

Furui Cheng

HUMAN-AI INTERACTION · DATA VISUALIZATION · EXPLAINABLE MACHINE LEARNING

Stampfenbachstrasse 114, 8092 Zürich, Switzerland

✉ furui.cheng@inf.ethz.ch | 🌐 www.furuicheng.tech | 📷 chengfr

Summary

I am a Postdoc at ETH Zürich, working with **Prof. Mennatallah El-Assady** and other amazing people in the IVIA Lab. Before joining ETH Zürich, I received my Ph.D. at the Hong Kong University of Science and Technology (HKUST) in 2022, advised by **Prof. Huamin Qu**.

I aim to bridge the gaps between humans and AI systems to enable a two-way knowledge transfer (i.e., human learning from and steering ML models). I develop and integrate **explainable machine learning** techniques with **interactive visualizations** to help users *probe, understand, and steer* machine learning models.

Education

Hong Kong University of Science and Technology

Hong Kong, China

PH.D. IN COMPUTER SCIENCE AND ENGINEERING

2018 - 2022

- Thesis: Interactive Visual Analytics for Human-Centered AI (Advisor: **Prof. Huamin Qu**)

Beihang University

Beijing, China

B.ENG. IN COMPUTER SCIENCE AND ENGINEERING

2014 - 2018

Experience

ETH Zürich

Zürich, Switzerland

POSTDOC WITH **PROF. MENNATALLAH EL-ASSADY**

Dec 2022 - present

- Research: improving the transparency of Large Language Models (LLMs) using visualization.
- Teaching: co-lecturer of Interactive Machine Learning: Visualization & Explainability

Hong Kong University of Science and Technology

Hong Kong, China

PHD STUDENT @VISLAB, ADVISED BY **PROF. HUAMIN QU**

Aug 2018 - Nov 2022

- Ph.D. Thesis Topic: Interactive Visual Analytics for Human-Centered AI

Harvard Medical School

MA, USA

VISITING STUDENT @HIDIVE LAB, ADVISED BY **PROF. NILS GEHLENBORG**

Feb 2022 - May 2022

- Worked with molecular biologists to help them improve the existing workflows in reference-based single-cell data analysis
- Proposed an interactive transfer learning framework, **Polyphony** that facilitates biologists to incorporate their knowledge into single-cell data integration and annotation models

Massachusetts Institute of Technology

MA, USA

VISITING STUDENT @DATA TO AI GROUP, ADVISED BY **DR. KALYAN VEERAMACHANENI**

Sep 2020 - Mar 2021

- Worked with clinicians from a children's hospital to understand clinicians' perspectives and behaviors of using ML model predictions and explanations for clinical decision-making
- Developed an AI-assisted clinical decision support system, **VBridge**, for post-operative complication predictions
- Contributed to **Cardea**, an open source AutoML library for using ML with Electronic Health Records

Microsoft Research Asia

Beijing, China

RESEARCH INTERN ADVISED BY **DR. WEIWEI CUI**

Apr 2017 - Jul 2017

- Contributed to developing a visual analytics system for probing, analyzing, and monitoring Convolutional Neural Networks

Featured Projects

RELIC: Investigating Large Language Model Responses using Self-Consistency

CHI 2024

- Researched how to support LLM (e.g., ChatGPT) users to identify and mitigate hallucinations in the generated text
- Proposed **RELIC**, an interactive system that enables users to verify semantic-level variations in multiple long-form responses. This allows users to recognize potentially inaccurate information in the generated text and make necessary corrections.

