# YANG CHEN

yang.chen.csphd@gmail.com | 😵 www.yangchen.info

#### EDUCATION

· Ph.D. in Computer Science	
University of Auckland, Auckland, New Zealand	November 2018 – September 2022
· First Class Honours in Computer Science	
University of Auckland, Auckland, New Zealand	$July \ 2017 - July \ 2018$
· Bsc in Computer Science & Technology	
Beijing Institute of Technology, Beijing, China	August 2013 – June 2017
EXPERIENCE	
Appointment	
· University of New South Wales	Sydney, Australia
Senior Research Associate Advisor: Flora Salim	October 2024 – Present
· Shanghai Artificial Intelligence Lab	Shanghai, China (Remote)
External Researcher	January 2025 – Present
· University of Auckland	Auckland, New Zealand
Research Fellow Advisor: Michael Witbrock	June 2021 – September 2024
Intern	

· DAMO Academy, Alibaba Research Intern Beijing, China September 2020 – January 2021

# **RESEARCH INTERESTS**

My research interests span a spectrum of reinforcement learning, large language models, multiagent systems and game theory. I have focused on addressing challenges modeled as multi-agent systems through reinforcement learning and game-theoretic approaches, with the ultimate goal of advancing theoretical and application insights into their intersection. Lately, my research has shifted toward leveraging reinforcement learning to improve the performance of large language models (LLMs) and use LLMs for AI-driven scientific discovery, with applications in embodied AI, ethical AI, and responsible AI.

# TEACHING

My lecture videos and original teaching materials (lecture and tutorial slides, assignment and exam questions with answers, codes and computer programmes) are archived and publicly available at https://www.yangchen.info/#teaching.

# **Postgraduate Courses**

· COMPSCI 713: Artificial Intelligence Fundamentals Instructor, University of Auckland

•	COMPSCI 761: Advanced Topics in Artificial Intelligence	
	Instructor, University of Auckland	Semester 2, $2022 - 2023$
•	COMPSCI 769: Natural Language Processing	
	Guest Lecturer, University of Auckland	Semester 2, 2024
	Undergraduate Courses	
•	COMPSCI 367: Artificial Intelligence	
	Guest Lecturer, University of Auckland	Semester 2, 2024
•	COMPSCI 220: Algorithms and Data Structures	
	Guest Lecturer, University of Auckland	Semester 1, 2022

# SUPERVISION

One PhD Student (mentoring); Four master students; One BSc (Honours) student (graduated).

# GRANTS

•	AI-based behavioural analytics for live sports broadcast	2024 - 2026
	Associate Investigator	
	Funding Agency: Ministry of Business, Innovation & Employment (MBIE) Endea	wour Fund
	Amount Awarded: 1 Million NZD	

# EVENT ORGANISING

- $\cdot$  Local co-chair of AAMAS 2024 (a core A\* conference).
- $\cdot\,$  Co-chair of AAMAS 2024 AAAI track.

# ACADEMIC SERVICES

- $\cdot$  Conference Reviewer: ICML'25, IJCAI'25, ICLR'25, ACL'24–25, EMNLP'24, ECAI'24, AAMAS'23–25.
- $\cdot\,$  Journal Reviewer: JMLR since 2024.

# AWARDS

· AAMAS 2022 Scholarship	April 2022
$\cdot$ Google Global PhD Fellowship Nomination (Australia & New Zealand)	August 2020
· Best Paper Award, BSCI 2019	July 2019
$\cdot$ Summer Scholarship Funding from PDH Research Partnership	November 2018
· University of Auckland Doctoral Scholarship	October 2018

# SELECTED PUBLICATIONS

#### Preprints

- Everyone You Meet Knows Something You Don't: Efficiently Scaling Test-Time Compute through Multi-LLM Repeated Sampling Jianhao Chen, Zishuo Xun, Bocheng Zhou, Han Qi, Qiaosheng Zhang, <u>Yang Chen</u>, Wei Hu, Yuzhong Qu, Wanli Ouyang, Shuyue Hu. *Preprint. 2025*
- Proximal Inverse Reward Optimisation
- Yang Chen, Jiaqi Zhang, Libo Zhang, Ping Cheng, Menglin Zou, Yitan Zhang. Preprint. 2025
- $\cdot\,$  Adversarial Inverse Reward-Constraint Learning with Reward-Contrast Prior

Yang Chen. Yitan Zhang, Gaël Gendron, Mitchell Rogers, David Arturo Soriano Valdez, Mihailo Azhar, Shahrokh Heidari, Padriac Amato Tahua O'Leary, Kobe Knowles, Jiamou Liu, Patrice Jean Delmas, Michael Witbrock.

Preprint. 2025

 Physics-Informed Utility Learning for Mean-Field Games <u>Yang Chen</u>, Junyi Yang. <u>Preprint.</u> 2025

#### **Reinforcement Learning & Imitation Learning**

- Multi-Agent, Human-Agent and Beyond: A Survey on Cooperation in Social Dilemmas. Hao Guo, Chunjiang Mu, <u>Yang Chen</u>, Chen Shen, Shuyue Hu, Zhen Wang. *Neurocomputing Journal.* 2024.
- $\cdot$  Meta-Inverse Reinforcement Learning for Mean Field Games with Probabilistic Context Variables.

Yang Chen, Xiao Lin, Bo Yan, Libo Zhang, Jiamou Liu, Neset Özkan Tan, Michael Witbrock. AAAI 2024. Thirty-Eighth AAAI Conference on Artificial Intelligence. (Core A\*)

- Adversarial Inverse Reinforcement Learning for Mean Field Games
  Yang Chen, Libo Zhang, Zhenyun Deng, Neset Özkan Tan, Jiamou Liu, Michael Witbrock.
  AAMAS 2023. The 22nd International Conference on Autonomous Agents and Multi-agent Systems. (Core A\*)
- Density-based Correlated Equilibrium for Markov Games
  Libo Zhang, Yang Chen (co-first & contact), Toru Takisaka, Bakh Khoussainov, Michael Witbrock, Jiamou Liu.

AAMAS 2023. The 22nd International Conference on Autonomous Agents and Multi-agent Systems. (Core  $A^*$ )

· Individual-Level Inverse Reinforcement Learning for Mean Field Games Yang Chen, Libo Zhang, Jiamou Liu, Shuyue Hu.

AAMAS 2022. The 21st International Conference on Autonomous Agents and Multi-agent Systems. (Core  $A^*$ )

- Interconnected Neural Linear Contextual Bandits with Upper Confidence Bound Exploration <u>Yang Chen</u>, Miao Xie, Jiamou Liu, Kaiqi Zhao. *PAKDD 2022. 26th Pacific-Asia Conference on Knowledge Discovery and Data Mining.*
- $\cdot$  Social Structure Emergence: A Multi-agent Reinforcement Learning Framework for Relationship Building

Yang Chen, Jiamou Liu, He Zhao, Hongyi Su.

AAMAS 2020. The 19th International Conference on Autonomous Agents and Multi-agent Systems. (Core  $A^*$ )

 Social Capital Games as A Framework for Social Structural Pattern Emergence Yang Chen, Jiamou Liu. ASONAM 2020. IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining.

### Multi-agent Behaviour Modelling and Simulation

- $\cdot$ Behaviour Modelling of Social Animals via Causal Structure Discovery and Graph Neural Networks
- Gaël Gendron (co-first), Yang Chen (co-first), Mitchell Rogers, Yiping Liu, Mihailo Azhar, Shahrokh Heidari, David Arturo Soriano Valdez, Kobe Knowles, Padriac O'Leary, Simon Eyre, Michael Witbrock, Gillian Dobbie, Jiamou Liu and Patrice Delmas.
- AAMAS 2024. The 23rd International Conference on Autonomous Agents and Multi-agent Systems. (Core A\*)
- $\cdot$  Meerkat Behaviour Recognition Dataset
- Mitchell Rogers, Gaël Gendron, David Soriano Valdez, Mihailo Azhar, <u>Yang Chen</u>, Shahrokh Heidari, Caleb Perelini, Padriac O'leary, Kobe Knowles, Izak Tait, Simon Eyre, Michael Witbrock, Patrice Delmas.
- 3rd Workshop on CV4Animals in conjunction with CVPR 2023): Computer Vision for Animal Behavior Tracking and Modeling.
- MSDC: Non-intrusive Load Monitoring with a Dual-CNN Model Jialing He, Jiamou Liu, Zijian Zhang, <u>Yang Chen</u>, Yiwei Liu, Bakh Khoussainov, Liehuang Zhu. AAAI 2023. Thirty-Seventh AAAI Conference on Artificial Intelligence. (Core A\*)

#### Natural Language Processing and Reasoning

- $\cdot$  Assessing and Enhancing the Robustness of Large Language Models with Task Structure Variations for Logical Reasoning
- Qiming Bao, Gaël Gendron, Alex Peng, Wanjun Zhong, Neset Tan, <u>Yang Chen</u>, Michael Witbrock, Jiamou Liu.
- ICONIP 2024. The 31st International Conference on Neural Information Processing.
- · Abstract Meaning Representation-Based Logic-Driven Data Augmentation for Logical Reasoning
- Qiming Bao, Alex Yuxuan Peng, Zhenyun Deng, Wanjun Zhong, Gael Gendron, Timothy Pistotti, Neset Tan, Nathan Young, <u>Yang Chen</u>, Yonghua Zhu, Paul Denny, Michael Witbrock, Jiamou Liu.
- ACL 2024 Findings. The 62nd Annual Meeting of the Association for Computational Linguistics. (Core  $A^*$ )
- $\cdot$  Neuromodulation Gated Transformer
- Kobe Knowles, Joshua Bensemann, Diana Benavides Prado, Vithya Yogarajan, Michael Witbrock, Gillian Dobbie, Yang Chen.
- ICLR 2023 Tiny Papers. The Eleventh International Conference on Learning Representations.

· Multi2Claim: Generating Scientific Claims from Multi-Choice Questions for Scientific Fact-Checking

Neset Tan, Trung Nguyen, Josh Bensemann, Alex Peng, Qiming Bao, <u>Yang Chen</u>, Mark Gahegan, Michael Witbrock. *EACL 2023. The 17th Conference of the European Chapter of the Association for Computational Linguistics.* 

- Interpretable AMR-Based Question Decomposition for Multi-hop Question Answering.
  Zhenyun Deng, Yonghua Zhu, Yang Chen, Michael Witbrock, Patricia Riddle.
  IJCAI 2022. The 31st International Joint Conference on Artificial Intelligence. (Core A\*)
- Prompt-based Conservation Learning for Multi-hop Question Answering Zhenyun Deng, Yonghua Zhu, <u>Yang Chen</u>, Qianqian Qi, Michael Witbrock, Patricia Riddle. COLING 2022. The 29th International Conference on Computational Linguistics.
- $\cdot$  An explainability analysis of a sentiment prediction task using a transformer-based attention filter

Neset Özkan Tan, Joshua Bensemann, Diana Benavides-Prado, <u>Yang Chen</u>, Mark Gahegan, Lia Lee, Alex Yuxuan Peng, Patricia Riddle, Michael Witbrock. *The Ninth Annual Conference on Advances in Cognitive Systems*.

# Graph Theory & Graph Neural Networks

- Graph Transformer against Graph Perturbation by Flexible-pass Filter Yonghua Zhu, Jincheng Huang, <u>Yang Chen</u>, Robert Amor, Michael Witbrock. Journal of Information Fusion. 2024.
- Robust Node Classification on Graph Data with Graph and Label Noise Yonghua Zhu, Lei Feng, Zhenyun Deng, Yang Chen, Robert Amor, Michael Witbrock. AAAI 2024. Thirty-Eighth AAAI Conference on Artificial Intelligence. (Core A\*)
- Efficient Size-Prescribed k-Core Search Yiping Liu, Bo Yan, Bo Zhao, Hongyi Su, Yang Chen, Michael Witbrock. ASONAM 2023. The 2023 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining.
- Chain of Propagation Prompting for Node Classification
  Yonghua Zhu, Zhenyun Deng, <u>Yang Chen</u>, Robert Amor, Michael Witbrock.
  ACM MultiMedia 2023. (Core A\*)
- A Reinforcement Learning Approach to Gaining Social Capital with Partial Observation He Zhao, Hongyi Su, <u>Yang Chen</u>, Jiamou Liu, Hong Zheng, Bo Yan.
   PRICAI 2019. The 16th Pacific Rim International Conference on Artificial Intelligence.
- Dynamic Relationship Building: Exploitation Versus Exploration on a Social Network Bo Yan, Yang Chen, Jiamou Liu. WISE 2017. The 18th International Conference on Web Information Systems Engineering.

- Can Reinforcement Learning Enhance Social Capital?
  He Zhao, Hongyi Su, <u>Yang Chen</u>, Jiamou Liu, Bo Yan, Hong Zheng.
  The International Workshop on Web Information Systems in the Era of AI. 2019.
- Distributed Community Detection over Blockchain Networks Based on Structural Entropy Yang Chen, Jiamou Liu. The 2019 ACM International Symposium on Blockchain and Secure Critical Infrastructure. (Best Paper Award)
- · Becoming Gatekeepers Together with Allies: Collaborative Brokerage over Social Networks Yang Chen, Jiamou Liu.

ASONAM 2019. The 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining.