

# Grace Guo

ML Interpretability | Explainable AI | Visual Analytics  
[gguo31@g.harvard.edu](mailto:gguo31@g.harvard.edu) | [Google Scholar](#) | [Personal Website](#) | [LinkedIn](#)

## Education

- Ph.D. in Human-centered Computing** 2024  
Georgia Institute of Technology  
Faculty Advisor: Alex Endert, PhD
- B.S. in Human-Computer Interaction and Cognitive Science** 2018  
Carnegie Mellon University

## Awards

- IBM PhD Fellowship** (1 of 10 worldwide) 2023

## Experience

- Postdoctoral Fellow**, SEAS, Harvard University, Cambridge, MA Jul 2024 -  
  - Develop novel interpretability methods for quantifying and evaluating concept distributions in generative AI outputs
  - Collaborate with physicians at the Harvard Medical School to develop new AI and visual analytics systems to analyze highly multiplexed spatial data of tissues
  - Mentor and advise students in the Doctoral Program
- Research Intern**, IBM Research, Cambridge, MA May 2023 - Aug 2023  
  - Developed a novel framework for counterfactual explanation of AI image and video classification models in biomedical domains
  - Built MiMICRI, an open-source visualization tool in JavaScript, D3, and React
  - Built a Python embedding of MiMICRI to support the generation of in-domain counterfactual cardiac MRI images in JupyterLab
  - MiMICRI was published and presented at ACM FAccT 2024
- Research Intern**, IBM Research, Cambridge, MA May 2022 - Aug 2022  
  - Collaborated with the IBM Healthcare Analytics team on causal inference problems
  - Built Causalvis, an open-source visualization package in JavaScript, D3, and React with a Python embedding to support causal inference analysis in JupyterLab
  - Causalvis was published and presented at ACM CHI 2023

**Research Intern**, Pacific Northwest National Laboratory, Richland, WA May 2020 - Aug 2020

- Designed and built VAINE, a visualization system for interactively estimating causal effects in natural experiments
- VAINE was published and presented at IEEE VIS 2021

## Skills

**Languages & Frameworks:** Python (Pandas, NumPy, scikit-learn, PyTorch), JavaScript (React, Vue, Svelte, D3.js, WebGL), Git

**Data Analytics & Visualization:** SQL, matplotlib, seaborn, Tableau, Looker, D3.js

**Research & Communication:** A/B testing, causal inference, experimental design, model evaluation, stakeholder storytelling, quantitative research methods, qualitative research methods

**Generative AI & NLP:** Explainable AI (XAI), retrieval-augmented generation (RAG), document parsing, prompt engineering, web scraping (Selenium, BeautifulSoup), text summarization

## Publications

[1] Salma Abdel Magid, Weiwei Pan, Simon Warchol, **Grace Guo**, Junsik Kim, Mahia Rahman, and Hanspeter Pfister. 2025. Is What You Ask For What You Get? Investigating Concept Associations in Text-to-Image Models. Transactions on Machine Learning Research (TMLR) (2025).

[2] Simon Warchol, **Grace Guo**, Johannes Knittel, Dan Freeman, Usha Bhalla, Jeremy L Muhlich, Peter K Sorger, and Hanspeter Pfister. 2025. SEAL: Spatially-resolved Embedding Analysis with Linked Imaging Data. bioRxiv (2025), 2025–07. (To appear in IEEE VIS 2025)

[3] **Grace Guo**, Subhajit Das, Jian Zhao, and Alex Endert. 2025. More Like Vis, Less Like Vis: Comparing Interactions for Integrating User Preferences Into Partial Specification Recommenders. IEEE Transactions on Visualization and Computer Graphics (2025).

[4] Mark S Keller, Eric Mörth, Thomas C Smits, Simon Warchol, **Grace Guo**, Qianwen Wang, Robert Krueger, Hanspeter Pfister, and Nils Gehlenborg. 2025. The State of Single-Cell Atlas Data Visualization in the Biological Literature. IEEE Computer Graphics and Applications (2025).

[5] **Grace Guo**, Lifu Deng, Animesh Tandon, Alex Endert, and Bum Chul Kwon. 2024. MiMICRI: Towards Domain-centered Counterfactual Explanations of Cardiovascular

Image Classification Models. In Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency (FAccT). 1–14.

[6] **Grace Guo**, Aishwarya Mudgal Sunil Kumar, Adit Gupta, Adam Coscia, Chris MacLellan, and Alex Endert. 2024. Visualizing Intelligent Tutor Interactions for Responsive Pedagogy. In Proceedings of the 2024 International Conference on Advanced Visual Interfaces (AVI). 1–9.

[7] **Grace Guo**, John Stasko, and Alex Endert. 2024. What We Augment When We Augment Visualizations: A Design Elicitation Study of How We Visually Express Data Relationships. In Proceedings of the 2024 International Conference on Advanced Visual Interfaces (AVI). 1–6.