Polyphony: an Interactive Transfer Learning Framework for Single-Cell Data Analysis

IEEE VIS'22, PRESENTED AT BioVis@ISMB WITH **BEST ABSTRACT AWARD**

- Researched *how to incorporate expert knowledge into single-cell RNA sequencing (scRNA-seq) data annotation models*
- Proposed **Polyphony**, an interactive transfer learning framework that integrates interactive visualization and anchor-based batch-effect removal methods to support biologists in annotating scRNA-seq data
- Conducted user simulation studies and expert interviews that proved the effectiveness and usability of the system

VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models

IEEE VIS'21, **BEST PAPER HONORABLE MENTION AWARD**

- Researched *how to inform clinicians' decision-making with ML models*
- Worked with pediatric clinicians to understand the challenges in adapting existing explainable ML techniques (i.e., feature attributions) in making critical clinical decisions
- Developed **VBridge**, a visual analytics system that connects the dots between features and data to support human-AI collaborations in making clinical decisions

DECE: Decision Explorer with Counterfactual Explanations for Machine Learning Models

IEEE VIS'20

- Researched *counterfactual explanations* (i.e., *how to alter an ML model prediction with minimal changes to the data input*) as probes to help humans understand the ML models' decision boundaries (DBs)
- Designed an analysis workflow of mentally approximating model's DBs with iterative hypothesizing (i.e., *what the users think the DBs should be*) and counterfactuals-guided refinements (i.e., *understanding the difference from the actual cases*)
- Developed **DECE**, a visual analytics system with novel visualization designs for visually summarizing counterfactual explanations

Selected Publications

Peer-reviewed Conference and Journal Publications

- **RELIC: Investigating Large Language Model Responses using Self-Consistency**
Furui Cheng, Vilém Zouhar, Simran Arora, Mrinmaya Sachan, Hendrik Strobelt, Mennatallah El-Assady
Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI24), 2024.
- **Leveraging Historical Medical Records as a Proxy via Multimodal Modeling and Visualization to Enrich Medical Diagnostic Learning**
Yang Ouyang, Yuchen Wu, He Wang, Chenyang Zhang, Furui Cheng, Chang Jiang, Lixia Jin, Yuanwu Cao, Quan Li
IEEE Transactions on Visualization and Computer Graphics (IEEE VIS'23), 2023.
- **ShortcutLens: A Visual Analytics Approach for Exploring Shortcuts in Natural Language Understanding Dataset**
Zhijia Jin, Xingbo Wang, Furui Cheng, Chunhui Sun, Qun Liu, Huamin Qu
IEEE Transactions on Visualization and Computer Graphics, 2023.
- **Polyphony: an Interactive Transfer Learning Framework for Single-Cell Data Analysis**
Furui Cheng, Mark S Keller, Huamin Qu, Nils Gehlenborg, Qianwen Wang
IEEE Transactions on Visualization and Computer Graphics (IEEE VIS'22), 2022. **Best Abstract Award at BioVis@ISMB**
- **In Defence of Visual Analytics Systems: Replies to Critics**
Aoyu Wu, Dazhen Deng, Furui Cheng, Yingcai Wu, Shixia Liu, Huamin Qu
IEEE Transactions on Visualization and Computer Graphics (IEEE VIS'22), 2022.
- **VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models**
Furui Cheng, Dongyu Liu, Fan Du, Yanna Lin, Alexandra Zytek, Haomin Li, Huamin Qu, Kalyan Veeramachaneni
IEEE Transactions on Visualization and Computer Graphics (IEEE VIS'21), 2021. **Best Paper Honorable Mention Award**
- **DECE: Decision Explorer with Counterfactual Explanations for Machine Learning Models**
Furui Cheng, Yao Ming, Huamin Qu
IEEE Transactions on Visualization and Computer Graphics (IEEE VIS'20), 2020.
- **ProtoSteer: Steering deep sequence model with prototypes**
Yao Ming, Panpan Xu, Furui Cheng, Huamin Qu, Liu Ren
IEEE Transactions on Visualization and Computer Graphics (IEEE VIS'19), 2019.

Pre-print Papers

- **Interactive Analysis of LLMs using Meaningful Counterfactuals**
Furui Cheng, Vilém Zouhar, Robin Shing Moon Chan, Daniel Fürst, Hendrik Strobelt, Mennatallah El-Assady
under review and available through *arXiv preprint*, 2024.

- **Interactive Data Analysis with Next-step Natural Language Query Recommendation**

Xingbo Wang, [Furui Cheng](#), Yong Wang, Ke Xu, Jiang Long, Hong Lu, Huamin Qu
arXiv preprint, 2022.

