Vishal Anand

vishal.anand@brown.edu · vishalanand@fas.harvard.edu · vishalanand.net

EDUCATION

- **Brown University**, PhD, Computer Science 2026 - Ongoing
  - AI Superlab: NLP and ML: Foundation Models, Agentic AI
  - Advisor: Stephen Bach
- **Columbia University**, MS, Computer Science, 3.84/4 2017 - 2019
  - NLP Lab: Low Resource learning. Conversational AI / Summary. Machine Reading
  - Class Marshal, PDL Fellow, CA Fellow Tuition Waiver, Student Body Career Chair
  - Advisor: Kathleen McKeown
- **IIT Guwahati**, BTech, Computer Science, 8.1/10 (3.66/4) 2011 - 2015
  - Thesis: Multi-task NLP: word-vector generation, causality evaluation

INVITED TALKS

- 'Interpretable Foundation Models', Pals of Autonomous Agents Meeting, Microsoft, Jan 2026
- 'Very Large Language Models: Self-Updating Inferences On-The-Fly', Bing, Nov 2023
- 'Semantic Understanding: Interaction Tracking in long form data', NIST, Dec 2022
- Panel on Industry Research, Columbia University, Jun 2022
- 'NLP and Localization', Custom MT, Jun 2022
- 'Semantic Inference on Multimodal long form datasets', NIST, Dec 2021
- 'On skipping dense nets: Explainability in DeepNets', AI4Retail Conference, Jun 2019
- 'ChatBots - Context sensitive NLP models in business decisions', eBay, Apr 2018

RESEARCH & EMPLOYMENT

- **Microsoft - Generative AI**, Senior Applied Scientist, 2019 - 2026
  - Four promotions in three years
  - The Microsoft Copilot project: Long-tailed Conversational LLMs
  - Smart-reply: conversational AI, MSR, saves - 294k hrs/year, $3.6 MM/year
  - Understanding text with temporal cues, Microsoft Research MSR program
  - Target extraction, explanation generation - internally published, US patent
  - Value Guidance: Intent Understanding; AutoComplete; Foundation Models
  - Sixteen papers/patents with MSR, Columbia collaborations, multiple promotions
- **NLP Lab, Columbia University**, Research Assistant, 2017 - 2019
  - Multi-lingual low resource transfer learning: Swahili-Tagalog-Somali, German
  - Conversational Summarization with primary text and secondary sound guidance
- **Uber**, Machine Learning Intern, Summer 2018
  - Publicly featured first of 11 selected among 200 interns in Uber's blog [link]
  - 8% CTR improved, Deep-Wide Tensorflow Models to debias under-sampled elasticities
- **Visa**, Senior Software Engineer, 2015 - 2017
  - Anomaly detection in transactions, 2 patents and 1 trade secret filed
  - CFO Vasant Prabhu directly met and recognized my work; fast-tracked promotion
- **Google Summer of Code**, Fellow, Summer 2015
  - Architected and developed open-source Kriging Interpolation libraries, Eclipse, [code]
- **Goldman Sachs London**, Equity Quantitative Strategist, Summer 2014
  - 40 million USD per trade: Hedged risks for Exchange Traded Funds' Algorithmic trading
- **University of Manchester**, Research Assistant, Summer 2013
  - Improved 30% efficiency with Queen's Swarm Optimization variant in ML models

ADVISORY APPOINTMENT

- Advisory board member, NGO: Serving Nicely Foundation, 2021 - Current
  - AI Assisted Education: 6 slums - 10,000 huts covered, 1750 students screened
  - Adopted 3 schools; building equitable mentoring systems across social segments [link]
  - Government: 15 schools approved by India's Government for pilot [link]

US PATENTS

- Xiaofeng Zhu, **Vishal Anand**, Thomas Lin, Matthew Calderwood, Eric Clausen-Brown, Gord Lueck, Wen-Wai Yim, Anuj Jain, Andres D'Elia, and Cheng Wu.
  Training a Learning-to-Rank Model using a Linear Difference Vector, Microsoft.
  GRANTED US 12,632,482 B2, May, 2026

- **Vishal Anand**, Pramodith Ballapuram, Ananya Mishra, and Cheng Wu.
  Building and Using Target-Based Sentiment Models, Microsoft.
  GRANTED US 12,321,701 B2, Jun, 2025

- **Vishal Anand**. Systems & methods for preventing fraud, Visa Inc., 15/674,187, Feb, 2019
- **Vishal Anand**. Security approaches for virtual reality, Visa Inc., 15/184,759, Dec 2017
- **Vishal Anand**. NLP based pre-event occurence prediction, Visa Inc., Trade Secret, Feb 2016

SELECTED PUBLICATIONS

- **Vishal Anand**, and David Rothschild. Autocomplete: A Hidden Query Modifier?
  In Preparation of submission. Experiments completed and analyzed over multiple year time-frame.

