

YIFAN HOU

yifan.hou@inf.ethz.ch | [Homepage](#) | [Google Scholar](#) | [GitHub](#) | [LinkedIn](#)

Research Profile

Research focus: reasoning in language models—its mechanisms, grounding, and reliability.

I study when language models reason from grounded evidence rather than shortcuts, how they represent and compose that evidence, and where these processes fail internally. I develop behavioral diagnostics, mechanistic analyses, and training and inference-time interventions for more reliable and interpretable reasoning. Multimodal perception provides a controlled testbed for core NLP questions about grounding, compositionality, and generalization.

Education

ETH Zürich Ph.D. in Computer Science; thesis defended Aug. 2026 Advisors: Prof. Mrinmaya Sachan and Prof. Antoine Bosselut	Oct. 2020–Sep. 2026
The Chinese University of Hong Kong (CUHK) M.Phil. in Computer Science and Engineering Advisors: Prof. James Cheng and Prof. Ming-Chang Yang	Sep. 2018–Jun. 2020
Huazhong University of Science and Technology (HUST) B.Eng. in Information Science, Advanced Class in Mathematics and Physics	Sep. 2014–Jun. 2018
University of Illinois Urbana-Champaign (UIUC) Summer Exchange	Jul.–Aug. 2017

Academic Research Experience

ETH Zürich — Doctoral Researcher	Sep. 2020–Sep. 2026
EPFL — Visiting Researcher Host: Prof. Antoine Bosselut	Jun.–Sep. 2022; Jun.–Sep. 2023
Toyota Technological Institute at Chicago (TTIC) — Remote Research Intern Host: Prof. Mrinmaya Sachan	Jun.–Sep. 2020
National University of Singapore (NUS) — Visiting Researcher Host: Prof. Richard T. B. Ma	Jun.–Sep. 2019

Industry Research Experience

Meta AI — Research Intern	Oct. 2024–Feb. 2025
Tencent — Research Intern	Jun.–Sep. 2021

Publications

*Notation: * equal contribution; underlined names denote students I mentored.*

Peer-Reviewed Publications

- [1] *KnowVis: Knowledge-Centric Visual Summarization for Video Lectures.*
Yi Xu, **Yifan Hou**, Xiaoyu Zhang.
Findings of EMNLP, 2026.
- [2] *Unveiling the Visual Counting Bottleneck in Vision-Language Models.*
Xingzhou Pang*, **Yifan Hou***, Junling Wang, Mrinmaya Sachan.
ICML, 2026.
- [3] *Diversity Matters: Revisiting Test-Time Compute in Vision-Language Models.*
Yijie Tong*, **Yifan Hou***, Shaobo Cui, Antoine Bosselut, Mrinmaya Sachan.
ICML, 2026.

- [4] *Compose and Fuse: Revisiting the Foundational Bottlenecks in Multimodal Reasoning.*
Yucheng Wang*, **Yifan Hou***, Aydin Javadov, Mubashara Akhtar, Mrinmaya Sachan.
ICLR, 2026.
- [5] *Can Vision-Language Models Solve Visual Math Equations?*
Monjoy Narayan Choudhury*, Junling Wang*, **Yifan Hou**, Mrinmaya Sachan.
EMNLP, 2025.
- [6] *HealthCards: Exploring Text-to-Image Generation as Visual Aids for Healthcare Knowledge Democratizing and Education.*
Qian Wu, Zheyao Gao, Longfei Gou, **Yifan Hou**, Ann Sin Nga Lau, Qi Dou.
EMNLP, 2025.
- [7] *Do Vision-Language Models Really Understand Visual Language?*
Yifan Hou, Buse Giledereli, Yilei Tu, Mrinmaya Sachan.
ICML, 2025.
- [8] *What Do Language Models Learn in Context? The Structured Task Hypothesis.*
Jiaoda Li*, **Yifan Hou***, Mrinmaya Sachan, Ryan Cotterell.
ACL, 2024.
- [9] *Unveiling the Art of Heading Design: A Harmonious Blend of Summarization, Neology, and Algorithm.*
Shaobo Cui*, Yiyang Feng*, Yisong Mao, **Yifan Hou**, Boi Faltings.
Findings of ACL, 2024.
- [10] *Towards a Mechanistic Interpretation of Multi-Step Reasoning Capabilities of Language Models.*
Yifan Hou, Jiaoda Li, Yu Fei, Alessandro Stolfo, Wangchunshu Zhou, Guangtao Zeng, Antoine Bosselut, Mrinmaya Sachan.
EMNLP, 2023.
- [11] *Mitigating Label Biases for In-context Learning.*
Yu Fei, **Yifan Hou**, Zeming Chen, Antoine Bosselut.
ACL, 2023.
- [12] *Adapters for Enhanced Modeling of Multilingual Knowledge and Text.*
Yifan Hou, Wenxiang Jiao, Meizhen Liu, Carl Allen, Zhaopeng Tu, Mrinmaya Sachan.
Findings of EMNLP, 2022.
- [13] *What Has Been Enhanced in my Knowledge-Enhanced Language Model?*
Yifan Hou, Guoji Fu, Mrinmaya Sachan.
Findings of EMNLP, 2022.
- [14] *Bird’s Eye: Probing for Linguistic Graph Structures with a Simple Information-Theoretic Approach.*
Yifan Hou, Mrinmaya Sachan.
ACL, 2021.
- [15] *Measuring and Improving the Use of Graph Information in Graph Neural Networks.*
Yifan Hou, Jian Zhang, James Cheng, Kaili Ma, Richard T. B. Ma, Hongzhi Chen, Ming-Chang Yang.
ICLR, 2020.
- [16] *A Representation Learning Framework for Property Graphs.*
Yifan Hou, Hongzhi Chen, Changji Li, James Cheng, Ming-Chang Yang.
KDD, 2019.
- [17] *Grasper: A High Performance Distributed System for OLAP on Property Graphs.*
Hongzhi Chen, Changji Li, Juncheng Fang, Chenghuan Huang, James Cheng, Jie Zhang, **Yifan Hou**, Xiao Yan.
SoCC, 2019.
- [18] *Large Scale Graph Mining with G-Miner.*
Hongzhi Chen, Xiaoxi Wang, Chenghuan Huang, Juncheng Fang, **Yifan Hou**, Changji Li, James Cheng.
SIGMOD (Demo), 2019.
- [19] *Course Recommendation of MOOC with Big Data Support: A Contextual Online Learning Approach.*
Yifan Hou, Pan Zhou, Jie Xu, Dapeng Oliver Wu.
IEEE INFOCOM Workshop, 2018.

Manuscripts and arXiv Preprints

- [M1] *Stage-Wise Visual Component Decomposition for Diagram Understanding Enhancement.*
Junling Wang, **Yifan Hou**, Mrinmaya Sachan.
Manuscript, 2026.

- [M2] *LLMs Find Their Way in Multi-Turn Conversation.*
Xingzhou Pang, Jiaoda Li, **Yifan Hou**, Ryan Cotterell, Mrinmaya Sachan.
Manuscript, 2026.
- [M3] *Chimera: Diagnosing Shortcut Learning in Visual-Language Understanding.*
Ziheng Chi*, **Yifan Hou***, Chenxi Pang, Shaobo Cui, Mubashara Akhtar, Mrinmaya Sachan.
Manuscript, 2026.
- [M4] *RecurrentGPT: Interactive Generation of (Arbitrarily) Long Text.*
Wangchunshu Zhou, Yuchen Eleanor Jiang, Peng Cui, Tiannan Wang, Zhenxin Xiao, **Yifan Hou**, Ryan Cotterell, Mrinmaya Sachan.
arXiv, 2023.
- [M5] *Understanding Graph Neural Networks from Graph Signal Denoising Perspectives.*
Guoji Fu, **Yifan Hou**, Jian Zhang, Kaili Ma, Barakeel Fansu Kamhoua, James Cheng.
arXiv, 2020.

Selected Awards

- | | |
|---|------|
| • Best Paper Award, Multilingual Representation Learning Workshop at EMNLP | 2022 |
| • Swiss Data Science Center Ph.D. Fellowship | 2022 |
| • Outstanding Reviewer, ICML | 2022 |
| • Faculty Outstanding M.Phil. Thesis Award, Faculty of Engineering, CUHK (sole recipient) | 2020 |
| • KDD Student Travel Award | 2019 |
| • Outstanding Graduate Award, HUST | 2018 |
| • National Scholarship | 2016 |

Student Mentoring

Yi Xu	Research project	2025–present	Ph.D., CityU of Hong Kong
Chengxi Li	Research project	2025–present	B.Sc., HKUST
Xingzhou Pang	Master’s thesis	2025–present	M.Sc., ETH Zürich
Yucheng Wang	Research project	2025–2026	M.Sc., ETH Zürich
Yijie Tong	Master’s thesis	2025–2026	M.Sc., University of Zürich
Ziheng Chi	Master’s thesis	2024–2025	ETH Zürich → Tencent
Nayanika Debnath	Research project	2024–2025	ETH Zürich → IBM
Monjoy Narayan Choudhury	Research project	2024–2025	IIIT Bangalore → SE, Google
Yilei Tu	Research project	2023–present	ETH Zürich → Ph.D., UBC
Buse Giledereli	Master’s thesis	2023–2024	ETH Zürich → Data Scientist, Prometeia
Yu Fei	Research project	2022–2023	ETH Zürich → Ph.D., UCI

Teaching Experience

- | | |
|--|-------------------|
| • TA , Large-Scale AI Engineering (LSAIE), ETH Zürich | Autumn 2026 |
| • Head TA , Computational Semantics for Natural Language Processing, ETH Zürich | Spring 2021–2025 |
| • TA , Large Language Models, ETH Zürich | Spring 2023, 2026 |
| • TA , Deep Learning, ETH Zürich | Autumn 2025 |
| • TA , Introduction to Social Networks, CUHK | Spring 2020 |

Academic Service

Area Chair: NeurIPS; ACL Rolling Review (ARR).
Reviewer: ICML, ICLR, and COLM.