Jun Luo

www.junluo.me | ljaiverson365@gmail.com | linkedin.com/in/junluo1 | Google Scholar | San José, California

EDUCATION

University of Pittsburgh — **Pittsburgh, PA**
Ph.D. in Intelligent Systems (in progress) — *Aug., 2019 – Present*

University of Southern California — **Los Angeles, CA**
*M.S. in Electrical Engineering (**Honors Fellow**), GPA: 3.92/4.0* — *Aug., 2017 – May, 2019*

Huazhong University of Science and Technology — **Wuhan, China**
B.E. in Electrical Engineering and Automation, GPA: 3.72/ 4.0 — *Sep., 2013 – Jun., 2017*

University of California, Berkeley — **Berkeley, CA**
Exchange program with selective admission: Berkeley International Study Program — *Aug., 2016 – Dec., 2016*

Core courses
Advanced Topics in Machine Learning, Found. of AI, Computer Vision, LLM Engineering, Intro. to NLP, Learning-Based Image Synthesis, Optimization, Mathematical Pattern Recognition, Time Series, ML with Large Datasets, Applied Statistical Methods

EXPERIENCES

Sony AI, Privacy-Preserving Machine Learning Team — **Tokyo, Japan (remote in US)**
Research Intern, mentored by Dr. Weiming Zhuang and Dr. Chen Chen — *May, 2023 – Jul., 2023*
Efficiently leveraging foundation models in federated learning
- Led research that empowers heterogeneous resource-constrained devices with knowledge from foundation models.
- Proposed a knowledge distillation-based communication- and computation-efficient personalized federated learning algorithm, leveraging large foundation models for clients with limited and heterogeneous resources.
- Implemented the proposed idea for Sony AI's federated learning framework and composed related documentation.

University of Pittsburgh, ICCI Lab — **Pittsburgh, PA**
Graduate Student Researcher, advised by Dr. Shandong Wu@Pitt and Dr. Chen Chen@UCF — *May, 2020 – Present*
***Federated learning with Vision-Language Models and general deep learning models** [Project page]*
- Proposed federated learning and its personalization algorithms under the scopes of Mixture of Experts (MoE), Vision-Language Model (VLM), prompt tuning, first-order approximation, and adaptive personalization.
- Proposed algorithms achieve high performance, accelerated convergence, and reduced overhead *(with up to 15.47% accuracy increase, 4.2× convergence speed-up, one magnitude ($O(N^2) \rightarrow O(N)$) of overhead reduction over baselines).*
- Related papers published in ICLR 2025, ICCV 2023 (including 1 oral, top 1.8 %), IJCAI 2022, and ISBI 2022.

***Curriculum learning with domain knowledge for medical image classification** [Project page]*
- Proposed data level and task level curriculum learning strategies for multi-view cancer and fracture diagnoses.
- Collaborated with radiologists on designing domain knowledge-guided ML pipelines with boosted performance *(up to 3.6% accuracy increase, 4.8% AUC increase, and 6.2% F1 score increase over baselines on medical datasets).*
- Related papers published in ISBI 2023, SPIE Medical Imaging 2022 & 2021, MICCAI-W 2021, and ARRS 2021.

University of Pittsburgh, Dr. Hauskrecht's Lab — **Pittsburgh, PA**
Graduate Student Researcher, advised by Dr. Milos Hauskrecht@Pitt — *Jan., 2020 – Apr., 2020*
***Time series modeling of Ordinary Differential Equation (ODE) -based biomedical system** [Project page]*
- Developed an ODE-based generative model for time series data synthesis for wound healing inflammation process.
- Adopted Markov chain Monte Carlo (MCMC) for ODE coefficients posterior estimation from synthesized data.
- Built a Seq2Seq model with LSTM to predict the future time series trajectories of the biomedical signals.

University of North Carolina at Charlotte, Computer Vision Lab — **Charlotte, NC**
Research Intern, advised by Dr. Chen Chen@UCF — *Jul., 2018 – Dec., 2018*
***Video anomaly detection pipeline development for the elderly** [Project page]*
- Fine-tuned Inflated 3D ConvNet (I3D) on fall detection video datasets and post-processed the results with YOLO-V3.
- Employed OpenPose and Spatial-Temporal Graph ConvNet (GCN) for fall detection from pose estimation in videos.

SELECTED PUBLICATIONS

- **Luo, J.**, Chen, C., & Wu, S. "Mixture of Experts Made Personalized: Federated Prompt Learning for Vision-Language Models". In *Proceedings of the International Conference on Learning Representations (**ICLR**)*, 2025. [PDF]

