

Yuntian DENG

yuntiandeng.com • Google Scholar
yuntian@uwaterloo.ca • Updated August 2026

Employment

University of Waterloo	Waterloo, Canada
Assistant Professor, Cheriton School of Computer Science	Aug 2024–Present
Vector Institute	Toronto, Canada
Faculty Affiliate	Aug 2024–Present
Harvard University	Remote
Associate, Computer Science	Jun 2024–Present
NVIDIA	Toronto, Canada
Visiting Professor, Yejin Choi's team	Nov 2024–Nov 2025
Allen Institute for Artificial Intelligence	Seattle, WA
Postdoc (advised by Yejin Choi)	Jul 2023–Jul 2024

Research Interests

Program-as-Weights and neural programs; neural world models and interactive systems; implicit chain-of-thought reasoning and internalized computation; real-world LLM use and open research infrastructure.

Education

Harvard University	Cambridge, MA
Ph.D. in Computer Science	May 2023
Advisors: Alexander Rush and Stuart Shieber	
Carnegie Mellon University	Pittsburgh, PA
Master in Language Technologies	Aug 2016
Advisor: Eric Xing	
Tsinghua University	Beijing, China
Bachelor of Engineering, Department of Automation	Jul 2014

Awards

Laude Institute Slingshots (NeuralOS)	Jun 2026
Google Research Award (co-PI)	Apr 2026
Argonne National Lab Impact Award	Jan 2023
University of Chicago Rising Stars in Data Science	Nov 2022
ACM Gordon Bell Special Prize for COVID-19 Research	Nov 2022
NVIDIA Fellowship	Dec 2021
Microsoft Turing Academic Program, Selected for Improving LM Reasoning	Dec 2021
Harvard Certificates of Distinction in Teaching	Spring 2019, Fall 2020, Fall 2021
Twitch Fellowship Finalist	Jan 2021
DAC 2020 Best Paper Award	Jul 2020
Baidu Fellowship	2019
French American Doctoral Exchange Program (French Embassy, 10 US laureates)	Jun 2018

Research Contributions

Program-as-Weights and neural programs. Compiles English function specifications into compact neural programs that run locally; its public hub contains 3,800+ community-built programs.

Neural world models and interactive systems. Uses a neural generative model to simulate an interactive operating system end to end; its public deployment has 50,000+ visits and USD 45,000 in Laude Institute support.

Implicit chain-of-thought reasoning and internalized computation. Introduced reasoning in continuous hidden states and helped catalyze the area by organizing the Latent & Implicit Thinking Workshop at ICLR 2026, which accepted 78 papers.

Real-world LLM use and open research infrastructure. WildChat provides 4.8 million real user-ChatGPT conversations and has been used in published work by OpenAI and Anthropic; the research was featured in *The Washington Post* and *Neue Zürcher Zeitung*.

Publications and Preprints

Representative Publications

- 48 Wentao Zhang*, Liliana Hotsko*, Woojeong Kim*, Pengyu Nie, Stuart Shieber, **Yuntian Deng**. Program-as-Weights: A Programming Paradigm for Fuzzy Functions. *arXiv 2026*
- 47 Luke Rivard, Sun Sun, Hongyu Guo, Wenhui Chen, **Yuntian Deng**. NeuralOS: Towards Simulating Operating Systems via Neural Generative Models. *ICLR 2026*
- 46 **Yuntian Deng**, Yejin Choi, Stuart Shieber. From Explicit CoT to Implicit CoT: Learning to Internalize CoT Step by Step. *arXiv 2024*
- 45 **Yuntian Deng**, Kiran Prasad, Roland Fernandez, Paul Smolensky, Vishrav Chaudhary, Stuart Shieber. Implicit Chain of Thought Reasoning via Knowledge Distillation. *arXiv 2023*
- 44 Wenting Zhao, Xiang Ren, Jack Hessel, Claire Cardie, Yejin Choi, **Yuntian Deng**. WildChat: 1M ChatGPT Interaction Logs In The Wild. *ICLR 2024 Spotlight*

