

Zhou Yu

Email: zy2461@columbia.edu
Cell: 412-944-4815
Addr: Schapiro CEPSR 723
530 West 120th Street, New York, NY 10027
Website: <https://www.cs.columbia.edu/~zhouyu/>

Professional Experience

2021/01- Assistant Professor in Computer Science Department, Columbia University
2017/07 – 2020/12 Assistant Professor in Computer Science Department, University of California, Davis

Education

2011/08 – 2017/05 Ph.D. in School of Computer Science, Carnegie Mellon University
2007/09 – 2011/06 Honors Program, Chu Kochen Honors College, Zhejiang University
Dual Degree: Bachelor of Science in Computer Science and Arts in Linguistics

Honors and Awards

2020/06 Sony Faculty Award
2019/10 Google GCP Credit Award
2019/08 **Best paper honorable mention in ACL 2019**
2019/05 Best GGCS Research Forum Presentation, UC Davis
2018/11 **Amazon Alexa Prize Champion with \$500,000 Award**
2018/04 Cisco Faculty Research Gift
2017/12 & 2019/08 Bosch Faculty Research Award
2017/11 **2018 Forbes 30 Under 30 in Science**
2017/11 & 2019/05 Tencent Faculty Research Award
2017/11 & 2018/11 Intel Faculty Research Award
2015/12 Rising Stars in EECS
2012/09 Quality of Life Fellowship
2010/12 Best Paper Award in International Conferences on Internet Multimedia Computing and Service
2010/05 Microsoft Young Fellowship from MSRA

Research Interests

Speech and Natural Language Processing

[\[AAAI 2020a\]](#)[\[AAAI 2020b\]](#)[\[AAAI 2020c\]](#)[\[AAAI 2020d\]](#)[\[AAAI 2020e\]](#)[\[EMNLP 2019a\]](#)[\[EMNLP 2019b\]](#)[\[EMNLP 2019c\]](#)[\[EMNLP 2019d\]](#)[\[SIGDIAL 2019\]](#)[\[ACL 2019a\]](#)[\[ACL 2019b\]](#)[\[NAACL 2019\]](#)[\[AAAI 2019\]](#)[\[EMNLP 2018a\]](#)[\[ACL 2018\]](#)[\[SIGDIAL 2018\]](#)[\[ICASSP 2018\]](#)[\[IWSIDS 2017\]](#)[\[IJCAI 2017\]](#)[\[SIGDIAL 2016a\]](#)[\[SIGDIAL 2016b\]](#)[\[IVA 2016b\]](#)[\[LREC 2016a\]](#)[\[LREC 2016b\]](#)[\[IWSIDS 2016a\]](#)[\[IWSIDS 2016b\]](#)[\[SIGDIAL 2015\]](#)[\[AAAI Symposium 2015\]](#)[\[ASRU 2015\]](#)

Multimodal Machine Learning

[\[EMNLP 2019a\]](#)[\[SIGDIAL2018\]](#)[\[EMNLP 2018a\]](#)[\[EMNLP 2018b\]](#)[\[IWSIDS 2016a\]](#)[\[IVA2016a\]](#)[\[SIGDIAL2015\]](#)[\[SIGDIAL2013\]](#)[\[SEMDIAL 2013\]](#)

Human-Computer Interaction

[\[CHI 2020\]](#)[\[IMWUT 2018 \(UbiComp 2019\)\]](#)[\[IVA 2016a\]](#)[\[HRI 2016\]](#)

Publications

Book Chapters

Zhou Yu, Vikram Ramanarayanan, Robert Mundkowsky, Patrick Lange, Alan Black, Alexei Ivanov, David Suendermann-Oeft, Multimodal HALEF: An Open-Source Modular Web-Based Multimodal Dialog Framework, in: Dialogues with Social Robots, Springer, 2017

Vikram Ramanarayanan, David Suendermann-Oeft, Patrick Lange, Robert Mundkowsky, Alexei V. Ivanov, [Zhou Yu](#), Yao Qian and Keelan Evanini (2016, in press), Assembling the jigsaw: How multiple open standards are synergistically combined in the HALEF multimodal dialog system, in: Multimodal Interaction with W3C Standards: Towards Natural User Interfaces to Everything, D. A. Dahl, Ed., ed New York: Springer, 2016.

