

KUMAIL ALHAMOUD

Website: <https://kumailalhamoud.netlify.app>

✉ kumail@mit.edu

Current Research Summary

Currently working on evaluation and post-training for multimodal reasoning models. I built [NegBench](#), which reveals why vision-language models (e.g. CLIP, LLaVA, AIMv2) misinterpret negation, and it provides synthetic datasets that measurably improve negation understanding after finetuning. I also worked on applications like [FedMedICL](#), which stress-tests medical AI under realistic shifts (e.g. pandemics). My first-authored work has been published at diverse venues (e.g., CVPR, TMLR, MICCAI, Computers & ChemE) and featured by [MIT News](#), [MIT CSAIL](#), and [New Scientist](#). I've collaborated with research teams at Google, Samsung, and Foxconn, and I've consulted for PwC.

Education

Massachusetts Institute of Technology (MIT)

PhD, Computer Science and Artificial Intelligence Lab (CSAIL); GPA: 5.0/5.0 (A)

September 2023 – Present

Cambridge, MA

- PhD Supervisor: Prof. Marzyeh Ghassemi
- **Relevant Courses:** Large Language Models, Hardware Architecture of Deep Learning
- **Main Programming Language and Framework:** Python and PyTorch

King Abdullah University for Science and Technology (KAUST)

M.S, Computer Science; GPA: 3.96/4.00 (A)

September 2021 – July 2023

Thuwal, Saudi Arabia

- Thesis Advisor: Prof. Bernard Ghanem
- **Relevant Courses:** Deep Generative Modeling, Low-resource Deep Learning, Algorithms in Bioinformatics

Cornell University

B.S, Electrical and Computer Engineering; GPA: 4.13/4.30 (A+)

May 2021

Ithaca, NY

- *Summa Cum Laude*
- Dean's List, all semesters
- **Relevant Courses:** Computer Vision, Natural Language Processing, Digital Signal Processing, Deep Learning

Publications

(* indicates equal contribution; clicking on paper title links to paper)

- **Kumail Alhamoud***, Sepehr Ranjbar*, Marzyeh Ghassemi. “[SpaceVLM: Sub-Space Modeling of Negation in Vision-Language Models](#)”, *Preprint 2025*.
- Hasan Hammoud, **Kumail Alhamoud**, Elie Bou-Zeid, Marzyeh Ghassemi, Bernard Ghanem. “[Train long, think short: Curriculum learning for efficient reasoning](#)”, *Preprint 2025*.
- Olawale Salaudeen, Haoran Zhang, **Kumail Alhamoud**, Sara Beery, Marzyeh Ghassemi. “[Aggregation Hides OOD Generalization Failures from Spurious Correlations](#)”, *NeurIPS 2025*. [Spotlight Paper]
- **Kumail Alhamoud**, Shaden Alshammari, Yonglong Tian, Guohao Li, Philip Torr, Yoon Kim, Marzyeh Ghassemi. “[Vision-Language Models Do Not Understand Negation](#)”, *CVPR 2025*.
- **Kumail Alhamoud***, Yasir Ghunaim*, Motasem Alfarra, Thomas Hartvigsen, Philip Torr, Bernard Ghanem, Adel Bibi, Marzyeh Ghassemi. “[FedMedICL: Towards Holistic Evaluation of Distribution Shifts in Federated Medical Imaging](#)”, *MICCAI 2024*.
- **Kumail Alhamoud***, Yasir Ghunaim*, Abdulelah Alshehri, Guohao Li, Bernard Ghanem, Fengqi You. “[Leveraging 2D Molecular Graph Pretraining for Improved 3D Conformer Generation with Graph Neural Networks](#)”, *Computers and Chemical Engineering 2024*.
- **Kumail Alhamoud***, Hasan Abed Al Kader Hammoud*, Motasem Alfarra, Bernard Ghanem. “[Generalizability of Adversarial Robustness Under Distribution Shifts](#)”, *Transactions on Machine Learning Research (TMLR) 2023*. [awarded a Featured Certification → presented in ICLR 2024]

- Yasir Ghunaim*, Adel Bibi*, **Kumail Alhamoud**, Motasem Alfarra, Hasan Abed Al Kader Hammoud, Ameya Prabhu, Philip Torr, Bernard Ghanem. “Real-Time Evaluation in Online Continual Learning: A New Hope”, *CVPR* 2023. [Highlight Paper, top 2.5%]
- Andrés Villa, Juan C. Alcazar, Motasem Alfarra, **Kumail Alhamoud**, Julio Hurtado, Fabian Caba, Alvaro Soto, Bernard Ghanem. “PIVOT: Prompting for Video Continual Learning”, *CVPR* 2023.
- Andrés Villa, **Kumail Alhamoud**, Juan C. Alcazar, Fabian Caba, Victor Escoria, Bernard Ghanem. “vCLIMB: A Novel Video Class Incremental Learning Benchmark”, *CVPR* 2022. [Oral Paper, top 2.5%]
- Akshay Ajagekar, **Kumail Alhamoud**, Fengqi You. “Hybrid Classical-Quantum Optimization Techniques for Solving Mixed-Integer Programming Problems in Production Scheduling”, *IEEE Transactions on Quantum Engineering* 2022.
- Sherry Liang, Khalid Alanazi, **Kumail Alhamoud**. “Set covering problem”, *Cornell University Computational Optimization Open Textbook Chapter* 2020.

