

HAO YU

<http://cs-people.bu.edu/haoyu/>

Department of Computer Science, Boston University, Boston, MA, 02215

+1-617-372-1630 ◊ haoyu@bu.edu

EDUCATION

Boston University, Boston, MA, USA

PhD Candidate in Computer Science

Advisor: Prof. Margrit Betke

Sep 2019 - Present

GPA: 3.98 / 4.0

Zhejiang University, Hangzhou, Zhejiang, China

B.E. in Computer Science and B.A. in English Language and Literature

Sep 2015 - Jun 2019

GPA: 3.95 / 4.0

RELEVANT COURSEWORK

Deep Learning, Machine Learning, Big Data Systems (Hadoop, MapReduce), Optimization Theory, Randomness in Computing, Computer Vision, Computer Graphics, Artificial Intelligence, Information Retrieval.

SELECTED RESEARCH PROJECTS

KenyanFace Mitigates Data Bias in Generative Face Models

Supervisor: Prof. Margrit Betke and Prof. Sarah Adel Bargal

Oct 2023 - May 2024

Boston University

- Proposed KenyanFace, a large-scale dataset consisting of face images of darker skins to mitigate the existing racial bias in large-scale face datasets, and KenyanFaceHQ, a high-quality subset for generative modeling.
- Demonstrated that incorporating KenyanFaceHQ into existing face datasets ensures that the model generates a balanced representation of White and Black faces, reducing bias while maintaining high image quality.
- Showed how KenyanFace improves the accuracy of gender classification models, especially for subgroups for which these models previously underperformed (e.g. Black individuals).

Affect Behavior Prediction: Using Transformers and Timing Information to Make Early Predictions of Student Exercise Outcome

Supervisor: Prof. Margrit Betke

Jan 2023 - Sep 2023

Boston University

- Proposed the first affective model that can predict student exercise outcomes at an *early* stage by analyzing only the first 5–20 seconds of student data.
- Proposed a multimodal system that augments state-of-the-art video representation with timing information obtained from students' learning log data.
- Designed an affect-aware Transformer for *early* student exercise outcome prediction.
- Presented extensive experimental results showing our model's superior performance over state-of-the-art.

COVES: A Cognitive-Affective Deep Model that Personalizes Math Problem Difficulty in Real Time

Supervisor: Prof. Margrit Betke

Jan 2022 - Oct 2023

Boston University

- Proposed a computer vision enhanced problem selector (COVES), a deep model to select a personalized difficulty level for each student based on predicting student-problem interactions through a combination of visual information and log data.
- Integrated COVES into MathSpring, an open-sourced online tutor by applying model quantization.
- Conducted real-time classroom experiments with twenty-two seventh-grade students.
- Results indicate that COVES leads to higher mastery of math concepts, better timing, and higher scores, thus providing a positive learning experience for the participants.

Leveraging Affect Transfer Learning for Student Behavior Prediction in Intelligent Tutoring Systems

Jan 2021 - Aug 2022

Supervisor: Prof. Margrit Betke

Graduate Research Fellow at Boston University

- Proposed a video-based transfer learning approach for predicting problem outcomes of students working with an intelligent tutoring system (ITS) by analyzing their faces and gestures.
- Created a large labeled dataset of student interactions with an intelligent online math tutor.
- Our model achieved a 50% relative increase in mean F-score over the previous state-of-the-art method.

Measurements and Interventions to Improve Student Engagement and Affect in Tutoring Systems

Sep 2020 - Sep 2021

Supervisor: Prof. Margrit Betke

Boston University

- Conducted two studies using computer vision techniques to measure students' engagement and affective states from their head pose and facial expressions, as they use an online tutoring system.
- Presented and integrated engagement strategies into the online tutoring system to help students re-engage once their attention wanders.
- A usability study was conducted and initial results indicate that students exposed to our re-engagement strategies were more confident and more persistent, responding positively to our learning companion.