Additional Publications and Preprints

- 43 Wentao Zhang*, Xuanhe Pan*, Han Zhou*, Yang Lu[†], **Yuntian Deng**[†]. Interactive Training 2: Auditable Control Plane for Live Model Training. *arXiv 2026*
- 42 Liliana Hotsko, Yinxi Li, **Yuntian Deng**, Pengyu Nie. Code2LoRA: Hypernetwork-Generated Adapters for Code Language Models under Software Evolution. *arXiv 2026*
- 41 Sajad Ebrahimi, Nima Jamali, Bardia Shirsalimian, Kelly McConvey, Wentao Zhang, Jalehsadat Mahdavi Moghadam, Maksym Taranukhin, Maura R. Grossman, Vered Shwartz, **Yuntian Deng**, Ebrahim Bagheri. DetectZoo: A Unified Toolkit for AI-Generated Content Detection Across Text, Audio, and Image Modalities. *arXiv 2026*
- 40 Kelly McConvey, Jalehsadat Mahdavi Moghadam, Nima Jamali, Maksym Taranukhin, Sajad Ebrahimi, Wentao Zhang, Karen Eltis, **Yuntian Deng**, Maura R. Grossman, Vered Shwartz, Ebrahim Bagheri. The CIFAR Synthetic Evidence Corpus for Detecting AI-Generated Evidence. *arXiv 2026*
- 39 Shayne Longpre, Anka Reuel, Dayeon Ki, Zhiping Zhang, Cathy Mengying Fang, Chuanyang Jin, Jennifer Mickel, Cedric Deslandes Whitney, Niloofar Mireshghallah, Victor Ojewale, Megan Richards, Ariel N. Lee, Alex Pentland, Tianshi Li, **Yuntian Deng**, Sara Hooker, Mykel Kochenderfer, Sanmi Koyejo. The AI Observatory: A Public Measure of Real-World AI Use. *In Submission*
- 38 David Acuna, Chao-Han Huck Yang, **Yuntian Deng**, Jaehun Jung, Ximing Lu, Prithviraj Ammanabrolu, Hyunwoo Kim, Yuan-Hong Liao, Yejin Choi. Long Grounded Thoughts: Synthesizing Visual Problems and Reasoning Chains at Scale. *ICML 2026*

