

Zhongtian (Falcon) Dai

Website: falcond.ai Email: dai@ttic.edu Address: Berkeley, CA.

Education

Toyota Technological Institute at Chicago, Chicago, IL

Ph.D. in Computer Science (defended in September 2022). September 2015 - September 2023.

- > Advised by Professor [Matthew R. Walter](#).
- > Thesis committee: Matthew R. Walter, [David McAllester](#), [Avrim Blum](#).
- > Thesis title: [On Reward Structures in Markov Decision Processes](#).
- > Select courses: learning theory, natural language processing, computer vision, dynamical systems.

The University of Chicago, Chicago, IL

B.S. with Honors in Mathematics and B.A. with Honors in Physics. September 2008 - June 2012.

Cumulative GPA: **3.76/4.00**.

- > Student Marshal of Class 2012 (top University distinction).
- > James Franck Institute Summer Undergraduate Research Fellowship, 2011.
- > Advanced courses: mathematical logic, graduate quantum mechanics, graduate general relativity.

Publications

- > Yunis D, Jung J, -, Walter MR. [Subwords as Skills: Tokenization for Sparse-Reward RL](#). Neural Information Processing Systems (NeurIPS), 2024.
- > -, Walter MR. [Loop Estimator for Discounted Values in Markov Reward Processes](#). Association for the Advancement of Artificial Intelligence conference (AAAI), 2021.
- > Zheng W, Minama Reddy GK, -, Chandramani A, Brang D, Hunter S, Kohrman MH, Rose S, Rossi M, Tao J, Wu S, Byrne R, Frim DM, Warnke P, Towle VL. [Chasing Language Through the Brain: Successive Parallel Networks](#). Clinical Neurophysiology, 2020.
- > -, Walter MR. [Maximum Expected Hitting Cost of a Markov Decision Process and Informativeness of Rewards](#). Neural Information Processing Systems (NeurIPS), 2019.
- > -, Cai Z. [Towards Near-imperceptible Steganographic Text](#). Association for Computational Linguistics (ACL), 2019. [[oral presentation](#), [nominated for best paper awards](#)]
- > Gehrmann S, -, Elder H, Rush AM. [End-to-End Content and Plan Selection for Natural Language Generation](#). International Conference on Natural Language Generation (INLG), 2018.
- > -*, Cai Z*. [Glyph-aware Embedding of Chinese Characters](#). Subword and Character level models in NLP workshop at Empirical Methods in Natural Language Processing conference (EMNLP), 2017.
- > Towle VL, -, Zheng W, Issa N. "[Mapping Cortical Function with Event-Related Electrocorticography](#)," in *Functional Mapping of the Cerebral Cortex*, ed. Richard W. Byrne. (Springer, 2016), 91-104.
- > Brang D, -, Zhang W, Towle VL. [Registering Imaged ECoG Electrodes to Human Cortex: A Geometry-based Technique](#). Journal of Neuroscience Methods, 64-73. 2016.
- > Brang D, Towle VL, Suzuki S, Hillyard SA, Di Tusa S, -, Wu S, Tao J, Grabowecky M. [Peripheral sounds rapidly activate visual cortex: evidence from electrocorticography](#). Journal of Neurophysiology, 3023-3028.

2015.

> Towle VL, Minama Reddy GK, -, Zhang W, Brang D, Hunter S, Kohrman MH, Marcucilli CJ, Tao J, Rossi MA, Frim DM, Byrne RW. Chasing Language Through the Brain: Three Successive Parallel Networks. Society for the Neurobiology of Language Conference, 2014.

