

Alyssa Hwang

ahwang16@seas.upenn.edu · alyssahwang.com

Modified: February 23, 2024

Research Interests

I study how AI assistants can be more efficient, interactive, and fluid communicators. I am particularly interested in leveraging large language models to help users manage many details at once through voice, text, AR/VR, or embedded assistants.

Education

- Sept 2020–Present **University of Pennsylvania** – Philadelphia, PA
Expected May 2026 *PhD* in Computer and Information Science
Dec 2023 *MS* in Computer and Information Science
Advisors: [Andrew Head](#), [Chris Callison-Burch](#)
Thesis topic: AI assistants that help users manage many details at once
- May 2020 **Columbia University** – New York, NY
BS in Computer Science
Advisor: [Kathleen McKeown](#)
Thesis: [Towards Augmenting Lexical Resources for Slang and African American English](#)

Work Experience

- Summer 2022 **AWS Comprehend** (Applied Scientist Intern) – New York, NY
NewsQs: *Multi-Source Question Generation for the Inquiring Mind*
Hosts: [Kathleen McKeown](#), Kalpit Dixit
- Summer 2021 **Google** (Software Engineering Intern) – Pittsburgh, PA
Image Ranking for Shopping
Hosts: Wei Wan, [Tao Chen](#)

Summer 2020 **Google** (Software Engineering Intern) – Mountain View, CA
Dialogue Act Recognition for Dialogflow
Hosts: Chuan Wang, [Igor Poletaev](#)

Honors

- 2023 **OpenAI**, DevDay Invited Attendee
- 2023 **Penn Engineering**, Outstanding Service Award
- 2022 **Amazon Alexa Prize Taskbot Competition**, Finalist Team
- 2022 **CRA-WP Grad Cohort for Women**, Invited Attendee
- 2020 **Columbia University**, Theodore R. Bashkow Award
Presented to a senior in computer science who has excelled in independent projects.
- 2020 **National Science Foundation**, Graduate Research Fellowship Program
- 2019 **Google**, Grace Hopper Celebration Travel Scholarship
- 2018 **Johnson & Johnson**, Summer Research Grant
- 2018 **Columbia University**, Data Science Institute Scholars Program and Grant

Publications, Articles, and Preprints

** Equal contribution*

- In submission at
ACL 2024 **NewsQs: Multi-Source Question Generation for the Inquiring Mind**
Alyssa Hwang, Kalpit Dixit, Miguel Ballesteros, Yassine Benajiba, Vittorio Castelli,
Markus Dreyer, Mohit Bansal, Kathleen McKeown
Association for Computational Linguistics (in submission).
Work conducted during an internship at AWS AI Labs.
- In submission at
ACL 2024 **RAID: A Shared Benchmark for Robust Evaluation of
Machine-Generated Text Detectors**
Liam Dugan, **Alyssa Hwang**, Filip Trhлік, Josh Magnus Ludan, Andrew Zhu, Hainiu Xu,
Daphne Ippolito, Chris Callison-Burch.
Association for Computational Linguistics (in submission).
- In submission at
ACL 2024 **FanOutQA: Multi-Hop, Multi-Document Question Answering for
Large Language Models**
Andrew Zhu, **Alyssa Hwang**, Liam Dugan, Chris Callison-Burch.
Association for Computational Linguistics (in submission). [arXiv preprint.](#)

- To appear at CHI 2024 **Ivie: Lightweight Anchored Explanations of Just-Generated Code**
Litao Yan, **Alyssa Hwang**, Zhiyuan Wu, Andrew Head
ACM Conference on Human Factors in Computing Systems (to appear).
- arXiv 2023 **Grounded Intuition of GPT-Vision’s Abilities with Scientific Figures**
Alyssa Hwang, Andrew Head, Chris Callison-Burch
arXiv preprint.
- NLP-OSS 2023 **Kani: A Lightweight and Highly Hackable Framework for Building Language Model Applications**
Andrew Zhu*, Liam Dugan*, **Alyssa Hwang**, Chris Callison-Burch
Workshop for Natural Language Processing Open-Source Software (co-located with EMNLP).
- arXiv 2023 **Large Language Models as Sous Chefs: Revising Recipes with GPT-3**
Alyssa Hwang*, Bryan Li*, Zhaoyi Hou*, Dan Roth
arXiv preprint.
- DIS 2023 **Rewriting the Script: Adapting Text Instructions for Voice Interaction**
Alyssa Hwang, Natasha Oza, Chris Callison-Burch, Andrew Head
ACM Conference on Designing Interactive Systems.
- Amazon 2022 **QuakerBot: A Household Dialog System Powered by Large Language Models**
Artemis Panagopoulou, Manni Arora, Li Zhang, Dimitri Cugini, Weiqiu You, Yue Yang, Liyang Zhou, Yuxuan Wang, Zhaoyi Hou, **Alyssa Hwang**, Lara Martin, Sherry Shi, Chris Callison-Burch, Mark Yatskar
Alexa Prize Taskbot Proceedings.
- VarDial 2020 **Towards Augmenting Lexical Resources for Slang and African American English**
Alyssa Hwang, William R. Frey, and Kathleen McKeown
Workshop on NLP for Similar Languages, Varieties and Dialects (co-located with COLING).
- EMNLP 2019 **AMPERSAND: Argument Mining for PERSuAsive oNline Discussions**
Tuhin Chakrabarty, Christopher Hidey, Smaranda Muresan, Kathy McKeown, **Alyssa Hwang**
2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing.
- MWE-WN 2019 **Confirming the Non-compositionality of Idioms for Sentiment Analysis**
Alyssa Hwang and Christopher Hidey
Joint Workshop on Multiword Expressions and WordNet (co-located with ACL).
- ArgMining 2017 **Analyzing the Semantic Types of Claims and Premises in an Online Persuasive Forum**
Christopher Hidey, Elena Musi, **Alyssa Hwang**, Smaranda Muresan, and Kathleen McKeown
Workshop on Argument Mining (co-located with ACL).

