

Bryan Wang

Web: <http://www.dgp.toronto.edu/~bryanw>

Mail: bryanw@adobe.com

Current Position

- 2024- **Research Scientist**, Adobe Research, Seattle, WA.
Building next-gen AI tools for creativity and storytelling.
Affiliated with STORIE Lab and CAVA research org.

Education

- 2020-2024 **Ph.D. in Computer Science.**
Department of Computer Science, University of Toronto.
Advisor: Tovi Grossman.
- 2018-2020 **MSc in Computer Science.**
Department of Computer Science, University of Toronto.
Advisor: Tovi Grossman.
- 2013-2018 **BSc in Computer Science.**
Department of Computer Science, National Taiwan University.
Advisor: Mike Y. Chen.

Publications

FULL CONFERENCE PAPERS

- 2023 [C. 15] **Bryan Wang**, Yuliang Li, Zhaoyang Lv, Haijun Xia, Yan Xu, Raj Sodhi. LAVE: LLM-Powered Agent Assistance and Language Augmentation for Video Editing. To Appear at *the 29th ACM Conference on Intelligent User Interfaces (IUI '24)*.
- 2023 [C. 14] Stephen Brade, **Bryan Wang**, Mauricio Sousa, Gregory Lee Newsome, Sageev Oore, Tovi Grossman. SynthScribe: Deep Multimodal Tools for Synthesizer Sound Retrieval and Exploration. To Appear at *the 29th ACM Conference on Intelligent User Interfaces (IUI '24)*.
- 2023 [C.13] Tao Li, Gang Li, Zhiwei Deng, **Bryan Wang**, Yang Li. A Zero-Shot Language Agent for Computer Control with Structured Reflection. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP '23)*.
- 2023 [C.12] Stephen Brade, **Bryan Wang**, Mauricio Sousa, Sageev Oore, Tovi Grossman. Promptify: Text-to-Image Generation through Interactive Prompt Exploration with Large Language Models. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '23)*.

- 2023 [C.11] **Bryan Wang**, Gang Li, Yang Li. Enabling Conversational Interaction with Mobile UI using Large Language Models. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '23)*
- 2023 [C.10] Jiannan Li, Mauricio Sousa, Karthik Mahadevan, **Bryan Wang**, Paula Akemi Aoyagui, Nicole Yu, Angela Yang, Ravin Balakrishnan, Anthony Tang, Tovi Grossman. Stargazer: An Interactive Camera Robot for Capturing How-To Videos Based on Subtle Instructor Cues. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '23)*
- 2023 [C.9] Rachel Phinnemore, Mohi Reza, Blaine Lewis, Karthik Mahadevan, **Bryan Wang**, Michelle Annett, Daniel Wigdor. Creepy Assistant: Development and Validation of a Scale to Measure the Perceived Creepiness of Voice Assistants. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '23)*
- 2022 [C.8] **Bryan Wang**, Zeyu Jin, Gautham J. Mysore. Record Once, Post Everywhere: Automatic Shortening of Audio Stories for Social Media. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '22)*.
- 2021 [C.7] **Bryan Wang**, Gang Li, Xin Zhou, Zhourong Chen, Tovi Grossman, Yang Li. Screen2Words: Automatic Mobile UI Summarization with Multimodal Learning. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '21)*.
- 2021 [C.6] **Bryan Wang**, Mengyu Yang, Tovi Grossman. Soloist: Generating Mixed-Initiative Tutorials from Existing Music Instructional Videos Through Audio Processing. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '21)*
- 2020 [C.5] **Bryan Wang**, Tovi Grossman. BlyncSync: Enabling Multimodal Smartwatch Gestures with Synchronous Touch and Blink. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '20)*
- 2019 [C.4] **Bryan Wang**, Yi-Hsuan Yang. PerformanceNet: Score-to-Audio Music Generation with Multi-Band Convolutional Residual Network. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI '19)*. Oral Presentation
- 2018 [C.3] Yu-Chian Wu, Te-Yen Wu, Paul Taele, **Bryan Wang**, Jun-You Liu, Po-En Lai, Pin-Sung Ku, Mike Y. Chen. ActiveErgo: Automatic and Personalized Ergonomics using Self-actuating Furniture. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '18)*.
- 2017 [C.2] Te-Yen Wu, **Bryan Wang**, Jiun-Yu Lee, Hao-Ping Shen, Yu-Chian Wu, Yu-An Chen, Pin-sung Ku, Ming-Wei Hsu, Yu-Chih Lin, Mike Y. Chen. CircuitSense: Automatic Sensing of Physical Circuits and Generation of Virtual Circuits to Support Software Tools. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '17)*.
- 2016 [C.1] Chiuang Wang, Hsuan-Ming Yeh, **Bryan Wang**, Te-Yen Wu, Hsin-Ruey Tsai, Rong-Hao Liang, Yi-Ping Hung, Mike Y. Chen. CircuitStack: Supporting Rapid Prototyping and Evolution of Electronic Circuits. In *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '16)*.

EXTENDED ABSTRACTS, DEMO, WORKSHOP, DOCTORAL SYMPOSIUM

- 2023 [D.3] Stephen Brade, **Bryan Wang**, Mauricio Sousa, Gregory Lee Newsome, Sageev Oore, Tovi Grossman. SynthScribe: Deep Multimodal Tools for Synthesizer Sound Retrieval and Exploration. *In NeurIPS Workshop on ML for Creativity and Design 2023*
- 2023 [D.4] **Bryan Wang**, Democratizing Content Creation and Consumption through Human-AI Copiloted Systems. *In Doctoral Symposium at the ACM Symposium on User Interface Software and Technology (UIST '23)*.
- 2023 [D.3] Yang Li, Ranjay Krishna, Helena Vasconcelos, **Bryan Wang**, and Forrest Huang. Workshop on Artificial Intelligence and Human-Computer Interaction. *In International Conference on Machine Learning (ICML), 2023*.
- 2021 [D.2] Zhuoyue Lyu, Jiannan Li, **Bryan Wang**, Alive: Interactive Visualization and Sonification of Neural Network in Virtual Reality. *In IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR '21)*
- 2019 [D.1] Yu-Hua Chen, **Bryan Wang**, Yi-Hsuan Yang. Demonstration of PerformanceNet: A Convolutional Neural Network Model for Score-to-Audio Music Generation. *In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI '19)*.

PRE-PRINTS

- 2024 [P.1] Yuwen Lu, Ziang Tong, Qinyi Zhao, Yewon Oh, **Bryan Wang**, Toby Jia-Jun Li. Flowy: Supporting UX Design Decisions Through AI-Driven Pattern Annotation in Multi-Screen User Flows. arXiv 2024
- 2024 [P.2] Hsin-Ruey Tsai, Shih-Kang Chiu, **Bryan Wang** GazeNoter: Co-Piloted AR Note-Taking via Gaze Selection of LLM Suggestions to Match Users' Intentions. arXiv 2024

DISSERTATION

- 2024 **Bryan Wang**. Human-AI Systems for Creating, Consuming, and Interacting with Digital Content. *Department of Computer Science, University of Toronto*.

Research Experiences

- 2018 - 2024 **University of Toronto**, Graduate Student Researcher. Toronto, ON. Advisor: Tovi Grossman. Driving research agenda on pushing the frontier of human-AI interaction by developing AI systems that empower human creativity. Supervised 10+ students at both graduate and undergraduate levels. Publishing at top-tier HCI conferences.
- Summer 2023 **Meta Reality Labs - Research**, Research Intern. Redmond, WA. Host: Raj Sodhi. Developed LAVE, a video editing tool that provides LLM-powered agent assistance and editing to assist users in the editing process. Conducted user studies to understand users' perception and adoption of an LLM agent for creative workflows. Work led to publication at IUI 2024 [C. 15].

- Summer 2022 **Google Research**, Student Researcher. Mountain View, CA. Hosts: Yang Li and Gang Li.
Designed and evaluated novel prompting techniques for enabling conversational interaction with mobile UIs using LLMs, laying the groundwork for foundation model-based GUI research. Work led to publication at CHI 2023 [C.9] and open-source code on Google’s Github [O.3].
- Summer 2021, Winter 2022 **Adobe Research**, Research Intern. Remote. Hosts: Gautham J. Mysore and Zeyu Jin.
Developed ROPE, an audio story editing tool capable of automatically shortening raw voice recordings through an NLP pipeline. Conducted studies to analyze audio story structures, the limitations of existing methods for shortening audio stories, and the effectiveness of ROPE in assisting users with audio editing. Work led to publication at UIST 2022 [C.8].
- Summer 2020 -Winter 2021 **Google Research**, Research Intern. Remote. Hosts: Yang Li and Gang Li.
Collected the first large-scale UI-language dataset for screen summarization, containing 110k+ summaries for 22k+ mobile UIs. Designed and trained multimodal deep neural networks to enable automatic UI summarization based on the dataset. Work led to a publication at UIST 2021 [C. 7] and the open sourcing of the dataset and model code [O. 2].
- Summer 2018 **Academia Sinica**, Research Assistant. Taipei, Taiwan. Host: Yi-Hsuan Yang.
Developed a deep convolutional network for audio music generation based on musical score input. Work led to publication at AAAI 2019 [C.4] and opensource code [O. 1].

Honours & Awards

- 2022-2024 **Special Recognition for Outstanding Reviews**, CHI ’22-’24, UIST ’22-’24. (11 recognitions)
- 2023-2024 **Wolfond Scholarship (\$5000 CAD)**, University of Toronto.
- 2022-2023 **DiDi Graduate Student Award in Computer Science (\$10000 CAD)**, University of Toronto.
- 2022 **Snap Creative Challenge Award**, Snap Inc.
- 2021-2022 **DiDi Graduate Student Award in Computer Science (\$10000 CAD)**, University of Toronto.
- 2021 **Best Presentation Award**, IEEE AIVR 2022
- 2016 **Best Talk Award**, ACM UIST 2016.
- 2016 **The James Dyson Award World Finalist**, James Dyson Foundation.
- 2017, 2018 **AI and CS Research Excellence Scholarship**, Appier Inc.

Services and Volunteering

PROGRAM COMMITTEE

- 2025 **Program Committee, Associate Chair** ACM CHI 2025 Full Papers.
- 2024 **Program Committee, Associate Chair** ACM UIST 2024 Full Papers.
- 2024 **Proposal Committee, Reviewer** ICML 2024 Workshops.
- 2023 **Program Committee, Associate Chair** ACM UIST 2023 Full Papers.
- 2023 **Program Committee, Area Chair**. ICML 2023 Workshop on AI and HCI.
- 2022-2023 **Program Committee, Associate Chair**. ACM CHI 2022-23 Workshop on Computational UI.
- 2021 **Program Committee, Associate Chair**. ACM CHI 2021 Late-Breaking Work.

ORGANIZING

- 2023 **Co-chair, Student Innovation Contest** ACM UIST 2023.
- 2023 **Co-organizer**, ICML 2023 Workshop on AI and HCI.
- 2023 **Session Chair**. ACM CHI 2023.
- 2021 **Session Chair**. ACM UIST 2021.
- 2019 **Student Volunteer**. ACM CHI 2019.

REVIEWING

- 2020-Now **Reviewer**. The ACM Conference on Human Factors in Computing Systems (CHI)
- 2020-Now **Reviewer**. The ACM Symposium on User Interface Software and Technology (UIST)
- 2023 **Reviewer**. The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)
- 2021 **Reviewer**. The ACM Designing Interactive Systems conference (DIS)
- 2021 **Reviewer**. The ACM conference on Creativity & Cognition (C&C)
- 2021 **Reviewer**. IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)
- 2021 **Reviewer**. Transactions of the International Society for Music Information Retrieval (TISMIR)

SCHOOL SERVICES

- 2021 **Student Member**. Working Group on Fostering a Safe and Inclusive Workplace, DGP Lab.
- 2021 **Graduate Admission Triager**. Department of Computer Science, University of Toronto.
- 2021 **Application Mentor**. Graduate Application Assistance Program, University of Toronto.

COMMUNITY SERVICES

- 2023-2024 **Host of HCI Pro Bono Office Hours**
In my last year of Ph.D, I dedicated 1-2 hours weekly for office hours to help junior students navigate the challenges of Ph.D. studies and HCI research, emphasizing supporting those from underrepresented backgrounds.
- 2022 **Technical Mentor for Refugee Professionals**, Upwardly Global.
While at Google, I volunteered for the *IT Automation with Python* course with Upwardly Global, an organization that helps eliminate employment barriers for immigrant and refugee professionals into the U.S. economy. Led Q&A sessions to support students toward certification completion.
- 2021-2022 **Co-President**, Taiwanese Graduate Student Association in Toronto.
Led a team of 5 graduate students to build a community for Taiwanese students coming to pursue graduate studies in Toronto, Canada. Hosted orientations and monthly social events.

Invited Talks

- 2024 **Adobe Research**. Remote. Host: Mira Dontcheva.
Augmenting Human Capabilities in Content Creation and Interaction through Generative AI-infused Systems.
- 2024

- Apple Research.** Remote. Host: Jeff Nichols.
Augmenting Human Capabilities in Content Creation and Interaction through Generative AI-infused Systems.
- 2024 **Google Research.** Remote. Host: Yang Li.
Augmenting Human Capabilities in Content Creation and Interaction through Generative AI-infused Systems.
- 2024 **ByteDance Research.** Remote. Host: Yu Tian
Augmenting Human Capabilities in Content Creation and Interaction through Generative AI-infused Systems.
- 2024 **Twelve Labs, Multimodal Weekly.** Remote. Host: James Le
Towards a Future of Co-Creating Videos with AI Agents.
- 2024 **Huawei Research.** Remote. Host: Yuanhao Yu
Augmenting Human Capabilities in Content Creation and Interaction through Generative AI-infused Systems.
- 2024 **University of Notre Dame.** Remote. Host: Toby (Jia-Jun) Li.
Augmenting Human Capabilities in Content Creation and Interaction through Generative AI-infused Systems.
- 2024 **Microsoft Research.** Remote. Host: Saleema Amershi
Augmenting Human Capabilities in Content Creation and Interaction through Generative AI-infused Systems.
- 2023 **HCI Seminar, University of Washington.** Seattle, Washington. Host: Chu Li.
Making GUIs Interactable with Natural Language.
- 2023 **Huawei Research.** Remote. Host: Zhen Li.
Making GUIs Interactable with Natural Language.
- 2023 **Meta Reality Labs - Research.** Remote. Host: Kashyap Todi.
Making GUIs Interactable with Natural Language.
- 2023 **University of Tokyo.** Tokyo, Japan. Host: Takeo Igarashi.
AI-Facilitated Human-Computer Interactivity.
- 2021 **National Taipei University of Technology.** Taipei, Taiwan.
Opportunities and Challenges in AI-infused Interactive Systems.
- 2020 **ACM SIGCHI, Toronto Chapter.** Toronto, ON.
BlyncSync: Enabling Multimodal Smartwatch Gestures with Synchronous Touch and Blink.
- 2017 **ACM SIGGRAPH, UIST Reprise.** Los Angeles, CA.
CircuitStack: Supporting Rapid Prototyping and Evolution of Electronic Circuits.
- 2017 **Appier Inc.** Taipei, Taiwan.
CircuitStack: Supporting Rapid Prototyping and Evolution of Electronic Circuits.

Student Mentoring

PHD STUDENTS:

- 2023-2024 **Yuwen Lu**, Computer Science, University of Notre Dame.
Project: Flowy: Supporting UX Design Decisions through AI-Driven Pattern Annotation in Multi-Screen User Flows.

MASTER'S STUDENTS:

- 2022-2024 **Stephen Brade**, Computer Science, University of Toronto.
Thesis: Promptify: Text-to-Image Generation through Interactive Prompt Exploration with Large Language Models [C.12]. Additional Publication [C. 14].
- 2021-2023 **Rachel Phinnemore**, Computer Science, University of Toronto.
Thesis: Creepy Assistant: Development and Validation of a Scale to Measure the Perceived Creepiness of Voice Assistants [C. 9].

