

## STATEMENT

“My interest in people lead me to do so much with computers.” During past two years, I researched at Future Laboratory at **Tsinghua University**, in the Natural Language Computing Group at **Microsoft Research Asia**, and at **Moonshot.ai**. My research interest lies in **human-computer interaction (HCI)** with **multimodal agents** and **alignment** to enhance human performance in productivity, cognition, and well-being. I hold publications on CHI [7], UIST [3], IEEE VR [2], IROS [1], HRI [6], and TEI [5].

**Research in Agent.** It began when I collaborated with the **Institute for AI Industry Research (AIR)** at Tsinghua University on competitive driver agent simulation, which was published by **IROS [1]**. At about the same time, when a Ph.D. student in **Future Laboratory** was trying to produce VR contents with smell experience, I proposed to apply agents. Previous methods relied on heavy manual labeling by perfumers to match objects with scent sequences. We first narrowed down to a VR movie scenario, and I proposed a **multimodal agent system** for image-odor matching by combining LLM with the vision-language model (VLM) and smaller expert models, with metrics of main object, emotion, space, and time. Frames of movies were turned into descriptions by VLM for LLM, and LLM with perfumer knowledge through in-context learning could request expert image segmentation model and storage to output JSON, to call functions of communication module for odor display hardware. The ablation experiment and user study that I designed led to a full mark from a reviewer and the paper was accepted by **IEEE VR [2]** and won the **Amazon GenAI Hackathon top 10**.

**Research in LLM Reasoning.** My knowledge of agents led me to the **Natural Language Computing Group** at **Microsoft Research Asia (MSRA)**. There, I found that different from the experimental environment, LLM in real scenarios will face heterogeneous inference tasks from users (math, code, writing, etc.), which tests the adaptive ability of LLM to the diversity. I proposed a **meta-reasoning scheme for LLM** to dynamically choose suitable reasoning methods before conducting actions, which was inspired by the human thinking pathway. On comprehensive datasets (a randomly sampled set of seven orthogonal popular benchmarks), it achieved a 12% improvement in accuracy than Chain-of-Thoughts, and is 5% higher than the most powerful Tree-of-Thoughts, with significant advantages in less interactions and more efficiency. Though I absorbed the most cutting-edge knowledge from MSRA, there is still a gap between academic benchmarks and user needs. I desired to get in touch with the real world.

**Engineering & Research in LLM Alignment.** **Moonshot.ai (Kimi)** is a unicorn company of LLM with **36 million monthly active users** in China, maintaining an enormous real data pool. I was involved in LLM post-training in live phone call scenarios to equip models with "high emotional quotient" - a vague metric without benchmarks. We started by focusing on an urgent issue of "abrupt change of subject". I established **data pipelines** for a joint group of alignment, product and data team, from auto-analysis, evaluation, rolling and annotation, which improved efficiency. However, the post-trained model by DPO with 10 thousand refined data improved by 10% within goals but weakened other performance. This lesson taught me how to cope with the features of the model on the balance between generic and specialized issues. So we processed user data by thematic coding analysis and grounding 1M real QA pairs to build values metrics. I was working on the **reward model for RLHF**. However, human values are non-transitive in real situations, and methods to minimize collective dissatisfaction, such as applying the Nash Game, encounter bottlenecks. Therefore, we also **aligned for individuals** – long-term and short-term **memory systems**, using large and small model collaboration to maintain dynamically personalized user bios. It functions by a rigorous state-discriminating smaller model that routes memory actions to a fine-tuned memory expert model to generate message bodies. The outcome systems have **served tens of millions of users** every day.

