Yichen “William” Huang

wh4@andrew.cmu.edu | yichenwilliamhuang.com

ACADEMIC EXPERIENCE

Carnegie Mellon University **Pittsburgh, USA**
Research Assistant, Computer Science Department *2024 - present*
Supervisor: Prof. Chris Donahue.

Mohamed bin Zayed University of Artificial Intelligence **Abu Dhabi, UAE**
Master of Science in Natural Language Processing (GPA: 3.77/4.0) *2022 - 2024*
Supervisors: Prof. Timothy Baldwin and Prof. Gus Xia.
- Thesis: Robust Automatic Evaluation for Natural Language Generation.

New York University Shanghai **Shanghai, China**
Bachelor of Science *2018 - 2022*
Double major with honors in Computer Science and Data Science (AI track) (GPA: 3.88/4.0)
- Honors: Magna Cum Laude, University Honors Scholar, Dean’s List for Academic Years 2018, 2019, and 2020.

PUBLICATIONS

- **Y. Huang**, & E. Kochmar (2024). REFeREE: A REference-FREE Model-Based Metric for Text Simplification. Accepted in *the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation*.
- **Y. Huang**, & T. Baldwin (2023). Robustness Tests for Automatic Machine Translation Metrics with Adversarial Attacks. In *Findings of the Association for Computational Linguistics: EMNLP 2023*.
- X. Liu, D. Chin, **Y. Huang**, & G. Xia (2023). Learning Interpretable Low-dimensional Representation via Physical Symmetry. In *The Thirty-seventh Conference on Neural Information Processing Systems*.
- **Y. Huang**, Y. Wang, & Y. Tam (2022). UNITER-Based Situated Coreference Resolution with Rich Multimodal Input. In *The Tenth Dialog System Technology Challenge workshop at AAAI 2022*.

RESEARCH

Improved Automatic Evaluation for Music Generation
Supervised by Prof. Chris Donahue and Prof. John Thickstun. *Ongoing*
- Proposed a new metric for open-ended music generation based on generative models and divergence frontiers, which overall outperforms FAD on sets of synthesized meta-evaluation data targeting fidelity, musicality, context, and diversity.
- Conducting a large-scale listening test to collect human preference data for further meta-evaluation.

Improving MIR with Generative Models
Supervised by Prof. Chris Donahue. *Ongoing*
- Proposed recasting music understanding task as generative audio-to-audio tasks.
- Benchmarked various strategies to apply generative models for MIR tasks.

REFeREE: A REference-FREE Model-Based Metric for Text Simplification [arXiv][Github]
Supervised by Prof. Ekaterina Kochmar. Accepted in LREC-COLING 2024. *2024*
- Proposed and implemented a three-stage training curriculum enabling arbitrarily scalable pre-training of supervised text amplification metrics.
- Analyzed the performance of the proposed metric on multiple datasets, which suggests SoTA performance in predicting overall ratings and competitive performance in predicting specific ratings.

Robustness Tests for Automatic Machine Translation Metrics with Adversarial Attacks [arXiv][GitHub]
Supervised by Prof. Timothy Baldwin. Accepted in Findings of EMNLP 2023. *2023*

- Proposed and applied adversarial attacks on machine translation metrics exposing overpenalization and self-inconsistency.
- Designed and coordinated a 10-person annotation run to validate the proposed attacks, gathering 10K quality-controlled quality annotations.

Learning Interpretable Low-dimensional Representation via Physical Symmetry [arXiv][GitHub]
Supervised by Prof. Gus Xia. Accepted in NeurIPS 2023. *2023*

- Participated in the development of a representation learning method capable of learning low-dimensional concepts (e.g. pitch and Cartesian coordinates) with minimal domain knowledge based on physical symmetry.
- Designed and implemented experiments on learning from arbitrary natural melodies and melodies rendered with varied timbers.

JukeControl: Enhancing Jukebox for Audio-to-Audio Music Generation with ControlNet [PDF][Github]
Supervised by Prof. Gus Xia. *2023*

- Extended ControlNet to Jukebox's transformer prior model to enable efficient fine-tuning for audio-conditioned audio generation.
- Proposed and experimented with recasting non-audio conditions and supervision signals as audio, expanding the use cases of the proposed model.

Causal Discovery on the Capacities and Specifications of Large Language Models [PDF][Github]
Course project for NLP-702: Advanced NLP. Instructed by Prof. Muhammad Abdul-Mageed. *2023*

- Applied causal discovery algorithms on MMLU performances of language models to analyze the causal relationships between task-specific performances (e.g. Math and Physics) and between performances and model specifications (e.g. model size and Math).

BabyLM v.s. OpenWebText: How Does Child-Oriented Language Affect Syntax Acquisition for Language Models? [PDF][Github]
Course project for NLP-705: Current Topics in NLP. Instructed by Prof. Ted Briscoe. *2023*

- Pretrained small-scale language models while controlling for properties of the pretraining corpus.
- Evaluated syntactic performance with acceptability judgements, studying the effect of child-oriented conversational corpus and, in particular, vocabulary and sentence length.

Faster Sequence-to-Sequence Symbolic Music Generation with Rule-Augmented Edit-Based Models and Knowledge Distillation [Slides][Github]
Supervised by Prof. Gus Xia. *2022*

- Proposed and implemented a fast conditional symbolic music generation method based on rule-based transformations and semi-autoregressive edit operations capable of inference four times faster than the autoregressive baseline.

An AI-Empowered Piano Performance Interface for Non-Pianists [PubPub][Github][Demo]
Supervised by Prof. Gus Xia. *2022*

- Developed a set of software interfaces for learning and performing contextualized harmonic progressions.
- Utilized harmonic style transfer methods to generate music samples based on a curated set of texture samples conditioned on user-input chord progressions in real time.
- Conducted quantitative and qualitative user studies with eight participants.

UNITER-Based Situated Coreference Resolution with Rich Multimodal Input [arXiv][GitHub]
Supervised by Prof. Yik-Cheung Tam. Accepted in the DSTC 10 workshop at AAAI 2022. *2021*

- Extended the UNITER model for effective multimodal coreference resolution in conversations in fashion and furniture shopping domains.
- Developed methods to obtain and fuse embeddings for multimodal inputs including dialogue history, scene images, structured knowledge base entries and scene graphs.
- Achieved an F1 score of 0.733, significantly outperforming the baseline (F1=0.366) and ranking second in the official evaluation at the DSTC 10 challenge.

PRACTICAL EXPERIENCE

Learning Assistant for CSCI-SHU 360: Machine Learning
Academic Resource Center, New York University Shanghai. *2021*

- Prepared and led weekly review sessions on fundamental concepts in Machine Learning and workshops on relevant libraries and packages.
- Provided one-on-one and small-group tutoring on coursework and projects.

Data Analysis Intern
AI R&D, Trip.com Group. *2020*

- Participated in the development of a recommender system combining points of interest into travel routes tailored to users' profiles.
- Scrapped and cleaned data for model training. Performed exploratory data analysis and model evaluation.

Music Club President
New York University Shanghai. *2019 - 2020*

- Led a team in planning and hosting club activities, including open mics, music workshops at local communities, and university-wide songwriting competitions, maintaining over 100 active club members.

LANGUAGES & SKILLS

- Natural languages: Mandarin Chinese (native), English (bilingual), Japanese (proficient, JLPT N1 Certified), Shanghainese (passive understanding).
- Programming languages: Python, Javascript, C#, C.
- Packages: PyTorch, Huggingface, Sci-kit Learn, Pandas, Numpy, Matplotlib.
- Web development: Vue, Bootstrap, D3, Flask, ExpressJS.
- Music instrument: Electric guitar.