- **Vishal Anand**, Milad Alshomary, and Kathleen McKeown.
  iBERT (Interpretable BERT): Interpretable Embeddings via Sense Decomposition.
  Foundational general-purpose model (9700 lines of code), improves SOTA by 8%
  In Proceedings of EACL (Main), Rabat, Morocco, 2026. **Oral Talk**. [paper] [web: iBERT.io]

- Milad Alshomary *, Nikhil Varimalla *, **Vishal Anand**, Smaranda Muresan, and Kathleen McKeown.
  Layered Insights: Generalizable Analysis of Human Style by Leveraging All Transformer Layers.
  In Proceedings of EMNLP (Main), Suzhou, China, 2025. [paper]

- Efsun Sarigolu Kayi *, **Vishal Anand** *, and Smaranda Muresan. MultiSeg: Parallel Data and Subword Information for Learning Bilingual Embeddings in Low Resource Scenarios
  [paper] [code] [visualization]

  In Proceedings of Language Resources and Evaluation Conference (LREC), Marseille, France
  Full paper: LREC - SLTU-CCURL, 2020

***** denotes equal contributing author

ALL PUBLICATIONS

- **Vishal Anand**, and David Rothschild. Autocomplete: A Hidden Query Modifier?
  In Preparation of submission. Experiments completed and analyzed over multiple year time-frame.

- Cheng Wu *, **Vishal Anand** * , Jaya Mandivarapu, Xiya Liu, and Rui Zhuang.
  DA-RAC: Distance-Aware Calibration of LLM Judges for Trustworthy AI Auditing.
  In Proceedings of Culture x AI: Evaluating AI as a Cultural Technology, ICML 2026.

- **Vishal Anand**, Milad Alshomary, and Kathleen McKeown.
  iBERT (Interpretable BERT): Interpretable Embeddings via Sense Decomposition.
  Foundational general-purpose model (9700 lines of code), improves SOTA by 8%
  Oral Talk. In Proceedings of EACL (Main), Rabat, Morocco, 2026. [paper] [web: iBERT.io]

- Milad Alshomary *, Nikhil Varimalla *, **Vishal Anand**, Smaranda Muresan, and Kathleen McKeown.
  Layered Insights: Generalizable Analysis of Human Style by Leveraging All Transformer Layers.
  In Proceedings of EMNLP (Main), Suzhou, China, 2025. [paper]

- Xiaofeng Zhu, Thomas Lin, **Vishal Anand**, Matthew Calderwood, Eric Clausen-Brown, Gord Lueck, Wen-Wai Yim and Cheng Wu. Explicit and Implicit Semantic Ranking Framework. [Acceptance 19.8%] In Proceedings of the ACM Web Conference (WWW), Texas, USA, 2023.
- **Vishal Anand**, Yifei Dong, Raksha Ramesh, Zifan Chen, Yun Chen, Linquan Li and Ching-Yung Lin. Semantic Understanding and Evolving Interaction Tracking in Long-form Multimodal Datasets. [**Leaderboard #1**] In Proceedings of TRECVID 2022, NIST, USA, 2022.
- Raksha Ramesh *, **Vishal Anand** * †, Zifan Chen, Yifei Dong, Yun Chen, and Ching-Yung Lin. Leveraging Text Representation and Face-head Tracking for Long-form Multimodal Semantic Relation Understanding. [**Leaderboard #2**] In Proceedings of the 30th ACM Multimedia Conference, Lisbon, Portugal, 2022.
- **Vishal Anand**, Raksha Ramesh, Boshen Jin, Ziyin Wang, Xiaoxiao Lei, and Ching-Yung Lin. Multi-Modal Language Modelling on Knowledge Graphs for Deep Video Understanding. [**Leaderboard #1**] In Proceedings of the 29th ACM Multimedia Conference, Chengdu, China, 2021.
- **Vishal Anand**, Yuki Miura. PREDISM: Pre-Disaster Modelling With CNN Ensembles for At-Risk Communities. In NeurIPS Workshop: Tackling Climate Change with Machine Learning, 2021
- Raksha Ramesh, **Vishal Anand**, Ziyin Wang, Tianle Zhu, Wenfend Lyu, Serena Yuan, and Ching-Yung Lin. Kinetics and Scene features for intent detection. International Workshop on Deep Video Understanding (DVU2020), Co-located with ACM ICMI 2020, Utrecht, The Netherlands, 2020. **Oral Presentation**.
- **Vishal Anand**, Raksha Ramesh, Ziyin Wang, Yijing Feng, Jiana Feng, Wenfeng Lyu, Tianle Zhu, Serena Yuan, and Ching-Yung Lin. Story Semantic Relationships from Multimodal Cognitions. [**Leaderboard #3**] In Proceedings of the 28th ACM Multimedia Conference, Seattle, USA, 2020.
- Efsun Sarigolu Kayi *, **Vishal Anand** *, and Smaranda Muresan. MultiSeg: Parallel Data and Subword Information for Learning Bilingual Embeddings in Low Resource Scenarios [paper] [code] [visualization]

