

Minsuk Kahng 강민석

 Seoul Nat'l Univ. Bldg 220, Room 320
 minsuk.kahng@snu.ac.kr
 <https://minsuk.com>

My research advances the **responsible development of AI systems** through a **human-centered, data-centric** approach. I design and develop interactive tools and methods that enable people to systematically analyze model failures, bias, and safety risks.

Areas of expertise: Human-Computer Interaction (HCI), Responsible AI, Data Visualization, Visual Analytics, Human-AI Interaction, Explainable AI, Data Science, Human-Centered AI

Employment

- Sept. 2026 – present **Seoul National University**, Seoul, South Korea
Assistant Professor, SNU College and Graduate School of Data Science (Joint Appointment)
- Mar. 2025 – Aug. 2026 **Yonsei University**, Seoul, South Korea
Assistant Professor, Department of Computer Science and Engineering, School of Computing
- 2022–2025 **Google DeepMind**, Atlanta, GA
Senior Research Scientist, People+AI Research (PAIR) Team
Foundational Research Unit; Responsible AI and Human-Centered Technology
- 2019–2022 **Oregon State University**, Corvallis, OR
Assistant Professor of Computer Science, School of Electrical Eng. and Computer Science
- Summer 2017 **Google**, Cambridge, MA
Software Engineering Intern, People+AI Research (PAIR) Team, Google Brain
- Summer 2016 **Facebook**, Menlo Park, CA
Research Intern, Applied Machine Learning Research Group
- Summer 2015 **Facebook**, Menlo Park, CA
Research Intern, Applied Machine Learning Research Group

Education

- 2013–2019 **Georgia Institute of Technology**, Atlanta, GA, USA
Ph.D. in Computer Science
Thesis: *Human-Centered AI through Scalable Visual Data Analytics*
Committee: Polo Chau (Advisor), Sham Navathe, Alex Endert, Martin Wattenberg & Fernanda Viégas
- 2009–2011 **Seoul National University**, Seoul, South Korea
M.S. in Computer Science and Engineering
Thesis: *Context-Aware Recommendation using Learning-to-Rank*
Advisor: Sang-goo Lee
- 2005–2009 **Seoul National University**, Seoul, South Korea
B.S. in Electrical and Computer Engineering

Awards & Honors

- 2021 **Dissertation Award**, College of Computing, Georgia Tech
- 2021 **Finalist, Research Award** on Transparency in Fairness, Meta
- 2020 **Best Paper, Honorable Mention**, ACM Transactions on Interactive Intelligent Systems
- 2018 **Google PhD Fellowship**, Google
Full tuition and stipend (\$35,000) for 2 years
- 2018 **Graduate Teaching Assistant of the Year**, Center for Teaching and Learning, Georgia Tech
- 2014 **NSF Graduate Research Fellowship**, National Science Foundation
Full tuition and stipend (\$34,000) for 3 years
- 2016 **Student Travel Award**, HILDA Workshop at SIGMOD Conference
- 2011 **Best Paper Award**, Ph.D. Workshop at CIKM Conference
- 2005 **National Scholarship** for Science and Engineering, Korea Student Aid Foundation

Selected Publications

This section lists **latest** and **highly-cited** publications. Please see Google Scholar for citations: <https://scholar.google.com/citations?user=vieI1GYAAAAJ> (total 6,000, h-index 29).

*: corresponding authors;

- [S18] Who Defines "Best"? Towards Interactive, User-Defined Evaluation of LLM Leaderboards. Minji Jung, Minjae Lee, Yejin Kim, Sarang Choi, and **Minsuk Kahng***. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2026.
- [S17] COMPASS: A Framework for Evaluating Organization-Specific Policy Alignment in LLMs. Dasol Choi, DongGeon Lee, Brigitta Jesica Kartono, Helena Berndt, Taeyoun Kwon, Joonwon Jang, Haon Park, Hwanjo Yu*, and **Minsuk Kahng***. *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2026.
- [S16] Data-Prompt Co-Evolution: Growing Test Sets to Refine LLM Behavior. Minjae Lee and **Minsuk Kahng***. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2026.
- [S15] Transformer Explainer: Learning LLM Transformers with Interactive Visual Explanation and Experimentation. Aeree Cho, Grace C. Kim, Alexander Karpekov, Seongmin Lee, Alec Helbling, Benjamin Hoover, Zijie J. Wang, **Minsuk Kahng***, and Duen Horng Chau*. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2026.
- [S14] Multi-turn Evaluation of Anthropomorphic Behaviours in Large Language Models. Lujain Ibrahim, Canfer Akbulut, Rasmi Elasmr, Charvi Rastogi, **Minsuk Kahng**, Meredith Ringel Morris, Kevin R. McKee, Verena Rieser, Murray Shanahan, and Laura Weidinger. *International Conference on Learning Representations (ICLR)*, 2026.
- [S13] SafePath: Preventing Harmful Reasoning in Chain-of-Thought via Early Alignment. Wonje Jeung, Sangyeon Yoon, **Minsuk Kahng***, and Albert No*. *Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2025.
- [S12] LLM Comparator: Interactive Analysis of Side-by-Side Evaluation of Large Language Models.

