Kuo-Han **Hung**

+1 (650) 283-6084 | khhung@stanford.edu | khhung-906.github.io | Kuo-Han-Hung

Summary

First-year MSCS student at Stanford specializing in AI and robotics, with 4 years experience in machine learning, including large language models, vision-language models, and policy learning. Published at ICLR 2025, NAACL 2025, and NeurIPS 2024. Interned at IBM Research and Microsoft Bing, where I worked on LLM systems and applied AI for real-world tasks.

Education

Stanford University — *California, USA*
M.S. in Computer Science — *Sept. 2025 – June 2027 (Expected)*
- Specializing in Artificial Intelligence and Robotics.

National Taiwan University (NTU) — *Taipei, Taiwan*
B.S. in Computer Science and Information Engineering — *Sept. 2020 – Dec. 2024*
- GPA: 4.24/4.3 overall; Ranking: 5/187 (top 2.6%); Program: Creativity and Entrepreneurship.
- Awards: Bachelor Thesis Dean's Award, 3 times Dean's List Awards, 2 times Presidential Awards.

Work Experience

IBM Research — *New York, USA*
Research Intern, Advisor: Dr. Pin-Yu Chen — *June 2024 – Sept. 2024*
- Proposed a novel perspective on prompt injection attacks in LLM agents through attention. *[NAACL'25]*
- Developed an attention-based detector, achieving a 31% AUROC improvement over prior best baselines.

Microsoft, Bing Team — *Remote (team in California, USA)*
Applied Scientist Intern — *July 2023 – June 2024*
- Integrated Bing Maps with Microsoft Copilot to create an enriched experience for conversational map queries.
- Developed a novel pathfinding algorithm that learns from history, achieving a 10x speedup. *[SIGSPATIAL'24]*
- Invented and evaluated a context-aware map search system using LLMs. *[SIGSPATIAL'24]*

Cinnamon AI — *Taipei, Taiwan*
Deep Learning Intern — *June 2022 – Aug. 2022*
- Developed and implemented the first image driven recipe retrieval model with adaptable ingredients, along with a cross modal retrieval system featuring a vision-ingredients seq2seq architecture.

Research Experience

Embodied AI Lab — *Taipei, Taiwan*
Researcher, Advisor: Prof. Tsung-Wei Ke — *June 2025 – Sept. 2025*
- Researching on *generalized dexterous manipulation policies.*
  - Developing a pipeline to convert in-the-wild videos into trainable digital twins, enabling scalable policy learning.
  - Training a dexterous manipulation policy that generalized to unseen demos with zero test-time tuning.

Robot Learning Lab — *Taipei, Taiwan*
Undergraduate Researcher, Advisor: Prof. Shao-Hua Sun — *Oct. 2024 – May 2025*
- Researched on *generalized policy by learning from videos*. *[Under Review @ NeurIPS'25]*
  - Developed a framework to learn reusable skills from action-free videos using optical flow as an action surrogate.
  - Enabled multi-task and long-horizon planning by translating flow-based skill plans into executable actions.

Communication and Multimedia Lab — *Taipei, Taiwan*
Undergraduate Researcher, Advisors: Prof. Winston H. Hsu & Prof. Yi-Ting Chen — *Feb. 2022 – May 2024*
- Researched on *vision-instruction correlation rewards for long-horizon manipulation.* *[ICLR'25]*
  - Developed a reward model leveraging LLMs and VLMs to infer stage-aware progress from videos and instructions.
  - Achieved a 43% improvement in success rate over state-of-the-art baselines on long-horizon task.
- Researched a behavior error detector in few-shot imitation policies using pattern-based feature analysis. *[NeurIPS'24]*

Machine Intelligence and Understanding Lab *Taipei, Taiwan*

Undergraduate Researcher, Advisor: Prof. Yun-Nung Chen *Feb. 2022 – Dec. 2023*

- Benchmarked *multi-source retrieval and reasoning in visual question answering.*
- Researched on *open-domain conversational questions, answering using historical answers.* *[AACL'22]*
  - Employed knowledge distillation to enhance the efficacy of retrieving passages using historical replies.
  - Achieved state-of-the-art performance in open-domain conversational retrieval with 77.9 (R@5).