UNDERGRADUATE STUDENTS:

- 2024 **Yu-Kai Hung**, Computer Science, National Taiwan University.
Project: SimTube: Simulating Audience Feedback on Videos using Vision-Language Models and Persona Dataset.
- 2023-2024 **Kieran Sasha**, Engineering Science, University of Toronto.
Thesis: ReciproCast: Enhancing Podcast Interactivity Through Conversational Avatars.
- 2023-2024 **Issam Arabi**, Computer Science, University of Toronto.
Thesis: TBD
- 2022-2023 **Yviel Castillejos**, Engineering Science, University of Toronto.
Thesis: Comic Generation with Foundation Models.
Now: Software Engineer at AMD.
- 2022-2023 **Shawn Zhang**, Engineering Science, University of Toronto.
Thesis: Generating Audio Stories from Workouts using Large Language Models.
Now: Machine Learning Engineer at StratumAI.
- 2021-2022 **Stephen Brade**, Engineering Science, University of Toronto.
Thesis: Automatic Musical Chord Suggesting System for Music Composition.
Now: Master's Student at University of Toronto.
- 2020-2021 **Qian (Philip) Chen**, Engineering Science, University of Toronto.
Thesis: AI-assisted Sketching Tutoring System.
Now: Machine Learning Scientist at Chisel AI.
- 2021 **Zhouyue Lyu**, Computer Science, University of Toronto.
Topic: Interactive Visualization and Sonification of Neural Networks. [D.2]
Now: Master's Student at Harvard University -> PhD student at Cambridge University.
- 2020 **Mengyu Yang**, Engineering Science, University of Toronto.
Topic: AI-assisted Musical Instrument Learning System [C.6].
Now: PhD student at Georgia Tech.

Teaching Experiences

- Winter 2024 **Seminar Moderator**. Topics in Interactive Computing: Large Language Models for Intelligent User Interfaces, CSC 2524. Graduate CS seminar course, University of Toronto.
Responsibilities: Design seminar materials, moderate class discussions, and provide feedback on students' term projects.
- 2019-2021 **Teaching Assistant**. Human-Computer Interaction, CSC 428-2514. Graduate CS course, University of Toronto.
Responsibilities: Lead weekly tutorials, mark assignments, and organize final presentations.

- Winter 2019 **Teaching Assistant.** Intro to Computer Science, CSC 148. Undergraduate CS course, University of Toronto.
Responsibilities: Lead weekly tutorials and mark exams.
- Fall 2018 **Teaching Assistant.** Intro to Computer Programming, CSC 108. Undergraduate CS course, University of Toronto.
Responsibilities: Hold office hours, and mark exams.

Open-Source Code and Datasets

- 2023 [O. 3] Enabling Conversational Interaction with Mobile UI using Large Language Models. [\[Code\]](#)
- 2021 [O. 2] Screen2Words: Automatic Mobile UI Summarization with Multimodal Learning. [\[Code\]](#) [\[Dataset\]](#)
- 2019 [O. 1] PerformanceNet: Score-to-Audio Music Generation with Multi-Band Convolutional Residual Network. [\[Code\]](#)

Selected Press Coverage

- 2024 [University of Toronto Magazine](#), "Tuning into Tomorrow: AI can help musicians compose and create new sounds. Is it just another music-making tool – or something else?"
- 2024 [BNN Breaking](#), "Revolutionizing Digital Creativity: LAVE's AI Merges with Video Editing for Novice-Friendly Innovation"
- 2024 [MarkTechPost](#), "Revolutionizing Video Editing: How LAVE and AI are Democratizing Creative Expression"
- 2024 [Heart of Machine](#), "When Sora's explosive video generation emerged, Meta began using an agent to automatically edit videos. (In Chinese)".
- 2023 [Google AI Blog](#), "Enabling conversational interaction on mobile with LLMs"
- 2023 [MarkTechPost](#), "This Artificial Intelligence (AI) Research Shows The Feasibility Of Enabling Conversational Interactions With Mobile UIs Using Large Language Models LLMs"
- 2021 [ITmedia](#), "Developed by the University of Toronto, Soloist turns YouTube musical instrument lesson videos into personalized instruction." (in Japanese)

Skills

I am proficient in building interactive AI systems with full-stack web development, including both frontend and backend. I am also experienced in designing, training, and evaluating deep learning models. I am familiar with the following programming languages, frameworks, and research methodologies.

Programming: Python, HTML/CSS/JavaScript

Development Framework: Pytorch, Tensorflow, Flask, React.js.

Research Methodologies: User Studies (quantitative/qualitative), ML experiments.

References

Available upon requests.

Last updated: October 2, 2024 • <http://www.dgp.toronto.edu/bryanw>