Minsuk Kahng, Ian Tenney, Mahima Pushkarna, Michael Xieyang Liu, James Wexler, Emily Reif, Krystal Kallarackal, Minsuk Chang, Michael Terry, and Lucas Dixon.
IEEE Transactions on Visualization and Computer Graphics (VIS), 2024.
[Deployed on Google’s LLM evaluation platforms. Featured at Google I/O 2024 as part of Responsible GenAI Toolkit for Gemini open models](#)

- [S11] Adversarial Nibbler: An Open Red-Teaming Method for Identifying Diverse Harms in Text-to-Image Generation. Jessica Quaye, Alicia Parrish, Oana Inel, Charvi Rastogi, Hannah Rose Kirk, **Minsuk Kahng**, Erin van Liemt, Max Bartolo, Jess Tsang, Justin White, Nathan Clement, Rafael Mosquera, Juan Ciro, Vijay Janapa Reddi, and Lora Aroyo.
ACM Conference on Fairness, Accountability, and Transparency (FAccT), 2024.
[Featured in Google Research blog](#)
- [S10] Gemma 2: Improving Open Language Models at a Practical Size.
 Google DeepMind Gemma Team. *arXiv*, 2024. [Released by Google; 2,000 citations](#)
- [S9] Visualizing Linguistic Diversity of Text Datasets Synthesized by Large Language Models.
 Emily Reif, **Minsuk Kahng***, and Savvas Petridis.
IEEE Visualization and Visual Analytics Conference (VIS Short), 2023.
- [S8] VLSlice: Interactive Vision-and-Language Slice Discovery.
 Eric Slyman, **Minsuk Kahng***, and Stefan Lee*.
International Conference on Computer Vision (ICCV), 2023.
- [S7] DendroMap: Visual Exploration of Large-Scale Image Datasets for Machine Learning with Treemaps. Donald Bertucci, Md Montaser Hamid, Yashwanthi Anand, Anita Ruangrotsakun, Delyar Tabatabai, Melissa Perez, and **Minsuk Kahng***.
IEEE Transactions on Visualization and Computer Graphics (VIS), 2022.
- [S6] One Explanation is Not Enough: Structured Attention Graphs for Image Classification.
 Vivswan Shitole, Li Fuxin, **Minsuk Kahng**, Prasad Tadepalli, and Alan Fern.
Annual Conference on Neural Information Processing Systems, (NeurIPS), 2021.
- [S5] FairVis: Visual Analytics for Discovering Intersectional Bias in Machine Learning. Ángel Alexander Cabrera, Will Epperson, Fred Hohman, **Minsuk Kahng**, Jamie Morgenstern, and Polo Chau. *IEEE Conference on Visual Analytics Science and Technology (VIS)*, 2019.
- [S4] Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers.
 Fred Hohman, **Minsuk Kahng**, Robert Pienta, and Polo Chau.
IEEE Transactions on Visualization and Computer Graphics, 25(8), 2019. [800 citations: 2nd most cited paper among all papers published in VIS and TVCG in 2018-2023 \[ref\]](#).
- [S3] GAN Lab: Understanding Complex Deep Generative Models using Interactive Visual Experimentation.
Minsuk Kahng, Nikhil Thorat, Polo Chau, Fernanda Viégas, and Martin Wattenberg.
IEEE Transactions on Visualization and Computer Graphics (VIS), 2018. [Open-sourced with Google Brain; Live system used by more than 200,000 people from over 190 countries.](#)
- [S2] ActiVis: Visual Exploration of Industry-Scale Deep Neural Network Models.
Minsuk Kahng, Pierre Y. Andrews, Aditya Kalro, and Polo Chau.
IEEE Transactions on Visualization and Computer Graphics (VIS), 2017. [Deployed by Facebook; Invited to SIGGRAPH; 500 citations: 8th most cited in IEEE TVCG in 2018–2023 \[ref\]](#).
- [S1] Interactive Browsing and Navigation in Relational Databases.
Minsuk Kahng, Shamkant B. Navathe, John T. Stasko, and Polo Chau.
International Conference on Very Large Data Bases (VLDB), 2016.

Full List of Publications

Conference papers that were published in journals (e.g., VIS papers published in TVCG) are listed in the Referred Conference papers. Please see Google Scholar for citations (total 6,000, h-index 29).