- 37 Yinxi Li, **Yuntian Deng**, Pengyu Nie. TokDrift: When LLM Speaks in Subwords but Code Speaks in Grammar. *ACL 2026*
- 36 Wei Liu, Ruochen Zhou, Yiyun Deng, Yuzhen Huang, Junteng Liu, **Yuntian Deng**, Yizhe Zhang, Junxian He. Learn to Reason Efficiently with Adaptive Length-Based Reward Shaping. *ICLR 2026*
- 35 Wentao Zhang, Yang Young Lu, **Yuntian Deng**. Interactive Training: Feedback-Driven Neural Network Optimization. *EMNLP 2025 Demo*
- 34 Wentao Zhang, Woojeong Kim, **Yuntian Deng**. From Chat Logs to Collective Insights: Aggregative Question Answering. *EMNLP 2025 Main Oral*
- 33 Bill Yuchen Lin, **Yuntian Deng**, Khyathi Chandu, Faeze Brahman, Abhilasha Ravichander, Valentina Pyatkin, Nouha Dziri, Ronan Le Bras, Yejin Choi. WildBench: Benchmarking LLMs with Challenging Tasks from Real Users in the Wild. *ICLR 2025 Spotlight*
- 32 Xiaoyan Bai, Itamar Pres, **Yuntian Deng**, Chenhao Tan, Stuart Shieber, Fernanda Viégas, Martin Wattenberg, Andrew Lee. Why Can't Transformers Learn Multiplication? Reverse-Engineering Reveals Long-Range Dependency Pitfalls. *arXiv 2025*
- 31 Shivalika Singh, Yiyang Nan, Alex Wang, Daniel D'Souza, Sayash Kapoor, Ahmet Üstün, Sanmi Koyejo, **Yuntian Deng**, Shayne Longpre, Noah A. Smith, Beyza Ermis, Marzieh Fadaee, Sara Hooker. The Leaderboard Illusion. *NeurIPS D&B 2025*
- 30 Yifan Zong, **Yuntian Deng**, Pengyu Nie. Mix-of-Language-Experts Architecture for Multilingual Programming. *LLM4Code 2025*
- 29 Zhangchen Xu, Fengqing Jiang, Luyao Niu, **Yuntian Deng**, Radha Poovendran, Yejin Choi, Bill Yuchen Lin. Magpie: Alignment Data Synthesis from Scratch by Prompting Aligned LLMs with Nothing. *ICLR 2025*
- 28 **Yuntian Deng**, Wenting Zhao, Jack Hessel, Xiang Ren, Claire Cardie, Yejin Choi. WildVis: Open Source Visualizer for Million-Scale Chat Logs in the Wild. *EMNLP 2024 Demo*
- 27 Wenting Zhao, Tanya Goyal, Yu Ying Chiu, Liwei Jiang, Benjamin Newman, Abhilasha Ravichander, Khyathi Chandu, Ronan Le Bras, Claire Cardie, **Yuntian Deng**, Yejin Choi. WildHallucinations: Evaluating Long-form Factuality in LLMs with Real-World Entity Queries. *arXiv 2024*
- 26 Jinjie Ni, Fuzhao Xue, Xiang Yue, **Yuntian Deng**, Mahir Shah, Kabir Jain, Graham Neubig, Yang You. MixEval: Deriving Wisdom of the Crowd from LLM Benchmark Mixtures. *NeurIPS 2024*
- 25 John X. Morris, Chandan Singh, Alexander Rush, Jianfeng Gao, **Yuntian Deng**. Tree Prompting: Efficient Task Adaptation without Fine-Tuning. *EMNLP 2023*
- 24 **Yuntian Deng**, Noriyuki Kojima, Alexander Rush. Markup-to-Image Diffusion Models with Scheduled Sampling. *ICLR 2023 (Demo: huggingface.co/spaces/yuntian-deng/latex2im)*
- 23 Richa Rastogi, Yair Schiff, Alon Hachohen, Zhaozhi Li, Ian Lee, **Yuntian Deng**, Mert R. Sabuncu, Volodymyr Kuleshov. Semi-Parametric Inducing Point Networks and Neural Processes. *ICLR 2023*
- 22 Maxim Zvyagin*, Alexander Brace*, Kyle Hippe*, **Yuntian Deng***, Bin Zhang, Cindy Orozco Bohorquez, Austin Clyde, Bharat Kale, Danilo Perez-Rivera, Heng Ma, Carla M. Mann, Michael Irvin, J. Gregory Pauloski, Logan Ward, Valerie Hayot-Sasson, Murali Emani, Sam Foreman, Zhen Xie, Diangen Lin, Maulik Shukla, Weili Nie, Josh Romero, Christian Dallago, Arash Vahdat, Chaowei Xiao, Thomas Gibbs, Ian Foster, James J. Davis, Michael E. Papka, Thomas Brettin, Rick Stevens, Anima Anandkumar, Venkatram Vishwanath, Arvind Ramanathan. GenSLMs: Genome-scale language models reveal SARS-CoV-2 evolutionary dynamics. *International Journal of High Performance Computing Applications*, 2023. **ACM Gordon Bell Special Prize for COVID-19 Research**
- 21 **Yuntian Deng**, Volodymyr Kuleshov, Alexander Rush. Model Criticism for Long-Form Text Generation. *EMNLP 2022*
- 20 Anton Bakhtin*, **Yuntian Deng***, Sam Gross, Myle Ott, Marc'Aurelio Ranzato, Arthur Szlam. Residual Energy-Based Models for Text. *JMLR 2021*