**Research in HCI.** It began two years ago when my first paper was accepted by **ACM TEI**, a software tool to assist digital fabrication. At this year's **ACM UIST**, my paper, MulO, was presented, showcasing a development tool of hardware and software that enables semi-professionals to create task-oriented olfactory innovation. The PCB, which I independently designed, contains 144 components and is integrated into an AirPods volume. Inspired by Event-Condition-Action rules, for software, I also developed the visualization design tools for set odor triggers, outputs, devices, and so on. I led software development, database, and APIs for three platforms. Additionally, I served as a **reviewer for ACM CHI, IUI, DIS, and TEI**. When reviewing, I reflected on myself and learned from other experienced reviewers, which deepened my understanding of academic writing and research tastes. I explored a lot and finally devoted to agent and alignment for HCI.

**Education & Teaching Experience & Entrepreneurship.** I always try to go beyond the environment. During my undergraduate study at Tongji University, I ranked **1st out of 77** students, having achieved honors of outstanding graduates and outstanding bachelor's thesis. The skills in **design, object-oriented programming, embedded programming, electronic engineering, and mechanical engineering** equipped me with the technical skills and interdisciplinary knowledge to begin my HCI research. As a master's student at Tsinghua University, I honed my skills in machine learning, generative models, and signal processing, and I have **received research grants of totally 100,000 RMB from Tsinghua**. In addition, I enjoy mentoring younger students. During a 15-day HCI summer hackathon at Tsinghua University, I **supervised 7 group projects for 30 students**. At the Future Lab, I helped to **supervise three research interns** in six months, one went to CMU for an MS in entertainment technology, and another got the offer of a Master of Computing in NUS. I also supervised a student with bachelor's thesis, through which he got an opportunity for an internship at Microsoft Research. During this year, I **founded a startup and serve as a CEO** in end-cloud AI hardware. I received **2,000,000 RMB of investment** from Miracle Plus and **200,000 USD of Google GPU resources**. The AI education product we developed is on the eve of mass production and is expected to be sold in the middle of 2025. And we use GPU resources with Ph.D. of Stanford to train a forthcoming open-source video generation world model. Even though it is coming to an end, this valuable experience enable me to taste failure and realize the both the limit and the potential of personal power.

**Future Work.** My research experience at Microsoft and Tsinghua has equipped me with a solid research foundation, interdisciplinary perspective, and teamwork ability, while Moonshot.ai cultivated my engineering skills and user awareness. I hope to continue working on **multimodal agents and alignment for individuals**. Specifically, I would like to focus on applying post-training, large-small model collaboration, and reinforcement learning to enable user intention understanding and personalization, thereby enhancing human performance through OS-level Agentic AI on wearables or software canvas. My dream is to witness AI understanding us humans better than we understand ourselves and help individuals achieve self-fulfillment.

## References

- [1] Jin Y, Shen X, Peng H, et al. Surrealdriver: Designing generative driver agent simulation framework in urban contexts based on large language model[J]. arXiv preprint arXiv:2309.13193, 2023.
- [2] Zhang Y, Gao P, Kang F, et al. OdorAgent: Generate Odor Sequences for Movies Based on Large Language Model[C]//2024 IEEE Conference Virtual Reality and 3D User Interfaces (VR). IEEE, 2024: 105-114.
- [3] Gao P, Liu F, Wen D, et al. Mul-O: Encouraging Olfactory Innovation in Various Scenarios Through a Task-Oriented Development Platform[C]//Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology. 2024: 1-17.
- [4] Gao P, Xie A, Mao S, et al. Meta Reasoning for Large Language Models[J]. arXiv preprint

arXiv:2406.11698, 2024.

[5] Gao P, Gao T, Yang Y, et al. Bamboo Agents: Exploring the Potentiality of Digital Craft by Decoding and Recoding Process[C]//Proceedings of the Seventeenth International Conference on Tangible, Embedded, and Embodied Interaction. 2023: 1-13.

[6] Wang X, Cui X, Xiong R, et al. Gobot: A Novel Shoe-Integrated Robot for Enriching Walking Experiences[C]//Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction. 2024: 1265-1268.

[7] Lu Q, Zhang Y, Zhang Y, et al. Atmospheror: Towards an olfactory interactive system for enhancing social presence and interaction in synchronous online classes[C]//Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems. 2023: 1-8.