*: corresponding authors

Refereed Conference Papers

- [C33] Dasol Choi, Jihwan Lee, Minjae Lee, and **Minsuk Kahng***. When Cars Have Stereotypes: Auditing Demographic Bias in Objects from Text-to-Image Models. *European Conference on Computer Vision (ECCV'26)*, Malmö, Sweden, Sept. 2026.
- [C32] Dasol Choi, DongGeon Lee, Brigitta Jesica Kartono, Helena Berndt, Taeyoun Kwon, Joonwon Jang, Haon Park, Hwanjo Yu*, and **Minsuk Kahng***. COMPASS: A Framework for Evaluating Organization-Specific Policy Alignment in LLMs. *Annual Meeting of the Association for Computational Linguistics (ACL'26)*, San Diego, CA, July 2026.
- [C31] Minji Jung, Minjae Lee, Yejin Kim, Sarang Choi, and **Minsuk Kahng***. Who Defines "Best"? Towards Interactive, User-Defined Evaluation of LLM Leaderboards. *ACM Conference on Fairness, Accountability, and Transparency (FAccT'26)*, Montréal, Canada, June 2026.
- [C30] Charvi Rastogi, Mukul Bhutani, **Minsuk Kahng**, Shamsuddeen Hassan Muhammad, Evgeniia Razumovskaia, Priyanka Suresh, Ibrahim Said Ahmad, Charu Kalia, Yaaseen Mahomed, Madhurima Maji, Minjae Lee, Alicia Parrish, Jessica Quaye, Vijay Janapa Reddi, Aishwarya Verma, and Lora Aroyo. Going PLACES: Participatory Localized Red Teaming for Text-to-Image Safety in the Global South. *ACM Conference on Fairness, Accountability, and Transparency (FAccT'26)*, Montréal, Canada, June 2026.
- [C29] Jessica Quaye, Alicia Parrish, Charvi Rastogi, **Minsuk Kahng**, Oana Inel, Lora Aroyo, and Vijay Janapa Reddi. Bridging the Scale Gap: Augmenting Human Red-Teaming to Uncover Latent Risks in T2I Models. *ACM Conference on Fairness, Accountability, and Transparency (FAccT'26)*, Montréal, Canada, June 2026.
- [C28] Minjae Lee and **Minsuk Kahng***. Data-Prompt Co-Evolution: Growing Test Sets to Refine LLM Behavior. *ACM Conference on Human Factors in Computing Systems (CHI'26)*, Barcelona, Spain, Apr. 2026. [25.3% acceptance rate]
- [C27] Aeree Cho, Grace C. Kim, Alexander Karpekov, Seongmin Lee, Alec Helbling, Benjamin Hoover, Zijie J. Wang, **Minsuk Kahng***, and Duen Horng Chau*. Transformer Explainer: Learning LLM Transformers with Interactive Visual Explanation and Experimentation. *ACM Conference on Human Factors in Computing Systems (CHI'26)*, Barcelona, Spain, Apr. 2026. [25.3% acceptance rate]
- [C26] Lujain Ibrahim, Canfer Akbulut, Rasmi Elasmara, Charvi Rastogi, **Minsuk Kahng**, Meredith Ringel Morris, Kevin R. McKee, Verena Rieser, Murray Shanahan, and Laura Weidinger. Multi-turn Evaluation of Anthropomorphic Behaviours in Large Language Models. *International Conference on Learning Representations (ICLR'26)*, Rio de Janeiro, Brazil, Apr. 2026. [27.4% acceptance rate]
- [C25] Wonje Jeung, Sangyeon Yoon, **Minsuk Kahng***, and Albert No*. SafePath: Preventing Harmful Reasoning in Chain-of-Thought via Early Alignment. *Annual Conference on Neural Information Processing Systems (NeurIPS'25)*, San Diego, CA, Dec. 2025. [24.5% acceptance rate]

- [C24] Yixuan Li, Emery D. Berger, **Minsuk Kahng***, and Cindy Xiong Bearfield. From Perception to Decision: Assessing the Role of Chart Types Affordances in High-Level Decision Tasks. *IEEE Visualization and Visual Analytics Conference (VIS'25)*, Short, Vienna, Austria, Oct. 2025.
- [C23] **Minsuk Kahng**, Ian Tenney, Mahima Pushkarna, Michael Xieyang Liu, James Wexler, Emily Reif, Krystal Kallarackal, Minsuk Chang, Michael Terry, and Lucas Dixon. LLM Comparator: Interactive Analysis of Side-by-Side Evaluation of Large Language Models. *IEEE Transactions on Visualization and Computer Graphics*, 31(1) (**VIS'24**), Jan. 2025. [22.3% acceptance rate] [Deployed on Google's LLM platforms. Featured at Google I/O 2024 as part of Responsible GenAI Toolkit for Gemini open models.](#)
- [C22] Jessica Quaye, Alicia Parrish, Oana Inel, Charvi Rastogi, Hannah Rose Kirk, **Minsuk Kahng**, Erin van Liemt, Max Bartolo, Jess Tsang, Justin White, Nathan Clement, Rafael Mosquera, Juan Ciro, Vijay Janapa Reddi, and Lora Aroyo. Adversarial Nibbler: An Open Red-Teaming Method for Identifying Diverse Harms in Text-to-Image Generation. *ACM Conference on Fairness, Accountability, and Transparency (FAccT'24)*, Rio de Janeiro, Brazil, June 2024. [24.1% acceptance rate] [Featured in Google Research blog](#) [[link](#)].
- [C21] Eric Slyman, **Minsuk Kahng***, and Stefan Lee*. VLSlice: Interactive Vision-and-Language Slice Discovery. *International Conference on Computer Vision (ICCV'23)*, Paris, France, Oct. 2023. [26.2% acceptance rate]
- [C20] Emily Reif, **Minsuk Kahng***, and Savvas Petridis. Visualizing Linguistic Diversity of Text Datasets Synthesized by Large Language Models. *IEEE Visualization Conference (VIS'23)*, Short, Melbourne, Australia, Oct. 2023. [33.7% acceptance rate]
- [C19] Donald Bertucci, Md Montaser Hamid, Yashwanthi Anand, Anita Ruangrotsakun, Delyar Tabatabai, Melissa Perez, and **Minsuk Kahng***. DendroMap: Visual Exploration of Large-Scale Image Datasets for Machine Learning with Treemaps. *IEEE Transactions on Visualization and Computer Graphics*, 29(1) (**VIS'22**), Oklahoma City, OK, Oct. 2022. [26.5% acceptance rate]
- [C18] Kin-Ho Lam, Delyar Tabatabai, Jed Irvine, Donald Bertucci, Anita Ruangrotsakun, **Minsuk Kahng**, and Alan Fern. Beyond Value: CheckList for Testing Inferences in Planning-Based RL. *32nd International Conference on Automated Planning and Scheduling (ICAPS'22)*, Virtual, June 2022. [30.7% acceptance rate]
- [C17] Jeongyeon Kim, Yubin Choi, **Minsuk Kahng**, and Juho Kim*. FitVid: Responsive and Flexible Video Content Adaptation. *ACM CHI Conference on Human Factors in Computing Systems (CHI'22)*, New Orleans, LA, Apr. 2022. [24.7% acceptance rate]
- [C16] Vivswan Shitole, Li Fuxin, **Minsuk Kahng**, Prasad Tadepalli, and Alan Fern. One Explanation is Not Enough: Structured Attention Graphs for Image Classification. *35th Conference on Neural Information Processing Systems (NeurIPS'21)*, Virtual, Dec. 2021. [26% acceptance rate]
- [C15] Matthew L. Olson, Thuy-Vy Nguyen, Gaurav Dixit, Neale Ratzlaff, Weng-Keen Wong, and **Minsuk Kahng***. Contrastive Identification of Covariate Shift in Image Data. *IEEE Visualization Conference (VIS'21)*, Short, Virtual, Oct. 2021. [29.0% acceptance rate]
- [C14] Delyar Tabatabai, Anita Ruangrotsakun, Jed Irvine, Jonathan Dodge, Zeyad Shureih, Kin-Ho Lam, Margaret Burnett, Alan Fern, and **Minsuk Kahng***. "Why did my AI agent lose?": Visual Analytics for Scaling Up After-Action Review. *IEEE Visualization Conference (VIS'21)*, Short, Virtual, Oct. 2021. [29.0% acceptance rate] [Featured in DARPA XAI Workshop.](#)