Presentations

- > Yunis D, Jung J, -, Walter MR. [Subwords as Skills: Tokenization for Sparse-Reward RL](#). Generalization in Planning workshop at NeurIPS, 2023.
- > -, Walter MR. Loop Estimator for Discounted Values in Markov Reward Processes. Poster session at Algorithmic Learning Theory (ALT), 2020.
- > -, Walter MR. Maximum Expected Hitting Cost of a Markov Decision Process and Informativeness of Rewards. Poster session at Algorithmic Learning Theory (ALT), 2020.
- > -, Walter MR. Finite Time Analysis of Potential-based Reward Shaping. Reinforcement Learning and Decision Making (RLDM), 2019. **[Student travel fellowship]**
- > Vasiljevic I, Kolkin N, Luo R, Wang H, -, Daniele AF, Mostajabi M, Basart S, Walter MR, Shakhnarovich G. [DIODE: A Dense Indoor and Outdoor DEpth Dataset](#). 3D Scene Understanding for Vision, Graphics, and Robotics workshop at Computer Vision and Pattern Recognition (CVPR), 2019.
- > -, Walter MR. Finite Time Analysis of Potential-based Reward Shaping. Midwest Machine Learning Symposium (MMLS), 2019.
- > -, Cai Z. Towards Near-imperceptible Steganographic Text. Midwest Machine Learning Symposium (MMLS), 2019.
- > -, Cai Z. Towards Near-imperceptible Steganographic Text. Midwest Speech and Language Days (MSLD), 2019.
- > - and others at RIPL @ TTIC. Rubik's cube solving robot. National robotics week special exhibit at the Museum of Science and Industry, 2019.
- > -, Walter MR. Reward-adjusted Diameters and Their Conditioning by Potential-based Reward Shaping. Learning by Instruction workshop at Neural Information Processing Systems (NeurIPS), 2018.
- > - and others at RIPL @ TTIC. Checkers-playing robot. National robotics week special exhibit at the Museum of Science and Industry, 2018.
- > Schaff C*, -*, Walter MR. [Towards Active Imitation Learning](#). Learning from Demonstrations in High-Dimensional Feature Spaces workshop at Robotics: Science and Systems conference (RSS), 2017. **[Student travel grant award]**
- > -, Walter MR. Notepad-Augmented Environments in Reinforcement Learning. Midwest Machine Learning Symposium, 2017.
- > -, Walter MR. Notepad-Augmented Environments in Reinforcement Learning. Midwest Robotics workshop, 2017.
- > -, Cai Z. Glyph-based Visual Chinese Character Embedding. Midwest Speech and Language Days, 2017.
- > -, Nettsheim G. Simulation and Modeling of the Anode of the Proposed Large-Area Picosecond

Photo-Detector. Chicago Area Undergraduate Research Symposium, 2011.

Working papers

- > Mason W, Brenken D, -, Cruz Castillo RG, St-Martin Cormier O, Sedal A. [Acoustic Tactile Sensing for Mobile Robot Wheels](#). In submission, 2024.
- > -. [Word2vec Conjecture and A Limitative Result](#). In submission, 2019.

Service to the community

- > Primary reviewer. IEEE Transactions on Robotics (T-RO), 2023.
- > Primary reviewer. International Conference on Learning Representations (ICLR), 2023.
- > Primary reviewer. International Conference on Machine Learning (ICML), 2021.
- > Primary reviewer. The Association for Computational Linguistics conference (ACL), 2021.
- > Primary reviewer. North American Chapter of the Association for Computational Linguistics conference (NAACL), 2021.
- > Primary reviewer. European Chapter of the Association for Computational Linguistics conference (EACL), 2021.
- > Primary reviewer. International Conference on Learning Representations conference (ICLR), 2021.
- > Primary reviewer. Association for the Advancement of Artificial Intelligence conference (AAAI), 2021.
- > Primary reviewer. Neural Information Processing Systems conference (NeurIPS), 2020.
- > Primary reviewer. Empirical Methods in Natural Language Processing conference (EMNLP), 2020.
- > Primary reviewer. International Conference on Learning Representations conference (ICLR), 2020.
- > Student volunteer. Symposium on Theory of Computing (STOC), 2020.
- > Primary reviewer. Association for Computational Linguistics conference (ACL), 2020.
- > Primary reviewer. International Joint Conference on Artificial Intelligence (IJCAI), 2020.
- > Secondary reviewer. Artificial Intelligence and Statistics (AISTATS), 2020.
- > Primary reviewer. International Conference on Learning Representations (ICLR), 2019.
- > Secondary reviewer. Neural Information Processing Systems conference (NeurIPS), 2019.
- > Student volunteer. Reinforcement Learning and Decision Making (RLDM), 2019.
- > Primary reviewer. International Journal of Robotics Research (IJRR), 2018.
- > Primary reviewer. International Conference on Learning Representations conference (ICLR), 2018.
- > Primary reviewer. International Symposium on Robotics Research (ISRR), 2017.
- > Secondary reviewer. Neural Information Processing Systems conference (NIPS), 2017.
- > Primary reviewer. Spatial-Semantic Representations in Robotics workshop at Robotics: Science and Systems conference (RSS), 2017.

Experience

Machine learning research, Symbolica AI, San Francisco, CA.

Machine learning researcher, April 2024 - Present

- > Automated theorem proving with large language models and reinforcement learning.
- > Representation learning and latent diffusion modeling with transformers.

Software engineering, Apple Special Project Group, Santa Clara, CA.

Machine learning engineer (contractor), October 2023 - March 2024

- > Support perception, annotation tooling, annotation ops teams at the special project group.
- > Make data ingestion and annotation ETL workflows more reliable, scalable and maintainable.

Teaching assistant to *An Introduction to Machine Learning Theory*, Toyota Technological Institute at Chicago, Chicago, IL

Teaching assistant (to Professor Avrim Blum), March 2022 - June 2022

- > Provided detailed technical explanations at office hours and over emails.
- > Automated the returning of graded homeworks via email with Gmail APIs.

Technical consulting, Waymark Inc, Detroit, MI

Independent contractor, January 2018 - February 2018

- > Prototyped and advised the CEO on an abstractive summarization system.

