Urvashi Khandelwal

CONTACT
INFORMATION

urvashik@google.com
https://urvashik.github.io/

EDUCATION

Stanford University, CA
Ph.D. Computer Science
Advisor: Dan Jurafsky
2015 - 2021

University of Illinois Urbana-Champaign, IL
B.S. Computer Science, Minor in Mathematics
Advisor: Jiawei Han
2011 - 2015

RESEARCH
INTERESTS

Understanding and improving the generalization capabilities of generative models, particularly language models; Model calibration and uncertainty estimation.

WORK
EXPERIENCE

Google
Research Scientist - Google DeepMind
August, 2021-Present

Facebook
Research Intern - Facebook AI Research
Summer, 2020

Facebook
Research Intern - Facebook AI Research
Summer, 2019

Google
Research Intern - Google Brain
Summer-Fall, 2018

Facebook
Software Engineering Intern - Data and Targeting, Ads
Summer, 2014

Google
Software Engineering Intern - Gmail Backend
Summer, 2013

Qualcomm
Engineering Intern - Qualcomm CDMA Technologies
Summer, 2012

AWARDS

Microsoft Research Dissertation Grant
Winner
2020

CRA Outstanding Undergraduate Researchers Award
National Winner
2015

C.W. Gear Outstanding Undergraduate Award
Winner, University of Illinois Urbana-Champaign
2015

PUBLICATIONS

Few-Shot Recalibration of Language Models.
Xiang Lisa Li, **Urvashi Khandelwal**, Kelvin Guu.
ArXiv, 2024.

From Pixels to UI Actions: Learning to Follow Instructions via Graphical User Interfaces.
Peter Shaw, Mandar Joshi, James Cohan, Jonathan Berant, Panupong Pasupat, Hexiang Hu, **Urvashi Khandelwal**, Kenton Lee, Kristina Toutanova.
Neural Information Processing Systems (NeurIPS), 2023.

Open-domain Visual Entity Recognition: Towards Recognizing Millions of Wikipedia Entities.
Hexiang Hu, Yi Luan, Yang Chen, **Urvashi Khandelwal**, Mandar Joshi, Kenton Lee, Kristina Toutanova, Ming-Wei Chang.
International Conference on Computer Vision (ICCV), 2023.

Pix2Struct: Screenshot Parsing as Pretraining for Visual Language Understanding.
Kenton Lee, Mandar Joshi, Iulia Turc, Hexiang Hu, Fangyu Liu, Julian Eisenschlos, **Urvashi Khandelwal**, Peter Shaw, Ming-Wei Chang, Kristina Toutanova.
International Conference on Machine Learning (ICML), 2023.

Nearest Neighbor Machine Translation.
Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer and Mike Lewis.
International Conference on Learning Representations (ICLR), 2021.

With Little Power Comes Great Responsibility.
Dallas Card, Peter Henderson, **Urvashi Khandelwal**, Robin Jia, Kyle Mahowald and Dan Jurafsky.
Empirical Methods in Natural Language Processing (EMNLP), 2020.

Emergent Linguistic Structure in Artificial Neural Networks Trained by Self-Supervision.
Chris Manning, Kevin Clark, John Hewitt, **Urvashi Khandelwal** and Omer Levy.
Proceedings of the National Academy of Sciences (PNAS), 2020.

Generalization through Memorization: Nearest Neighbor Language Models.
Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer and Mike Lewis.
International Conference on Learning Representations (ICLR), 2020.

What does BERT look at? An Analysis of BERT's Attention.
Kevin Clark, **Urvashi Khandelwal**, Omer Levy and Christopher D. Manning.
BlackboxNLP, 2019.
(Best Paper Award)

BAM! Born-Again Multi-Task Networks for Natural Language Understanding.
Kevin Clark, Minh-Thang Luong, **Urvashi Khandelwal**, Christopher D. Manning and Quoc V. Le.
Association for Computational Linguistics (ACL), 2019.

Sample Efficient Text Summarization Using a Single Pre-Trained Transformer.
Urvashi Khandelwal, Kevin Clark, Dan Jurafsky, Lukasz Kaiser.
ArXiv Preprint, 2019. Presented at WestCoast NLP, 2019.