Recent Work Experience

Foxconn <i>AI Research Intern</i>	July 2025 – Present <i>Taipei, Taiwan</i>
<i>Enhancing Reasoning in Large Language Models (LLMs)</i>	
• Designed and implemented post-training strategies to strengthen the reasoning capabilities of FoxBrain, Foxconn’s proprietary LLM, leveraging reinforcement learning methods such as DPO, PPO, GRPO, and Dr. GRPO.	
MIT, Healthy ML Group <i>PhD Student Researcher</i>	
<i>Enhancing Negation Understanding of Vision-Language Models with Synthetic Data</i>	
• Proposed a benchmark to evaluate negation understanding in VLMs, which involves recognizing what is <i>not</i> present in an image; highlighted a severe limitation in current models and innovated a synthetic data approach to mitigate this issue.	
PwC <i>AI Research Consultant (Part-Time)</i>	July 2024 – July 2025 <i>Remote</i>
<i>Strategic AI Advisory for Government Transformation</i>	
• Authored strategic studies on Human-Centered AI, Enterprise RAG, and AI Agents as Government Partners, providing clear explanations and actionable recommendations tailored to decision makers.	
• Designed roadmaps for implementing these technologies by benchmarking 20 global case studies across 7 sectors and embedding best practices such as UNESCO AI Ethics Guidelines.	
University of Oxford, Torr Vision Group <i>Visiting PhD Student Researcher</i>	May 2024 – September 2024 <i>Oxford, UK</i>
<i>Advancing Reliable Decision-Making with Large Language Models (LLMs)</i>	
• Developed innovative uncertainty quantification techniques to guide decision-makers on when to trust LLM outputs, improving decision-making processes under conditions of uncertainty.	

Awards & Recognition

CVPR Outstanding Reviewer Award	2025
SACM Fellowship for PhD Studies (up to four years of PhD Funding)	2023-2027
MIT Jameel Clinic PhD Fellowship for ML and Health (\$100,000 award for one year)	2023
KAUST Fellowship for MSc Studies	2021
The MISK Fellowship	2021
Outstanding Teaching Assistant Award at Cornell University	2021
• Awarded by the Information Science Department; nominated by professor for taking lead of all in-class activities	
2021 Rhodes Scholarship Finalist (only 8 finalists from my constituency)	2021

Academic and Teaching Experience

Academic Reviewer for Top-Tier Research Publication Venues

2022 – Present

- Including CVPR, ECCV, NeurIPS, ICLR, MICCAI, ML4H, JAMIA, and TPAMI

Teaching Assistant

2019 – 2022

- CS 323 Deep Learning for Visual Computing, Prof. Bernard Ghanem, KAUST
- CS 4300 Language and Information, Prof. Cristian Danescu-Niculescu-Mizil, Cornell (Received the Best TA Award)
- CS 2800 Discrete Structures, Prof. Anke van Zuylen, Cornell
- ECE 3250 Mathematics of Signals and System Analysis, Prof. David Delchamps, Cornell

Research Presentations

ECCV24 Workshop on Emergent Visual Abilities and Limits of Foundation Models

September 2024

- **Vision-Language Models Do Not Understand Negation**

Kumail Alhamoud, Shaden Alshammari, Yonglong Tian, Guohao Li, Philip Torr, Yoon Kim, Marzyeh Ghassemi

KAUST Conference on Scientific Computing and Machine Learning [Best Poster Award]

November 2022

- **Chemistry-informed Graph Representation Learning for Molecular Conformation Generation and Beyond**

Kumail Alhamoud, Yasir Ghunaim, Guohao Li, Bernard Ghanem

Cornell Engineering Learning Initiatives Summer Presentation

August 2020

- **Harnessing Quantum Computing to Improve the SOTA in Solving Industrial-Scale Scheduling Problems**

Kumail Alhamoud, Akshay Ajagekar, Fengqi You

Boston University 22nd Annual UROP Undergraduate Research Symposium

October 2019

- **DNA Sequence Alignment Framework for Sequence Pathogenicity Screening**

Kumail Alhamoud, Samuel M.D. Oliveira, Douglas Densmore