

# Jiacheng Pang

+1-857-400-1382 | [✉ pangj@usc.edu](mailto:pangj@usc.edu) | [🌐 LinkedIn](#) | [🐙 Github](#) | [🌐 Personal Website](#)  
Los Angeles, California 90007, USA

## AREAS OF INTEREST

---

Natural Language Processing, Multimodality, Multi-agent systems, Time Series Analysis, Large Language Models

## EDUCATION

---

- **University of Southern California** May. 2026 - Dec. 2030  
*Ph.D. in Computer Science*  
◦ GPA: 4.00/4.00, supervised by Dr. Yan Liu at the USC Melady Lab  
Los Angeles, USA
- **University of Southern California** Jan. 2024 - Dec. 2025  
*M.S. in Computer Science (Artificial Intelligence)*  
◦ GPA: 3.92/4.00  
Los Angeles, USA
- **Northeastern University** Sept. 2018 - May 2022  
*B.S. in Computer Science, Minor in Mathematics*  
◦ GPA: 3.74/4.00  
Boston, USA  
◦ Honors: magna cum laude, Dean's List of all semesters

## RESEARCH EXPERIENCE

---

- **Melady Lab, USC Viterbi School of Engineering** [🌐](#) June 2025 - Present  
*Graduate Researcher*  
Los Angeles, USA  
◦ Co-built **HILA (ICLR 2026)**: a human-in-the-loop multi-agent LLM system that learns when to defer to a human expert, combining cost-aware GRPO with continual learning from expert interventions, improving baselines by **7.5%** accuracy on average  
◦ Co-developed **Maestro**, a multi-agent LLM framework with parallel execution agents for divergent exploration and a central agent for convergent synthesis, trained via Conditional Listwise Policy Optimization (CLPO), outperforming SOTA MAS baselines by **+6%** on average  
◦ Co-authored **two survey papers** on 1) time series reasoning with LLMs and 2) LLM based agentic systems with multi-agent reinforcement learning (**under review**)
- **Intelligent Human Perception Lab, USC Institute for Creative Technologies** [🌐](#) June 2025 - May 2026  
*Graduate Researcher*  
Los Angeles, USA  
◦ Designed **VoxParadox (ICML 2026)**, an adversarial paralinguistic benchmark to diagnose multimodal hallucination in speech LLMs, performed **linear probing** to examine the cause, and proposed **PCLM**, a novel layer fusion method to alleviate the issue, coupled with **DPO**, improving Qwen2-Audio baselines by **57% accuracy**  
◦ Contributed to **AVERE (ICLR 2026)**, a multimodal preference optimization method for audiovisual emotion reasoning; AVEm-DPO achieved **6–19%** relative gains over reference baselines in zero-shot (13%↑ UAR on DFEW over EmotionLLaMA); built the extensive evaluation pipeline  
◦ Contributed to **MoD-DPO (CVPR 2026)**, a modality-decoupled preference optimization framework that mitigates cross-modal hallucinations in omni LLMs via modality-aware regularization and language-prior debiasing; outperformed DPO and OmniDPO baselines on AVHBench and CMM; designed comprehensive analysis suite





## PUBLICATIONS

---

- **Do Audio LLMs Listen or Read? Analyzing and Mitigating Paralinguistic Failures with VoxParadox.** Jiacheng Pang, Ashutosh Chaubey, Mohammad Soleymani, ICML 2026. [\[pdf\]](#) [\[project page\]](#)
- **MoD-DPO: Towards Mitigating Cross-modal Hallucinations in Omni LLMs using Modality Decoupled Preference Optimization.** Ashutosh Chaubey, Jiacheng Pang, Mohammad Soleymani, CVPR 2026. [\[pdf\]](#) [\[project page\]](#)
- **AVERE: Improving Audiovisual Emotion Reasoning with Preference Optimization.** Ashutosh Chaubey, Jiacheng Pang, Maksim Siniukov, Mohammad Soleymani, ICLR 2026. [\[pdf\]](#) [\[project page\]](#)
- **Adaptive Collaboration with Humans: Metacognitive Policy Optimization for Multi-Agent LLMs with Continual Learning.** Wei Yang, Defu Cao, Jiacheng Pang, Muyan Weng, Yan Liu, ICLR 2026. [\[pdf\]](#) [\[code\]](#)
- **Maestro: Learning to Collaborate via Conditional Listwise Policy Optimization for Multi-Agent LLMs.** Wei Yang, Jiacheng Pang, Shixuan Li, Paul Bogdan, Stephen Tu, Jesse Thomason. [\[pdf\]](#)
- **Toward Evolutionary Intelligence: LLM-based Agentic Systems with Multi-Agent Reinforcement Learning.** Wei Yang, Muyan Weng, Jiacheng Pang, Defu Cao, Heng Ping, Peiyu Zhang, Shixuan Li, Yue Zhao, Qiang Yang, Mengdi Wang, Yan Liu. 2025. AI Open [\[pdf\]](#).
- **A Survey of Reasoning and Agentic Systems in Time Series with Large Language Models.** Ching Chang, Yidan Shi, Defu Cao, Wei Yang, Jeehyun Hwang, Haixin Wang, Jiacheng Pang, Wei Wang, Yan Liu, Wen-Chih Peng, Tien-Fu Chen. 2025. Under review at TMLR. [\[pdf\]](#).



