

WENHU CHEN

Email: (hustchenwenhu@gmail.com, wenhuchen@uwaterloo.ca), **Website:** <https://wenhuchen.github.io/>

ABOUT ME

I'm a tenure-track assistant professor at Computer Science of University of Waterloo and Vector Institute since Aug 2022. I became a Canadian CIFAR AI chair since Jan 2023. I mainly work on different aspects of foundation models. My research can be summarized as the following aspects:

- 1) **large language models:** post-training, reasoning, coding, etc.
- 2) **multimodal language models:** post-training, multi-image/video understanding, etc.
- 3) **diffusion models:** text-to-image generation, image editing, text-to-video generation, video editing, etc.
- 4) **retrieval augmentation:** information retrieval, retrieval-augmented generation, etc.
- 5) **evaluation:** benchmarks, automatic metrics, etc.

WORKING EXPERIENCE

Vector Institute

Aug 15th, 2022-Now

Location: Toronto, Canada

Role: Faculty Member

Duties: 1. Conducting research in natural language processing. 2. Service in postdoc and faculty review committee. 3. Hosting research events in Vector Institute.

University of Waterloo

Aug 15th, 2022-Now

Location: Waterloo, Canada

Role: Assistant Professor in Cheriton School of Computer Science

Duties: 1. Leading research projects and supervising graduate students. 2. Teaching 3 units of courses per year. 3. Administrative service in the department (like graduate admission committee, etc).

Google Deepmind

Aug 15th, 2022-Now

Location: Remote

Role: Senior Research Scientist (20% Part-Time)

Manager: William W. Cohen & Ming-Wei Chang

Duties: Conducting Research. We aim at building more robust image generation and editing systems. We publish our research results on [CVPR 2024], [NeurIPS2 2023] and [ICLR 2023].

Google Research

Aug 16th, 2021-Aug 12th, 2022

Location: Pittsburgh, USA

Role: Research Scientist (Full time)

Manager: William W. Cohen

Duties: Conducting Research. We aim at integrating search results to enhance language generative models. We publishe our research results in [EACL 2023] and [EMNLP 2022].

EDUCATION

University of California, Santa Barbara

Sep 2017-Jun 2021

Location: Santa Barbara, USA

Ph.D in Computer Science

Advisor: William (Yang) Wang and Xifeng Yan

RWTH Aachen University, Germany

Sep 2014-Dec 2016

Location: Aachen, Germany

M.Sc. in Electronics Engineering

Huazhong University of Science and Technology, China

Sep 2010-June 2014

Location: Wuhan, China

B.Sc. in Electronics Engineering

INTERNSHIP EXPERIENCE

Google Research

June 15th, 2020 - Dec 18th, 2020

Location: Pittsburgh, PA, USA

Mentors: Ming-Wei Chang, Eva Schlinger, William W. Cohen

Duties: Work as research intern to explore research topics. We Build question answering system that can understand both structured and unstructured information. Our work was published in [ICLR 2021].

Microsoft AI & Cloud

June 17th, 2019 - Sep 6th, 2019

Location: Seattle, WA, USA

Mentors: Zhe Gan, Jingjing Liu

Duties: Work as research intern to explore research topics. We build explainable visual question answering system that can ground on images to answer simple questions. Our work was published in [WACV 2021].

Samsung Research America

June 18th, 2018 - Sep 14th, 2018

Location: Mountain View, CA, USA

Mentors: Yi-Lin Shen, Hongxia Jin

Duties: Work as research intern to explore research topics. We investigate different compression algorithm in natural language systems to save computation and storage. Our work was published in [NAACL 2019].

Microsoft Research Asia

Feb 13th, 2017 - Aug 11th, 2017

Location: Beijing, China

Mentors: Shujie Liu, Ming Zhou

Duties: Work as research intern to explore research topics. We enhance the machine translation systems by utilizing data augmentation strategies. Our work was published in [NAACL 2018].

eBay Research

Dec 7th, 2015 - May 27th, 2016

Location: Aachen, Germany

Mentors: Matusov Evgency, Shahram Khadivi

Duties: Work as research intern to explore research topics. We enhance machine translation with statistical alignment algorithm. Our work was published in [AMTA 2016].

