

Education

University of California, Los Angeles	LOS ANGELES, CA, US
Ph.D degree in Computer Science (Data Mining and Natural Language Processing)	2017.9 – 2021.7
University of California, Los Angeles	LOS ANGELES, CA, US
MS degree in Electrical Engineering (Signals & Systems)	2015.9 – 2017.6
Southeast University	NANJING, JIANGSU, P.R.CHINA
BS degree in Electronics Engineering	2011.8 – 2015.6

Experiences

Google Research	MOUNTAIN VIEW, CA, US
Senior Software Engineer (ML Research)	2021.8 – Now
Working on LLM-based structured summarization and comparison solutions.	
Google Research	MOUNTAIN VIEW, CA, US
Research Intern	2020 Summer
Working on information extraction from semi-structured data (web page; email).	
YAHOO Research	SUNNYVALE, CA, US
Research Intern	2018, 2019 Summer
Working on ads targeting problems such as understanding user activity trails and ad advertisement theme recommendation.	
HELLA Co.	LIPPSTADT, NRW, GERMANY
Optical Design Intern for Undergraduate Thesis	2015 Winter and Spring
Working on high-resolution automotive headlight design and test.	

Publications and Patents

Publications

740+ Google Scholar citations in total and published at ACL, KDD, AAAI, WWW, EMNLP, CIKM, ICDE, etc.

1. **Yichao Zhou**, James B. Wendt, Navneet Potti, Jing Xie, Sandeep Tata, Selective Labeling: How to Radically Lower Data-Labeling Costs for Document Extraction Models, *The 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP2023*, 6-10 December 2023, Singapore.
2. Zilong Wang, **Yichao Zhou**, Wei Wei, Chen-Yu Lee, Sandeep Tata, A Benchmark for Structured Extractions from Complex Documents, *SIGKDD Conference on Knowledge Discovery and Data Mining, KDD2023*, 6-10 August, 2023, Long Beach, California, USA.
3. Xiusi Chen, Jyun-Yu Jiang, Kun Jin, **Yichao Zhou**, Mingyan Liu, Jefferey Brantingham, Wei Wang, ReLiable: Offline Reinforcement Learning for Tactical Strategies in Professional Basketball Games, *The 31th ACM International Conference on Information & Knowledge Management, CIKM2022*, 17-21 October 2022, Atlanta, Georgia, USA.
4. **Yichao Zhou**, Ying Sheng, Nguyen Vo, Nick Edmonds, Sandeep Tata, Learning Transferable Node Representations for Attribute Extraction from Web Documents, *The Fifteenth International Conference on Web Search and Data Mining, WSDM2022*, 21-25 February 2022, Phoenix, Arizona, USA.
5. Jyun-Yu Jiang, **Yichao Zhou**, Xiusi Chen, Yan-Ru Jhou, Liqi Zhao, Sabrina Liu, Po-Chun Yang, Jule Ahmar, Wei Wang, COVID-19 Surveiller: Toward a Robust and Effective Pandemic Surveillance System based on Social Media Mining, *Philosophical Transactions A*, issue: Data Science Approach to Infectious Diseases Surveillance, November 2021.
6. **Yichao Zhou**, Learning Robust Representations for Low-resource Information Extraction, *Doctoral Dissertation 2021*.

7. **Yichao Zhou**, Jyun-yu Jiang, Xiusi Chen, Wei Wang, #StayHome or #Marathon? Social Media Enhanced Pandemic Surveillance on Spatial-temporal Graphs, *The 30th ACM International Conference on Information & Knowledge Management, CIKM2021*, 1-5 November 2021, Online.
8. **Yichao Zhou**, Yu Yan, J. Harry Caufield, Kai-Wei Chang, Yizhou Sun, PeiPei Ping, Wei Wang, Clinical Temporal Relation Extraction with Probabilistic Soft Logic Regularization and Global Inference, *The 35th AAAI Conference on Artificial Intelligence, AAAI2021*, 2-9 February 2021, Online.
9. **Yichao Zhou**, Wei-Ting Chen, Bowen Zhang, David Lee, J. Harry Caufield, Kai-Wei Chang, Yizhou Sun, PeiPei Ping, Wei Wang, CREATE: Clinical Report Extraction and Annotation Technology, *IEEE 37th International Conference on Data Engineering, (Demo), ICDE2021*, 19-22 April 2021, Online.
10. Rujun Han, **Yichao Zhou**, Nanyun Peng, Domain Knowledge Empowered Structured Neural Net for End-to-End Event Temporal Relation Extraction, *The 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP2020*, 16-20 November 2020, Online.
11. **Yichao Zhou**, Jyun-yu Jiang, Jieyu Zhao, Kai-Wei Chang, Wei Wang, "The Boating Store Had Its Best Sail Ever": Pronunciation-attentive Contextualized Pun Recognition", *2020 Annual Conference of the Association for Computational Linguistics, ACL2020*, 5-10 July 2020, Online.
12. **Yichao Zhou**, Shaunak Mishra, Manisha Verma, Narayan Bhamidipati, Wei Wang, "Recommending Themes for Ad Creative Design via Visual-Linguistic Representations", *The Web Conference, (Short), WWW2020*, 20-24 April 2020, Online.
13. Shaunak Mishra, Manisha Verma, **Yichao Zhou**, Kapil Thadani, Wei Wang, Learning to Create Better Ads: Generation and Ranking Approaches for Ad Creative Refinement, *The 29th ACM International Conference on Information & Knowledge Management, CIKM2020*, 19-23 October 2020, Online.
14. Cheng Zheng, Jyun-Yu Jiang, **Yichao Zhou**, Sean D. Young, Wei Wang, "Social Media User Geolocation via Hybrid Attention", *The 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR2020*, 25-30 July 2020, Online.
15. J. Harry Caufield, **Yichao Zhou**, Yunsheng Bai, David Liem, Garlid, A., Kai-Wei Chang, Yizhou Sun, Wei Wang, Peipei Ping, "Resources for comprehensive information extraction and data integration in biomedicine", *28th Conference on Intelligent Systems for Molecular Biology, (Poster), ISMB2020*, 13-16 July 2020, Online.
16. **Yichao Zhou**, Shaunak Mishra, Jelena Gligorijevic, T. Bhatia, Narayan Bhamidipati, "Understanding Consumer Journey using Attention based Recurrent Neural Networks", *SIGKDD Conference on Knowledge Discovery and Data Mining, KDD2019*, 4-8 August 2019, Anchorage, Alaska, USA.
17. **Yichao Zhou***, Jyun-yu Jiang*, Kai-Wei Chang, Wei Wang, "Learning to Discriminate Perturbations for Blocking Adversarial Attacks in Text Classification", *2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP2019*, 3-7 November 2019 Hong Kong, China. (* equal contribution)
18. **Yichao Zhou**, J. Harry Caufield, Yizhou Sun, Peipei Ping, David Liem, D. Sigdel, Kai-Wei Chang, Wei Wang, "Learning Structured Knowledge from Clinical Case Reports", *27th Conference on Intelligent Systems for Molecular Biology and 18th European Conference on Computational Biology, (Poster), ISMB/ECCB2019*, 21-25 July 2019, Basel, Switzerland.
19. J. Harry Caufield, **Yichao Zhou**, Yunsheng Bai, David Liem, Anders O. Garlid, Kai-Wei Chang, Yizhou Sun, Peipei Ping, Wei Wang, A Comprehensive Typing System for Information Extraction from Clinical Narratives, (Pre-print), *medRxiv2019*.
20. **Yichao Zhou**, Chelsea Ju, J. Harry Caufield, Kevin Shih, Calvin Chen, Yizhou Sun, Kai-Wei Chang, Peipei Ping, Wei Wang. Clinical Named Entity Recognition using Contextualized Token Representations, (Pre-print), *arXiv2019*.
21. Jieyu Zhao, **Yichao Zhou**, Zeyu Li, Wei Wang, Kai-Wei Chang, "Learning Gender-Neutral Word Embedding", *Conference on Empirical Methods in Natural Language Processing, (Short), EMNLP2018*, October 31–November 4 2018, Brussels, Belgium.
22. Tan Patrick, **Yichao Zhou**, Xinxin Huang, Wei Wang, "AZTEC: A Cloud-based Computational Platform to Integrate Biomedical Resources", *IEEE 33rd International Conference on Data Engineering, (Demo), ICDE17*, 19-22 April 2017, San Diego, California, USA.