- [C13] Rupika Dikkala, Roli Khanna, Caleb Matthews, Jonathan Dodge, Sai Raja, Catherine Hu, Jed Irvine, Zeyad Shureih, Kin-Ho Lam, Andrew Anderson, **Minsuk Kahng**, Alan Fern, and Margaret Burnett. Doing Remote Controlled Studies with Humans: Tales from the COVID Trenches. *ACM/IEEE 14th International Conference on Cooperative and Human Aspects of Software Engineering (CHASE'21)*, Virtual, May 2021. [28.0% acceptance rate]
- [C12] Zijie J. Wang, Robert Turko, Omar Shaikh, Haekyu Park, Nilaksh Das, Fred Hohman, **Minsuk Kahng**, and Duen Horng (Polo) Chau. CNN Explainer: Learning Convolutional Neural Networks with Interactive Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 27(2) (**VIS'20**), Virtual, Oct. 2020. [24.8% acceptance rate] 200 citations; Invited to present at SIGGRAPH'21 as one of the selected top papers.
- [C11] **Minsuk Kahng** and Duen Horng (Polo) Chau. How Does Visualization Help People Learn Deep Learning? Evaluating GAN Lab with Observational Study and Log Analysis. *IEEE Visualization Conference (VIS'20)*, Short, Virtual, Oct. 2020. [36.0% acceptance rate]
- [C10] Ángel Alexander Cabrera, Will Epperson, Fred Hohman, **Minsuk Kahng**, Jamie Morgenstern, and Duen Horng (Polo) Chau. FairVis: Visual Analytics for Discovering Intersectional Bias in Machine Learning. *IEEE Conference on Visual Analytics Science and Technology (VIS'19)*, Vancouver, Canada, Oct. 2019. [30.0% acceptance rate] 190 citations; Open-sourced.
- [C9] **Minsuk Kahng**, Nikhil Thorat, Duen Horng (Polo) Chau, Fernanda Viégas, and Martin Wattenberg. GAN Lab: Understanding Complex Deep Generative Models using Interactive Visual Experimentation. *IEEE Transactions on Visualization and Computer Graphics*, 25(1) (**VIS'18**), Berlin, Germany, Oct. 2018. [25.0% acceptance rate] Open-sourced with Google Brain; Live system used by more than 200,000 people from over 190 countries.
- [C8] **Minsuk Kahng**, Pierre Y. Andrews, Aditya Kalro, and Duen Horng (Polo) Chau. ActiVis: Visual Exploration of Industry-Scale Deep Neural Network Models. *IEEE Transactions on Visualization and Computer Graphics*, 24(1) (**VIS'17**), Phoenix, AZ, Oct. 2017. [21.4% acceptance rate] Deployed on Facebook ML platform; Invited to present at SIGGRAPH'18 as one of the selected top papers. 500 citations: top 8th most cited in IEEE TVCG in 2018–2023.
- [C7] Robert Pienta, **Minsuk Kahng**, Zhiyuan Lin, Jilles Vreeken, Partha Talukdar, James Abello, Ganesh Parameswaran, and Duen Horng (Polo) Chau. FACETS: Adaptive Local Exploration of Large Graphs. *SIAM International Conference on Data Mining (SDM'17)*, Houston, TX, Apr 2017. [26.0% acceptance rate]
- [C6] **Minsuk Kahng**, Shamkant B. Navathe, John T. Stasko, and Duen Horng (Polo) Chau. Interactive Browsing and Navigation in Relational Databases. *Proceedings of the VLDB Endowment*, 9(12) (**VLDB'16**), New Delhi, India, Sept. 2016. [15.6% acceptance rate]
- [C5] Hugo Gualdrón, Robson Cordeiro, Jose Rodrigues Jr., Duen Horng (Polo) Chau, **Minsuk Kahng**, and U Kang. M-Flash: Fast Billion-scale Graph Computation Using a Bimodal Block Processing Model. *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD'16)*, Riva del Garda, Italy, Sept. 2016. [28.3% acceptance rate]
- [C4] Charles D. Stolper, **Minsuk Kahng**, Zhiyuan Lin, Florian Foerster, Aakash Goel, John Stasko, and Duen Horng (Polo) Chau. GLO-STIX: Graph-Level Operations for Specifying Techniques and Interactive eXploration. *IEEE Transactions on Visualization and Computer Graphics*, 20(12) (**VIS'14** (formerly InfoVis)), Paris, France, Nov. 2014. [23.0% acceptance rate]
- [C3] Zhiyuan Lin, **Minsuk Kahng**, Kaeser Md. Sabrin, Duen Horng (Polo) Chau, Ho Lee, and U