TEACHING EXPERIENCE

University of Waterloo, 2023 Winter : CS 486/686 Introduction to Artificial Intelligence

University of Waterloo, 2024 Winter : CS 886 Recent Advances on Foundation Models

University of Waterloo, 2024 Spring : CS 486/686 Introduction to Artificial Intelligence

SELECTED PUBLICATIONS (AS FIRST & LAST AUTHOR)

[ICLR 2025] OmniEdit: Building Image Editing Generalist Models Through Specialist Supervision
Cong Wei, Zheyang Xiong, Weiming Ren, Xinrun Du, Ge Zhang, Wenhu Chen

[ICLR 2025] VLM2Vec: Training Vision-Language Models for Massive Multimodal Embedding Tasks
Ziyang Jiang, Rui Meng, Xinyi Yang, Semih Yavuz, Yingbo Zhou, Wenhu Chen

[ICLR 2025] MEGA-Bench: Scaling Multimodal Evaluation to over 500 Real-World Tasks
Jiacheng Chen, Tianhao Liang, Sherman Siu, Zhengqing Wang, Kai Wang, Yubo Wang, Yuansheng Ni, Wang Zhu, Ziyang Jiang, Bohan Lyu, Dongfu Jiang, Xuan He, Yuan Liu, Hexiang Hu, Xiang Yue, Wenhu Chen

[TMLR 2024] MANTIS: Interleaved Multi-Image Instruction Tuning
Dongfu Jiang, Xuan He, Huaye Zeng, Cong Wei, Max Ku, Qian Liu, Wenhu Chen

[TMLR 2024] AnyV2V: A Plug-and-Play Framework For Any Video-to-Video Editing Tasks
Max Ku, Cong Wei, Weiming Ren, Huan Yang, Wenhu Chen

(Reproducibility Certification)

[NeurIPS 2024] MMLU-Pro: A More Robust and Challenging Multi-Task Language Understanding Benchmark
Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj,

Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, Wenhui Chen
(Spotlight, Top 10% of accepted papers)

[NeurIPS 2024] GenAI Arena: An Open Evaluation Platform for Generative Models
Dongfu Jiang, Max Ku, Tianle Li, Yuansheng Ni, Shizhuo Sun, Rongqi Fan, Wenhui Chen

[NeurIPS 2024] MAMMO-TH2: Scaling Instructions from the Web
Xiang Yue, Tuney Zheng, Ge Zhang, Wenhui Chen

[EMNLP 2024] Unifying Multimodal Retrieval via Document Screenshot Embedding
Xueguang Ma, Sheng-Chieh Lin, Minghan Li, Wenhui Chen, Jimmy Lin

[EMNLP 2024] VideoScore: Building Automatic Metrics to Simulate Fine-grained Human Feedback for Video Generation
Xuan He, Dongfu Jiang, Ge Zhang, Max Ku, Achint Soni, Sherman Siu, Haonan Chen, Abhramil Chandra, Ziyang Jiang, Aaran Arulraj, Kai Wang, Quy Duc Do, Yuansheng Ni, Bohan Lyu, Yaswanth Narsupalli, Rongqi Fan, Zhiheng Lyu, Yuchen Lin, Wenhui Chen

[Findings of EMNLP 2024] Augmenting Black-box LLMs with Medical Textbooks for Clinical Question Answering
Yubo Wang, Xueguang Ma, Wenhui Chen

[COLM 2024] StructLM: Towards Building Generalist Models for Structured Knowledge Grounding
Alex Zhuang, Ge Zhang, Tianyu Zheng, Xinrun Du, Junjie Wang, Weiming Ren, Stephen W. Huang, Jie Fu, Xiang Yue, Wenhui Chen

[ACL 2024] VIEScore: Towards Explainable Metrics for Conditional Image Synthesis Evaluation
Max Ku, Dongfu Jiang, Cong Wei, Xiang Yue, Wenhui Chen

[Findings of ACL 2024] OpenCodeInterpreter: Integrating Code Generation with Execution and Refinement
Tianyu Zheng, Ge Zhang, Tianhao Shen, Xueling Liu, Bill Yuchen Lin, Jie Fu, Wenhui Chen, Xiang Yue