23. **Yichao Zhou**, Zai-Fa Zhou, Nihuan Wang, "A novel residual stress test structure for MEMs thin films", *12th IEEE International Conference on Solid-State and Integrated Circuit Technology, ICSICT2014*, 28-31 October 2014, Gulin, China.

Patents

1. S. Mishra, R.H Tang, **Yichao Zhou**, V.R.Raithatha, L.N. Buamidipati, M. Abasi, S. Ravindran, "Stage-based content item selection and transmission", 2020
-

Professional Activities

- Conference Chair
 1. NAACL 2021 Publication Chair (*Annual Conference of the North American Chapter of the Association for Computational Linguistics*)
 2. KDD 2021 Session Chair (*SIGKDD Conference on Knowledge Discovery and Data Mining*)
 - Program Committee Member
 1. ICML 2023 (*International Conference on Machine Learning*)
 2. EMNLP 2023 (*Empirical Methods in Natural Language Processing*)
 3. WWW 2022 (*The Web Conference*)
 4. ACL 2023 2022 Rolling Review
 5. AAAI 2022 2021 (*AAAI Conference on Artificial Intelligence*)
 6. CIKM 2022 2021 (*ACM International Conference on Information and Knowledge Management*)
 7. AACL 2023 2022 2020 (*Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*)
 8. KDD 2022 Workshop on Content Understanding and Generation for E-commerce
 9. KDD 2021 2020 (*SIGKDD Conference on Knowledge Discovery and Data Mining*)
 10. ACL 2021 (*Annual Conference of the Association for Computational Linguistics*)
 - Journal Reviewer
 1. Big Data 2021
 2. TKDE 2019 (*Transactions on Knowledge and Data Engineering*)
 - Conference Sub-Reviewer:
 1. ICDM 2019 2017 (*IEEE International Conference on Data Mining*)
 2. HealthNLP 2018 (*International Workshop on Health Natural Language Processing*)
 3. KDD 2017 (*SIGKDD Conference on Knowledge Discovery and Data Mining*)
-

Honors, Awards, and Scholarships

- Student Volunteer/Travel Award
 1. EMNLP 2019 (*Conference on Empirical Methods in Natural Language Processing*)
 2. SIGKDD 2019 (*SIGKDD Conference on Knowledge Discovery and Data Mining*)
 - UCLA Teaching Fellow, 2021
 - UCLA Graduate Division Fellowships, 2018 2019 2020
 - Excellent Undergraduate Thesis (2.7%), Southeast University, 2015
-

Teaching Assistant Experiences

Taught 4 courses in 7 quarters and promoted as a teaching fellow (top title) at UCLA.

1. UCLA 2021 Spring & 2018 Winter, Prof. Wei Wang, **Data Mining**
 - This course introduces concepts, algorithms, and techniques of data mining on different types of datasets, which covers basic data mining algorithms, as well as advanced topics on text mining, recommendation systems, and graph/network mining.
 - This course guides students to understand, develop, and use data mining systems to analyze large amounts of data.
2. UCLA 2020 Winter & Spring, Prof. David Smallberg, **Data structures and Algorithms** (C++)
 - This course teaches Object-oriented software development and abstract data type definition and use.
 - This course introduces Object-oriented view of data structures: stacks, queues, lists and algorithm analysis (e.g. Trees, graphs, and associated algorithms. Searching and sorting). Case studies and exercises from computer science applications.
3. UCLA 2019 & 2020 Fall, Prof. Todd Millstein, **Principles & Practices of Computing** (Python)
 - This course teaches students how to use computers as a tool for problem solving, creativity, and exploration via the design and implementation of computer programs.
 - This course covers key concepts including functional decomposition, usage of common data types, and usage of common control structures.
4. UCLA 2018 Fall, Prof. Kai-Wei Chang, **Introduction to Machine Learning**
 - This class introduces the fundamental concepts and algorithms in machine learning (supervised/unsupervised learning) and best practices in applying machine learning to practical problems.
 - This class covers a number of problem domains including computer vision, natural language processing, computational biology and robotics.