Selected Projects

Vulnerabilities in VLM-Powered Policies *[Slide]*

- Developed a text and image-based adversarial attack method to manipulate embodied AI behavior, achieving a 40% increase in targeted attack success on VLM-powered policy networks.

Zero-shot Text Behavior Retrieval *[Report]*

- Proposed a zero-shot behavior retrieval system using text-guided object detection and CLIP to retrieve task-relevant data from offline datasets, enabling policy training without expert data.

Honors & Awards

Student Research Scholarship, issued by Taiwan National Science Council	Sept. 2024
Lin Hsiung Chen Memorial Scholarship (for top 1 CS student)	Nov. 2023
Jason International Fund Scholarship, issued by Acer Inc.	Sept. 2022
Dean's Award (for top two theses in EECS dept.), NTU Bachelor's Thesis Award	June 2024
Presidential Award/Dean's List (for top 2%/5% students), NTU	Fall'22, '23; Spring'21, '22, '23
Second prize, NTU CSIE Undergraduate Thesis Exhibition	June 2023
Creativity and Entrepreneurship Excellence Award, NTU D-school	June 2023
Honorable Award, NTU CSIE Undergraduate Thesis Exhibition	June 2022
Special Award (for top 5/300+ teams), LINE Fresh Hackathon	Dec. 2021

Teaching Experience

Teacher Assistant, CSIE1000 Introduction to Computer Science *Sept. 2024 – Dec. 2024*

- Developed assignments and led weekly TA sessions to reinforce students' understanding across major CS fields.

Teacher Assistant, CSIE5043 Machine Learning (graduate level) *Jan. 2023 – June 2023*

- Designed and implemented advanced machine learning assignments on VC dimension and SVM.
- Organized a final project competition focused on music recommendation systems using tabular data.

Publications (*co-authorship)

[1] **K. H. Hung***, P. C. Lo*, J. F. Yeh*, H. Y. Hsu, Y. T. Chen, and W. H. Hsu, "*VICtoR: Learning Hierarchical Vision-Instruction Correlation Rewards for Long-horizon Manipulation*." The 13th International Conference on Learning Representations (**ICLR**), 2025. *[PDF]*

[2] J. F. Yeh, **K. H. Hung***, P. C. Lo*, C. M. Chung, T. H. Wu, H. T. Su, Y. T. Chen, and W. H. Hsu, "*AED: Adaptable Error Detection for Few-shot Imitation Policy*." The 38th Conference on Neural Information Processing Systems (**NeurIPS**), 2024. *[PDF]*

[3] **K. H. Hung**, C. Y. Ko, A. Rawat, I. H. Chung, W. H. Hsu, and P. Y. Chen, "*Attention Tracker: Detecting Prompt Injection Attacks in LLMs*." Annual Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics (**NAACL**), 2025. *[PDF]*

[4] **K. H. Hung**, C. Zhang, and D. Yankov, "*Customizable Routing with Learning from Past Recommendations*." The 32nd International Conference on Advances in Geographic Information Systems (**SIGSPATIAL**), 2024. *[PDF]*

[5] C. Zhang, A. Sriram, **K. H. Hung**, R. Wang, and D. Yankov, "*Context-aware conversational map search with LLM*." The 32nd International Conference on Advances in Geographic Information Systems (**SIGSPATIAL**), 2024. *[PDF]*

[6] H. C. Fang*, **K. H. Hung***, C. W. Huang, and Y. N. Chen, "*Open-Domain Conversational Question Answering with Historical Answers*." The 2nd Asian Chapter of the Association for Computational Linguistics (**AACL**), 2022. *[PDF]*