Ángel Alexander Cabrera

I am a PhD candidate in the Human-Computer Interaction Institute (HCII) at Carnegie Mellon University, advised by Adam Perer and Jason Hong. I work on human-centered AI, specifically in applying techniques from HCI and visualization to help people better understand and improve their AI systems. I am supported by an NSF Graduate Research Fellowship and have spent time at Apple AI/ML, Microsoft Research, and Google.

- A cabreraalex.com
- 🔽 cabrera@cmu.edu
- 🞓 Google Scholar
- G GitHub

Education

2019 - Present **Ph.D. in Human-Computer Interaction (HCI)** Carnegie Mellon University Advised by Adam Perer and Jason Hong.

- 2022 M.S. in Human-Computer Interaction Carnegie Mellon University
- 2019 **B.S. in Computer Science** Georgia Institute of Technology Concentration in intelligence and modeling/simulation. Minor in economics.
- Fall 2017Sciences Po Paris, FranceExchange program with a focus on economics and political science.

Work Experience

- Summer 2021 Apple AI/ML Research Intern Modular machine learning interfaces, see Symphony.
- Summer 2020 Microsoft Research Research Intern Behavioral model analysis, see AlFinnity.
- Summer 2018 **Google** Software Engineering Intern Automated driver assistance and hyperlocal weather prediction for Android Auto.
- Summer 2017 Google Software Engineering Intern

Anomaly detection and regression analysis system for Google's data processing pipelines.

Summer 2016 Google Engineering Practicum Intern Analytics platform for monitoring and detecting erroneous edits to Google Maps.

Awards

- 2023 Mozilla Technology Fund \$50,000 grant to develop Zeno as an auditing tool for AI.
- 2023 Stanford HAI Audit Challenge Zeno was a finalist for the HAI challenge for designing better AI auditing tools.
- 2019 2022 National Science Foundation Graduate Research Fellowship (NSF GRFP) Three-year graduate fellowship for independent research. Full tuition with an annual stipend of \$34,000.
 - 2019 Love Family Foundation Scholarship Co-awarded the \$10,000 scholarship for the undergraduate with the most outstanding scholastic record.
- 2015 2019 Stamps President's Georgia Tech and the Stamps Family Charitable Foundation

Full ride scholarship with \$15,000 in extracurricular funding awarded to 10 incoming students.

2018 **The Data Open Datathon** Correlation One and Citadel Securities Placed third and won \$2,500 for creating a ML system to predict dangerous road areas.

Refereed Publications

[11] Where Does My Model Underperform? A Human Evaluation of Slice Discovery Algorithms

Nari Johnson, Ángel Alexander Cabrera, Gregory Plumb, Ameet Talwalkar AAAI Conference on Human Computation and Crowdsourcing (HCOMP). Delft, Netherlands, 2023.

Details PDF 🚯 Details

[10] Towards a More Rigorous Science of Blindspot Discovery in Image Classification Models

Gregory Plumb*, Nari Johnson*, Ángel Alexander Cabrera, Ameet Talwalkar

Transactions on Machine Learning Research (TMLR). 2023.

🛃 PDF 🌎 Code 🚯 Details

[9] Zeno: An Interactive Framework for Behavioral Evaluation of Machine Learning

Ángel Alexander Cabrera, Erica Fu, Donald Bertucci, Kenneth Holstein, Ameet Talwalkar, Jason I. Hong, Adam Perer

ACM Conference on Conference on Human Factors in Computing Systems (CHI). Hamburg, Germany, 2023.

Details PDF 🗐 BibTex 🌐 Website

[8] Improving Human-AI Collaboration with Descriptions of AI Behavior

Angel Alexander Cabrera, Adam Perer, Jason I. Hong ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW). Minneapolis, 2023.

🔀 PDF 📒 BibTex 🚯 Details

[7] What Did My Al Learn? How Data Scientists Make Sense of Model Behavior Ángel Alexander Cabrera, Marco Tulio Ribeiro, Bongshin Lee, Rob DeLine, Adam Perer, Steven M. Drucker

ACM Transactions on Computer-Human Interaction (TOCHI). 2023.

🔀 PDF 🗧 BibTex 🚯 Details

[6] Symphony: Composing Interactive Interfaces for Machine Learning

Ángel Alexander Cabrera*, Alex Bäuerle*, Fred Hohman, Megan Maher, David Koski, Xavier Suau, Titus Barik, Dominik Moritz

ACM Conference on Conference on Human Factors in Computing Systems (CHI). New Orleans, 2022.