## PROJECTS

---

- **VoxParadox (ICML 2026)** June 2025 - April 2026  
Advisor: *Prof. Mohammad Soleymani (Research Associate Professor, USC)* 
  - Created an **adversarial benchmark** to evaluate language dominance on Audio LLMs' paralinguistic understanding, with controlled linguistic-acoustic contradiction (2k clips over 10 tasks)
  - Provided analysis of failure modes through **layer-wise probing** and **attention scaling**, revealing text-prior reliance when answer is grounded in acoustic signals in major Audio LLM models (Audio Flamingo series, Qwen2-Audio, GPT-4o-Audio, and Gemini 2.5, etc.)
  - Proposed **Prompt-Conditioned Layer Mixer (PCLM)**, a novel layer fusion strategy that improved Audio Flamingo 3 baselines by **48%** and Qwen2-Audio baselines by **57%**, coupled with **DPO**
- **MSECap: Fusion Based Audio-textual Speech Emotion Captioning** Jan. 2025 - June 2025  
Advisor: *Prof. Mohammad Soleymani (Research Associate Professor, USC)* 
  - Designed **multimodal feature extractors** with three **fusion** strategies: early, late, and **X-Norm** fusion
  - Designed **feature adapters** for representation alignment and prefix-conditioned decoders with pretrained LLM
  - Implemented **training and evaluation pipelines** with **IEMOCAP**, X-Norm achieving **20%↑ match rate** compared with early/late baselines, showing high potential for improvement with more data
- **PoisonedGraphRAG: Knowledge Corruption Attacks Against GraphRAG** Sept. 2024 - Dec. 2024  
Advisor: *Prof. Yan Liu (Full Professor, USC)* 
  - Developed a **benchmark dataset** with a highly inter-connected graphical structure of named entities
  - Crafted poison texts based on generation condition and retrieval condition (introduced in **PoisonedRAG**) and achieved up to **70% attack success rate** against GraphRAG (with GPT-4o-mini)
  - Designed extensive per-question indexing **testing pipeline** to prevent inter-poison interference, and conducted experiments under various hyper-parameter settings
- **From Gold Labels to Trap Labels: Investigating the Reliability of SNLI** Sept. 2024 - Dec. 2024  
Advisor: *Prof. Swabha Swayamdipta (Assistant Professor, USC)* 
  - Proposed a **new method** for detecting benchmark leakage - exposing model overfitting to mislabeled or low-confidence examples - and conducted experiments with the SNLI corpus
  - Identified dataset outliers with unanimous aggregation of four BERT based model and extensive human annotation
  - Verified outliers with clustered Borda ranking, utilizing model prediction confidence as an additional safeguard for identification of mislabeled examples


## INTERNSHIP EXPERIENCE

---

- **RUIS Software Co., Ltd.**  Sept. 2022 - Nov. 2023  
*Software Engineer* Yinchuan, China
  - Extended the massive text processing capability of RUIS information system backend with BERT-based language models finetuned with proprietary datasets, achieving **91% accuracy** on document classification and rating, speeding up the pipeline by **over 20x**
  - Participated in the development and maintenance of RUIS internal management platform using Express.js
- **YUSUR Technology Co., Ltd.**  Jan. 2021 - July 2021  
*Data Science Researcher Intern* Beijing, China
  - Contributed to YUSUR's DPU research project, aimed to accelerate database operations using specially designed external computational devices
  - Benchmarked and tuned DPU-accelerated cluster computation on Apache Spark and XGBoost
  - Researched performance bottlenecks for Apache Spark (especially in shuffling) and formulated solutions for better performance on targeted conditions, achieving up to a **10× speedup** during proof-of-concept testing for clients

## SERVICE

---

- **Khoury College of Computer Sciences**  Jan. 2020 - June 2020  
*TA for Computer Science Fundamentals II (with Prof. Amit Shesh)* Boston, USA

## SKILLS

---

- **Programming Languages:** Python, C++, Java, SQL, JavaScript
- **ML Frameworks:** PyTorch, Hugging Face Transformers, TRL, PEFT/LoRA, DeepSpeed, Accelerate
- **LLM Inference & Serving:** vLLM, OpenAI SDK, Hugging Face Hub
- **Experiment Tools:** Weights & Biases, SLURM, Docker, Git
- **Data & Web:** NumPy, Pandas, Apache Spark, Node.js, Express.js, Angular