Chris Yuhao Liu

Contact Information

Engineering 2 Room 489
1156 High Street
Santa Cruz, CA 95060

{yliu298,chrisliu298}@ucsc.edu
chrisliu298.ai

Research Interests

My current research focuses on *machine unlearning*, which aims to remove information from a trained model. Previously, I worked on fundamental problems of modern deep learning. These problems include topics such as double descent, data scaling laws, and structural risk minimization.

Education

University of California, Santa Cruz, Santa Cruz, CA
Ph.D. in Computer Science and Engineering — September 2023 - present
Advisors: Yang Liu and Jeffrey Flanigan

University of California, Santa Cruz, Santa Cruz, CA
M.S. in Computer Science and Engineering — September 2021 - June 2023
Advisor: Jeffrey Flanigan

University of California, Santa Cruz, Santa Cruz, CA
B.S. in Computer Sciences and Engineering — September 2017 - March 2021

Research Experience

REAL, University of California, Santa Cruz, Santa Cruz, CA, USA
Graduate Student Researcher — September 2023 - Present
Advised by Yang Liu and Jeffrey Flanigan

JLab, University of California, Santa Cruz, Santa Cruz, CA, USA
Student Researcher — April 2020 - August 2023
Advised by Jeffrey Flanigan

REAL, University of California, Santa Cruz, Santa Cruz, CA, USA
Research Intern — June 2022 - June 2023
Worked with Yang Liu

Research Projects

Understanding the Role of Optimization and Loss Function in Double Descent [Paper]
Overfitted models do not exhibit the double descent phenomenon due to 1) weak optimizers struggling to land at a low-loss local minimum and the 2) presence of an exponential tail in the shape of the loss function.

Structural risk minimization for deep neural networks
We propose a new regularization based on structural risk minimization that directly minimizes the generalization gap.

What Determines Sample Complexity Rate in Practice?
We empirically estimate the power-law exponents of various model architectures and study how they are altered by a wide range of training conditions for classification.

Toward Disentangling Double Descent and Information Flow in Deep Neural Networks [Paper], [Code], [Slides]
We study the relationship between the amount of mutual information compression and generalization given the double descent phenomenon.

Learning to Extract Compact Vector Representations from Weight Matrices [Paper], [Code], [Slides]
We study the problem of learning to construct compact representations of neural network weight matrices by projecting them into a smaller space.

Sample Complexity Scaling Laws For Adversarial Training [Paper], [Code]
We show that adversarially training (Fast Gradient Sign Method and Projected Gradient Descent) reduces the empirically sample complexity rate for MLP and a variety of CNN architectures on MNIST and CIFAR-10.

Faster Sample Complexity Rates With Ensemble Filtering
We present a dataset filtering approach that uses sets of classifiers, similar to ensembling, to estimate noisy (or non-realizable) examples and exclude them so a faster sample complexity rate is achievable in practice.

Course Projects

Conditional Generation of Research Paper Abstracts [Code]

- Fine-tuned a GPT-2 model using all research paper titles and abstracts under cs.AI, cs.LG, cs.CL, and cs.CV on arXiv.
- This project was the winner of the Generative Modeling Competition for the course CSE142 Machine Learning in Spring 2020.

Text Augmentation Using Pre-Trained Transformers With Reinforcement Learning [Code]

- Trained distilled RoBERTa model as a text classifier and a GPT-2 as a text generator using proximal policy optimization synchronously to generate augmented text for text classification tasks.

Sentiment Analysis with Transformers [Code]

- Fine-tuned a RoBERTa model on the IMDb dataset for sentiment analysis.
- This project was the winner of the Sentiment Analysis Competition for the course CSE142 Machine Learning in Spring 2020.

Teaching Experience

University of California, Santa Cruz, Santa Cruz, CA, USA
Teaching Assistant September 2021 - June 2023

- Courses: CSE 20 Introduction to Python (Fall 2021, Spring 2022, Fall 2022), CSE 30 Programming Abstractions: Python (Spring 2023), CSE 144 Applied Machine Learning (Winter 2022).

University of California, Santa Cruz, Santa Cruz, CA, USA
Tutor and Grader September 2020 - December 2020

- Course: CSE 142 Machine Learning (Fall 2020).

Honors and Awards

Honors in the Major, Cum Laude, University of California, Santa Cruz 2021

Services

Volunteer, Thirty-eighth International Conference on Machine Learning 2021
Volunteer, Ninth International Conference on Learning Representations 2021