[8] Anh-Ton Tran, **Grace Guo**, Jordan Taylor, Katsuki Chan, Elora Raymond, and Carl DiSalvo. 2024. Situating Data Sets: Making Public Data Actionable for Housing Justice. In Proceedings of the 2024 CHI conference on human factors in computing systems (CHI). 1–16.

[9] **Grace Guo**, Ehud Karavani, Alex Endert, and Bum Chul Kwon. 2023. Causalvis: Visualizations for Causal Inference. In Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI). 1–20.

[10] **Grace Guo**, Maria Glenski, ZhuanYi Shaw, Emily Saldanha, Alex Endert, Svitlana Volkova, and Dustin Arendt. 2021. Vaine: Visualization and AI for natural experiments. In Proceedings of the 2021 IEEE Visualization Conference (VIS). IEEE, 21–25.

[11] Fabian Sperrle, Mennatallah El-Assady, **Grace Guo**, Rita Borgo, D Horng Chau, Alex Endert, and Daniel Keim. 2021. A Survey of Human-Centered Evaluations in Human-Centered Machine Learning. In Computer Graphics Forum, Vol. 40. Wiley Online Library, 543–568.

[12] Austin P Wright, Zijie J Wang, Haekyu Park, **Grace Guo**, Fabian Sperrle, Mennatallah El-Assady, Alex Endert, Daniel Keim, and Duen Horng Chau. 2020. A comparative analysis of industry human-AI interaction guidelines. arXiv preprint arXiv:2010.11761 (2020).

[13] **Grace Guo**, Bianchi Dy, Nazim Ibrahim, Sam Conrad Joyce, and Ate Poorthuis. 2020. Examining Design-Centric Test Participants in Graphical Perception Experiments. In EuroVis (Short Papers). 43–47.

[14] Ate Poorthuis, Lucas van der Zee, **Grace Guo**, Jo Hsi Keong, and Bianchi Dy. 2020. Florence: a Web-based Grammar of Graphics for Making Maps and Learning Cartography. Cartographic Perspectives 96 (2020), 32–50

[15] Sam C Joyce, **Grace Guo**, Bianchi Dy, Nazim Ibrahim, and Ate Poorthuis. 2019. Seeing numbers: Considering the effect of presentation of engineering data in design. In

Proceedings of IASS Annual Symposia, Vol. 2019. International Association for Shell and Spatial Structures (IASS), 1–8.

## Presentations

**Explainable AI** ACM FAccT, 2024  
MiMICRI: Towards Domain-centered Counterfactual Explanations of Cardiovascular Image Classification Models

**Visual Tools for Education** ACM AVI, 2024  
Visualizing Intelligent Tutor Interactions for Responsive Pedagogy

**Visualization II** ACM AVI, 2024  
When We Augment Visualizations: A Design Elicitation Study of How We Visually Express Data Relationships

**Politics of Datasets** ACM CHI, 2024  
Situating Data Sets: Making Public Eviction Data Actionable for Housing Justice

**Invited Talk @ Tableau Research** Salesforce, 2024  
Visualizations in Context: Toolkits for Expressive Visualization Augmentation

**Making Sense & Decisions with Visualization** ACM CHI, 2023  
Causalvis: Visualizations for Causal Inference

**Doctoral Colloquium** IEEE VIS, 2022  
Flexible and Expressive Augmentation of Domain-Specific Visualizations

**AI+VIS** IEEE VIS, 2021  
VAINE: Visualization and AI for Natural Experiments

**State-of-the-Art Reports (STARs)** EuroVis, 2021  
Survey of Evaluations in Human-Centered Machine Learning

## Teaching

**CS 1710: Visualization** Fall 2025  
Undergraduate course, 90 students  
Harvard SEAS

**EC 2135: Data Visualization for Analysis and Communication** Spring 2025  
Graduate course, 20 students  
Harvard Business School

**CS4460: Introduction to Information Visualization** Spring 2023  
Undergraduate course, 100 students  
Georgia Institute of Technology

**CS7455: Issues in Human-Centered Computing** Spring 2022  
Graduate course, 20 students  
Georgia Institute of Technology

**CS4873: Computing, Society and Professionalism** Summer 2021  
Undergraduate course, 10 students  
Georgia Institute of Technology

**CS7450: Information Visualization** Fall 2020  
Graduate course, 100 students  
Georgia Institute of Technology

**85-211: Cognitive Psychology** Fall 2016, Fall 2017  
Undergraduate course  
Carnegie Mellon University

**15-112: Fundamentals of Programming and CS** Fall 2015, Spring 2016  
Undergraduate course  
Carnegie Mellon University

## Professional Service

### **Conference Organizing Committee**

Information+, IEEE Vis Posters Track  
2025

### **Reviewing**

CHI Papers, EuroVis Full Papers, CG&A Special Issue, TVCG Journal Papers  
2025

CHI Papers, VIS Full Papers, VIS Short Papers, TVCG Journal Papers  
2024

VIS Full Papers  
2023