. Patent Application No. 14/865,665, filed 2015.
- [4] Detecting that a mobile device is riding with a vehicle.
Leonard Henry Gropop and Bhuwan Dhingra.
U.S. Patent Application No. 13/619,143, filed 2012.

Invited Talks & Posters

- *Text as a Virtual Knowledge Base (talk)*.
Stanford NLP Seminar, October 2019.
UC Irvine AI/ML Seminar, October 2019.
IIT Delhi, November 2019.
UMass Amherst ML & Friends Lunch, December 2019.
USC AI Rising Stars Symposium, December 2019.
- *Differentiable Reasoning over a Virtual Knowledge Base (poster)*.
Microsoft AI Breakthroughs Workshop, September 2019.
- *Learning Neural Knowledge Representations (talk)*.
Siemens Research, Princeton, April 2019.
- *Reading Graphs of Facts and Text for Question Answering (talk)*.
Intuit Research, San Diego, July 2018.
- *Improving the Statistical Efficiency of Machine Reading Comprehension (talk)*.
ML Lunch, Carnegie Mellon University, April 2018.
- *Neural Architectures for Reading and Reasoning over Documents (talk)*.
AI Seminar sponsored by Apple, Carnegie Mellon University, September 2017.
- *Gated Attention Readers for Text Comprehension (poster)*.
IBM Cognitive Colloquium, September 2016.

Professional Activities

- Area Chair
2021: ACL-IJCNLP (Interpretability and Analysis of Models)
2020: NAACL (Information Extraction)
- Reviewing
2020: ICML, EMNLP, ACL, NAACL, AAAI, ICLR, JMLR

2019: NAACL (Outstanding Reviewer), ICML, AKBC, ACL, EMNLP.

2018: NeurIPS (Top 30% Reviewer), ACL, EMNLP, SIGIR, CoLing, NLP-CC English.

2017: IJCNLP.

References

Available on request.