Peer-reviewed Conferences and Journals

2021

Derek Chen, Howard Chen, Yi Yang, Alexander Lin and [Zhou Yu](#), Action-Based Conversations Dataset: A Corpus for Building More In-Depth Task-Oriented Dialogue Systems, NAACL 2021

Zhaojiang Lin, Bing Liu, Seungwhan Moon, Paul Crook, Zhenpeng Zhou, Zhiguang Wang, [Zhou Yu](#), Andrea Madotto, Eunjoon Cho, Rajen Subba, Leveraging Slot Descriptions for Zero-Shot Cross-Domain Dialogue State Tracking, NAACL 2021

Mingyang Zhou, Luwei Zhou, Yu Cheng, Linjie Li, [Zhou Yu](#), Jingjing Liu, CVPR 2021

Qingyang Wu, Lei Li and [Zhou Yu](#), TextGAIL: Generative Adversarial Imitation Learning for Text Generation, AAAI 2021

Jing Gu, Qingyang Wu and [Zhou Yu](#), Perception Score: A Learned Metric for Open-ended Text Generation Evaluation, AAAI 2021

Kun Qian, Wei Wei and [Zhou Yu](#), A Student-Teacher Architecture for Dialog Domain Adaptation under the Meta-Learning Setting, AAAI 2021

Jing Gu, Mostafa Mirshekari, [Zhou Yu](#) and Aaron Sisto, ChainCQG: Flow-Aware Conversational Question Generation, EACL 2021

Qingyang Wu, Yichi Zhang, Yu Li and [Zhou Yu](#), Alternating Recurrent Dialog Model with Large-scale Pre-trained Language Models, EACL 2021

Dian Yu and [Zhou Yu](#), MIDAS: A Dialog Act Annotation Scheme for Open Domain Human Machine Spoken Conversations, EACL 2021

2020

David Gros, Hariharan Sezhiyan, Prem Devanbu, [Zhou Yu](#), Code to Comment "Translation": Data, Metrics, Baseline & Evaluation, ASE 2020

Liang Qiu, Yizhou Zhao, Weiyan Shi, Yuan Liang, Feng Shi, Tao Yuan, [Zhou Yu](#), Song-Chun Zhu, Structured Attention for Unsupervised Dialogue Structure Induction, EMNLP 2020

Shirley Anugrah Hayati, Dongyeop Kang, Qingxiaoyang Zhu, Weiyan Shi and [Zhou Yu](#), INSPIRED: Toward Sociable Recommendation Dialog Systems. EMNLP 2020

Weixin Liang, James Zou and [Zhou Yu](#), ALICE: Active Learning with Contrastive Natural Language Explanations, EMNLP 2020

Weixin Liang, James Zou and [Zhou Yu](#), Beyond User Self-Reported Likert Scale Ratings: A Comparison Model for Automatic Dialog Evaluation, ACL 2020

Silin Gao, Yichi Zhang, Zhijian Ou and [Zhou Yu](#), Paraphrase Augmented Task-Oriented Dialog Generation, ACL 2020

Jiaying Hu, Yan Yang, Chencai Chen, Liang He and [Zhou Yu](#), SAS: Dialogue State Tracking via Slot Attention and Slot Information Sharing, ACL 2020

Yiheng Zhou, Yulia Tsvetkov, Alan Black, and [Zhou Yu](#), Augmenting Non-Collaborative Dialog Systems with Explicit Semantic and Strategic Dialog History, ICLR 2020

Weiyan Shi, Yoo Jung Oh, Xuwei Wang, Saurav Sahay, Jingwen Zhang, and [Zhou Yu](#), Effects of Persuasive Dialogues: Testing Bot Identities and Inquiry Strategies, CHI 2020

Qingyang Wu, Lei Li, Hao Zhou, Ying Zeng and Zhou Yu, Importance-Aware Learning for Neural Headline Editing, AAAI 2020

Weixin Liang, Youzhi Tian, Chengcai Chen, and Zhou Yu, MOSS: End-to-End Dialog System Framework with Modular Supervision, AAAI 2020

Xiyuan Zhang, Chenxi Li, Sam Davidson, Dian Yu and Zhou Yu, Filling Conversation Ellipsis for Better Social Dialog Understanding, AAAI 2020