Workshop Publications

- **ExpLIMEable: An exploratory framework for LIME**

Sonia Laguna, Julian Heidenreich, Jiugeng Sun, Nilüfer Cetin, Ibrahim Al Hazwani, Udo Schlegel, [Furui Cheng](#), Mennatallah El-Assady
XAI in Action: Past, Present, and Future Applications @NeurIPS 23, 2023.

- **Pulse: Toward a Smart Campus by Communicating Real-time Wi-Fi Access Data**

Aoyu Wu, Bon Kyung Ku, [Furui Cheng](#), Xinhuan Shu, Abishek Puri, Yifang Wang, and Huamin Qu
Workshop on Visualization for Communication, the IEEE Visualization Conference (IEEE VIS'18), 2018.

Selected Awards

- 2022 **Best Abstract Award**, BioVis@ISMB 2022
For the paper “Polyphony: an Interactive Transfer Learning Framework for Single-Cell Data Analysis”
- 2022 **RedBird Academic Excellence Award**, Hong Kong University of Science and Technology
- 2021 **Best Paper Honorable Mention Award**, IEEE VIS 2021 (Top 5%)
For the paper “VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models”
- 2016 **Academic Excellence**, School of Computer Science and Engineering (Top 10%)
- 2015 **Academic Excellence**, School of Mathematical Sciences (Top 10%)
- 2013 **First Price (3rd Place)**, 30th Chinese Physics Olympiad (CPhO) Regional (Shannxi, China) with over 15,000 participants. Invited to the 30th CPhO Final and got **Bronze Award**.

Presentation

ExpLIMEable: a Visual Analytics Approach for Exploring LIME

CONFERENCE PRESENTATION AT **IEEE VIS 2023**

[Melbourne, Australia](#)

Oct. 2023

From Human-Data Interaction to Human-AI Interaction

PRESENTATION AT THE ETH SYSTEMS GROUP SEMINAR

[Zürich, Switzerland](#)

May. 2023

Polyphony: an Interactive Transfer Learning Framework for Single-Cell Data Analysis

CONFERENCE PRESENTATION AT **IEEE VIS 2022**

[Virtual](#)

Oct. 2022

VBridge: Connecting the Dots Between Features and Data to Explain Healthcare Models

CONFERENCE PRESENTATION AT **IEEE VIS 2021**

[Virtual](#)

Oct. 2021

INVITED TALK AT **CHINAVis 2021**

Jul. 2021

Visual Analytics on Explainable Machine Learning for Informed Decision Making

INVITED TALK AT **STATE KEY LAB OF CAD&CG, ZHEJIANG UNIVERSITY**

[Hangzhou, China](#)

Jul. 2021

DECE: Decision Explorer with Counterfactual Explanations for Machine Learning Models

CONFERENCE PRESENTATION AT **IEEE VAST 2020**

[Virtual](#)

Oct. 2020

Teaching

Interactive Machine Learning: Visualization & Explainability

CO-LECTURER AND HEAD TEACHING ASSISTANT

[ETH Zürich](#)

2023, 2024 Spring

Fundamentals of Web Engineering

HEAD TEACHING ASSISTANT

[ETH Zürich](#)

2023 Autumn

Exploring and Visualizing Data

TEACHING ASSISTANT

[HKUST](#)

2020 Spring

Services

PROGRAM COMMITTEE MEMBERSHIPS

- ACM Conference on Intelligent User Interfaces (IUI), 2023 - 2024
- IEEE VIS Workshop on Visualization in Biomedical AI, 2022

PAPER REVIEW

- IEEE Transactions on Visualization and Computer Graphics (TVCG)
- IEEE VIS: Visualization & Visual Analytics, 2020 - 2024
- ACM Conference on Human Factors in Computing Systems (CHI), 2023 - 2024
- ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW), 2023 - 2024
- ACM Conference on Intelligent User Interfaces (IUI), 2022
- IEEE Pacific Visualization Symposium (PacificVIS), 2022 - 2023
- China Visualization and Visual Analytics Conference (ChinaVis), 2021 - 2023
- Visual Informatics