岸 PDF 📃 BibTex 🕟 Video 🚯 Details

[5] An open repository of real-time COVID-19 indicators

Alex Reinhart, Logan Brooks, Maria Jahja, Aaron Rumack, Jingjing Tang, [et al, including Ángel Alexander Cabrera]

Proceedings of the National Academy of Sciences (PNAS). 2021.

上 PDF 📃 BibTex 🌐 Website 🌎 Code 🚯 Details

[4] Discovering and Validating AI Errors With Crowdsourced Failure Reports Ángel Alexander Cabrera, Abraham Druck, Jason I. Hong, Adam Perer ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW). Virtual, 2021.

📩 PDF 🗧 BibTex 🚯 Details

[3] Regularizing Black-box Models for Improved Interpretability

Gregory Plumb, Maruan Al-Shedivat, Ángel Alexander Cabrera, Adam Perer, Eric Xing, Ameet Talwalkar

Conference on Neural Information Processing Systems (NeurIPS). Vancouver, 2020.

🔓 PDF 🗧 BibTex 🌎 Code 🚯 Details

[2] Designing Alternative Representations of Confusion Matrices to Support Non-Expert Public Understanding of Algorithm Performance

Hong Shen, Haojian Jin, Ángel Alexander Cabrera, Adam Perer, Haiyi Zhu, Jason I. Hong

ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW). Virtual, 2020.

📩 PDF 📃 BibTex 🚯 Details

[1] FairVis: Visual Analytics for Discovering Intersectional Bias in Machine Learning

Ángel Alexander Cabrera, Will Epperson, Fred Hohman, Minsuk Kahng, Jamie Morgenstern, Duen Horng (Polo) Chau

IEEE Conference on Visual Analytics Science and Technology (VAST). Vancouver, Canada, 2019.



Workshops, Demos, Posters, and Preprints

[4] Evaluating Systemic Error Detection Methods using Synthetic Images Gregory Plumb, Nari Johnson, Ángel Alexander Cabrera, Marco Tulio Ribeiro, Ameet Talwalkar

ICML - Workshop on Spurious Correlations, Invariance and Stability. Baltimore, MD, 2022.

PDF 1 Details

[3] "Public(s)-in-the-Loop": Facilitating Deliberation of Algorithmic Decisions in Contentious Public Policy Domains

Hong Shen, Ángel Alexander Cabrera, Adam Perer, Jason I. Hong CHI - Fair & Responsible Al Workshop. Hawaii, USA, 2020.

👌 PDF 🌐 Workshop 🚯 Details

[2] Discovery of Intersectional Bias in Machine Learning Using Automatic Subgroup Generation

Ángel Alexander Cabrera, Minsuk Kahng, Fred Hohman, Jamie Morgenstern, Duen Horng (Polo) Chau ICLR - Debugging Machine Learning Models Workshop (Debug ML). New Orleans, Louisiana, USA, 2019.

PDF 🌐 Workshop 🚯 Details

[1] Interactive Classification for Deep Learning Interpretation

Ángel Alexander Cabrera, Fred Hohman, Jason Lin, Duen Horng (Polo) Chau

CVPR - Demo. Salt Lake City, Utah, USA, 2018.

🔁 PDF 🌐 Website 🕟 Video 🌎 Code 🚯 Details

Talks

2023	Authoring Interactive AI Evaluation Reports with Zeno LLMs in Production - Part III
2022, 2023	Evaluating Machine Learning CMU 05-618/318: Human-AI Interaction CMU 10-605/805: ML with Large Datasets CMU 17-634: Applied Machine Learning
2022	Visualization and Machine Learning CMU 17-428/728: ML and Sensing
2022	Designing Large Web Applications CMU 05-431/631: Software Structures for User Interfaces (SSUI)
2022	Modern Web Frameworks CMU 05-431/631: Software Structures for User Interfaces (SSUI)
2021	Ethics in Data Visualization CMU 05-899: Data Visualization
2021	D3 Deep Dive CMU 05-899: Data Visualization
2021	Data Science Widgets with Svelte and Jupyter Svelte Summit 2021

Teaching

Fall 2022	05-431/631: Software Structures for User	Graduate Teaching Assistant @
	Interfaces (SSUI)	Carnegie Mellon
	To a de sur a lub state a a stiene a sur ale terste ana	l la aura accoración

Teach weekly lab sections, grade tests and homeworks.

Fall 2021 **05-499:C: Data Visualization** Graduate Teaching Assistant @ Carnegie Mellon Taught a D3 course and led an ethics workshop in addition to grading and course management. 2016 - 2018 CS1332: Data Structures and Algorithms

Undergraduate Teaching Assistant @ Georgia Tech

Taught a weekly recitation, graded tests and homework, and helped create assignments.