Yu Li, Kun Qian, Weiyang Shi and Zhou Yu, End-to-End Trainable Non-Collaborative Dialog System, AAAI 2020

Yichi Zhang, Zhijian Ou and Zhou Yu, Improving Task-Oriented Dialog Systems by Considering Multiple Appropriate Responses with Different Dialog Policies, AAAI 2020

2019

Yuan Lin & Zhou Yu, Abstractive Dialog Summarization with Semantic Scaffolds, arXiv, 2019

Qingyang Wu, Yichi Zhang, Yu Li, and Zhou Yu, Alternating Recurrent Dialog Model with Large-scale Pre-trained Language Models, arXiv, 2019

Dian Yu & Zhou Yu, MIDAS: A Dialog Act Annotation Scheme for Open Domain Human Machine Spoken Conversations, arXiv, 2019

Mingyang Zhou, Joshua Arnold and Zhou Yu, Building Task-Oriented Visual Dialog Systems Through Alternative Optimization Between Dialog Policy and Language Generation, EMNLP 2019

Weiyang Shi, Kun Qian, Xuewei Wang and Zhou Yu, How to Build User Simulators to Train RL-based Dialog Systems, EMNLP 2019

Sam Davidson, Dian Yu and Zhou Yu, Dependency Parsing for Spoken Dialog Systems, EMNLP 2019

Michelle Cohn, Chun-Yen Chen and Zhou Yu, A Large-Scale User Study of an Alexa Prize Chatbot: Effect of TTS Dynamism on Perceived Quality of Social Dialog, SIGDIAL 2019

Kun Qian & Zhou Yu, Domain Adaptive Dialog Generation via Meta Learning, ACL 2019

Xuewei Wang, Weiyang Shi, Yoo Jung Oh, Jingwen Zhang, and Zhou Yu, Persuasion for Good: Towards A Personalized Persuasive Dialogue System for Social Good, ACL 2019 ***Best Paper honorable mention**

Weiyang Shi, Tiancheng Zhao and Zhou Yu, Unsupervised Dialog Structure Learning, NAACL 2019

Jiaao Chen, Jianshu Chen and Zhou Yu, Incorporating Structured Commonsense Knowledge in Story Completion, AAAI 2019

2018

Mingyang Zhou, Runxiang Cheng, Yong Jae Lee and Zhou Yu, A Visual Attention Grounding Neural Model for Multimodal Machine Translation, EMNLP 2018

Weiming Wen, Songwen Su and Zhou Yu, Cross-Lingual Cross-Platform Rumor Verification Pivoting on Multimedia Content, EMNLP 2018

Zijie Zhu, Xuewei Wang, Aakaash Kapoor, Tingrui Pan and Zhou Yu, EIS: A Wearable Device for Epidermal American Sign Language Recognition, IMWUT 2018 (UbiComp 2019)

Weiyang Shi and Zhou Yu, Sentiment Adaptive End-to-End Dialog Systems, ACL 2018

Jiaping Zhang, Tiancheng Zhao and Zhou Yu, Multimodal Hierarchical Reinforcement Learning Policy for Task-Oriented Visual Dialog, SIGDIAL 2018

Neville Ryant, Erika Bergelson, Kenneth Church, Alejandrina Cristia, Jun Du, Sriram Ganapathy, Sanjeev Khudanpur, Diana Kowalski, Mahesh Krishnamoorthy, Rajat Kulshreshta, Mark Liberman, Yu-Ding Lu, Matthew Maciejewski, Florian Metzger, Jan Profant, Lei Sun, Yu Tsao, [Zhou Yu](#), Enhancement and Analysis of Conversational Speech, ICASSP 2018

2017

[Zhou Yu](#), Alan W Black and Alexander I. Rudnicky, Learning Conversational Systems that Interleave Task and Non-Task Content, IJCAI 2017

[Zhou Yu](#), Vikram Ramanarayanan, Patrick Lange and David Suendermann-Oeft, "An open-source multimodal dialog system with real-time engagement tracking for job interview training applications," in IWSDS 2017.

2016

[Zhou Yu](#), Ziyu Xu, Alan W Black and Alexander I. Rudnicky, Strategy and Policy Learning for Non-Task-Oriented Conversational Systems, ACM SIG Meetings on Discourse and Dialogue (SIGDIAL 2016).

[Zhou Yu](#), Leah Nicolich-Henkin, Alan W Black and Alexander I. Rudnicky, A Wizard-of-Oz Study on A Non-Task-Oriented Dialog Systems that Reacts to User Engagement, ACM SIG Meetings on Discourse and Dialogue (SIGDIAL 2016).

[Zhou Yu](#), Xinrui He, Alan W Black and Alexander I. Rudnicky, User Engagement Modeling in Virtual Agents Under Different Cultural Contexts, International Conference on Intelligent Virtual Agents (IVA 2016).

Luis F. D'Haro, Bayan Abu Shawar and [Zhou Yu](#), 2nd WOCHAT Shared Task Data Collection and Annotation, International Conference on Intelligent Virtual Agents (IVA 2016).

Sean Andrist, Dan Bohus, [Zhou Yu](#) and Eric Horvitz, Are You Messing with Me? : Querying about the Sincerity of Interactions in the Open World, International Conference on Human-Robot Interaction (HRI 2016).

[Zhou Yu](#), Ziyu Xu, Alan W Black and Alexander I. Rudnicky, Chatbot evaluation and database expansion via crowdsourcing, In Proceedings of the RE-WOCHAT workshop in Language Resources and Evaluation Conference (LREC 2016).

Luis F. D'Haro, Bayan Abu Shawar and [Zhou Yu](#), 1st WOCHAT Shared Task Data Collection and Annotation, Language Resources and Evaluation Conference (LREC 2016).

[Zhou Yu](#), Vikram Ramanarayanan, Robert Mundkowsky, Patrick Lange, Alan Black, Alexei Ivanov, David Suendermann-Oeft, Multimodal HALEF: An Open-Source Modular Web-Based Multimodal Dialog Framework, International Workshop for Spoken Dialog Systems (IWSDS 2016).

Alexei Ivanov, Patrick Lange, David Suendermann-Oeft, Vikram Ramanarayanan, Yao Qian, [Zhou Yu](#) and Jidong Tao, Speed vs. Accuracy: Designing an Optimal ASR System for Spontaneous Non-Native Speech in a Real-Time Application, International Workshop for Spoken Dialog Systems (IWSDS 2016).

2015

[Zhou Yu](#), Dan Bohus and Eric Horvitz, Incremental Coordination: Attention-Centric Speech Production in a Physically Situated Conversational Agent, SIG Meetings on Discourse and Dialogue (SIGDIAL 2015).

[Zhou Yu](#), Vikram Ramanarayanan, David Suendermann-Oeft, Xinhao Wang, Klaus Zechner, Lei Chen, Jidong Tao and Yao Qian, Using Bidirectional LSTM Recurrent Neural Networks to Learn High-Level Abstractions of Sequential Features for Automated Scoring of Non-Native Spontaneous Speech, IEEE Automatic Speech Recognition and Understand Workshop (ASRU 2015).

[Zhou Yu](#), Alexandros Papangelis and Alexander I. Rudnicky, TickTock: Engagement Awareness in a non-Goal-Oriented Multimodal Dialogue System, AAAI Spring Symposium on Turn-taking and Coordination in Human-Machine Interaction (AAAI Symposium 2015).

2013

[Zhou Yu](#), David Gerritsen, Amy Ogan, Alan W Black and Justine Cassell, Automatic Prediction of Friendship via Multi-model Dyadic Features, SIG Meetings on Discourse and Dialogue (SIGDIAL 2013).

Zhou Yu, Stefan Scherer, David Devault, Jonathan Gratch, Giota Stratou, Louis-Philippe Morency and Justine Cassell, Multimodal Prediction of Psychological Disorder: Learning Verbal and Nonverbal Commonality in Adjacency Pairs, Workshop on Semantics and Pragmatics of Dialog (SEMDIAL 2013).

2010

Zhou Yu, Deng Cai and Xiaofei He, Error-correcting Output Hashing in Fast Similar Search, in International Conference on Internet Multimedia Computing and Service, International Conferences on Internet Multimedia Computing and Service (ICIMCS 2010, **best paper award**).

