

How Sampling Shapes LLM Alignment: From One-Shot Optima to Iterative Dynamics

SPEAKER **Yurong Chen**

Postdoc
SIERRA-team, INRIA Paris

DATE 26 Feb, 2026 (Thu)

TIME 2:00 PM - 3:00 PM

VENUE G7315, 7th Floor Green Zone, Yeung Kin Man Academic Building, City University of Hong Kong, 83 Tat Chee Avenue, Kowloon Tong

ABSTRACT

Standard methods for aligning large language models with human preferences learn from pairwise comparisons among sampled candidate responses and regularize toward a reference policy. Despite their effectiveness, the effects of sampling and reference choices are poorly understood theoretically. We investigate these effects through Identity Preference Optimization, a widely used preference alignment framework and show that proper instance-dependent sampling can yield stronger ranking guarantees, while skewed on-policy sampling can induce excessive concentration under structured preferences. We then analyze iterative alignment dynamics in which the learned policy feeds back into future sampling and reference policies, reflecting a common practice of model-generated preference data. We prove that these dynamics can exhibit persistent oscillations or entropy collapse for certain parameter choices, and characterize regimes that guarantee stability. Our theoretical insights extend to Direct Preference Optimization, indicating the phenomena we captured are common to a broader class of preference-alignment methods. Experiments on real-world preference data validate our findings.

BIOGRAPHY

Yurong Chen is currently a postdoc at SIERRA-team, INRIA Paris, working with Michael I. Jordan, and was recently awarded a Marie Skłodowska-Curie Actions (MSCA) Postdoctoral Fellowship to be jointly hosted by Francis Bach and Michael I. Jordan. She earned her PhD in Computer Science at Peking University, where she was advised by Xiaotie Deng. Her research focuses on the intersection of learning, economics, and game theory, especially on how strategic agents exploit information advantage during interactions with learning agents. She is a recipient of the Best Student Paper Award at WINE 2022.

All are welcome!



In case of questions, please contact Prof. Li Minming at minming.li@cityu.edu.hk, or visit the CS Departmental Seminar Web at <https://www.cs.cityu.edu.hk/events/cs-seminars/recent-cs-colloquiums>.