Kang. MMap: Fast Billion-Scale Graph Computation on a PC via Memory Mapping. *IEEE International Conference on Big Data (BigData'14)*, Washington, D.C., Oct. 2014. [39.4% acceptance rate]

- [C2] Sangkeun Lee, Sungchan Park, **Minsuk Kahng**, and Sang-goo Lee. PathRank: A Novel Node Ranking Measure on a Heterogeneous Graph for Recommender Systems. *ACM Conference on Information and Knowledge Management (CIKM'12)*, Maui, HI, Oct. 2012. [27.8% acceptance rate]
- [C1] Sangkeun Lee, Sang-il Song, **Minsuk Kahng**, Dongjoo Lee, and Sang-goo Lee. Random Walk based Entity Ranking on Graph for Multidimensional Recommendation. *ACM Conference on Recommender Systems (RecSys'11)*, Chicago, IL, Oct. 2011. [20.0% acceptance rate]

Journal Articles

This list does not include conference papers published in journal issues (e.g., VIS papers published in IEEE TVCG).

- [J7] Roli Khanna, Jonathan Dodge, Andrew Anderson, Rupika Dikkala, Jed Irvine, Zeyad Shureih, Kin-Ho Lam, Caleb R. Matthews, Zhengxian Lin, **Minsuk Kahng**, Alan Fern, and Margaret Burnett. Finding AI's Faults with AAR/AI: An Empirical Study. *ACM Transactions on Interactive Intelligent Systems*, 12(1), 2022.
- [J6] Li Fuxin, Zhongang Qi, Saeed Khorram, Vivswan Shitole, Prasad Tadepalli, **Minsuk Kahng**, and Alan Fern. From Heatmaps to Structural Explanations of Image Classifiers. *Applied AI Letters*, 2021.
- [J5] Jonathan Dodge, Andrew Anderson, Roli Khanna, Jed Irvine, Rupika Dikkala, Kin-Ho Lam, Delyar Tabatabai, Anita Ruangrotsakun, Zeyad Shureih, **Minsuk Kahng**, Alan Fern, and Margaret Burnett. From “No Clear Winner” to an Effective XAI Process: An Empirical Journey. *Applied AI Letters*, 2021.
- [J4] Fred Hohman, **Minsuk Kahng**, Robert Pienta, and Duen Horng (Polo) Chau. Visual Analytics in Deep Learning: An Interrogative Survey for the Next Frontiers. *IEEE Transactions on Visualization and Computer Graphics*, 25(8), 2019. 800 citations: 2nd most cited paper published in the IEEE TVCG journal in 2018–2023 [ref].
- [J3] Peter J. Polack Jr., Shang-Tse Chen, **Minsuk Kahng**, Kaya De Barbaro, Rahul Basole, Moushumi Sharmin, and Duen Horng (Polo) Chau. Chronodes: Interactive Multifocus Exploration of Event Sequences. *ACM Transactions on Interactive Intelligent Systems*, 8(1), 2018. Best Paper, Honorable Mention.
- [J2] Sangkeun Lee, **Minsuk Kahng**, and Sang-goo Lee. Constructing Compact and Effective Graphs for Recommender Systems via Node and Edge Aggregations. *Expert Systems with Applications*, 42(7), 2015.
- [J1] Rahul C. Basole, Mark Braunstein, Vikas Kumar, Hyunwoo Park, **Minsuk Kahng**, Duen Horng (Polo) Chau, Acar Tamersoy, Daniel A. Hirsh, Nicoleta Serban, James Bost, Burton Lesnick, Beth L. Schissel, and Michael Thompson. Understanding Variations in Pediatric Asthma Care Processes in the Emergency Department using Visual Analytics. *Journal of the American Medical Informatics Association (JAMIA)*, 22(2), 2015.

