

# Shaomu Tan

✉ shaomu.tan.edu@gmail.com 🏠 HomePage 🌐 shaomutan 🎓 G-Scholar

## Professional Summary

---

Shaomu is a researcher specializing in machine translation, multilingual LLMs, reasoning systems, and reward modeling. He has extensive experience in post-training 7B–72B LLMs using distributed systems, reward modeling from human preferences, agentic evaluate-and-improve pipelines. He also published 7 papers at ACL, EMNLP, and NeurIPS, and earned Kaggle gold/silver medals.

- Expected to complete Ph.D. in 2026 and currently seeking Research Scientist roles in Large Language Models training, pre- and post-training, evaluation and alignment.
- Research expertise and interests: Multilingual LLMs, LLM Reasoning, Evaluation and reward modeling, Reinforcement Learning from Human Feedback, Vision-Language LLMs, World Models.

## Education

---

**University of Amsterdam, Netherlands** 2022 – 2026  
*Ph.D. in Computer Science. Focus on LLMs, Machine Translation, supervised by [Christof Monz](#).*

**Utrecht University, Netherlands** 2020 – 2022  
*Master of Science in Artificial Intelligence. GPA: 4.0/4.0*

**Shandong University, China** 2016 – 2020  
*Bachelor of Science in Information System.*

## Professional Experience

---

**Applied Scientist Research Intern, Amazon AGI, Berlin, Germany** Oct 2025 – Jan 2026

- Conducted research on test-time scaling for long-context (book-level) machine translation, bridging research and production LLM systems ([Amazon Nova](#)).
- Designed and led a large-scale human evaluation pipeline (annotated 1M words, costs \$90K USD) to assess long-context generation quality. Research output has been accepted at ACL 2026.

**Research Intern, Sony (Speech & Language AI Lab), Tokyo, Japan** Jun 2025 – Sep 2025

- Developed reasoning-based LLM evaluation models for translation quality estimation and refinement.
- Applied reinforcement learning (RLVR) to improve evaluation reliability and alignment with human preferences, based on 7-32B LLMs.
- Built a multi-role reasoning agentic LLM framework for evaluation + refinement pipelines. Research output has been accepted at ACL 2026.

## Competitions

---

**Kaggle Gold Medal (Ranking: 11/1,315 teams)** — in Google’s 2023 American Sign Language Fingerspelling Recognition. I design and implement Conformer-Transformer architecture, use sign language landmarks for training, and fuse CTC with cross-entropy for learning objectives, [\[link\]](#).

**Kaggle Silver Medal (Ranking: 19/950 teams)** — in 2024 WSDM Cup: Multilingual Chatbot Arena Human Preference Reward Modeling. I train 7-14B reward models for multilingual human preferences, optimized with DeepSpeed, VLLM, quantization, pseudo labeling, and model merging, [\[link\]](#).

**WMT23 1st place** — in General Machine Translation Competition: English-Hebrew. We train small Encoder-Decoder system (0.5B) that performs on-par with GPT-4-5shot, [\[link\]](#).

## Selected Publications

---

**Shaomu Tan**, Dawei Zhu, Ke Tran, Michael Denkowski, Sony Trenous, Bill Byrne, Leonardo Ribeiro, Felix Hieber. "What Does LLM Refinement Actually Improve? A Systematic Study on Document-Level Literary Translation", **ACL 2026**.

**Shaomu Tan**, Ryosuke Mitani, Ritvik Choudhary, Qiyu Wu, Toshiyuki Sekiya, Christof Monz. "Remedy-R: Generative Reasoning for Machine Translation Evaluation without Error Annotations", **ACL 2026**, [\[link\]](#).

**Shaomu Tan** and Christof Monz. "Remedy: Learning Machine Translation Evaluation from Human Preferences with Reward Modeling", **EMNLP 2025**, [\[link\]](#).

**Shaomu Tan**, Di Wu, and Christof Monz. "Neuron Specialization: Leveraging intrinsic task modularity for multilingual machine translation.", **EMNLP 2024**, [\[link\]](#).

Di Wu, **Shaomu Tan**, and Christof Monz. "How Far Can 100 Samples Go? Unlocking Overall Zero-Shot Multilingual Translation via Tiny Multi-Parallel Data.", **ACL 2024**, [\[link\]](#).

**Shaomu Tan**, and Christof Monz. "Towards a Better Understanding of Variations in Zero-Shot Neural Machine Translation Performance.", **EMNLP 2023**, [\[link\]](#).

Baohao Liao, **Shaomu Tan**, and Christof Monz. "Make Pre-trained Model Reversible: From Parameter to Memory Efficient Fine-Tuning.", **NeurIPS 2023**, [\[link\]](#).

## Service & Teaching

---

### Invited Talks

- "Toward Explainable, Robust, and Actionable Translation Quality Estimation" at **EACL 2026**. [MME workshop](#), March 2026.
- "Remedy-R: Large Reasoning Models for Machine Translation Evaluation" at **University of Tokyo**, invited by Prof. [Yoshimasa Tsuruoka](#), August 2025.
- "The Second Half of Machine Translation" at Nara Institute of Science and Technology (**NAIST**), invited by Prof. [Taro Watanabe](#), July 2025.

### Lectures & Education

- "[Interference and Knowledge Transfer in Multilingual Translation Models](#)" at **University of Amsterdam** for AI Master course "Deep Learning for NLP", co-presented with Prof. Christof Monz, October 2024.
- "[A Journey on Multilingual Neural Machine Translation](#)" at **Utrecht University** for AI Bachelor course "Models for Language Processing", co-presented with Prof. Denis Paperno, June 2024.

### Reviewing

- **2023-2026**: ACL Rolling Review (ACL, EMNLP, EACL, NAACL), Transactions on Audio, Speech, and Language Processing (TASLP), WMT, Multilingual Representation Learning (MRL).

## Technical Skills

---

- **Programming related**: Python, Bash and Linux, Git, Docker, Slurm, AWS.
- **Framework**: PyTorch, HF (Transformers, TRL, PEFT), SentencePiece, Fairseq.
- **LLMs**: Verl, OpenRLHF, vLLMs, Megatron-LM, TRL, Sentencepiece, Llama-Recipes, NeMo.
- **Scaling**: DeepSpeed, FSDP, Flash-Attention, Megatron-LM.
- **Large-Scale Experience**: Distributed training and inference for up to 72B LLMs.
- **Languages**: English (full professional proficiency). Chinese (native speaker).