- 19 Justin Chiu*, **Yuntian Deng***, Alexander Rush. Low-Rank Constraints for Fast Inference in Structured Models. *NeurIPS 2021*
- 18 **Yuntian Deng**, Alexander Rush. Sequence-to-Lattice Models for Fast Translation. *EMNLP 2021 Findings*
- 17 Keyon Vafa, **Yuntian Deng**, David Blei, Alexander Rush. Rationales for Sequential Predictions. *EMNLP 2021 Oral*
- 16 **Yuntian Deng**, Anton Bakhtin, Myle Ott, Arthur Szlam, Marc'Aurelio Ranzato. Residual Energy-Based Models for Text Generation. *ICLR 2020*
- 15 **Yuntian Deng**, Alexander Rush. Cascaded Text Generation with Markov Transformers. *NeurIPS 2020*
- 14 Thierry Tambe, En-Yu Yang, Zishen Wan, **Yuntian Deng**, Vijay Janapa Reddi, Alexander Rush, David Brooks, Gu-Yeon Wei. Algorithm-Hardware Co-Design of Adaptive Floating-Point Encodings for Resilient Deep Learning Inference. *Design Automation Conference (DAC) 2020 Best Paper Award*
- 13 Zachary Ziegler*, **Yuntian Deng***, Alexander Rush. Neural Linguistic Steganography. *EMNLP 2019 Oral (Demo: steganography.live)*
- 12 Anton Bakhtin, Sam Gross, Myle Ott, **Yuntian Deng**, Marc'Aurelio Ranzato, Arthur Szlam. Real or Fake? Learning to Discriminate Machine from Human Generated Text. *arXiv 1906.03351*
- 11 **Yuntian Deng**, David Rosenberg, Gideon Mann. Challenges in End-to-End Neural Scientific Table Recognition. *ICDAR 2019*
- 10 **Yuntian Deng***, Yoon Kim*, Justin Chiu, Demi Guo, Alexander Rush. Latent Alignment and Variational Attention. *NeurIPS 2018*
- 9 Sebastian Gehrmann, **Yuntian Deng**, Alexander Rush. Bottom-Up Abstractive Summarization. *EMNLP 2018*
- 8 **Yuntian Deng**, Anssi Kanervisto, Jeffrey Ling, Alexander Rush. Image-to-Markup Generation with Coarse-to-Fine Attention. *ICML 2017 (Demo: im2markup.yuntiangdeng.com)*
- 7 Pengtao Xie, **Yuntian Deng**, Yi Zhou, Abhimanu Kumar, Yaoliang Yu, James Zou, Eric P Xing. Learning Latent Space Models with Angular Constraints. *ICML 2017*
- 6 Guillaume Klein, Yoon Kim, **Yuntian Deng**, Jean Senellart, Alexander M Rush. OpenNMT: Open-Source Toolkit for Neural Machine Translation. *ACL 2017 Best Demo Paper Award Runner-Up*
- 5 Zichao Yang, Zhiting Hu, **Yuntian Deng**, Chris Dyer, Alex Smola. Neural Machine Translation with Recurrent Attention Modeling. *EACL 2017*
- 4 Xuezhe Ma, Yingkai Gao, Zhiting Hu, Yaoliang Yu, **Yuntian Deng**, Eduard Hovy. Dropout with Expectation-linear Regularization. *ICLR 2017*
- 3 Hao Zhang, Zhiting Hu, **Yuntian Deng**, Mrinmaya Sachan, Zhicheng Yan, Eric P. Xing. Learning Concept Taxonomies from Multi-modal Data. *ACL 2016*
- 2 Pengtao Xie, **Yuntian Deng**, Eric P. Xing. Diversifying Restricted Boltzmann Machine for Document Modeling. *KDD 2015*
- 1 Zhiting Hu, Poyao Huang, **Yuntian Deng**, Yingkai Gao, Eric P. Xing. Entity Hierarchy Embedding. *ACL 2015*

Media Coverage

Interviews and Direct Features

The Canadian Press: Real or fake? Researchers to develop tool that would help courts spot AI evidence	Dec 5, 2025
NZZ: What people are really asking ChatGPT? (English translation)	Nov 30, 2024
TechCrunch: Why is ChatGPT so bad at math?	Oct 2, 2024
The Washington Post: What do people really ask chatbots?	Aug 4, 2024

Coverage of Coauthored Research

Tech Xplore: AI models stumble on basic multiplication without special training methods, study finds	Dec 29, 2025
Harvard Business School AI Institute: When Giants Stumble: What Multiplication Reveals about AI's Capabilities	Oct 8, 2025
Computerworld: Leaderboard illusion: How big tech skewed AI rankings on Chatbot Arena	May 2, 2025