[TMLR 2024] TIGERScore: Towards Building Explainable Metric for All Text Generation Tasks
Dongfu Jiang*, Yishan Li*, Ge Zhang, Wenhao Huang, Bill Yuchen Lin, Wenhui Chen

[ECCV 2024] UniIR: Training and Benchmarking Universal Multimodal Information Retrievers
Cong Wei, Yang Chen, Haonan Chen, Hexiang Hu, Ge Zhang, Jie Fu, Alan Ritter, Wenhui Chen
(Oral, Top 5% of accepted papers)

[CVPR 2024] Instruct-Imagen: Image Generation with Multi-modal Instruction
Hexiang Hu*, Kelvin C.K. Chan*, Yu-Chuan Su*, Wenhui Chen*, Yandong Li, Kihyuk Sohn, Yang Zhao, Xue Ben, Boqing Gong, William Cohen, Ming-Wei Chang, Xuhui Jia
(Oral, Top 3% of accepted papers)

[CVPR 2024] MMMU: A Massive Multi-discipline Multimodal Understanding and Reasoning Benchmark for Expert AGI
Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, Wenhui Chen
(Best Paper Finalist, Top 1% of accepted papers)

[ICLR 2024] ImagenHub: Standardizing the evaluation of conditional image generation models
Max Ku, Tianle Li, Kai Zhang, Yujie Lu, Xingyu Fu, Wenwen Zhuang, Wenhui Chen

[ICLR 2024] MAMMO-TH: Building Math Generalist Models through Hybrid Instruction Tuning
Xiang Yue*, Xingwei Qu, Ge Zhang, Yao Fu, Wenhao Huang, Huan Sun, Yu Su, Wenhui Chen*
(Spotlight, Top 20% of accepted papers)

[WACV 2024] Synthesizing Coherent Story with Auto-Regressive Latent Diffusion Models
Xichen Pan, Pengda Qin, Yuhong Li, Hui Xue, Wenhui Chen
(Oral, Top 5% of accepted papers)

[TMLR 2023] DreamEdit: Subject-driven Image Editing
Tianle Li, Max Ku, Cong Wei, Wenhui Chen

[**TMLR 2023**] Program of Thoughts Prompting: Disentangling Computation from Reasoning for Numerical Reasoning Tasks

Wenhu Chen*, Xueguang Ma*, Xinyi Wang, William W. Cohen

[**EMNLP 2023**] TheoremQA: A Theorem-driven Question Answering dataset

Wenhu Chen, Ming Yin, Max Ku, Elaine Wan, Xueguang Ma, Jianyu Xu, Tony Xia, Xinyi Wang, Pan Lu

[**NeurIPS 2023**] Subject-driven Text-to-Image Generation via Apprenticeship Learning

Wenhu Chen, Hexiang Hu, Yandong Li, Nataniel Ruiz, Xuhui Jia, Ming-Wei Chang, William W. Cohen

[**ACL 2023**] Few-shot In-context Learning on Knowledge Base Question Answering

Tianle Li, Xueguang Ma, Alex Zhuang, Yu Gu, Yu Su and Wenhu Chen

[**Findings of EACL 2023**] Large Language Models are few(1)-shot Table Reasoners

Wenhu Chen

[**EACL 2023**] Augmenting Pre-trained Language Models with QA-Memory for Open-Domain Question Answering

Wenhu Chen, Pat Verga, Michiel de Jong, John Wieting, William Cohen

[**ICLR 2023**] Re-Imagen: Retrieval-Augmented Text-to-Image Generator

Wenhu Chen, Hexiang Hu, Chitwan Saharia, William W. Cohen

[**EMNLP 2022**] MuRAG: Multimodal Retrieval-Augmented Generator for Open Question Answering over Images and Text

Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, William W. Cohen

[**NeurIPS 2021**] A Dataset for Answering Time-Sensitive Questions

Wenhu Chen, Xinyi Wang, William Yang Wang

[**EMNLP 2020**] KGPT: Knowledge-Grounded Pre-Training for Data-to-Text Generation

Wenhu Chen, Yu Su, Xifeng Yan, William Yang Wang.