  In Proceedings of Language Resources and Evaluation Conference (LREC), Marseille, France
  Full paper: LREC - SLTU-CCURL, 2020

***** denotes equal contributing author, † denotes corresponding author

SELECTED HONORS

- Fourth Promotion in three years at Microsoft as Scientist, 2024
- Class Marshal in graduating Columbia class for demonstrating academic achievements, 2019
- CA Fellowship, Columbia: full-tuition waiver for exceptional performance, 2018
- Featured in Uber's blog [link], 11 out of 200 interns, 2018
- Silver Innovation Award: Pioneering with a purpose - session with CFO, Visa, 2016
- Asia Ranked 79, ACM ICPC Asia, Regionals Onsite, 2014
- National Selection - India, International Physics Olympiad, 2011
- All India Ranked 104, National Cyber Olympiad, 2011

SKILLS

- *Code* 8k+ production lines: Python, C++, C#, Java, Scala, C, Fortran
- *Frameworks*: Docker, Jinja, PyTorch, Kubernetes, Tensorflow, Jenkins, Azure
- *Music*: Violin, Harmonium, Hindustani Classical Vocalist
- *Extra-Curricular*: Salsa, Bachata, Shito-Ryu Genbu-Kai Karate, Marathon
- *Languages*: English (fluent), Hindi (native), Odiya (fluent), Mandarin (basic)

SERVICE

ACL 2023, EACL 2023, EMNLP 2022, EMNLP 2021.

MENTORING

- MSR PhD Intern's Mentor - Microsoft Research: Summer 2022
  - Screened, interviewed, hired and mentored the PhD intern
  - NLP: analysis of conversational text from a temporal lens
  - Discussions with MSR's Hao Cheng, and Cheng Wu
- Multimodal - After office hours, Prof. Ching-Yung Lin: 2020 - 2022
  - 12 MS students collaborated with me, and co-authored their first papers
  - Students continued to pursue research at GaTech, CMU, UW, Dartmouth, and industry
  - Led long-form multimodal research, leaderboard rank 1 across multiple years
- Head TA: Advanced Big Data Analytics, Columbia University, Fall 2018, Spring 2018, Fall 2017
  - Worked with Professor Ching-Yung Lin to design research projects.
  - TA award with full tuition waiver. Managed 6 TAs and 180 graduate students.
  - Bi-weekly one-on-one team mentoring on their research.
- Co-Lead TA: CS101, CS110, Introduction to Computing, IIT Guwahati, Fall 2014
  - Program to help under-performing students; undergrads usually not allowed to tutor

LINKS

- Website: https://vishalanand.net
- AI+Research Notes: https://blog.vishalanand.net
- CV: https://vishalanand.net/cv_vishalanand.pdf
- Scholar: https://scholar.google.com/citations?hl=enuser=YV3ZpKQAAAAJ

REFERENCES

Web - Stephen Bach, Eliot Horowitz Associate Professor, Computer Science, Brown University

Web - Kathleen R. McKeown, Henry and Gertrude Rothschild Professor of CS, Columbia University

Web - David M. Rothschild, PhD, Senior Principal Researcher, Microsoft Research NYC

Web - Cheng Wu, PhD, Partner Group Science Manager, Microsoft

Web - Ching-Yung Lin, Professor, IEEE Fellow, Columbia University

Web - Smaranda Muresan, Associate Professor, Computer Science, Barnard College, Columbia University

Web - Ashish Anand (undergrad advisor), Professor, Computer Science, IIT Guwahati