INDUSTRY EXPERIENCE

Meta, Menlo Park, CA

Student Researcher (Part-time)

Aug 2024 - Nov 2024

Manager: Hanchao Yu

Project: Leveraging Multimodal Large Language Models (LLMs) for Unifying Representation and Generation with Contrastive-Autoregressive Finetuning.

Software Engineer Intern (Full-time)

May 2024 - Aug 2024

Manager: Hanchao Yu

Project: Advancing multi-modal representations for social media posts using large vision-language models.

AWARDS AND HONORS

AI and Education Initiative Doctoral Fellowship, Rafik B. Hariri Institute for Computing and Computational Science & Engineering 2024

Best Poster Award (4% award rate), IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021) 2021

First-Class Scholarship for Research and Innovation, Zhejiang University 2017, 2018

TEACHING EXPERIENCE

Teaching Fellow, Data Science Tools and Applications (CS 506), Boston University 2024 Fall

Head Teaching Fellow, Image and Video Computing (CS 585), Boston University 2024 Spring

Teaching Fellow, Principles of Machine Learning (CS 542), Boston University 2023 Spring, 2023 Fall

Teaching Fellow, Image and Video Computing (CS 585), Boston University 2021 Spring, 2022 Spring

Teaching Fellow, Introduction to Computer Science II (CS 112), Boston University 2020 Spring

PROFESSIONAL ACTIVITIES

Reviewer for Annual AAAI Conference on Artificial Intelligence (AAAI) 2025

Reviewer for IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2022-2025

Reviewer for ACM Multimedia (ACM MM) 2023, 2024

Reviewer for European Conference on Computer Vision (ECCV) 2022, 2024

Reviewer for International Conference on Affective Computing & Intelligent Interaction (ACII) 2024

