

# Nikitha Rao

Research Scientist working on AI for Code  
Google DeepMind  
Mountain View, USA

✉ [raonikitha@google.com](mailto:raonikitha@google.com)  
🏠 <https://raonikitha.github.io/>  
🎓 Google Scholar

## EDUCATION

- **Carnegie Mellon University**, Pittsburgh, USA August 2021 - May 2025  
*Ph.D. in Software Engineering.*  
Advisors: Prof. Vincent Hellendoorn and Prof. Claire Le Goues  
Research Interests: Artificial Intelligence for Code, Large Language Models, Generative AI  
Thesis title: Navigating Challenges with LLM-based Code Generation using Software-specific Insights.  
Thesis committee: Daniel Fried (CMU - LTI), Andrew Begel (CMU - S3D), Tom Zimmermann (UC Irvine)  
Cumulative GPA: 4.08/4.0
- **PES University**, Bangalore, India 2015 - 2019  
*B.Tech in Computer Science and Engineering with a specialization in Data Science.*  
Advisor: Dr. Gowri Srinivasa  
Cumulative GPA: 9.48/10.0

## PUBLICATIONS

14. **DiffSpec: Differential Testing with LLMs using Natural Language Specifications and Code Artifacts**  
Nikitha Rao, Elizabeth Gilbert, Harrison Green, T. Ramananandro, N. Swamy, Claire Le Goues, Sarah Fakhoury  
*Under Submission, 2025.* [preprint]
13. **SegSub: Evaluating Robustness to Knowledge Conflicts and Hallucinations in Vision-Language Models**  
Peter Carragher, Nikitha Rao, Abhinand Jha, R Raghav, and Kathleen Carley  
*Workshop on Misinformation Detection in the Era of LLMs at ICWSM, 2025.* [MisD@ICWSM25]
12. **Prompts Are Programs Too! Understanding How Developers Build Software Containing Prompts**  
Jenny Liang, Melissa Lin\*, Nikitha Rao\*, and Brad Myers (\* equal contribution)  
*Foundations of Software Engineering, 2025.* [FSE 25]
11. **From Overload to Insight: Bridging Code Search and Code Review with LLMs** [FSE-IVR 25]  
Nikitha Rao, Bogdan Vasilescu, and Reid Homes  
*Foundations of Software Engineering - Ideas, Visions and Reflections Track, 2025.*
10. **AI for Low-Code for AI** [IUI 24]  
Nikitha Rao, Jason Tsay, Kiran Kate, Vincent Hellendoorn, and Martin Hirzel  
*Intelligent User Interfaces, 2024.*
9. **CAT-LM: Training Language Models on Aligned Code And Tests** [ASE 23]  
Nikitha Rao\*, Kush Jain\*, Uri Alon, Claire Le Goues, and Vincent Hellendoorn (\* equal contribution)  
*Automated Software Engineering, 2023.*
8. **Comments on Comments: Where Code Review and Documentation Meet** [MSR 22]  
Nikitha Rao, Jason Tsay, Martin Hirzel, and Vincent Hellendoorn  
*Mining Software Repositories, 2022.*
7. **SoftNER: Mining Knowledge Graphs From Cloud Incidents** [EMSE 22]  
Manish Shetty, Chetan Bansal, Sumit Kumar, Nikitha Rao, and Nachiappan Nagappan  
*Empirical Software Engineering (SEIP Special Issue), 2022.*
6. **Search4Code: Code Search Intent Classification Using Weak Supervision** [MSR 21]  
Nikitha Rao, Chetan Bansal, and Joe Guan  
*Mining Software Repositories, 2021.*
5. **Neural Knowledge Extraction from Cloud Service Incidents** [ICSE - SEIP 21]  
Manish Shetty, Chetan Bansal, Sumit Kumar, Nikitha Rao, Nachiappan Nagappan, and Thomas Zimmermann  
*International Conference on Software Engineering, 2021.*  
🏆 Nominated for the IEEE Software Distinguished Paper Award (5/41)  
🏆 Featured on VentureBeat: Microsoft's SoftNER AI uses unsupervised learning to help triage cloud service outages.  
🏆 Featured on Techzine: Microsoft's SoftNER AI evaluates disruptions in cloud services

4. **Handling Class Imbalance with POISE: pAUC Optimization in Supervised Experiments** [MLADS 20]  
 Nikitha Rao, and Sreangsu Acharyya  
 🏆 **Best Short Paper Award** at MLADS-SYNAPSE, 2020.  
*Microsoft internal Conference on Machine Learning and Data Science for Asia-Pacific region*  
 [Acceptance Rate  $\approx$  8%]
3. **Analyzing Web Search Behavior for Software Engineering Tasks** [IEEE BigData 20]  
 Nikitha Rao, Chetan Bansal, Thomas Zimmermann, Ahmed Hassan Awadallah, and Nachiappan Nagappan  
*IEEE International Conference on Big Data, 2020.*
2. **Product Insights: Analyzing Product Intents in Web Search** [CIKM 20]  
 Nikitha Rao, Chetan Bansal, Subhabrata Mukherjee, and Chandra Maddila  
*International Conference on Information and Knowledge Management, 2020.*
1. **Studying Ransomware Attacks Using Web Search Logs** [SIGIR 20]  
 Chetan Bansal, Pantazis Deligiannis, Chandra Maddila, and **Nikitha Rao** (*alphabetical order*)  
*International Conference on Research and Development in Information Retrieval, 2020.*

## PATENTS

---

- **Identification of Content Gaps based on Relative User-Selection Rates between Multiple Discrete Content Sources** filed with the USPTO. *October 16, 2020*  
 Co-inventors: Chetan Bansal, Junia George, Casey Gossard, Dung Nguyen, Dave Ludwig, and Curtis Anderson.
- **ExtraQuery Context-Aided Search Intent Detection** filed with the USPTO. *October 9, 2020*  
 Co-inventors: Chetan Bansal, Joe Guan, Mark Wilson-Thomas, Nachiappan Nagappan, and Thomas Zimmermann.
- **Automatic Recognition of Entities Related to Cloud Incidents** filed with the USPTO. *June 19, 2020*  
 Co-inventors: Manish Shetty, Chetan Bansal, Sumit Kumar, Nachiappan Nagappan, and Thomas Zimmermann.

## AWARDS AND HONORS

---

- Invited to give a **keynote** and serve as a **panelist at the VerifAI workshop at ICLR 2025.** *2025*
- **Hima and Jive Fellowship in Computer Science** at CMU, awarded \$40,000 in total. *2024*
- Invited to **Dagstuhl Seminar on Automated Programming and Program Repair** as a Young Researcher. *2024*
- Invited to **Dagstuhl Seminar on Code Search** as a Young Researcher. [report] *2024*
- Nominated for **IBM Ph.D. Fellowship**, 1 of 4 students from CMU-SCS. *2023*
- **Google Collab Ph.D. Fellowship**, awarded \$100,000 in total. *2021*
- **Graduate Dean's Scholar Award**, Computer Science, UCLA (*declined in favor of CMU*). *2021*
- **Computer Science Excellence Fellowship**, Computer Science, UIUC (*declined in favor of CMU*). *2021*
- **Dean's Distinguished Graduate Fellowship**, Computer Science, UC Davis (*declined in favor of CMU*). *2021*
- **Chair's Award**, Informatics, UC Irvine (*declined in favor of CMU*). *2021*
- **Best Short Paper Award** at MLADS-SYNAPSE. *2020*
- **Best Outgoing Student Award** for class of 2019 (360 students), Computer Science, PES University. *2019*
- Five time recipient of the **CNR Rao Scholarship**, Computer Science, PES University. *2016 - 2019*

## TEACHING

---

- **Neural Code Generation** (11891), CMU - Teaching Assistant *Spring 2024*
- **Applied Deep Learning** (17644), CMU - Teaching Assistant *Spring 2023*
- **Applied Machine Learning**(17634), CMU - Teaching Assistant *Spring 2023*

## ONGOING PROJECTS

---

- **Differential Testing with LLMs using Natural Language Specifications** *May, 2024 - Present*  
 Collaborators: *Sarah Fakhoury, Nikhil Swamy (MSR), and Claire Le Goues (CMU)*  
 Several real world systems like eBPF, WASM, network protocols, etc, have multiple implementations that need to conform to the same specifications, and should therefore have the same behaviour. However, there exists discrepancies in behavior that point to bugs. In this work, we make use of informal artifacts such as natural language specifications, code implementations, bug reports and so on to improve the quality of test suites by generating differential tests using LLMs.

The goal is to be able to generate tests that return different outputs and therefore point to discrepancies in the various implementations.

- **Teaching Large Language Models to Debug Code Collaboratively** *September, 2022 - Present*  
Collaborators: [Vincent Hellendoorn](#), and [Claire Le Goues](#) (CMU)  
Even tools such as ChatGPT or Copilot tend to generate code containing subtle bugs that are hard to find for inexperienced developers. In this work, we leverage the execution of code generated by these LLMs as a signal. Specifically, we employ several such models working in tandem: one observes the failing execution and generates debugging instructions, which other models use to repair the generated code before presenting it to the developer. By observing this interaction we can then teach these models to collaboratively debug the code they, or regular developers, generate.

## WORK EXPERIENCE

---

- **Google DeepMind, Mountain View - Research Scientist** *June, 2025 - present*  
Manager: [Dr. Daniel Tarlow](#)
- **Microsoft Research, Redmond - Research Intern** *May - August, 2024*  
Advisors: [Dr. Sarah Fakhoury](#) and [Dr. Nikhil Swamy](#)  
Project: Differential Testing with LLMs using Natural Language Specifications and Code Artifacts.
- **IBM T.J. Watson Research Center, Yorktown Heights, NY - AI Research Intern** *May - August, 2023*  
Advisor: [Anuradha Bhamidipaty](#)  
Project: Built a unified dialogue-based domain-specific question-answering system using LLMs.
- **IBM T.J. Watson Research Center, Yorktown Heights, NY - AI Research Intern** *May - August, 2022*  
Advisor: [Dr. Martin Hirzel](#)  
Project: AI for Low-Code for AI.
- **Microsoft Research, India - Research Fellow** *July, 2019 - July 2021*  
Advisors: [Chetan Bansal](#), [Dr. Subho Mukherjee](#), [Dr. Nachi Nagappan](#), and [Dr. Tom Zimmermann](#)  
Project Domains: Machine Learning for Software Engineering, Data Science, and Web Search
- **Microsoft Research, India - Research Intern** *January - June, 2019*  
Advisor: [Dr. Sreangsu Acharyya](#)  
Project: Partial AUC optimization for extreme class imbalance.
- **Carnegie Mellon University, Pittsburgh - Research Intern** *Summer 2018*  
Advisor: [Prof. Shawn Blanton](#)  
Project Domain: Machine Learning
- **Center for Pattern Recognition, PES University - Research Intern** *August, 2017 - December, 2019*  
Advisor: [Dr. Gowri Srinivasa](#)  
Project Domain: Data Science, Computer Vision
- **Indian Institute of Science, India - 5<sup>th</sup> Summer School Program** *July, 2017*  
Among the youngest students selected for the program organized by the Computer Science and Automation Department.

## TALKS

---

- **Navigating Challenges with LLM-based Code Generation using Software-specific Insights**  
PLSE seminar at National University of Singapore *April 2025*  
Google Deepmind, Mountain View *January 2025*  
PROSE at Microsoft, Redmond *November 2024*  
FAIR, Paris *November 2024*
- **Testing, Testing, 1-2-3: Test Generation with LLMs**  
Invited Keynote at VerifAI workshop at ICLR 25 *April 2025*  
Dagstuhl Seminar on Automated Programming and Program Repair *October 2024*
- **Test Generation with LLMs** *November 2024*  
Guest Lecture for Generative AI for Software Engineering, NCSU
- **Differential Testing with LLMs using Natural Language Specifications** *August 2024*  
Microsoft Research, Redmond
- **User Intent and Needs for Code Search** *April 2024*  
Dagstuhl Seminar on Code Search

- **Beyond Syntax: Navigating Challenges in AI-Generated Code** *December 2023*  
Microsoft Research, India
- **CAT-LM: Training Language Models on Aligned Code And Tests** *November 2023*  
JetBrains Research [\[Recording\]](#)
- **Code Generation and Alignment** *November 2023*  
Guest Lecture for Advanced NLP (11711), CMU [\[Website\]](#)
- **A Unified Dialogue Based Domain-Specific Question-Answering System Using LLMs** *August 2023*  
IBM T.J. Watson Research Center, Yorktown Heights, NY
- **Prompting and Tuning LLMs** *April 2023*  
Guest Lecture for Applied Deep Learning (17644), CMU
- **Introduction to Deep Learning** *March 2023*  
Guest Lecture for Applied Deep Learning (17644), CMU
- **AI for Low-Code for AI** *August 2022*  
IBM T.J. Watson Research Center, Yorktown Heights, NY
- **Search Insights: Analysing Web Search Behavior to Mine Insights** *July 2021*  
Microsoft Research, India
- **Partial-AUC Optimization to Handle Class Imbalance** *August 2020*  
Microsoft Research, India

## SERVICE

---

- **PC Research Track:** ICSE 26
- **Journal Reviewer:** TOSEM 25, JSERD 21
- **PC Industry Showcase Track:** ASE 25, ASE 24
- **PC Artifact Evaluation:** ICSE 24
- **Sub-reviewer:** FSE 23, ICLR 22
- **Shadow PC:** MSR 22
- **Ask Me Anything** session on Grad School Applications with Research Fellows at Microsoft Research India (2023).
- **DNI Representative at Microsoft Research India** (2019 - 2021)  
I was the Research Fellow representative in the Diversity and Inclusion (DNI) committee at Microsoft Research India. We took several initiatives that include workshops on LGBTQ+ sensitization, talks and workshops for women, panel discussions and an annual diversity and inclusion day dedicated to increasing awareness for all new interns and research fellows. I also started a virtual book club during the pandemic to help reduce isolation and to increase awareness on DNI topics, which received a lot of positive feedback.