Sharp Nearby, Fuzzy Far Away: How Neural Language Models Use Context.
Urvashi Khandelwal, He He, Peng Qi and Dan Jurafsky.
Association for Computational Linguistics (ACL), 2018.

ClusCite: Effective Citation Recommendation by Information Network Based Clustering.

Xiang Ren, Jialu Liu, Xiao Yu, **Urvashi Khandelwal**, Quanquan Gu, Lidan Wang and Jiawei Han.
International Conference on Knowledge Discovery and Data Mining (KDD), 2014.

Personalized Entity Recommendation in Heterogeneous Information Networks with Implicit User Feedback.
Xiao Yu, Xiang Ren, Yizhou Sun, Quanquan Gu, Bradley Sturt, **Urvashi Khandelwal**, Brandon Norick and Jiawei Han.
International Conference on Web Search and Data Mining, (WSDM), 2014.
(Test of Time Award, WSDM 2024)

HeteRec: Entity Recommendation in Heterogeneous Information Networks with Implicit User Feedback.
Xiao Yu, Xiang Ren, Yizhou Sun, Bradley Sturt, **Urvashi Khandelwal**, Quanquan Gu, Brandon Norick, and Jiawei Han.
International Conference on Recommender Systems (RecSys), 2013.

REPORTS

Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies.
Report submitted to the Administrative Conference of the United States (ACUS).
February 2020.
Served as the technical lead for two case studies:

- Informal Adjudication at the U.S. Patent and Trademark Office
  Daniel E. Ho, **Urvashi Khandelwal**, Alex Yu

- Formal Adjudication at the Social Security Administration
  Daniel E. Ho, Derin McLeod, **Urvashi Khandelwal**, Liza Starr, Emma Wang

INVITED TALKS

The Generalizability and Interpretability of Neural Language Models
Google Research	February, 2021
Carnegie Mellon University	March, 2021
New York University	March, 2021
Facebook AI Research	March, 2021
Microsoft Research	March, 2021
Square	May, 2021

Generalization through Memorization: Nearest Neighbor Language Models
Microsoft Research AI Breakthroughs	September 2020
Berkeley NLP Seminar	November, 2019

Sharp Nearby, Fuzzy Far Away: How Neural Language Models Use Context
Bay Area Research in NLP and ML Meetup	March, 2019

Media Portrayals of AI
Stanford AI Lab - AI Salon	April, 2017

Neural Text Summarization
Stanford Data Science Initiative	October, 2016

PRESS COVERAGE

Facebook's AI speeds up natural language processing without additional training.
VentureBeat, February 19, 2020.

Helpful Neighbors.
The Batch by Andrew Ng. January 29, 2019.

Stanford policy lab explores government use of artificial intelligence.
Stanford News Service, February 28, 2019.

Khandelwal receives CRA Outstanding Undergraduate Researcher award.
Department of Computer Science - CS@Illinois News, January 23, 2015.

SKILLS	Languages: Python, C++, C, MATLAB, LaTeX Frameworks: PyTorch, Tensorflow

TEACHING EXPERIENCE

CS124 - From Languages to Information
Head Teaching Assistant, Stanford   Winter 2019, 2020

CS225 - Data Structures
Teaching Assistant, UIUC   Spring 2013 - Fall 2014

CS173 - Discrete Mathematics
Teaching Assistant, UIUC   Fall 2013

ECE110 - Introduction to Electrical and Computer Engineering
Teaching Assistant, UIUC   Spring 2012

SERVICE

ICML 2022 Workshop on Knowledge Retrieval and Language Models
Co-Organizer   2022

NeuralGen 2019 - Workshop on Methods for Optimizing and Evaluating Neural Language Generation
Co-Organizer   2019

Stanford Computer Science PhD Admissions
Committee Member   2018-2019

Area Chair, Reviewer
NeurIPS, ICLR, ACL, NAACL, EMNLP, NeuralGen (Meta-reviewer), DeepGen, ACL-Student Research Workshop, KDD, RecSys