Research Funding

Laude Institute Slingshots: NeuralOS (USD 45,000 in-kind support)	Jun 2026
Google Research Award: "LLM-Guided Interactive Training of Biomedical Foundation Models on TPUs" (co-PI; USD 25,000 unrestricted gift and USD 75,000 in Google Cloud credits)	Apr 2026
CIFAR Canadian AI Safety Institute – AI Safety Solution Network: "Safeguarding Courts from Synthetic AI Content" (co-PI; CAD 700,000)	2025–2027
Canada Foundation for Innovation John R. Evans Leaders Fund: "Progressive Internalization of Skills for Advanced Language Model Reasoning" (PI; CAD 244,846 total project value, including CAD 96,000 CFI contribution)	2024–2028
Digital Research Alliance of Canada, Resources for Research Groups: "Neural OS: Simulating Operating Systems using Diffusion Models" (PI; RRG 2025 No. 5275, estimated in-kind value CAD 61,165; RRG 2026 ID 3797, estimated in-kind value CAD 28,772)	Jul 2025–Apr 2027
Mitacs Business Strategy Internship: "Decentralized Electric Microgrid Optimization by Reinforcement Learning" (PI; CAD 10,000)	2025
Manulife Research Grant: "End-to-End Optimization of LLM-Based Multi-Agent Systems via Text Differentiation" (PI; CAD 30,000)	2024–2026
NSERC Discovery Grant: "Multi-Agent Communication and Collaboration for Specialized Large Language Models" (PI; RGPIN-2024-05178; CAD 41,000/year + CAD 12,500 Discovery Launch Supplement)	Aug 2024–Apr 2029
University of Waterloo Starter Grant (CAD 150,000)	Aug 2024–Aug 2029

Workshop Organization

Organizer of Latent & Implicit Thinking – Going Beyond CoT Reasoning Workshop	Apr 27, 2026, ICLR
---	--------------------

Professional Service

Editorial and Leadership: Co-Chief Technical Officer, ACL Rolling Review (ARR), Dec 2025–Present (technical infrastructure supporting 17,087 submissions in the May 2026 cycle); Action Editor, TACL, 2025–Present; Co-Chair, AAAI Undergraduate Consortium, 2027; Mentor Chair, ICLR 2026.

Area Chair / Senior Program Committee: Area Chair, ICLR (2024–2026), ICML (2026), NeurIPS (2026), ACL Rolling Review (2024–2026), COLM (2025–2026), EMNLP Demo (2026), LREC (2026), EMNLP GEM Workshop (2023), and IJCNLP–AAACL (Language Modeling, 2025); Senior Program Committee Member, AAAI 2027.

Reviewing: Reviewer, *Nature* (2026); Reviewer, Schmidt Sciences Science of Trustworthy AI (2026); TACL, TMLR, Computational Linguistics, ACM AI Letters, ACM Computing Surveys, Foundations and Trends in Signal Processing, IEEE TNNLS, IEEE TIFS, ACM Transactions on Interactive Intelligent Systems, ACM Transactions on Information Systems, and Journal of Computer Science and Technology; Reviewer for ICML, NeurIPS, ICLR, ACL, ACL Rolling Review, EMNLP, AAAI, COLM, COLING, AACL, EACL, NAACL, ACL Demo, NAACL Demo, EMNLP Demo, and various workshops and reproducibility challenges; Proposal Reviewer, Mitacs; Selection Committee Member, Vector Scholarship in Artificial Intelligence (2026–2027).

Conference Service: Oral Session Moderator, ICLR (2024, 2026); Oral Session Chair, EMNLP 2024; Co-host, ICLR 2026 Mentorship Session, Apr 23, 2026; Oral Session Volunteer, ACL 2021; Volunteer, EMNLP 2021 and ICML (2016, 2020); Volunteer Moderator, ACL 2020.

Recognition: Top 10% Reviewer, NeurIPS 2020; Best Reviewer (Top 10%), ICML 2021.