Refereed Poster and Demo Papers

- [P15] Crystal Qian, Michael Xieyang Liu, Emily Reif, Grady Simon, Nada Hussein, Nathan Clement, James Wexler, Carrie J. Cai, Michael Terry, and **Minsuk Kahng***. LLM Adoption in Data Curation Workflows: Industry Practices and Insights. *ACM Conference on Human Factors in Computing Systems (CHI'25)* (Late Breaking Work Poster), Yokohama, Japan, May 2025. [32.8% acceptance rate].
- [P14] Seongmin Lee, Zijie J. Wang, Aishwarya Chakravarthy, Alec Helbling, ShengYun Peng, Mansi Phute, Duen Horng (Polo) Chau, and **Minsuk Kahng**. LLM Attributor: Interactive Visual Attribution for LLM Generation. *Annual AAAI Conference on Artificial Intelligence (AAAI'25)* (Demo), Philadelphia, PA, Feb. 2025.
- [P13] **Minsuk Kahng**, Ian Tenney, Mahima Pushkarna, Michael Xieyang Liu, James Wexler, Emily Reif, Krystal Kallarackal, Minsuk Chang, Michael Terry, and Lucas Dixon. LLM Comparator: Visual Analytics for Side-by-Side Evaluation of Large Language Models. *ACM Conference on Human Factors in Computing Systems (CHI'24)* (Late Breaking Work Poster), Honolulu, HI, 2024. [33.9% acceptance rate]. [Deployed on Google's internal LLM platforms. Featured at Google I/O 2024 as part of Responsible GenAI Toolkit for Gemini open models](#) [link].
- [P12] Emily Reif, Crystal Qian, James Wexler, and **Minsuk Kahng***. Automatic Histograms: Leveraging Language Models for Text Dataset Exploration. *ACM Conference on Human Factors in Computing Systems (CHI'24)* (Late Breaking Work Poster), Honolulu, HI, 2024. [33.9% acceptance rate].
- [P11] Crystal Qian, Emily Reif, and **Minsuk Kahng***. Understanding the Dataset Practitioners Behind Large Language Model Development. *ACM Conference on Human Factors in Computing Systems (CHI'24)* (Late Breaking Work Poster), Honolulu, HI, 2024. [33.9% acceptance rate].
- [P10] Anita Ruangrotsakun, Dayeon Oh, Thuy-Vy Nguyen, Kristina Lee, Mark Ser, Arthur Hiew, Rogers Ngo, Zeyad Shureih, Roli Khanna, and **Minsuk Kahng***. VIVA: Visual Exploration and Analysis of Videos with Interactive Annotation. *ACM International Conference on Intelligent User Interfaces (IUI'23)* (Demo), Sydney, Australia, Mar. 2023.
- [P9] Zijie J. Wang, Robert Turko, Omar Shaikh, Haekyu Park, Nilaksh Das, Fred Hohman, **Minsuk Kahng**, and Duen Horng (Polo) Chau. CNN 101: Interactive Visual Learning for Convolutional Neural Networks. *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'20)* (Late Breaking Work), Virtual, Apr. 2020.
- [P8] Dezhi Fang, Fred Hohman, Peter Polack, Hillol Sarker, **Minsuk Kahng**, Moushumi Sharmin, Mustafa al'Absi, and Duen Horng (Polo) Chau. mHealth Visual Discovery Dashboard. *ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp'17)* (Demo), Maui, HI, Sept. 2017.
- [P7] Robert Pienta, Leilei Xiong, Santiago Grijalva, Duen Horng (Polo) Chau, and **Minsuk Kahng**. STEPS: A Spatio-temporal Electric Power Systems Visualization. *ACM International Conference on Intelligent User Interfaces (IUI'16)* (Poster), Sonoma, CA, Mar. 2016.
- [P6] Peter J. Polack Jr., Shang-Tse Chen, **Minsuk Kahng**, Moushumi Sharmin, and Duen Horng (Polo) Chau. TimeStitch: Interactive Multi-focus Cohort Discovery and Comparison. *IEEE Conference on Visual Analytics Science and Technology (VAST as part of VIS'15)* (Poster), Chicago, IL, Oct. 2015.
- [P5] Robert Pienta, Zhiyuan Lin, **Minsuk Kahng**, Jilles Vreeken, Partha P Talukdar, James Abello,

Ganesh Parameswaran, Duen Horng (Polo) Chau. AdaptiveNav: Adaptive Discovery of Interesting and Surprising Nodes in Large Graphs. *IEEE VIS* (Poster), Chicago, IL, Oct. 2015.

- [P4] Charles D. Stolper, Florian Foerster, **Minsuk Kahng**, Zhiyuan Lin, Aakash Goel, John Stasko, and Duen Horng (Polo) Chau. GLOs: Graph-Level Operations for Exploratory Network Visualization. *ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'14)* (Poster), Toronto, Canada, Apr. 2014.
- [P3] **Minsuk Kahng** and Sang-goo Lee. Exploiting Paths for Entity Search in RDF Graphs. *ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'12)* (Poster), Portland, OR, Aug. 2012.
- [P2] **Minsuk Kahng**, Sangkeun Lee, and Sang-goo Lee. Ranking in Context-Aware Recommender Systems. *International Conference on World Wide Web (WWW'11)* (Poster), Hyderabad, India, Mar. 2011.
- [P1] Inbeom Hwang, **Minsuk Kahng**, Sung Eun Park, Jinwook Seo, and Sang-goo Lee. Si-Fi: Interactive Similar Item Finder. *ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'10)* (Demo), Geneva, Switzerland, Aug. 2010.

Refereed Workshop Papers

- [W11] Jessica Quaye, Alicia Parrish, Oana Inel, **Minsuk Kahng**, Charvi Rastogi, Hannah Rose Kirk, Jess Tsang, Nathan L Clement, Rafael Mosquera, Juan Manuel Ciro, Vijay Janapa Reddi, and Lora Aroyo. Lexically-Constrained Automated Prompt Augmentation: A Case Study using Adversarial T2I Data. *Safe Generative AI Workshop (at NeurIPS'24)*, Vancouver, Canada, Dec. 2024.
- [W10] Donald R. Bertucci and **Minsuk Kahng***. Backprop Explainer: Interactive Explanation of Backpropagation in Neural Network Training. *4th Workshop on Visualization for AI Explainability (at VIS'21)*, Virtual, Oct. 2021. (Interactive Article: <https://xnought.github.io/backprop-explainer/>)
- [W9] Kin-Ho Lam, Zhengxian Lin, Jed Irvine, Jonathan Dodge, Zeyad T. Shureih, Roli Khanna, **Minsuk Kahng**, and Alan Fern. Identifying Reasoning Flaws in Planning-Based RL Using Tree Explanations. *IJCAI-PRICAI 2020 Workshop on Explainable Artificial Intelligence (XAI)*, Virtual, Jan. 2021.
- [W8] **Minsuk Kahng** and Duen Horng (Polo) Chau. How Does Visualization Help People Learn Deep Learning? Evaluation of GAN Lab. *Workshop on Evaluation of Interactive Visual Machine Learning Systems (EVIVA-ML at VIS'19)*, Vancouver, Canada, Oct. 2019.
- [W7] Ángel Alexander Cabrera, **Minsuk Kahng**, Fred Hohman, Jamie Morgenstern, and Duen Horng (Polo) Chau. Discovery of Intersectional Bias in Machine Learning Using Automatic Subgroup Generation. *Debugging Machine Learning Models Workshop (at ICLR'19)*, New Orleans, LA, May 2019.
- [W6] **Minsuk Kahng**, Dezhi Fang, and Duen Horng (Polo) Chau. Visual Exploration of Machine Learning Results using Data Cube Analysis. *Workshop on Human-In-the-Loop Data Analytics (HILDA at SIGMOD'16)*, San Francisco, CA, June 2016. [Deployed on Facebook ML platform.](#)
- [W5] Rahul C. Basole, Hyunwoo Park, Mayank Gupta, Mark Braunstein, Duen Horng (Polo) Chau, Michael Thompson, Vikas Kumar, Robert Pienta, and **Minsuk Kahng**. A Visual Analytics Approach to Understanding Care Process Variation and Conformance. *Workshop on Visual*

Analytics in Healthcare (at VIS'15), Chicago, IL, Oct. 2015.

- [W4] Rahul C. Basole, Hyunwoo Park, Vikas Kumar, Mark Braunstein, James Bost, Duen Horng (Polo) Chau, and **Minsuk Kahng**. Bicentric Visualization of Pediatric Asthma Care Process Activities. *Workshop on Visualizing Electronic Health Record Data (at VIS'14)*, Paris, France, Nov. 2014.
- [W3] Yiqi Chen, Zhiyuan Lin, Robert Pienta, **Minsuk Kahng**, and Duen Horng (Polo) Chau. Towards Scalable Graph Computation on Mobile Devices. *Workshop on Scalable Machine Learning: Theory and Applications (at BigData'14)*, Washington, D.C., Oct. 2014.
- [W2] Vikas Kumar, Hyunwoo Park, Rahul C. Basole, Mark Braunstein, **Minsuk Kahng**, Duen Horng (Polo) Chau, Acar Tamersoy, Daniel A. Hirsh, Nicoleta Serban, James Bost, et al. Exploring Clinical Care Processes Using Visual and Data Analytics: Challenges and Opportunities. *Workshop on Data Science for Social Good (at KDD'14)*, New York, NY, Aug. 2014.
- [W1] **Minsuk Kahng**, Sangkeun Lee, and Sang-goo Lee. Ranking Objects by Following Paths in Entity-Relationship Graphs. *ACM Workshop for Ph.D. Students in Information and Knowledge Management (at CIKM'11)*, Glasgow, UK, Oct. 2011. [Best Paper Award](#).

Invited Papers

- [I1] Robert Pienta, James Abello, **Minsuk Kahng**, and Duen Horng (Polo) Chau. Scalable Graph Exploration and Visualization: Sensemaking Challenges and Opportunities. *International Conference on Big Data and Smart Computing*, Jeju, South Korea, Feb. 2015.

Book Chapters

- [B2] Peter J. Polack Jr., Moushumi Sharmin, Kaya de Barbaro, **Minsuk Kahng**, Shang-Tse Chen, and Duen Horng (Polo) Chau. Exploratory Visual Analytics of Mobile Health Data: Sense-making Challenges and Opportunities. *Mobile Health: Sensors, Analytic Methods, and Applications*. Springer, 2017.
- [B1] Dongjoo Lee, Sung Eun Park, **Minsuk Kahng**, Sangkeun Lee, and Sang-goo Lee. Exploiting Contextual Information from Event Logs for Personalized Recommendation. *Computer and Information Science 2010, Studies in Computational Intelligence*. Springer, 2010.

Preprints

- [A4] Alicia Parrish, Rajat Shinde, Sanket Badhe, Xinyi Bai, Sree Bhargavi Balija, Hua-Rong Chu, Emilio Ferrara, Armstrong Foundjem, Rajat Ghosh, Aakash Gupta, Xuanli He, Ong Chen Hui, Minji Jung, Madhangi Karimanal, Faiza Khan Khattak, Boryoung Kim, Eugenia Kim, Liliya Lavitas, Seok Min Lim, Victor Lu, Jim Moirangthem, Dhivya Nagasubramanian, Deepak Pandita, Sita Rajagopal, Geetha Raju, Evgeniia Razumovskaia, Aravind Reddy, Federico Ricciuti, Nobin Sarwar, Sungpil Shin, Sunayana Sitaram, Snehal Thorat, Tharindu Cyril Weerasooriya, Jasmijn Bastings, Joachim Baumann, Kongtao Chen, Murali Emani, Mariya Hendriksen, Jiho Jin, Jun Seong Kim, Younghoon Ko, Alicja Kwasniewska, Minjae Lee, Tom Wei-cyuan Lin, Kashyap Ramanandula Manjusha, Junho Myung, Junyeong Park, Roma Patel, Shyam Ratan, Sudarsun Santhiappan, Priyanka Suresh, Tuesday, Ksheeraj Sai Vepuri, Laura Amortegui-Ordóñez, Claire Dennis, **Minsuk Kahng**, Chris Knotz, Alice Oh, Balaraman Ravindran, Soojung Ryu, William Bartholomew, Hiwot Tesfaye, and Lora Aroyo. Pluralis v0.1: Towards a Multi-cultural, Multimodal, Multilingual Benchmark for AI Risk and Reliability. *arXiv preprint*, 2026.

- [A3] Crystal Qian, Michael Xieyang Liu, Emily Reif, Grady Simon, Nada Hussein, Nathan Clement, James Wexler, Carrie J. Cai, Michael Terry, and **Minsuk Kahng**. The Evolution of LLM Adoption in Industry Data Curation Practices. *arXiv preprint*, 2024.
- [A2] Google DeepMind Gemma Team. Gemma 2: Improving Open Language Models at a Practical Size. *arXiv preprint*, 2024. [Released by Google DeepMind; 2,000 citations.](#)
- [A1] Ian Tenney, Ryan Mullins, Bin Du, Shree Pandya, **Minsuk Kahng**, and Lucas Dixon. Interactive Prompt Debugging with Sequence Saliency. *arXiv preprint*, 2024. [Open-sourced as part of Google’s Responsible Toolkit for Gemini open models \[link\].](#)

Teaching Experience

Seoul National University

Fall 2026 991.306A. **Creative Fusion Project** (창의융합프로젝트)

Yonsei University

Spring 2026 CSI 6305. **Responsible AI**

Spring 2026 CAS 4150. **Introduction to Data Visualization**

Fall 2025 GCB 6114. **Data Visualization**

Fall 2025 CAS 4102. **AI Capstone Design**

Spring 2025 CAS 4150. **Introduction to Data Visualization**
<https://minsuk.com/courses/cas4150-spring2025/>

Spring 2025 CSI 7110. **Topics in Responsible AI**
<https://minsuk.com/courses/csi7110-spring2025/>

Oregon State University

Spring 2022 CS 565. **Human-Computer Interaction**

Winter 2022 CS 499/549. **Visual Analytics**
<https://minsuk.com/courses/cs499-549-winter2022/>

Spring 2021 CS 565. **Human-Computer Interaction**
<https://minsuk.com/courses/cs565-spring2021/>

Fall 2020 CS 539. **Selected Topics in AI: Data Visualization for Machine Learning**
<https://minsuk.com/courses/cs539-002-fall2020/>

Spring 2020 CS 565. **Human-Computer Interaction**

Others

Fall 2012 21003390. **Databases for Statistics**
Instructor. Sookmyung Women’s University

Fall 2017 CS 6400. **Database Systems Concepts and Design**
Teaching Assistant. Georgia Institute of Technology (Instructor: Sham Navathe)

Fall 2014 CSE 6242 / CX 4242. **Data and Visual Analytics**
Teaching Assistant. Georgia Institute of Technology (Instructor: Polo Chau)

Student Advising & Mentoring

Student Advising at Yonsei University

2025–present Minjae Lee, M.S. Student in Computer Science [c28, c30-31, c33]
2025–present Minji Jung, M.S. Student in Computer Science [c31]
2025–2026 Yejin Kim, Undergraduate Intern [c31]
2025–2026 Sarang Choi, Undergraduate Intern [c31]
2025–2026 Dasol Choi, Project Advising [c32-33]

Graduate Student Advising at Oregon State University

2021–2022 Eric Slyman, Ph.D. in Artificial Intelligence & Computer Science (co-advised with Stefan Lee) [c21] (Placed: Adobe Research)
2021–2022 Montaser Hamid, Ph.D. Student in Computer Science [c19]
2021–2022 Yashwanthi Anand, M.S./Ph.D. Student in Computer Science [c19]
2020–2022 Delyar Tabatabai, M.S. in Computer Science [c14,18-19, J5] (Placed: Apple)
2020–2022 Anita Ruangrotsakun, B.S./M.S. in Computer Science [c14,18-19, w20] (Placed: Momento)
2020–2022 Dayeon Oh, M.S. in Computer Science [w20] (Placed: LG)
2020–2021 Roli Khanna, M.S. in Computer Science [c13, J7, w18,20] (Placed: Microsoft)

Undergraduate Student Research Advising at Oregon State

2021–2022 Donald R. Bertucci, Computer Science (URSA Engage, Research Assistant) [c19, w19] (Placed: PhD Student at Georgia Tech)
2021–2022 Melissa Perez, Computer Science (Ecampus) (SRC Program) [c18-19]
2020–2021 Mark Ser, Computer Science (STEM Leaders) [w20] (Placed: Microsoft)
2020–2021 Kristina Lee, Computer Science (Research Assistant) [w20] (Placed: Oracle)
2020–2021 Thuy-Vy Nguyen, Computer Science (Research Assistant) [c15, w20] (Placed: Oracle)
2020–2020 Junhyeok “Derek” Jeong, Computer Science

Graduate Student Committee at Oregon State University

Eric Slyman, Matthew L. Olson, Saeed Khorram, Souti “Rini” Chattopadhyay, Amreeta Chatterjee, Andrew Anderson, Malinda Dilhara, Sanad Saha, Christopher Buss, Josiah Blaisdell, Brett Stoddard, Kin-Ho Lam, and Vaishnavi Rangarajan

Student Mentoring at Georgia Tech

2018–2019 Ángel Alexander Cabrera, B.S. in Computer Science, