

Kellen Cheng

✉ kellentan@princeton.edu 📠 4074094268 🏠 Website 🇺🇸 U.S. Citizen 📅 Available Aug. 2025

Education

Princeton University

Aug. 2022 - Nov. 2024

M.A. ELECTRICAL AND COMPUTER ENGINEERING, GPA: 4.00

PRINCETON, NJ

University of California, Los Angeles (UCLA)

Sept. 2018 - Jun. 2022

B.S. ELECTRICAL ENGINEERING, TECHNICAL BREADTH COMPUTER SCIENCE, GPA: 3.93

LOS ANGELES, CA

Publications

- **Kellen Tan Cheng**, Anna Lisa Gentile, Pengyuan Li, Chad DeLuca, Guang-Jie Ren. *Don't Be My Doctor! Recognizing Healthcare Advice in Large Language Models*. EMNLP 2024 Industry Track.
- **Kellen Tan Cheng**, Suma Bhat. *No Context Needed: Contextual Quandary In Idiomatic Reasoning With Pre-Trained Language Models*. NAACL 2024 Main.
- Ziheng Zeng, **Kellen Tan Cheng**, Srihari Venkat Nanniyur, Jianing Zhou, Suma Bhat. *IEKG: A Commonsense Knowledge Graph for Idiomatic Expressions*. EMNLP 2023 Main.
- **Kellen Tan Cheng**, Kunakorn Atchaneeyasakul, Zeid Barakat, David Liebeskind, Fabien Scalzo. *CT Perfusion Imaging of the Brain with Machine Learning*. ISVC 2021.

Experience

IBM Research: Research Scientist Intern

May. 2025 - Aug. 2025

MENTOR: DR. ANNA LISA GENTILE

SAN JOSE, CA

- Returning intern for Summer 2025.

Samsung Research America: Research Scientist Intern

Jan. 2025 - May. 2025

MENTOR: TBD

MOUNTAIN VIEW, CA

- NLP/LLM intern in the GenAI for Health team.

Princeton: Graduate Research Assistant

Nov. 2022 - Nov. 2024

ADVISOR: DR. SUMA BHAT

PRINCETON, NJ

- Advanced the performance and understanding of figurative language reasoning capabilities of pre-trained language models ranging from 110M to 7B parameters.
- Designed and implemented an end-to-end two-mask infilling fine-tuning framework for idiomatic knowledge injection with compact language models (e.g. BART) using the IEKG dataset.
- Coupled my fine-tuning framework with transfer learning from the MNLI dataset, achieving a new state-of-the-art performance of 83.75% on the IMPLI benchmark, an improvement of 12% compared to previous state-of-the-art.
- Validated my framework through an ablation study, demonstrating gains of 23.08% compared with off-the-shelf models and 5.79% compared with vanilla transfer learning.
- Formulated and conducted a comprehensive context perturbation study for off-the-shelf language models (ranging from BART to Mistral-7B), uncovering that they are negatively affected by the context (up to 3.89% better *without* context).
- Published my works as a first-author NAACL 2024 main conference paper and a second-author EMNLP 2023 main conference paper.

Princeton: Graduate Teaching Assistant

Jan. 2024 - Dec. 2024

SML/PHI 354: PROF. SARAH-JANE LESLIE

PRINCETON, NJ

- Sole assistant instructor (TA) responsible for 56 (Spring 2024) and 47 students (Fall 2024) enrolled in SML/PHI 354 (Artificial Intelligence: A Hands-On Introduction from Basics to ChatGPT).

- Lead 4 (Fall 2024) and 6 (Spring 2024) weekly precept sections (50 min. each) to review fundamental course material, spanning artificial neural networks, convolutional neural networks, and transformer models.
- Held weekly office hours and was responsible for handling assignments, final projects, and exams.

IBM Research: Research Scientist Intern

Jun. 2024 - Sept. 2024

MENTOR: DR. ANNA LISA GENTILE

SAN JOSE, CA

- Developed a lightweight BART-Large detector model for health advice as a summer intern in the OneShield AI safety guardrails team at IBM Research-Almaden.
- Manually synthesized a new, gold standard, evaluation benchmark (HeAL) for health advice guardrails, containing 402 samples from a mix of professional, medical, and conversational sources on Common Crawl.
- Designed an on-demand human-in-the-loop system based on grouped clustering for sparse human annotation of synthetic production data at scale.
- Created a new framework to generate synthetic production data for health advice guardrailing, resulting in state-of-the-art performance of 85.07% from a BART-Large detector model, an increase of 3.48% compared to GPT-4o.
- Validated my findings by running additional experiments on proprietary IBM data, demonstrating real world performance gains with a false positive rate reduction of up to 3.18%.
- Programmed and deployed an internal Rest API module to automate the loading and evaluation process of all internal use datasets for the OneShield team.
- Published my work as a first-author EMNLP 2024 industry track paper, with another first-author work in submission to ACL 2025.
- Defended and filed a patent detailing a continual learning framework with model version-control and knowledge distillation for health advice detector development (bar date of April 2025).

UCLA: ECE Course Grader

Sept. 2021 - Jun. 2022

ECE 102, ECE C143A: PROF. JONATHAN KAO

LOS ANGELES, CA

- Graded weekly assignments for over 100 students enrolled in ECE 102 systems and signals (Fall 2021) and ECE C143A neural signal processing (Spring 2022).
- Provided detailed feedback and additional notes for students on assignments.

PerfuseAI: Machine Learning Intern

Jun. 2021 - Sept. 2021

ADVISOR: PROF. FABIEN SCALZO

REMOTE

- Optimized random forest and kernel ridge regression models to predict CT perfusion blood flow parameters such as rCBF, rCBV, MTT, and TMax.
- Incorporated the end-to-end model pipeline from my ISVC 2021 paper into the PerfuseAI product, with a final objective of deployment for Thailand’s hospital network.
- Empirically tested the prototype arterial input function (AIF) image annotation tool designed to be used by physicians on ischemic stroke patients.

Awards & Organizations

Toby & Jack Wolf Travel Grant	2024
Bede Liu Travel Grant	2023
Princeton ECE Departmental Fellowship	2022
Tau Beta Pi	2020 - Present
IEEE Eta Kappa Nu (HKN)	2019 - Present
UCLA Dean’s Honor List	2019 - 2022