Teaching Experience

Waterloo

Instructor, CS 486/686: Introduction to Artificial Intelligence	Spring 2026
Instructor, CS 486/686: Introduction to Artificial Intelligence	Spring 2025
Guest Lecture, Computational Linguistics (Freda Shi)	Feb 11, 2025
Instructor, CS 886: Topics in Language Models	Fall 2024

Harvard University

Head TA, Intro to Computational Linguistics and NLP (Stuart Shieber)	Aug 2021–Dec 2021
Head TA, Intro to Computational Linguistics and NLP (Stuart Shieber)	Aug 2020–Dec 2020
TA, Natural Language Processing (Alexander Rush)	Jan 2019–May 2019
Lecture, Natural Language Processing – Translation (Alexander Rush)	Feb 13, 2019
TA, Advanced Machine Learning (Alexander Rush)	Aug 2017–Dec 2017

Cornell University

Lab Material Preparation, Break Through Tech AI (Alexander Rush)	Aug 2021
--	----------

Carnegie Mellon University

TA, Probabilistic Graphical Models (Matthew Gormley, Eric Xing)	Jan 2016–May 2016
TA, Graduate Machine Learning (Ziv Bar-Joseph, Eric Xing)	Aug 2015–Dec 2015

Community Service

Co-Treasurer, Board of Directors, Bright Starts Early Learning Centre	May 2025–May 2027
Childcare Chair, NAACL 2022 Diversity & Inclusion Committee	Apr–Jul 2022
Leader, Harvard English Language Table	Sep 2021–Jan 2023
Mentor, Harvard Women in CS NLP Reading Group	Sep 2020–Sep 2023
Volunteer Instructor of AP CS, Revere High School	Sep 2020–May 2021

Internships

NVIDIA (Anima Anandkumar, Weili Nie, Arash Vahdat, Chaowei Xiao)	May–Dec 2022
Facebook AI Research (Marc’Aurelio Ranzato, Arthur Szlam)	May–Dec 2019
Bloomberg CTO Office (David Rosenberg, Gideon Mann)	Jan–Aug 2017
UCSD (Charles Elkan)	Jul–Sep 2013

Talks

Apple MLR Internal Reading Group: NeuralOS	Aug 2026, Remote
Vector Endless Summer School: NeuralOS	Jun 2026, Remote
NSERC CREATE Responsible AI consortium: NeuralOS	Mar 2026, Remote
Mila/McGill NLP Reading Group: Implicit Reasoning in Language Models	Oct 2025, Remote
U Cambridge Language Technology Lab: Implicit Chain-of-Thought: Internalizing Reasoning in Language Models	Oct 2025, Remote
Google DeepMind: Implicit Reasoning in Language Models	May 2025, Remote
Colin Raffel Group: Implicit Reasoning in Language Models	Mar 2025, Toronto
U Oklahoma: Implicit Chain of Thought Reasoning	Jan 2025, Remote
NYU: Implicit Chain of Thought Reasoning	Nov 2024, New York, NY
Google DeepMind: Implicit Chain of Thought Reasoning	Jul 2024, Remote
Queen's University: Implicit Chain of Thought Reasoning	Feb 2024, Remote
UBC: Implicit Chain of Thought Reasoning	Jan 2024, Vancouver, Canada
AI2 Mosaic: WildChat	Sep 2023, Seattle, Washington
AI2 Semantic Scholar: Structure Modeling in Language Models	Aug 2023, Seattle, Washington
NYU AI School 2023: Natural Language Processing	Jun 2023, New York, NY
AI2 Mosaic: Structure Modeling in Language Models	Jun 2023, Seattle, Washington
UMass Amherst: Structural Coherence in Text Generation	May 2023, Remote
Georgia Tech: Structural Coherence in Text Generation	Apr 2023, Atlanta, Georgia
Harvard Kempner Institute: Structural Coherence in Text Generation	Mar 2023, Remote
U Alberta: Structural Coherence in Text Generation	Mar 2023, Edmonton, Canada
TTIC: Structural Coherence in Text Generation	Feb 2023, Chicago
U Waterloo ECE: Structural Coherence in Text Generation	Feb 2023, Remote
Cornell Seminar in NLU: Structural Coherence in Text Generation	Feb 2023, New York
U Waterloo CS: Structural Coherence in Text Generation	Jan 2023, Waterloo, Canada
U Chicago Rising Star: Model Criticism for Long-Form Text Generation	Nov 2022, Chicago
Princeton NLP: Model Criticism for Long-Form Text Generation	Nov 2022, Princeton
EMNLP 2021 Findings: Sequence-to-Lattice Models for Fast Translation	Nov 2021, Remote
OpenAI: Residual Energy-based Models for Text Generation	Apr 2021, Remote
NeurIPS 2020: Cascaded Text Generation with Markov Transformers	Dec 2020, Remote
Baidu: Cascaded Text Generation with Markov Transformers	Jun 2020, Remote
ICLR 2020: Residual Energy-based Models for Text Generation	Apr 2020, Remote
LinkedIn: Residual Energy-based Models for Text Generation	Mar 2020, Remote
Wayfair: Neural Encoder-Decoder Models	Nov 2019, Boston
EMNLP 2019: Neural Linguistic Steganography	Nov 2019, Hong Kong
FAIR NLP Meeting: Energy-Based Models for Text Generation	Aug 2019, New York City
NeurIPS 2018: Latent Alignment and Variational Attention	Dec 2018, Montreal, Canada
The French National Center for Scientific Research: Variational Attention	Jul 2018, France
Association for Machine Translation in the Americas: OpenNMT	Mar 2018, Boston
Open-Source Neural Machine Translation Workshop: Image-to-Text	Mar 2018, Paris, France
Tencent Social Network Group TSAIC: OpenNMT	Dec 2017, Shenzhen, China
ICML: Image-to-Markup Generation	Aug 2017, Sydney, Australia

Panels

KAIST at ICML

Jul 6, 2026, Seoul

Panelist alongside Jaehong Yoon and Young Jin Kim. Discussed Vision, Multimodal AI & Agents.

Vector Institute's Remarkable 2025

Mar 2025, Toronto

Panelist alongside Colin Raffel, Frank Rudzicz, Gerald Penn, Hongyang Zhang, Tracy Jenkin, and David Duvenaud. Discussed questions in NLP.

ACL Mentorship: Sharing Learnings & Identifying Research Directions

Dec 2023, Singapore

Panelist alongside Mohit Bansal, Zhijing Jin, and Rada Mihalcea. Discussed promising research directions.

ACL Mentorship: Preparing Your Grad School Application

Oct 2023, Remote

Participated in a mentorship session with Luke Zettlemoyer, Zhijing Jin, Nathan Schneider, and Oana Ignat, focused on preparing graduate school applications.

NAACL 2022: Balancing Professional and Family Commitments

Jul 2022, Remote

Co-hosted with Shirley Hayati, discussing the balance between professional and family commitments.

Thesis Committees

- MSc Thesis Committee for Guilherme Campos (Waterloo, advised by Wojciech Golab and Tahsin Reza) Aug 12, 2026
- MMath Thesis Committee for Benjamin Schneider (Waterloo, advised by Wenhui Chen and Florian Kerschbaum) Aug 11, 2026
- MMath Thesis Committee for Arthur Chen (Waterloo, advised by Victor Zhong) Aug 4, 2026
- MMath Thesis Committee for Ruotian Wu (Waterloo, advised by Pascal Poupart) Aug 4, 2026
- PhD Thesis Committee for William Loh (Waterloo, advised by Pascal Poupart) Jul 21, 2026
- PhD Thesis Committee for Xueguang Ma (Waterloo, advised by Jimmy Lin) Jul 15, 2026
- MMath Thesis Committee for Hosna Oyarhoseini (Waterloo, advised by Jimmy Lin and Amir-Hossein Karimi) Jul 14, 2026
- PhD Comprehensive-II Committee for Xiang Li (Waterloo, advised by Jian Zhao and Kate Larson) Jun 26, 2026
- MMath Thesis Committee for Justin Xu (Waterloo, advised by Pascal Poupart) May 2026
- MMath Thesis Committee for Saarang Agarwal (Waterloo, advised by Mei Nagappan and Pengyu Nie) Apr 2026
- PhD Comprehensive-II Committee for Weiming Ren (Waterloo, advised by Wenhui Chen and Jimmy Lin) Apr 2026
- PhD Comprehensive-II Committee for Dongfu Jiang (Waterloo, advised by Wenhui Chen) Mar 2026
- PhD Comprehensive-II Committee for Cong Wei (Waterloo, advised by Wenhui Chen) Mar 2026
- PhD Comprehensive-II Committee for Dongfu Jiang (Waterloo, advised by Wenhui Chen) Jan 2026
- MMath Thesis Defense Committee for Yuansheng Ni (Waterloo, advised by Wenhui Chen) Nov 2025
- PhD Comprehensive-II Committee for Xueguang Ma (Waterloo, advised by Jimmy Lin) Oct 2025
- PhD Seminar Committee for Ronak Pradeep (Waterloo, advised by Jimmy Lin) Oct 2025
- MMath Thesis Defense Committee for Allen Zhang (Waterloo, advised by Anita Layton) Aug 2025
- MMath Thesis Defense Committee for Carter Blair (Waterloo, advised by Kate Larson and Edith Law) Jun 2025

