

CHENMIEN TAN

Phone: +44 7536 478885 ◇ Email: chenmien.tan@ed.ac.uk

[Homepage](#) | [Google Scholar](#) | [GitHub](#) | [Kaggle](#)

EDUCATION

- **University of Edinburgh** 2023/09–present
M.Res. in Language Processing, Supervisor: [Shay Cohen](#)
- **University of Nottingham Ningbo China** 2018/09–2022/06
B.Sc. in Applied Mathematics
 - Selected modules: Vector Calculus (95), Linear Mathematics (94), Probability (97), Statistical Models (94), and Discrete Mathematics and Graph Theory (95).

INTERNSHIPS

- **Hong Kong University of Science and Technology** 2023/04–2023/08
Research Assistant, Supervisor: [Jie Fu](#)
 - Proposed LM editing hyper-network with scalability in synchronous editing operation amount, where the parameter shift aggregation is formulated as a least square problem to mitigate cancellation effect and the computation between the hyper-network and LM is delineated to accommodate multiple editing with limited memory budgets.
- **Shanghai Jiao Tong University** 2022/09–2022/12
Research Assistant, Supervisor: [Paul Weng](#)
 - Investigated on the tail risk of stochastic and adversarial multi-armed bandits, proposed (nearly) optimal algorithms with theoretical guarantees, and elucidated the characteristics of implicit exploration for achieving trade-off between the expected regret and the variability of regret.
 - Implemented online and offline deep reinforcement learning algorithms, *e.g.*, PPO and SAC, in dynamic control and game environments, *i.e.*, MuJoCo and Atari.

PUBLICATIONS

- [Massive Editing for Large Language Models via Meta Learning](#)
Chenmien Tan, Ge Zhang, and Jie Fu
In *ICLR'24*
- [Learning Rewards to Optimize Global Performance Metrics in Deep Reinforcement Learning](#)
Junqi Qian, Paul Weng, and **Chenmien Tan**
In *AAMAS'23*
- [CVaR-Regret Bounds for Multi-armed Bandits](#)
Chenmien Tan and Paul Weng
In *ACML'22*

COMPETITIONS

- Learning Equality – Curriculum Recommendations 2023/01–2023/03
Ranking: 17/1057 = 1.6%
 - Task: retrieve educational contexts for curriculum topics based on textual information.

- Retrieval: encoded topics and contexts separately and searched for the nearest neighbours for each topic, where Sentence Transformers were fine-tuned through contrastive learning using multiple negative ranking loss and hard negatives were retrieved from false positives.
 - Ranking: fine-tuned cross-encoders that classify the congruency of topic-context pairs with adversarial learning using the fast gradient method.
 - Re-ranking: aggregated the retrieval and ranking scores of candidates using gradient boosting decision tree, which is superior to merely applying threshold on ranker.
- [Google AI4Code – Understand Code in Python Notebooks](#) 2022/06–2022/08
Ranking: 25/1135 = 2.2%
 - Task: given the order of code cells, insert Markdown cells into Jupyter notebooks.
 - Modeling: pair-wise (a Markdown-code pair is fed into Transformer to predict whether the Markdown cell is adjacent to the code cell) and point-wise (a Markdown cell is fed with code cells in the same notebook to predict which code cell is adjacent to the Markdown cell).
 - Inference: to leverage the positional information of code cells, the Markdown cell is allocated to a position based on weighted probability, which is superior to merely assigning the Markdown cell to the code cell with the highest probability.
 - Ensemble: designed an entropy based confidence metric for adaptive ensemble.
 - [H&M Personalized Fashion Recommendations](#) 2022/03–2022/05
Ranking: 45/2952 = 1.5%
 - Task: predict customers' future purchase conditioned on the transaction history.
 - Data processing: retrieved positive samples using sliding windows, generated embedding for customers and products using TF-IDF and Word2Vec, embedded product images using pre-trained ResNet, and constructed and crossed features manually.
 - Retrieval: hand-crafted rules (*e.g.*, repurchase, popularity), collaborative filtering, matrix factorization, and two tower models (to aggregate representations from TF-IDF, Word2Vec, ResNet, and matrix factorization).
 - Ranking: gradient boosting decision tree and neural network.

SERVICES

- Reviewer: ICML'24
- TA: [NLU+](#) 23-24, [MLP](#) 23-24