Workshops, Reports and Demos

Dian Yu, Michelle Cohn, Yi Mang Yang, Chun-Yen Chen, Weiming Wen, Jiaping Zhang, Mingyang Zhou, Kevin Jesse, Austin Chau, Antara Bhowmick, Shreenath Iyer, Girithaja Sreenivasulu, Sam Davidson, Ashwin Bhandare, **Zhou Yu**, Gunrock: A Social Bot for Complex and Engaging Long Conversations, EMNLP 2019 demo

Zhou Yu, Ziyu Xu, Alan W Black and Alexander I. Rudnicky, A Chatbot for Film Promotion, in Proceedings of IEEE SLT 2016, Demo, San Diego, California, Dec. 13-16

Vikram Ramanarayanan, Patrick Lange, David Pautler, Zhou Yu and David Suendermann-Oeft, Interview With An Avatar: A Real-time Engagement Tracking-enabled Cloud-based Multimodal Dialog System For Learning And Assessment, in Proceedings of IEEE SLT 2016, Demo, San Diego, California, Dec. 13-16

Vikram Ramanarayanan, David Suendermann-Oeft, Patrick Lange, Alexei V. Ivanov, Keelan Evanini, **Zhou Yu**, Eugene Tsuprun, and Yao Qian (2016), Bootstrapping Development of a Cloud-Based Spoken Dialog System in the Educational Domain From Scratch Using Crowdsourced Data, in: *ETS Research Report Series*, Wiley. doi: 10.1002/ets2.12105.

Zhou Yu, Noisy Multimodal Sequence Data Modeling for Interactive Conversation Modeling, in Proceeding of Advances in Neural Information Processing Systems (NIPS), Women in Machine Learning Workshop 2015, Montreal, Canada, Dec. 7-12

Zhou Yu, Attention and Engagement Aware Multimodal Conversational Systems, in Proceeding of ACM International Conference on Multimodal Interaction (ICMI) Doctoral Consortium 2015, Seattle, Washington, Nov. 9-13

Zhou Yu, Multimodal Interactive System Design, Implement and Deployment, in Proceedings of the 11th Young Researchers' Roundtable on Spoken Dialog Systems (YRRSDS), 2015, Prague, Czech Republic, Aug. 31- Sep.1

Zhou Yu, Engaging Social Chatbot, in Proceedings of the 10th Young Researchers' Roundtable on Spoken Dialog Systems (YRRSDS), 2014, Philadelphia, Pennsylvania, June 16-17.

Patents

Computer-Implemented Systems and Methods for a Crowd Source-Bootstrapped Spoken Dialog System

A bootstrapping method in developing real-world deployable spoken dialog systems

Processor-Implemented Systems and Methods for Determining Sound Quality

A deep learning algorithm and end-to-end system for automatic non-native English proficiency scoring

Grants

2020/07	Amazon AWS ML research faculty gift, PI, \$40,000 and \$80,000 Amazon AWS credits
2020/06-2025/06	SemaFor, DARPA, Co-PI, \$600,000 (out of \$3,000,000) with Berkeley, UW and Boston University.
2020/06	Sony Faculty Research Award, Single PI, \$80,000
2020/01-2025/01	Facebook, Recommender Dialog Systems, Single PI, \$2,000,000
2020/01-2022/01	UCSF, Building a chatbot to deliver personalized physical activity interventions to reduce metabolic syndrome, Co-PI, \$75,000
2020/01-2022/01	NSF, Device-DS as a window into mechanisms of speech production: An investigation of adults and children, Co-PI, \$138,000
2019/05-2020/05	Amazon Alexa Prize, Towards An User-Adaptive Social Bot, Single PI, \$125,000.
2019/05	Tencent AI Lab Rhino-Bird Gift Fund, Personalized Second Language Conversation Skill Training Dialog Systems, Single PI, \$50,000