Invited Talks

- Apr 17, 2024 **Columbia University Natural Language Processing Seminar**
Human-centered Goals, Methods, and Evaluation for Large Language Models (tentative)
- Mar 21, 2024 **Google**
Human-centered Opportunities for Conversational Assistants (tentative)
- Feb 9, 2024 **Women in Data Science @ Penn**
Human-centered AI Research and Design
- Jan 29, 2024 **University of Pennsylvania Computational Linguistics Lunch (CLunch)**
Developing Grounded Intuition of Large Language Models
- Sept 18, 2023 **University of Pennsylvania Computational Linguistics Lunch (CLunch)**
Rewriting the Script: Adapting Text Instructions for Voice Interaction

Selected Media and Impact

- Nov 16, 2023 **A Peak into the Future of Visual Data Interpretation**
Nathi Magubane, [Penn Today](#).
- Nov 7, 2023 **ChatGPT-Maker OpenAI Hosts its First Big Tech Showcase as the AI Startup Faces Growing Competition**
Barbara Ortutay and Matt O'Brien, [Associated Press](#).
Reposted by 162 US news outlets, including The Independent, ABC News, Washington Post, U.S. News & World Report, The Business Journal, and CBS News.
- Nov 6, 2023 **As OpenAI's Multimodal API Launches Broadly, Research Shows It's Still Flawed**
Kyle Wiggers, [TechCrunch](#)
- Oct 2, 2023 **Rewriting the Script: Developing Effective AI Assistants**
Alyssa Hwang, [Penn Engineering Today](#).
- Sept 23, 2023 **Top Trending LLM Projects of the Week: Dive into the Future of Tech!**
AutoGPT Tutorials, [YouTube](#).
- Sept 20, 2023 **Kani, an Open-Source Framework from the University of Pennsylvania, is Reshaping Large Language Model Applications**
[Multiplatform.ai](#).

- Sept 19, 2023 **Kani: A Lightweight and Customizable Framework for Language Model Applications**
Michał Rogucki, [TS2 Space](#).
- Sept 19, 2023 **Unlocking AI Potential: Unveiling Kani, the Groundbreaking Open-Source Framework Revolutionizing Large Language Model Applications**
Casey Jones, [CJ&CO](#).
- Sept 18, 2023 **Researchers from the University of Pennsylvania Introduce Kani: A Lightweight, Flexible, and Model-Agnostic Open-Source AI Framework for Building Language Model Applications**
Tanya Malhotra, [MarkTechPost](#).
- Feb 19, 2021 **Doctorates do Valentines!**
Eboney Johnson, [Official Blog of the Penn Dept. of Computer and Information Science](#).

Teaching

- Fall 2022 **Head Teaching Assistant**, CIS 3990 Introduction to Human-Computer Interaction
Instructor: [Danaë Metaxa](#) (University of Pennsylvania)
- Fall 2021 **Teaching Assistant**, CIS 5300 Computational Linguistics
Instructor: [Mark Yatskar](#) (University of Pennsylvania)
- Spring 2020 **Teaching Assistant**, COMS 4705 Natural Language Processing
Spring 2019 Instructor: [Michael Collins](#) (Columbia University)
- Fall 2019 **Teaching Assistant**, COMS 4705 Natural Language Processing
Instructor: [Kathleen McKeown](#) (Columbia University)

Mentoring

- Spring 2024 Emma Jin (Master's Independent Study and Career Development)
- Fall 2023–
Spring 2024 Matthew Burke, Ryan Hsing-Smith, Bryan Yao, Joseph Koo,
Sebastian Lee, Zachary Arnold (Undergraduate Senior Design Project)
- Summer 2022 Juni Kim (High School Independent Study)
- Spring 2022 Jasmine Chen (Undergraduate Thesis)

Outreach

- Sept 2023 **University of Pennsylvania Department of Computer and Information Science**
Invited Panelist, PhD New Student Orientation
- Nov 2022 **MindCORE Diversity and Equity Initiative (DivE In)**
Invited Speaker, Overview of the PhD Application Process
- Nov 2021, 2020 **University of Pennsylvania School of Engineering and Applied Science**
Invited Panelist, PhD Open House
- Mar 2021 **ACM Capital Region Celebration of Women in Computing (CAPWIC)**
Invited Panelist, Graduate School: Is It Right for Me?
- Jan 2021 **University of Pennsylvania School of Engineering and Applied Science**
Invited Panelist, Graduate Recruitment Panel
- Oct 2020 **Columbia University Department of Computer Science**
Invited Speaker, Alumni-In-Residence Discussion

Leadership

- May 2022–Present **Penn CIS PhD Mentorship Program**
Founder
- Sept 2020–Present **Computer and Information Science Doctoral Association (CISDA)**
Co-Chair
Advocate for student interests and liaise with faculty of CIS and Penn Engineering.
- Sept 2022–
May 2023 **Department of Computer and Information Science Office Committee**
PhD Representative
- Sept 2021–
May 2022 **Penn Engineering Dean's Doctoral Advisory Board**
PhD Representative
- Sept–Dec 2021 **Penn NLP Computational Linguistics Lunch (CLunch)**
Co-organizer
- Mar–Oct 2021 **MindCORE Diversity and Equity Initiative (DivE In)**
Programming Development Committee Lead, Steering Committee Member
Collaborate across departments to establish inaugural DivE In, a fully funded weekend of workshops for underrepresented undergraduates in the mind sciences.