Yi-Chern Tan

yichern.tan[at]aya.yale.edu | linkedin.com/in/yicherntan | github.com/yichern | google scholar | website: yich.io

Experience

Cohere — **London, UK**

Post-Training Co-Lead, Generative Modeling — **June '24 – present**

- [Command A post-training lead] Led post-training across ~60 research scientists and engineers, designed and executed a merging-based post-training recipe, high-signal evaluation suites, new agent reasoning prompt template, alignment of agentic post-training across domains with sharing of scientific insights, collaboration with quantization and serving, and leadership reporting, which culminated in open-weights **Command A** (1313 LMArena elo, 13th at launch, ~DeepSeek v3) and **Command R7B** (technical report)
- [Technical lead, post-training recipes] Scientific co-advisor for developing post-training pipelines integrating vision, text, agentic and reasoning capabilities, with a focus on RLVR in multi-task multi-env settings and interaction effects between mid-training and RL

Senior Member of Technical Staff, Generative Modeling — **Sep '22 – June '24**

- [Core post-training research engineer] Led the implementation of research infrastructure such as **internal offline preference and RLHF algorithms** (in JAX, equivalent to SLiC-HF), prompt templating, post-training tokenizers, experiment orchestration framework, packing and masking of chat data, team processes for model training, unique identifiers for post-training data, data filtering
- Design and ran thousands of experiments for the entire post-training pipeline at large scale, released 2 models open-weights as **Command R** (35B, 1146 LMArena elo, 10th at launch) and **Command R+** (103B, 1194 LMArena elo, 6th at launch, ~Claude 3 Sonnet)
- [Technical lead, Cohere's internal annotation platform] Led a team of 8 across engineering, design and product to build the platform 0 to 1, reducing spend on external annotation vendors by ~75% and improving post-training human annotation quality via data engineering, in-house controls, data review, and bespoke annotation pipelines

Waymo — **Mountain View, CA**

Machine Learning Engineer, Simulation and Safety Evaluation — **Mar '22 – Sep '22**

- Improved automated data refresh, training and evaluation infrastructure (C++) for a classifier that scores the self-driving car's likelihood of a realistic collision with other road users in simulation, decreasing experiment iteration time from ~5 days to ~2 days
- Implemented feature engineering and architectural changes (mixture of experts, TensorFlow) to the collision likelihood model, improving P@R95 by ~4% overall and ~22% for vulnerable road users, reducing human-in-the-loop triage load by ~10%.

Facebook — **Menlo Park, CA**

Software Engineering Intern, Instagram Safety AI — **May '19 – Aug '19**

- Developed high-performant text classifiers for detecting violating comments (toxicity, bullying) with a new contextual encoder, distributed learning, and knowledge distillation; increased calibrated recall by ~70% at ~40% lower latency across 9 languages
- Implemented MLOps tooling for model training at scale (automated training pipelines, ~billions of comments), evaluation (score visualization) and deployment (publishing workflows) of classifiers; increased model iteration speed and robustness

Representative Research and Peer-Reviewed Publications

- **'23 – '25 (Lead):** Command A (technical report), Command R, Command R+
- **'23 – '25 (Contributor):** Aya 23 (report), Aya Expanse (report)
- **NeurIPS '19 (1st author, spotlight):** Assessing Social and Intersectional Biases in Contextualized Word Representations
- **ACL '19 (co 1st author, BlackBoxNLP workshop):** Open Sesame: Getting Inside BERT's Linguistic Knowledge
- **NeurIPS '25 (2nd author, spotlight):** Reverse Engineering Human Preferences with Reinforcement Learning
- **EMNLP '25:** No Need for Explanations: LLMs Can Implicitly Learn From Mistakes In-Context
- **EMNLP '19:** CoSQL: A Conversational Text-to-SQL Challenge Towards Cross-Domain Natural Language Interfaces to Databases
- **ACL '19:** SParC: Cross-Domain Semantic Parsing in Context
- **Reviewing:** EMNLP '21-'22, NeurIPS '21 (Outstanding Reviewer Award)-'25, ICLR '22-'25, ICML '22-24, COLM '24-'25

Selected Projects

- **RL for tool use by Baxter robot:** Built a sim2real learning pipeline with ROS and MuJoCo for a Baxter robot to learn time extended series of actions to use a physical tool using inverse kinematics solver with various policy gradient methods (A2C, TRPO, PPO)
- **needle library**: Developed a deep learning library using numpy, own NDArray implementation and CUDA kernels (DLSys course)
- **UNIX Shell in mCertiKOS:** Implemented a UNIX-like shell in the mCertiKOS operating system (C, QEMU)
- **Conversational personas**: Built and open-sourced conversant, a framework for customizable dialogue agents with system prompts

Education

Yale University — **New Haven, CT**

Bachelor of Science (CGPA: 3.98), Summa Cum Laude, Phi Beta Kappa (top 5% of cohort), College Valedictorian — **Aug '16 – May '20**

Double Major in Computer Science (GPA: 4.00) and Ethics, Politics, Economics (GPA: 3.97)

- **Selected Coursework:** NLP (PhD level), ML, Game Theory, Linear Algebra, Probability Theory, Statistical Learning Theory, Operating Systems, Classical AI, Robotics, Data Science Ethics, Normative Ethics, Distributive Justice, Cybersecurity Policy