Mentoring

Spring 2023	Josh Zhou B.S. in Computer Science, Carnegie Mellon	
- Present	Instance tagging for Zeno.	
Summer 2022	Erica Fu B.S. in Information Systems, Carnegie Mellon	
- Present	UX design for an ML evaluation platform. See Zeno.	
Spring 2023	Steven Huang Research Associate, Carnegie Mellon	
- Summer 2023	Chart builder for Zeno.	
Fall 2022	Tianqi Wu M.S. in Computer Science, Carnegie Mellon	
- Spring 2023	Interactive slice discovery for Zeno.	
Summer 2022 - Spring 2023	DonnyB.S. in Computer Science, Oregon State University. REU at CarnegieBertucciMellonInteractive model debugging. See Zeno.	
Summer 2022	Kan Sun B.S. in Math, Carnegie Mellon. Algorithmic discovery of ML errors.	
Fall 2021	Emily Guo B.S. in Statistics and Machine Learning, Carnegie Mellon	
- Spring 2022	Improving human-AI interaction with descriptions of model behavior.	
Spring 2020 - Spring 2021	Abraham DruckB.S. in Mathematical Sciences, Carnegie Mellon. Now: Technology Analyst at Morgan StanelyCrowdsourced discovery of ML failures for image captioning. See Deblinder.	
2020	CMU AI Mentoring Program	

Service

Organizer

2023 CSCW Workshop on Supporting User Engagement in Testing, Auditing, and Contesting Al

Program Committee

- 2022 2023 AC, ACM Conference on Human Factors in Computing Systems (CHI) Late Breaking Work
- 2022 2023 PC, ACM Fairness, Accountability, and Transparency (FAccT)
- 2022 2023 PC, IEEE VIS Workshop on Visualization for AI Explainability (VISxAI)
 - 2023 PC, CHI Workshop on Trust and Reliance in AI-Assisted Tasks (TRAIT)
 - 2022 AC, ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW) Posters

Reviewer

- 2021 2024 ACM Conference on Human Factors in Computing Systems (CHI)
- 2021 2023 ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)
- 2020 2023 IEEE VIS
- 2022 2023 ACM Fairness, Accountability, and Transparency (FAccT)
 - 2022 ACM Symposium on User Interface Software and Technology (UIST)
 - 2022 IEEE Computer Graphics and Applications (CGASI)
- 2019 2021 IEEE Transactions on Visualization and Computer Graphics (TVCG)
 - 2019 ACM Transactions on Interactive Intelligent Systems (TiiS)

Student Volunteer

- 2019 IEEE VIS
- 2019 ACM Fairness, Accountability, and Transparency (FAccT)

Department

- 2022 2023 REU application reviewer
- 2020 2021 Ph.D. student faculty representative

Press

- 2023 "Auditing AI: Announcing the 2023 Mozilla Technology Fund Cohort" Mozilla
- 2020 "New forecasting data could help public health officials prepare for what's next in the coronavirus pandemic" *CNN*
- 2020 "Facebook and Google Survey Data May Help Map Covid-19's Spread" -Wired
- 2020 "Carnegie Mellon Unveils Five Interactive COVID-19 Maps" *Carnegie Mellon*
- 2020 "Visualizing Fairness in Machine Learning" Data Stories Podcast
- 2019 "Alex Cabrera Wins Love Family Foundation Scholarship" GT SCS

- 2019 "Georgia Tech Satellite Successfully Launched Into Space " Georgia Tech
- 2018 "Datathon Challenges Students to Create Solutions to Real-World Problems" - GT SCS

Projects and Open Source

2023 ZenoAn interactive ML evaluation framework for any data or model.Website G GitHub

2021 Svelte + Vega

A Svelte component for reactively rendering Vega and Vega-Lite visualizations.

GitHub Demo

2021 Svelte + Jupyter Widgets

A framework for creating reactive data science widgets using Svelte JS.

M Blog 🜎 GitHub 🕞 Video

2020 COVIDCast Visualization of COVID-19 Indicators

Interactive visualization system of COVID-19 indicators gathered through >20,000,000 surveys on Facebook and Google by CMU Delphi.

Website GitHub

2015 - 2017 **PROX-1 Satellite** Flight Software Lead and Researcher Led a team of engineers in developing the software for a fully undergraduate-led satellite mission.

In space! In space!

2014 CTF Resources

Guide and resources for capture the flag (CTF) competitions with over 1.6k stars on GitHub.

Website

Selected Courses

- Ph.D. MultiModal Machine Learning Causality and Machine Learning Human Judgement and Decision Making Applied Research Methods
 - B.S. Deep Learning Data and Visual Analytics

Machine Learning Computer Simulation Honors Algorithms