2019/02	UC Davis AI Health Award, A Personalized Chatbot Driven by Artificial Intelligence Algorithms for Physical Activity Promotion, Co-PI, \$5,000 (out of \$10,000)
2019/01-2021/01	DARPA, ASED: Detecting and Responding to Social Engineering Attacks, Co-PI, \$500,000(out of \$2,000,000)
2018/11	Intel Research Gift, Persuasive Dialog Systems, Year 2, Single PI, \$100,000
2018/08	Cisco Gift, Multiple Reinforcement Learning, Single PI, \$70,000
2018/03	Bosch Gift, Sentiment Adaptive Dialog Systems, Single PI, \$30,000
2018/05	CITRIS, WeCare: WiFi enabled Device-free Activity Monitoring Platform for Elderly Healthcare and Smart Home Automation, co-PI, \$30,000 (out of \$60,000)
2018/02-2018/12	Amazon Alexa Prize, Robust Spoken Dialog Understanding, Single PI, \$125,000
2017/12	Tencent AI Lab Rhino-Bird Gift Fund, Personalized Recommendation Dialog Systems, Single PI, \$50,000
2017/11	Intel Research Gift, Persuasive Dialog Systems, Year 1, Single PI, \$100,000

Teaching

2018/01 & 2019/04	Undergraduate Introduction to Natural Language Processing (~130 students)
2019/09	Graduate Natural Language Processing (~40 students)

Current PhD Students

Mingyang Zhou, joined 2017
 Weiyang Shi, joined 2018
 Kun Qian, joined 2018
 Qingyang Wu, joined 2018
 David Gros, joined 2019, co-advised with Prem Devanbu
 Kai-hui Liang, joined 2019
 Yu Li, joined 2019

Invited Talks

“Seamless Natural Communication between Humans and Machines”

Columbia University, Carnegie Mellon University, University of Texas Austin, University of Maryland, UCLA, USC ICT. 2020

“Enabling Machines with Situational Awareness, Communication, and Decision-Making Abilities with Multimodal Information”

Johns Hopkins University Joint Seminar, Berkeley NLP Reading Group, 2019

“Building Dialog Systems with Less Supervision”

University of Minnesota Seminar, UC Davis Colloquium, Berkeley NLP Reading Group, Invited talk at WWW IR-ConvAI workshop, West Cost NLP (WeCNLP) workshop, 2019

“Casual Conversational Systems”

Facebook, Uber invited talk, 2019

“Grounding Reinforcement Learning with Real-World Applications”

Caltech ML Seminar, UCLA CS Seminar, USC CS Seminar, CMU LTI Colloquium, UT Austin CS Seminar, Amazon AWS invited talk, 2018-2019

“Multimodal Dialog Systems”

Airbnb, San Francisco, CA, Santa Cruz CS Seminar, Berkeley NLP Seminar, Google, Facebook FAIR, 2018-2019

“Situating Intelligent Interactive Systems”

UC Davis, Communication Department, Davis, Statistics Department, Huawei, UPenn, Intel, Stanford, Cambridge University, University of Edinburgh, Harriet-Watt University, UC Davis, University of Rochester, Rutgers University, University of Colorado Boulder, University of Washington, Allen Institute of Artificial Intelligence, Ohio State University, Amazon, 2018

“Engagement in Multimodal Interactive Conversational Systems”

Baidu Research, Honda Research Institute, Shanghai Jiao Tong University, 2018

“Engagement in Multimodal Interactive Conversational Systems”

Northeastern University and MIT CS Seminar, 2017

“Attention and Engagement Aware Multimodal Dialog Systems”

SRI, ICSI Berkeley, Nuance, 2017

Professional Services

Program Chair, SIGDAIL 2021

Area Chair: EMNLP 2018, ACL 2019, 2020, AAAI 2020

Faculty Advisor: Student Research Workshop ACL 2020

Faculty Mentor: Widening NLP workshop 2018 & 2019

Organizer: ACL 2020 Conversation AI workshop, International Conference on Multimodal interaction (ICMI) 2019 Publicity Chair, IUI Dialog System Workshop 2019, SIGDIAL Special Session 2018, AAAI Symposium Natural Communication for Human-Robot Collaboration 2017, Chatbots and Conversational Agents special session in International Workshop for Spoken Dialog Systems (IWSDS) 2017, RE-WOCHAT workshop share task in International Conference on Intelligent Virtual Agents (IVA) 2016, WOCHAT workshop share task in Language Resources and Evaluation Conference (LREC) 2016, 10th and 11th Young Researchers’ Roundtable on Spoken Dialog Systems (YRRSDS), 2015 & 2016

Volunteer lead for IEEE Automatic Speech Recognition and Understanding (ASRU) 2016 & 2017