- MMath Thesis Defense Committee for Achint Soni (Waterloo, advised by Sirisha Rambhatla and Charles Clarke) May 2025
- MMath Thesis Defense Committee for Chi-Chung Cheung (Waterloo, advised by Anita Layton) Dec 2024
- PhD Thesis Proposal Committee for William Loh (Waterloo, advised by Pascal Poupart) Aug 2024

PhD Seminar Attendance

- Faculty Attendee, PhD Seminar for Cong Wei (Waterloo, advised by Wenhui Chen) Jul 27, 2026
- Faculty Attendee, PhD Seminar for Cong Wei (Waterloo, advised by Wenhui Chen) Jul 24, 2026
- Faculty Attendee, PhD Seminar for Weiming Ren (Waterloo, advised by Wenhui Chen) Jul 22, 2026
- Faculty Attendee, PhD Seminar for Dongfu Jiang (Waterloo, advised by Wenhui Chen) Jul 15, 2026
- Faculty Attendee, PhD Seminar for Weiming Ren (Waterloo, advised by Wenhui Chen) Jul 15, 2026
- Faculty Attendee, PhD Seminar for William Loh (Waterloo, advised by Pascal Poupart) Jun 15, 2026
- Faculty Attendee, PhD Seminar for William Loh (Waterloo, advised by Pascal Poupart) Jun 5, 2026
- Faculty Attendee, PhD Seminar for William Loh (Waterloo, advised by Pascal Poupart) Mar 2026

Open Source Projects

ProgramAsWeights (230 GitHub stars; 3,856 community-built neural programs)	Python
NeuralOS (202 GitHub stars; 50,000+ demo visits)	Flask & WebSocket
WildChat-4.8M-Full (909 total Hugging Face likes across five AllenAI versions)	Dataset
OpenAIWatch (101,000 visits since 2025; 87,000 in 2026 YTD)	Python
WildVisualizer (3.1 million visits since 2025; 2.0 million in 2026 YTD)	Flask & D3.js
Internalize CoT Step by Step (208 GitHub stars)	PyTorch
Implicit Chain of Thought (149 GitHub stars)	PyTorch
WildChat o1 Chatbot (450 Hugging Face likes)	PyTorch
WildChat o1-mini Chatbot (258 Hugging Face likes)	PyTorch
WildChat GPT-4 Chatbot (492 Hugging Face likes)	PyTorch
ChatGPT Chatbot (416 Hugging Face likes)	PyTorch
Cascaded Generation (130 GitHub stars)	PyTorch
Neural Linguistic Steganography (14,200 demo visits since 2025; 214 GitHub stars)	PyTorch
Variational Attention (327 GitHub stars)	PyTorch
Image-to-Markup Generation (1.3k GitHub stars)	LuaTorch
OpenNMT-py (7k GitHub stars)	PyTorch
OpenNMT (archived) (2.4k GitHub stars)	LuaTorch
Attention OCR (1.1k GitHub stars)	TensorFlow