- **Luo, J.**, Mendieta, M., Chen, C., & Wu, S. "PGFed: Personalize Each Client's Global Objective for Federated Learning". In *Proceedings of the IEEE/CVF International Conference on Computer Vision (**ICCV**)*, 2023. (**Oral, top 1.8%**) [PDF]
- Sun, G., Mendieta, M., **Luo, J.**, Wu. S., & Chen. C. "FedPerfix: Towards Partial Model Personalization of Vision Transformers in Federated Learning". In *Proceedings of the IEEE/CVF International Conference on Computer Vision (**ICCV**)*, 2023. [PDF]
- **Luo, J.**, Arefan, D., Vasireddi, A., Wu, S., and Nguyen, N. "Potential Use of Artificial Intelligence in Sincalide-Stimulated Cholescintigraphy: A Pilot Study". In *Journal of Nuclear Medicine (**JNM**), P1119-P1119*, 2023. [URL]
- Zhou, Z., **Luo, J.**, Arefan, D., and Kitamura, G., & Wu, S. "Human Not in the Loop: Objective Sample Difficulty Measures for Curriculum Learning". In *2023 IEEE 20th International Symposium on Biomedical Imaging (**ISBI**), pp. 1-5. IEEE*, 2023. (**Oral**) [PDF]
- **Luo, J.**, & Wu, S. "Adapt to Adaptation: Learning Personalization for Cross-Silo Federated Learning". In *Proceedings of the International Joint Conferences on Artificial Intelligence (**IJCAI**)*, 2022. (**Acceptance rate: 14.9%**) [PDF]
- **Luo, J.**, & Wu, S. "FedSLD: Federated Learning with Shared Label Distribution for Medical Image Classification". In *2022 IEEE 19th International Symposium on Biomedical Imaging (**ISBI**), pp. 1-5. IEEE*, 2022. (**Oral**) [PDF]
- **Luo, J.**, Arefan, D., Zuley, M., Sumkin, J., & Wu, S. "Deep Curriculum Learning in Task Space for Multi-Class Based Mammography Diagnosis". In *Medical Imaging 2022: Computer-Aided Diagnosis (Vol. 12033, p. 71-76). International Society for Optics and Photonics (**SPIE**)*, 2022. (**Oral**) [PDF]
- **Luo, J.**, Kitamura, G., Arefan, D., Doganay, E., Panigrahy, A., & Wu, S. "Knowledge-Guided Multiview Deep Curriculum Learning for Elbow Fracture Classification". In *MICCAI 2021 Workshop on Machine Learning in Medical Imaging (**MICCAI-w**) (pp. 555-564). Springer, Cham*, 2021. (Poster) [PDF]
- **Luo, J.**, Kitamura, G., Doganay, E., Arefan, D., & Wu, S. "Medical Knowledge-Guided Deep Curriculum Learning for Elbow Fracture Diagnosis from X-Ray Images". In *Medical Imaging 2021: Computer-Aided Diagnosis (Vol. 11597, p. 1159712). International Society for Optics and Photonics (**SPIE**)*, 2021. (**Oral**) [PDF]

PROJECTS

LLM fine-tuning for Amazon item price prediction [GitHub] — Fall 2024

- Curated and published [dataset], trained baseline regression models with bag of words and engineered features.
- Full fine-tuned GPT-4o mini using OpenAI API, tracking loss and progress with wandb, and designed evaluation.
- Parameter-efficient fine-tuning with QLoRA on Llama 3.1 8B, published [trained model weights] on hugging face.

LLM-based agentic AI for best deals finder from deals RSS [GitHub] — Fall 2024

- Pipelined an agentic AI solution that pushes phone notifications for best deals with high-quality item descriptions.
- Designed a planner agent that sequentially calls an LLM-based scanner agent to acquire high-quality deals from RSS, an ensemble agent to estimate the discount, and a messenger agent to push cell phone notifications for the best deals.
- Implemented the ensemble price predictor agent with RAG empowered GPT-4o mini (Chroma knowledge base of SentenceTransformer features), fine-tuned Llama 3.1 deployed on cloud (Modal), and locally trained Random Forest.
- Designed advanced UI with Gradio integration of details for cached deals and real-time log visualization.

Deblurring ultra-low dose radiation dental CT images with pix2pix [GitHub] — Fall 2021

- Applied pix2pix to enhance and deblur the collected ultra-low dose radiation dental CT image data, acquired radiologist -qualified better sharpness, suggesting potential reduction of radiation exposure while maintaining diagnostic quality.
- Results has been published in AAOMR Annual Session, 2021, as a clinical abstract for oral presentation. [PDF] [Slides]

Spatio-temporal interpolation in fluid dynamics [GitHub] — Spring 2019

- Collaborated with domain experts on fluid dynamics data (synthesized with Mantaflow) spatio interpolation with Super Resolution CNN and Super Resolution GAN, and temporal interpolation with SuperSloMo.
- Implemented vanilla and upgraded SRCNN, [poster] was awarded deep learning course Best Poster Runner-up.

HONORS AND AWARDS

- **2024** Provost Fellowship of University of Pittsburgh
- **2023** Intelligent Systems Program Travel Grant Award
- **2021** Intelligent Systems Program Fellowship
- **2020** Intelligent Systems Program Fellowship
- **2019** Master of Science Honors Fellow graduate
- **2018** Top 10% of Kaggle TGS Salt segmentation Challenge
- **2016** First place (university-wise) in *Neoway* electronic contest

SKILLS

- **Programming:** Python (PyTorch, TensorFlow, Keras, Pandas), Java, C/C++, SQL, R, MATLAB, JavaScript, HTML/CSS
- **Tools:** Git, Anaconda, Hugging Face, Chroma, LangChain, Wandb, Gradio, Node.js, React.js, Visual Studio, IntelliJ
- **Cloud:** AWS (EC2, S3, EMR), GCP (Deep Learning VM), Spark, Hadoop, Databricks, Docker, Kubernetes, Modal