Reviewer for IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) 2019

PUBLICATIONS

- [1] Farid Karimli, **Hao Yu**, Srishti Jain, Emmanuel Sarpong Akosah, Margrit Betke, and Wenxin Feng. Demonstration of CameraMouseAI: A Head-Based Mouse-Control System for People with Severe Motor Disabilities. In Proceedings of the 26th International ACM SIGACCESS Conference on Computers and Accessibility, pp. 1-6. 2024. <https://doi.org/10.1145/3663548.3688499>
- [2] **Hao Yu**, Danielle A. Alessio, William Rebelsky, Tom Murray, John J. Magee, Ivon Arroyo, Beverly P. Woolf, Sarah Adel Bargal, Margrit Betke. “Affect Behavior Prediction: Using Transformers and Timing Information to Make Early Predictions of Student Exercise Outcome.” In 25th International Conference on Artificial Intelligence in Education (AIED), 2024. 14 pages. https://doi.org/10.1007/978-3-031-64299-9_14. **Nominated for Best Paper Award and Best Student Paper Award.**
- [3] Lei Guo, H. Denis Wu, **Hao Yu**, and Margrit Betke. “The emotional power of partisan media: A computer vision analysis of the 2020 Democratic Party presidential primaries.” In Mass Communication and Society, 1–23. 2024. <https://doi.org/10.1080/15205436.2024.2370562>.
- [4] **Hao Yu**, Danielle A. Alessio, Will Lee, William Rebelsky, Frank Sylvia, Tom Murray, John J. Magee, Ivon Arroyo, Beverly P. Woolf, Sarah Adel Bargal, and Margrit Betke. “COVES: A Cognitive-Affective Deep Model that Personalizes Math Problem Difficulty in Real Time and Improves Student Engagement with an Online Tutor.” In Proceedings of the 31st ACM International Conference on Multimedia (ACM MM), pp. 6152-6160. 2023. 9 pages. <https://doi.org/10.1145/3581783.3613965>.
- [5] Beverly Woolf, Margrit Betke, **Hao Yu**, Sarah Adel Bargal, Ivon Arroyo, John Magee, Danielle Alessio, William Rebelsky. “Face Readers: The Frontier of Computer Vision and Math Learning.” In 24th International Conference on Artificial Intelligence in Education Workshop (AIED Workshop), 2023. 13 pages. [pdf](#).
- [6] Maorong Wang, **Hao Yu**, Ling Xiao, and Toshihiko Yamasaki. “Bridging the Capacity Gap for Online Knowledge Distillation.” In 2023 IEEE 6th International Conference on Multimedia Information Processing and Retrieval (MIPR), pp. 1-4. IEEE, 2023. <https://doi.org/10.1109/MIPR59079.2023.00018>
- [7] Nataniel Ruiz*, **Hao Yu***, Danielle A. Alessio, Mona Jalal, Ajjen Joshi, Tom Murray, John J. Magee, Kevin Manuel Delgado, Vitaly Ablavsky, Stan Sclaroff, Ivon Arroyo, Beverly P. Woolf, Sarah Adel Bargal, and Margrit Betke. “ATL-BP: A Student Engagement Dataset and Model for Affect Transfer Learning for Behavior Prediction.” In IEEE Transactions on Biometrics, Behavior, and Identity Science, 2023. 14 pages. <https://doi.org/10.1109/TBIOM.2022.3210479> (*Equal contribution)
- [8] Will Lee, Danielle Alessio, William Rebelsky, Sai Satish Gattupalli, **Hao Yu**, Ivon Arroyo, Margrit Betke et al. “Measurements and Interventions to Improve Student Engagement Through Facial Expression Recognition.” In International Conference on Human-Computer Interaction (HCII), pp 286–301, 2022. 16 pages. https://doi.org/10.1007/978-3-031-05887-5_20
- [9] **Hao Yu**, Ankit Gupta, Will Lee, Ivon Arroyo, Margrit Betke, Danielle Alessio, Tom Murray, John Magee, and Beverly P. Woolf. “Measuring and Integrating Facial Expressions and Head Pose as Indicators of Engagement and Affect in Tutoring Systems.” In International Conference on Human-Computer Interaction (HCII), pp 219–233, 2021. 15 pages. https://doi.org/10.1007/978-3-030-77873-6_16.
- [10] Nataniel Ruiz, **Hao Yu**, Mona Jalal, Danielle Alessio, Ajjen Joshi, Tom Murray, Vitaly Ablavsky, John Magee, Jacob Whitehill, Ivon Arroyo, Beverly Woolf, Stan Sclaroff, Margrit Betke. “Leveraging Affect Transfer Learning for Behavior Prediction in an Intelligent Tutoring System.” In 2021 16th IEEE International Conference on Automatic Face and Gesture Recognition (FG), pp. 1-8, IEEE, 2021. 8 pages. <https://doi.org/10.1109/FG52635.2021.9667001>. **Best Poster Award (4% award rate).**
- [11] Kevin Delgado, Juan Manuel Origgi, Tania Hasanpoor, **Hao Yu**, Danielle Alessio, Ivon Arroyo, William Lee, Margrit Betke, Beverly Woolf, and Sarah Adel Bargal. “Student Engagement Dataset.”

In Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW), pp. 3628-3636, 2021. 9 pages. [pdf](#).

- [12] Krishna Kumar Singh, **Hao Yu**, Aron Sarmasi, Gautam Pradeep, and Yong Jae Lee. “Hide-and-seek: A data augmentation technique for weakly-supervised localization and beyond.” arXiv preprint [arXiv:1811.02545](#), 2018. 14 pages.

TECHNICAL SKILLS

Computer Languages	Python (Proficient), C/C++ (Proficient), Java, L ^A T _E X, JavaScript, SQL
DL Frameworks	PyTorch, TensorFlow and TensorFlow.js, Keras
Toolkits	OpenCV, Scikit-learn, SciPy, Pandas, Numpy, CUDA, ONNX Runtime