Ge Shi

(413) 425-3452 Github geshi@ucdavis.edu Personal Page

Education

University of California, Davis

Ph.D. in Computer Science; GPA: 3.8/4.0 University of Massachusetts, Amherst Master of Science in Computer Science; GPA: 3.9/4.0

Davis, CA

Linkedin

Zhejiang University

Bachelor of Engineering in Automation; GPA: 3.7/4.0

PROGRAMMING SKILLS

Coding: Python, Pytorch, Hugging Face, DeepSpeed, FSDP, Bash, Matlab, D3.js, Node.js, C/C++, Java, Flask. **Tools**: Docker, Git, LangChain, ROS, WANDB, RAY, SageMaker, AWS Capacity Block, VectorDB.

Selected AI Projects

Automatic Prompt Optimization for Social Science

- **Background**: Traditional text annotation is labor-intensive and relies heavily on fine-tuning language models. Few-shot learning with uses limited fixed examples with little human knowledge.
- Method: Developed a high-performance, cost-effective, three-stage (pre-process, in-process, post-process) framework with LangChain for automated text labeling in social science datasets.
- Enhanced Prompt Generation: Employed UMAP for feature reduction and K-Means for prototype selection, followed by MapReduce on the prototype pool, resulting in a 3.8% improvement in F1 score.
- Dynamic Few-Shot Learning: Implemented a dynamic few-shot example retriever utilizing a class-wise diverse MMR algorithm on pre-computed embeddings, enhancing F1 score by 1.4% over the first stage.
- Advanced LLM as a Judge: Leveraged GPT-4 to resolve discrepancies between two weaker LLMs through chain-of-thought reasoning, achieving a 1.1% increase in F1 score and reducing costs by 78%.

Large Language Model (LLM)

- **Background**: Pre-trained LLMs like Llama excel in natural language understanding and auto-completion but require supervised fine-tuning (**SFT**) to effectively follow human instructions.
- Method: Explored various parameter-efficient fine-tuning (PEFT) methods for instruction-following using Alpaca and HealthCareMagic Chat QA datasets on the Llama-3 8B model, utilizing 8 A100 GPUs.
- LoRA and QLoRA: Adopted LoRA (Low-Rank Adaptation) and quantized LoRA with DeepSpeed, which fine-tunes small low-rank matrices added to the model's weights, reducing the number of trainable parameters by up to 90% saving up to 75% of memory and 40% training time.
- Adapter and Llama-Adapter: Deployed Adapter (introduces small bottleneck layers between the existing layers) and Llama-Adapter (adds adaptation prompts prefixed to the input instruction tokens) outperforming the LoRA fine-tuned model (61% to 39%), as evaluated by GPT-4.

Vision Language Model (VLM)

- **Background**: Transformer-based visual and language encoders excel at feature extraction but struggle with collaboration on zero-shot classification and text generation from visual prompts.
- Prompt Tuning for Zero-Shot Classification: Implemented both general and class-specific prompt tuning for CLIP with a learnable prefix context on the MIT Indoor Scenes dataset, achieving an improved accuracy of 4.5% for zero-shot classification.
- Vision-Language Model Fine-tuning for Image Captioning: Fine-tuned BLIP-2 on the COCO 2017 dataset for image captioning, utilizing the validation set and test-dev set for training and testing. Outperformed the pre-trained model by 2.8 in BLEU score .
- Visual Instruction Tuning for Visual QA: Explored visual instruction tuning using CLIP and LLaMA-2 7B, following Llava. Distributed trained and evaluated on the ScienceQA benchmark and achieved a 79.3% accuracy.

Davis, CA Sep. 2024 Amherst, MA May. 2019

Hangzhou, China July. 2017

Davis, CA

Davis, CA

Davis, CA

Multi-view Learning on Neuroimaging Data

Supervisor: Prof. Ian Davidson

Davis, CA Apr. 2020 - Nov. 2021

Davis, CA

Dec. 2021 - Mar. 2023

- **Background**: Prognosis problem (predicting the outcome of a treatment) relies on costly clinical evaluation. We leveraged t-fMRI brain scans in early disease auto-prognosis problem to alleviate the burden.
- **Difficulty**: The scans are long and rare compared to multimedia data. Each scan is split into snippets given 4 imbalanced events and with only one paired label, which forms a unconventional machine learning problem.
- **MVMI Learning**: Proposed a novel multi-view multi-instance learning scheme employing 3D CNN, transfer learning, even sampling, ensemble stacking learning to reach SOTA accuracy 75.6%.
- Mixup Data Augmentation: Applied both intra-group and inter-group mixup data augmentation method to enhance the prediction ability of the model for both improvers (2.4%) and non-improvers (3.7%).

Explainable Artificial Intelligence

Supervisor: Prof. Ian Davidson

- **Background**: Post-hoc local explanation methods assign feature importances to unveil the decision making of black-box machine learning models to a single prediction. The popular methods include CAM, IG, SHAP.
- Feature Selection: Explored leveraging local XAI in feature selection which beats the classical wrapper method. Created a synthetic dataset and benchmark to evaluate local XAI in low signal-to-noise environments.
- Shapley IG: Proposed a novel attribution method called *Shapley Integrated Gradients* (SIG) which adapts Integrated Gradients to approximate Shapley Value which is faster than SHAP and more robust than IG.
- **Multi-level Explanations**: Proposed a method to aggregate local explanations at various granularities based on the closeness on both features and explanations and explored the its correlation with the model's local performance.

INTERNSHIPS

Lawrence Berkeley National Laboratory

Research Assistant in Michael Mahoney's Lab

Berkeley, CA Apr. 2023 - Mar. 2024

- **Background**: Loss landscapes visualization are essential to understand the model robustness, generalization, and optimization problems (e.g. model quantization, meta-optimization, fast ensembling).
- **Difficulty**: There is a lack of visual analytics tools to 1) visualize both local loss landscapes and global loss landscape statistics systematically 2) extend to more than 2D loss landscapes.
- LossLens: Created a visualization tool (Losslens) for ML . Primarily worked on the back-end implementation of pyHessian loss landscapes and mode-connectivity for ResNet, UNet, ViT, and PINN.
- **Topological Analysis**: Extended the 2D loss landscapes visualization to 3D and 4D subspaces defined by top eigenvectors of the Hessian matrix leveraging vertices sampling, neighborhood graph, and merge tree.

Selected Publications

KDD Deep learning for prognosis using task-fmri: A novel architecture and training scheme.
IJCAI HALE: Hierarchical Aggregation of Local Linear Additive Explanations.
TVCG Losslens: Diagnostics for Machine Learning Models through Loss Landscape Visual Analytics.
NeurIPS D&B ChaosMining: A Benchmark to Evaluate Post-Hoc Local Attribution Methods in Low SNR Environments.

PROFESSIONAL SERVICES

Academic Service: AAAI, IJCAI, KDD, VIS, CIKM, JMLR, MLIS, Frontiers in Neuroscience. Teaching Assistant: Introduction to Artificial Intelligence, Machine Learning and Discovery, Programming in C.

Awards

UC Davis Computer Science Department Dissertation Writing Award 2024 National Science Foundation (NSF) Traveling Award 2024 UC Davis Academic Traveling Award 2022, 2023 Meritorious Winner in The International Mathematical Contest in Modeling 2017 ZJUNlict 3rd Place in Robocup Kidsize Humanoid League 2016 Zhejiang University Excellent Academic Scholarship 15/100