[**Findings of EMNLP 2020**] HybridQA: A Dataset of Multi-Hop Question Answering over Tabular and Textual Data

Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, William Wang.

[**WACV 2021**] Meta Module Network for Compositional Visual Reasoning

Wenhu Chen, Zhe Gan, Linjie Li, Yu Cheng, William Wang, Jingjing Liu.

(Best Student Paper Honorable Mention, Top 1% of accepted papers)

[**Findings of EMNLP 2020**] HybridQA: A Dataset of Multi-Hop Question Answering over Tabular and Textual Data

Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, William Wang.

[**ACL 2020**] Logical Natural Language Generation from Open-Domain Tables

Wenhu Chen, Jianshu Chen, Yu Su, Zhiyu Chen and William Yang Wang.

[**ICLR 2020**] TabFact: A Large-scale Dataset for Table-based Fact Verification

Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou and William Yang Wang.

[**ACL 2019**] Semantically Conditioned Dialog Response Generation via Hierarchical Disentangled Self-Attention

Wenhu Chen, Jianshu Chen, Pengda Qin, Xifeng Yan and William Yang Wang.

[**NAACL 2019**] How Large A Vocabulary Does Text Classification Need? A Variational Approach on Vocabulary Selection

Wenhu Chen, Yu Su, Yilin Shen, Zhiyu Chen, Xifeng Yan and William Yang Wang.

[**EMNLP 2018**] XL-NBT: A Cross-lingual Neural Belief Tracking Framework

Wenhu Chen, Jianshu Chen, Yu Su, Xin Wang, Dong Yu, Xifeng Yan and William Yang Wang.

[**ACL 2018**] No Metrics Are Perfect: Adversarial Reward Learning for Visual Storytelling

Wenhu Chen*, Xin Wang*, Yuan-Fang Wang and William Yang Wang.

[**NAACL 2018**] Variational Knowledge Graph Reasoning

Wenhu Chen, Wenhan Xiong, William Yang Wang, Xifeng Yan.

[**NAACL 2018**] Generative Bridging Network in Neural Sequence Prediction

Wenhu Chen, Guanlin Li, Shuo Ren, Shujie Liu, Zhirui Zhang, Mu Li, Ming Zhou.

FUNDINGS

Waterloo Starter Faculty Fund: 2022 - 2027
CIFAR AI Chair Fund: 2022 - 2027
NSERC Discovery Fund: 2023 - 2028
Mitacs Accelerate Fund: 2023 - 2025
CIFAR AI Catalyst Fund: 2024 - 2026
Canada National Research Council Fund: 2024 - 2026

SERVICE

Program Committee: NAACL, ACL, EMNLP, ICML, NeurIPS, ICLR
Conference Session Chair: EMNLP 2020, NAACL 2021
Area Chair: AAAI 2022, AACL 2022, EMNLP 2022, EACL 2023, NLPCC 2023, EMNLP 2023, AACL 2023, EACL 2024, EMNLP 2024

STUDENTS

Max Ku (PhD, University of Waterloo)
Dongfu Jiang (PhD, University of Waterloo)
Cong Wei (PhD, University of Waterloo)
Weiming Ren (PhD, University of Waterloo)
Benjamin Schneider (MMATH, University of Waterloo) Yuansheng Ni (MMATH, University of Waterloo)
Zhiheng Lyu (MMATH, University of Waterloo)

MENTEES

Shahan Nedadahandeh (CS Undergrad, University of Waterloo)
Ritvik Bale (CS Undergrad, University of Waterloo)
Aaran Arulraj (CS Undergrad, University of Waterloo)
Sahel Sharifmoghaddam (CS Undergrad, University of Waterloo)

GRADUATED

Tianle Li (Master's, University of Waterloo)
Max Ku (Master's, University of Waterloo)

ACADEMIC AWARD

2024 CVPR Best Paper Finalist.
2023 AACL-IJCNLP Area Chair Award.
2022 CIFAR AI Chair.
2021 WACV Best Student Paper Honorable Mention.
2019 NeurIPS Outstanding Reviewer Award.