Teaching assistant to *Duckietown*, Toyota Technological Institute at Chicago, Chicago, IL

Teaching assistant (to Professor Matthew Walter), October 2017 - December 2017

- > Created material for the hands-on self-driving robotics course.
- > Provided technical assistance to students.

Research in Abstractive Summarization, Harvard University, Cambridge, MA

Visiting Research Intern (hosted by [Professor Alexander Rush](#)), July 2017 - September 2017

- > Re-implemented state-of-the-art methods in paragraph-to-sentence summarization.
- > Maintained the open-sourced OpenNMT-py repository.

Data Science and Analytics, Strikingly Inc, Shanghai, China

Data Scientist, February 2015 - August 2015

- > Designed an improved web analytics implementation workflow.
- > Recruited and managed a data engineer.
- > Analyzed user behaviors, user acquisition campaigns, user referral programs.
- > Defined business growth/health metrics and implemented monitoring dashboards.

Application of Artificial Neural Network in NLP, Toyota Technological Institute at Chicago, Chicago, IL

Student Visitor (of Professor Kevin Gimpel and Professor Mohit Bansal), June 2014 - December 2014

- > Read and reviewed relevant academic articles.
- > Implemented a neural network library in Python (optimized with NumPy).

Neurological Research, Towle Lab, University of Chicago, Chicago, IL

Research Assistant (to [Professor Vernon L. Towle](#)), November 2012 - December 2014

- > Studied language processing via electrocorticographic data.
- > Developed novel methods for registering intracranial electrodes.
- > Implemented state-of-the-art medical image analysis and visualization software.

Sociological Research, Knowledge Lab, University of Chicago, Chicago, IL

Research Assistant (to [Professor James Evans](#)), October 2012 – December 2013

- > Built a machine learning pipeline to predict sociological attributes from Google StreetView images.
- > Built a web application for collecting graph-structured information from users.
- > Analyzed author networks induced by co-authorship and citations.

TwIthinks, a startup project, Chicago, IL and Cambridge, MA

Co-founder, April 2011 – January, 2014

- > Won web track in MIT-CHIEF Business Plan Contest at Massachusetts Institute of Technology
- > Built a prototype twithinks.com to visualize the Twitter users' reactions to 2012 presidential election.
- > Initiated and completed #ivoted map on election day which received 20K pageviews within 6 hours.
- > Featured on MIT-CSAIL news and a Swiss national news outlet Tages-Anzeiger.

Summer Research Experience in Physics, University of Chicago, IL

Technical Support for *Quantum Computing Project* (Professor [David Schuster](#)) and *Photon Detector Project* (Professor [Henry Frisch](#)), June 2009 – September 2009, June 2010 – September 2010, June 2011 – September 2011

- > Modeled the secondary electron emission process, implemented simulation programs, and analyzed experiment signal data.
- > Learnt basic signal processing and printed circuit board design.
- > built a custom spectrum analyzer and coded a custom GUI program.

Computer Science Department and Mathematics Department, University of Chicago, IL

Teaching Assistant for *Mobile Computing*, *Introduction to Programming (C++)*, *Grader* for *Introduction to Scientific Computing* and *Honors Calculus*, October 2010 – November 2011

- > Graded more than 30 students' work per week with detailed corrections, comments and guidance.

Honors

- > [Best paper award finalist](#) at Association for Computational Linguistics conference (ACL), 2019.
- > Best app award at an invited hackathon in Shanghai, for a social chat app prototype, 2015.
- > 2nd place in [BattleHack hackathon in Chicago](#), for a facial recognition-assisted social payment app, 2014.
- > Ranked 164th on [Kaggle](#), 2013.
- > Co-founder of the Engineering Society at the University of Chicago, 2010.
- > 23rd team (out of 138 teams) at International Collegiate Programming Contest (ICPC) regional, 2010.
- > AP scholar after self-studying and scoring 5/5 on nine AP subjects, 2007.
- > 2nd place at Alamo regional Science & Engineering Fair; Honorable Mention at Texas state, 2007.

Skills

- > Programming languages: python, javascript, C, C++, CUDA, Java, prolog, SQL.
- > Robotics: ROS, RGBD data.
- > "Big data" software stack: Scikit-Learn, Spark, ElasticSearch, Hadoop, Pandas, IPython.
- > Web development: React, D3js, WebGL, HTML5, CSS3, Node.js, PostgreSQL, MongoDB.
- > Software: Atom editor, Tmux, Docker, Eclipse, Git, Mathematica, Octave/Matlab, LaTeX, Linux.
- > Web API's: Amazon Web Services, DigitalOcean, Google Cloud Platform, Twitter API, WeChat API.
- > Fluent in Mandarin and Cantonese.

Interests

- > Personal: skiing, climbing, automotive technologies, dog training, cooking, drawing, fixing electronics.
- > Academic: cognition, physics, quantum mechanics, logic.