

Khai-Nguyen Nguyen

khainguyen.nguyen@virginia.edu | +1 272-788-0512 | [Website](#) | [Google Scholar](#) | [LinkedIn](#)

Research Interests

I am broadly interested in multimodal, robustness and AI safety as well as their applications. To this end, I have worked on: **data curation** [C1][C5][C7], **evaluation** [C2][C3], **training** [C1][C5], and **post-hoc methods** [C4][C8].

Education

University of Virginia <i>Ph.D. student in Data Science</i>	2026 – 2029 (expected) <i>Advisor: Chirag Agarwal</i>
College of William and Mary <i>Research Assistant in Computer Science</i>	2023 – 2026 <i>Advisor: Antonio Mastropaolo</i>
Bucknell University <i>B.Sc. in Computer Science and Engineering; minor in Statistics</i> GPA: 3.94/4.00; <i>summa cum laude</i>	2019 – 2023

Relevant Experiences

Nguyen Lab, Auburn University <i>Remote Collaborator</i>	Mar 2025 – May 2026 <i>Advisor: Anh Totti Nguyen</i>
Research on bias in vision-language and language models, resulting in publications [C1] and [C3].	
CodaMetrix, Boston <i>Machine Learning Intern</i>	May – Aug 2025 <i>Advisor: Cheng Li</i>
- Developed a Gemini-based judge to evaluate and refine medical entity extractions in health records. - Designed & deployed medical entity extraction systems on Databricks using DSPy prompt optimization.	
CodaMetrix, Boston <i>Machine Learning Intern</i>	May – Aug 2024 <i>Advisor: Cheng Li</i>
- Trained a BERT-based model for multilabel ICD-10 classification. - Implemented distributed and parallel model training on Databricks.	
FPT Software AI Center, Vietnam <i>Remote Research Intern</i>	Jan – Sep 2024 <i>Advisor: Truong-Son Hy</i>
Research on trustworthy LLMs for healthcare applications, resulting in publications [C5], [C6], and [C7].	

Selected Publications

* Equal contribution.

- [C1] **Deep and shallow biases in language models** [Paper]
An Vo, Vy Tuong Dang, **Khai-Nguyen Nguyen**, Emilio Villa-Cueva, Thamar Solorio, Anh Totti Nguyen, Daeyoung Kim
EMNLP 2026
- [C2] **Pattern over Pixels: Measuring Pattern Completion Bias in Multimodal Code Generation** [Paper]
Khai-Nguyen Nguyen, Oscar Chaparro, Antonio Mastropaolo
ASE 2026
- [C3] **Vision Language Models are Biased** [Paper]
An Vo*, **Khai-Nguyen Nguyen***, Mohammad Reza Taesiri, Vy Tuong Dang, Anh Totti Nguyen, Daeyoung Kim
ICLR 2026
- [C4] **Toward Explaining Large Language Models in Software Engineering Tasks** [Paper]
Antonio Vitale*, **Khai-Nguyen Nguyen***, Denys Poshyvanyk, Rocco Oliveto, Simone Scalabrino, Antonio Mastropaolo
ACM Transactions on Software Engineering and Methodology (TOSEM), 2026

- [C5] **Sentiment Reasoning for Healthcare** [\[Paper\]](#)
Khai-Nguyen Nguyen*, Khai Le-Duc*, Bach Phan Tat, Duy Le, Long Vo-Dang, Truong-Son Hy
ACL 2025, Industry Track (Oral)
- [C6] **Medical Spoken Named Entity Recognition** [\[Paper\]](#)
 Khai Le-Duc, David Thulke, Hung-Phong Tran, Long Vo-Dang, **Khai-Nguyen Nguyen**, Truong-Son Hy, Ralf Schlüter
NAACL 2025, Industry Track
- [C7] **Real-time Speech Summarization for Medical Conversations** [\[Paper\]](#)
 Khai Le-Duc*, **Khai-Nguyen Nguyen***, Long Vo-Dang, Truong-Son Hy
Interspeech 2024 (Oral)
- [C8] **Getting Away with More Network Pruning: From Sparsity to Geometry and Linear Regions** [\[Paper\]](#)
 Jeffrey Cai*, **Khai-Nguyen Nguyen***, Nishant Shrestha, Aidan Good, Ruisen Tu, Xin Yu, Shandian Zhe, Thiago Serra
CPAIOR 2023; Sparsity in Neural Networks Workshop at ICLR 2023

Preprints

- [P1] **Prior Bias in Vision Language Models on UML Diagram Interpretation** [\[Paper\]](#)
 Zaiyu Cheng*, **Khai-Nguyen Nguyen***, Antonio Mastropaolo
arXiv preprint, 2026
- [P2] **S-Chain: Structured Visual Chain-of-Thought for Medicine** [\[Paper\]](#)
 Khai Le-Duc et al.
arXiv preprint, 2025

Awards and Honors

Computer Science Fellowship , College of William and Mary	2023
Dean's List , Bucknell University (every semester)	2019 – 2023
Program for Undergraduate Research Grant , Bucknell University (\$4,000)	2021
Honorable Mention , Mathematical Contest in Modeling	2020
Awarded Honorable Mention out of 13,753 teams worldwide (top 26%).	

Industry and Press Coverage

Industry Usage

2026: [Qwen used VLMsAreBiased to evaluate Qwen3.8-Max.](#)

2025: [Google DeepMind used VLMsAreBiased to evaluate Gemini 3 Pro.](#)

News Coverage

2026: Featured on Sky News' [video](#) and [article](#) on potential dangers of AI in modern warfare.

Teaching and Professional Services

University of Virginia: Computation for Data Science.

Teaching Assistant, William and Mary: Intro to Data Science; Computer Organization; Software Engineering.

Teaching Assistant, Bucknell University: Data Structures and Algorithms; Calculus I and II.

Reviewer: ACL Rolling Review (ARR), May 2026.