

# 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, Large Language Models, Benchmarking

## EDUCATION

- **University of Southern California** Jan. 2024 - Dec. 2025  
M.S. in Computer Science (Artificial Intelligence)  
◦ GPA: 3.91/4.00  
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

- **Intelligent Human Perception Lab, USC Institute for Creative Technologies** [🌐] June 2025 - Present  
Graduate Researcher  
Los Angeles, USA  
◦ Collaborated with [Prof. Mohammad Soleymani](#) on multimodal LLMs to benchmark and improve their affective understanding and multimodal representations  
◦ Designed an adversarial paralinguistic benchmark (**28k verified examples**) with systematically mismatched textual and acoustic cues, diagnosing multimodal hallucination in speech LLMs, and performed **linear probing** and **attention analysis** to examine the cause  
◦ Enhanced emotion reasoning with preference optimization (**13% ↑ UAR** on [DFEW](#) over [EmotionLLaMA](#)) and carried out **extensive baseline experiments** with multimodal LLMs ([AffectGPT](#), [VITA-1.5](#), and [VideoLLaMA 3](#), etc.) across various audiovisual datasets (**paper under review**)
- **Melady Lab, USC Viterbi School of Engineering** [🌐] June 2025 - Present  
Graduate Researcher  
Los Angeles, USA  
◦ Collaborated with [Prof. Yan Liu](#), and explored Conditional Listwise Policy Optimization (based on GRPO) for group decision-making vs. baselines, such as [GroupDebate](#), centralized arbiter (**3%↑ accuracy** on [MMLU](#) over [GroupDebate](#)) (**paper under review**)  
◦ Co-authored **two survey papers** on 1) time series reasoning with LLMs and 2) LLM based agentic systems with multi-agent reinforcement learning (**paper under review**)

## PUBLICATIONS

- **VoxParadox: Speech LLMs Hear the Words, Not the Voice.**  
[Jiacheng Pang](#), Ashutosh Chaubey, Mohammad Soleymani, preprint, submitting to ICML 2026.
- **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](#)].
- More under double-blind review at ICLR.

## PROJECTS

- **VoxParadox: Adversarial Paralinguistic Benchmark for Speech LLMs** June 2025 - Present  
Adviser: [Prof. Mohammad Soleymani](#) (Research Associate Professor, USC)  
◦ Implemented a content-style contradiction testbed to evaluate if speech LLMs hallucinate based on transcripts when paralinguistic cues carry the true label  
◦ Created a benchmark (**28k speech clips** and QA pairs) with **controlled text-to-speech** generation tools ([OpenAI TTS](#), [Azure](#), [ElevenLabs](#), and [CosyVoice 2](#)) for **14 paralinguistic tasks**, such as age, gender, pitch, speed  
◦ Found transcript overreliance in multiple major speech LLMs, and explored methods to diagnose and alleviate this issue, including **linear probing**, **projection of middle-layer** audio embeddings, and **attention scaling**
- **MSECap: Fusion Based Audio-textual Speech Emotion Captioning** Jan. 2025 - June 2025 [🌐]  
Adviser: [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%↑ GPT-evaluated 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  
*Adviser: 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  
*Adviser: 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
- **Artist Search: Full-Stack Web and Android Development** Jan. 2025 - May 2025  
*Adviser: Marco Papa (Senior Lecturer, USC)* [🌐]
  - Developed a **responsive web application** using Angular, Node.js/Express, Bootstrap and MongoDB, deployed via GCP App Engine, equipped with secure account management
  - Designed a native Android counterpart with Kotlin, adapting web functionalities to a **mobile-first UI**

## 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
  - Held weekly lab sessions and office hours to help **over 200 students** understand class materials and graded classwork and exams, averaging 10 - 20 hours of work time per week

## SKILLS

- **Programming Languages:** Python, Java, C++, SQL, JavaScript, HTML, CSS
- **AI/ML Tools & Frameworks:** PyTorch, Transformers, NumPy, Pandas
- **Engineering Tools & Frameworks:** Apache Spark, Angular, Node.js, Express.js, Bootstrap, MongoDB, GCP
- **Tools:** VSCode, Git, Anaconda
- **Tests:** GRE: 332 + 4, TOEFL: 115 (30, 29, 28, 28)