Jiayu Zheng

401-447-9534 | jiayu_zheng@brown.edu | Academic Homepage | github.com/JiayuZheng

LLMs SFT, PEFT, Instruction-tuning, Preference Optimization, VLM finetuning

Education

Brown University — Providence, RI
Sc.M. in Data Science — Sept. 2022 – Present

- GPA 4.0
- Relevant Courses:
  Data Science, Data Engineering, Statistical Learning, Machine Learning, Deep Learning, Operating Systems, Fairness in ML, Statistical Learning, Artificial Intelligence (TA), Computer Vision(TA)

Zhejiang University — Zhejiang, China
Sc.B. in Chemistry with Honors — Sept. 2018 – June 2022

- Graduated from the Cho Kochen Honors College
- Major GPA: 3.92/4.0, Cumulative GPA: 3.8/4.0, Rank: 2%
- Relevant Courses:
  C, Java, Object-Oriented Programming, Mathematical Modeling, Digital Logic Design, Linux Application Technology, Advanced Data Structure and Algorithm Analysis

Publications And Manuscripts

A Molecular Stereostructure Descriptor Based On Spherical Projection — 2020

- Licheng Xu, Xin Li, Miaojiong Tang, Luotian Yuan, **Jiayu Zheng**, Shuoqing Zhang, Xin Hong

Research Experience

Graduate Research Assistant — Providence, RI
BATS Lab, *Brown University* — Apr. 2023 – Present

- Advised by Prof. Stephen Bach
- Designed data-efficient, transferable soft prompts, called *reasoning function tokens*, trained in a bootstrapping manner to improve the out-of-domain (OOD) generalizability of LLMs and VLMs
- Measured the performance of Mistral-7b-instruct LLM on text classification datasets, like Ledgar legal documents classification and SMS Spam
- Collected "text relation entity" corpora by prompting Mistral-7b-instruct and bootstrap-finetuned itself on the corpora

Graduate Research Assistant — Providence, RI
Conversational AI Lab, *Brown University* — June 2023 – Jan 2024

- Co-advised by Prof. Shekhar Pradhan and Prof. Ritambhara Singh
- Conducted research on relation-aware CLIP with specialized training objectives and hard negative mining
- Collected a hard-negative dataset for MSCOCO Captions by prompting large language models (LLM)
- Designed a relation-aware image-text contrastive (ITC) training objective to explicitly encourage models to discriminate between highly similar images and captions
- Designed a weakly supervised, LLM-in-the-loop, labeling function to determine whether alternative captions retain original semantics

Undergraduate Research Assistant — Zhejiang, China
Zhejiang University — Apr. 2020 – July 2022

- Conducted research on representation learning of molecules, contributing to one published study and one thesis
- Designed a numeric descriptor to transform the van der Waals force field of molecules to a sequence of 2D matrices, collected a dataset using the descriptor, trained a CNN to predict the stereo-selectivity of reactions
- Designed a distance-aware message-passing mechanism inspired by the locality of atomic interaction, which greatly reduces the MSE of energy prediction using GNNs

Projects

JNeedle | *Python, C, CUDA* Dec. 2023 – Feb. 2024

- Implemented a PyTorch-like reverse-mode AutoDiff Deep Learning Library with a Numpy-like C/CUDA accelerated array operation backend
- Built trainable deep neural nets like ResNet, CNN, and Transformer using only the package

Tilt-Shift Filter | *C++, OpenMP, CUDA* Sept. 2023 – Dec. 2023

- Implemented a tilt-shift filter to convert real-life photos to miniature scenes by applying Gaussian blur that increases with the distance from the focal plane
- Achieved 10x speed-up (*c.f.* Python) by creating filters using OpenMP and implementing convolution using CUDA; 560x speed-up in total with further quantized blur scales

GourmAIt | *Python, PyTorch* July 2023 – Sept. 2023

- Performed image classification on a noisy dataset, Food101, with self-supervised learning (SSL)
- Implemented Google Research's Noisy Student Training, where pseudo labels generated by a teacher model are used to train a larger-or-equal-size student model, which is then used as the teacher in the next iteration
- Implemented ResNet with stochastic depth, step-wise unfreezing scheduling, and learning rate scheduling to maximize the performance gain in the fine-tuning phase

TransformerHub | *Python, PyTorch* Sept. 2022 – Aug. 2023

- Implemented various encoder-only, decoder-only, and encoder-decoder Transformer models including Transformer, BERT, GPT, ViT, and CLIP
- Implemented advanced features like sliding window attention, rotary position embedding, mixed precision training, and gradient accumulation

WeenixOS | *C, Unix kernel, X86_64 arch* Feb. 2023 – May 2023

- Developed a well-functioning, kernel-based, Unix operating system, the WeenixOS, as the semester-long project of the course Operating Systems, instructed by Prof. Thomas Doeppner
- Implemented a thread pool, device drivers, a virtual file system (VFS), an S5FS, and a virtual memory system
- Implemented various system calls, like do_waitpid, do_read, do_write, do_brk, do_fork, etc.

CyberBarista | *Python, TensorFlow, Scikit-learn* Sept. 2022 – Dec. 2022

- Implemented an ensemble of nine models, including an ElasticNet, a multi-layer perceptron (MLP), etc., to predict the quality score of coffee beans
- Interpreted the model output using feature-based methods: random perturbation, weight magnitude, and SHAP

Teaching and Mentorship

Teaching Assistant | *Computer Vision*, Instructor: Prof. Srinath Sridhar, Brown University Spring 2024
Teaching Assistant | *Artificial Intelligence*, Instructor: Dr. Thao Nguyen, Brown University Fall 2023
Mentor | *Data Science bootcamp*, Brown University Apr. 2023
Mentor | *Women in Data Science (WiDS)*, Brown University Feb. 2023
Head Teaching Assistant | *Structural Chemistry and Spectroscopy*, Zhejiang University Feb. 2020 – Dec. 2021

Honors and Awards

Outstanding Graduates of Zhejiang University 2022
Cho Kochen Honors College Scholarship for Innovation (publications or international contest awards) 2021
Second-Class Scholarship for Elite Students in Basic Sciences (top 20%) 2021
Outstanding Students of Zhejiang University 2019, 2021
Zhejiang University Scholarship, Third Prize (top 15%) 2020
Zhejiang University Scholarship, First Prize (top 3%) 2019, 2021

Technical Skills

Programming Language: Python, C/C++, Java, CUDA, SQL (SQLite), Latex, R
Deep Learning Packages: Tensorflow, PyTorch, Huggingface Transformers, LangChain, LlamaIndex, OpenAI Gym, Scikit-learn
Development Tools: Venv, Conda, Docker, Git, ssh, Jupyter, CMake, SLURM, gdb, wandb
Training/Inference Systems: Accelerate, PEFT, DDP/FSDP, DeepSpeed, ZeRO, Flexgen
LLMs/VLMs: GPT2-4, OPT, T0++, Llama1/2/3, T5, Flan-T5, Falcon, Mistral, Gemma; CLIP, BLIP, BLIP-2, BEIT-3, ALBEF, GPT-4v, Llava