

YI-LIN (PASCAL) TUAN

✉ pascaltuan@gmail.com ✉ ytuan@cs.ucsb.edu ✎ Pascalson ✎ pascalson.github.io

EDUCATION

- PhD, Computer Science, University of California Santa Barbara** 2019 - 2024
Advisor: William Yang Wang
- BS, Electrical Engineering, National Taiwan University** 2013 - 2017
Advisors: Lin-shan Lee, Hung-yi Lee, Yun-Nung Chen

WORK EXPERIENCE

- University of California Santa Barbara, CA, USA** Sep2019 - Jun2024
Graduate Student Researcher, Advisor: William Yang Wang
 - Collaborated with Google (led by Tania Bedrax-Weiss and William W. Cohen), UC Santa Cruz (Professor Lise Getoor), USC (Professor Jay Pujara) on dialogue research from 2020.
- Meta GenAI, NY, USA** Jun2023 - Nov2023
Research Intern, Manager: Dan Bikel
 - Controlling helpfulness and safety levels in Large Language Models. [Preprint'24]
- Meta AI, CA, USA** Jun2022 - Sep2022
Research Intern, Mentor: Wenhan Xiong; Manager: Yashar Mehdad
 - Developing alternative large language model architecture for inference efficiency.
- Amazon Alexa AI, MA, USA** Jun2021 - Sep2021
Applied Scientist Intern, Mentor: Sajjad Beygi; Manager: Maryam Fazel-Zarandi
 - Interpretable and scalable knowledge graph grounded transformer for conversation agent. [ACL'22]
- Facebook AI, CA, USA** Jun2020 - Sep2020
Research Intern, Mentor: Ahmed El-Kishky; Manager: Paco Guzmán
 - Machine translation quality estimation with no human annotations. [EACL'21]
- National Taiwan University, Taipei, Taiwan** Oct2017 - Jul2019
Full-time Research Assistant
 - Training efficient sequential generative adversarial networks and reinforcement learning frameworks [Journal'18; Preprint'18; Interspeech'19]. (Advised by Professor Hung-yi Lee)
 - Dynamic knowledge graph grounded dialogue generation model architecture [EMNLP'19]. (Co-advised by Professor Hung-yi Lee and Professor Yun-Nung Chen)
 - Benchmarking commercial songs lyrics recognition [ICASSP'18]. (Advised by Professor Lin-shan Lee)

PUBLICATIONS

* indicates equal contribution. P: pre-print or under submission. C: conference. W: workshop. J: journal.

[P6] [Yi-Lin Tuan](#) and William Yang Wang. “A Gradient Analysis Framework for Rewarding Good and Penalizing Bad Examples in Language Models”. Preprint, 2024.

[P5] [Yi-Lin Tuan](#), Xilun Chen, Eric Michael Smith, Louis Martin, Soumya Batra, Asli Celikyilmaz, William Yang Wang, Daniel M. Bikel. “Towards Safety and Helpfulness Balanced Responses via Controllable Large Language Models”. Preprint, 2024.

[P4] [Yi-Lin Tuan](#), Zih-Yun Chiu, William Yang Wang. “Dynamic Latent Separation for Deep Learning”. Preprint, 2024.

[C12] Zih-Yun Chiu*, [Yi-Lin Tuan](#)*, William Yang Wang, Michael Yip. “Flexible Attention-Based Multi-Policy Fusion for Efficient Deep Reinforcement Learning”. *NeurIPS*, 2023.

- [C11] [Yi-Lin Tuan](#), Alon Albalak, Wenda Xu, Michael Saxon, Connor Pryor, Lise Getoor, William Yang Wang. “*CausalDialogue: Modeling Utterance-level Causality in Conversations*”. **ACL Findings**, 2023.
- [C10] Alon Albalak, [Yi-Lin Tuan](#), Pegah Jandaghi, Connor Pryor, Luke Yoffe, Deepak Ramachandran, Lise Getoor, Jay Pujara, William Yang Wang. “*FETA: A Benchmark for Few-Sample Task Transfer in Open-Domain Dialogue*”. **EMNLP**, 2022.
- [C9] Wenda Xu, [Yi-Lin Tuan](#), Yujie Lu, Michael Saxon, Lei Li, William Yang Wang. “*Not All Errors are Equal: Learning Text Generation Metrics using Stratified Error Synthesis*”. **EMNLP Findings**, 2022.
- [C8] Kai Nakamura, Sharon Levy, [Yi-Lin Tuan](#), Wenhui Chen, William Yang Wang. “*HybridDialogue: An Information-Seeking Dialogue Dataset Grounded on Tabular and Textual Data*”. **ACL Findings**, 2022.
- [C7] [Yi-Lin Tuan](#), Sajjad Beygi, Maryam Fazel-Zarandi, Qiaozi Gao, Alessandra Cervone, and William Yang Wang. “*Towards Large-Scale Interpretable Knowledge Graph Reasoning for Dialogue Systems*”. **ACL Findings**, 2022.
- [W1] Alon Albalak, Varun Embar, [Yi-Lin Tuan](#), Lise Getoor, William Yang Wang. “*D-REX: Dialogue Relation Extraction with Explanations*”. **ACL NLP4ConvAI workshop**, 2022.
- [C6] Lucas Relic, Bowen Zhang, [Yi-Lin Tuan](#), Michael Beyeler. “*Deep Learning–Based Perceptual Stimulus Encoder for Bionic Vision*”. **ACM AHs**, 2022.
- [C5] [Yi-Lin Tuan](#), Connor Pryor, Wenhui Chen, Lise Getoor, and William Yang Wang. “*Local Explanation of Dialogue Response Generation*”. **NeurIPS**, 2021.
- [C4] [Yi-Lin Tuan](#), Ahmed El-Kishky, Adi Renduchintala, Vishrav Chaudhary, Francisco Guzman, and Lucia Specia. “*Quality Estimation without Human-labeled Data*”. **EACL**, 2021.
- [P3] Zih-Yun Chiu, [Yi-Lin Tuan](#), Hung-yi Lee, Li-Chen Fu. “*Parallelized Reverse Curriculum Generation*”. Preprint, 2021.
- [P2] [Yi-Lin Tuan](#), Wei Wei, and William Yang Wang. “*Knowledge Injection into Dialogue Generation via Language Models*”. Preprint, 2020.
- [C3] [Yi-Lin Tuan](#), Yun-Nung Chen, and Hung-yi Lee. “*DyKgChat: Benchmarking Dialogue Generation Grounding on Dynamic Knowledge Graphs*”. **EMNLP**, 2019.
- [C2] Feng-Guang Su*, Aliyah R Hsu*, [Yi-Lin Tuan](#), and Hung-Yi Lee. “*Personalized Dialogue Response Generation Learned from Monologues*”. **Interspeech**, 2019.
- [J1] [Yi-Lin Tuan](#) and Hung-yi Lee. “*Improving conditional sequence generative adversarial networks by stepwise evaluation*”. **IEEE/ACM Transactions on Audio, Speech, and Language Processing**, 2019.
- [P1] [Yi-Lin Tuan](#)*, Jinzhi Zhang*, Yujia Li, and Hung-yi Lee. “*Proximal Policy Optimization and its Dynamic Version for Sequence Generation*”. Preprint, 2018.
- [C1] Che-Ping Tsai*, [Yi-Lin Tuan](#)*, and Lin-shan Lee. “*Transcribing lyrics from commercial song audio: the first step towards singing content processing*”. **ICASSP**, 2018.

OPEN-SOURCE PROJECTS

Total Github stars over 100.

Knowledge-Grounded Reinforcement Learning (Pytorch) [\[link\]](#)

Training Language Models with Many Dialogue Branches (Pytorch) [\[link\]](#)

Training Language Models with Differentiable Knowledge Graphs (Pytorch) [\[link\]](#)

Explanation Extraction for Language Model Generation (Pytorch) [\[link\]](#)

Dynamic Knowledge Graph in Conversations (Tensorflow) [\[link\]](#)

Reinforcement Learning and Generative Adversarial Network for Text Generation (Tensorflow) [\[link\]](#)

ACADEMIC SERVICES

Reviewer at <i>ACL Rolling Review</i>	2024
Reviewer at <i>Conference on Neural Information Processing Systems (NeurIPS)</i>	2024
Reviewer at <i>ACL 6th Workshop on NLP for Conversational AI</i>	2024
Reviewer at <i>Annual Meeting of the Association for Computational Linguistics (ACL)</i>	2024
Reviewer at <i>International Conference on Learning Representations (ICLR)</i>	2024
Reviewer at <i>Conference on Empirical Methods in Natural Language Processing (EMNLP)</i>	2023
Reviewer at <i>Conference on Neural Information Processing Systems (NeurIPS)</i>	2023
Reviewer at <i>Annual Meeting of the Association for Computational Linguistics (ACL)</i>	2023
Reviewer at <i>International Conference on Learning Representations (ICLR)</i>	2023
Reviewer at <i>AAAI KnowledgeNLP workshop</i>	2022
Reviewer at <i>Conference on Empirical Methods in Natural Language Processing (EMNLP)</i>	2022
Reviewer at <i>Association for the Advancement of Artificial Intelligence (AAAI)</i>	2021

TALKS

<i>Human-Like Chatbot: Reasoning and Optimization</i> At Google Gemini	2024
<i>Modeling Causal-Effect in Conversations</i> At Heinrich Heine University Düsseldorf, Host: Milica Gasic	2024
<i>Local Explanation of Dialogue Response Generation</i> At Google AI, Host: Deepak Ramachandran	2021
<i>Fact Verification with Semi-Structured Knowledge</i> At Fact Verification and Extraction workshop (on behalf of William Wang)	2019
<i>Dynamic Knowledge-Grounded Dialogue Generation via Walking on the Graphs</i> At SoCal NLP symposium, CA, USA	2019
At ASUS, Taipei, Taiwan	2019

TEACHING EXPERIENCE

UCSB CS190I: <i>Natural Language Processing</i> Teaching assistant, Lecturer: William Wang Hold seminars in entry-level natural language processing concepts and implementation.	2022
NTU CommE 5045: <i>Machine Learning and Having it Deep and Structured</i> Teaching assistant, Lecturer: Hung-yi Lee Lectured in reinforcement learning and generative adversarial network for text generation.	2017

AWARDS

Best Poster Award at ACM Augmented Humans <i>Deep learning for bionic vision [AHs'22]</i>	2022
Best Paper Award at South California NLP Symposium (SoCalNLP) <i>Dynamic knowledge graph grounded dialogue generation via walking on the graphs [EMNLP'19]</i>	2019
Outstanding Achievement Award at Altera, Intel <i>Mobile 3D projector via FPGA</i>	2016
Presidential Award at National Taiwan University Electrical Engineering <i>Awarded to top 5% GPA and ranked 1/169.</i>	2014

SKILLS

Programming: Python, PyTorch, HuggingFace, TensorFlow, Matlab, C++
Language: English (fluent), Mandarin (native), Taiwanese (basic)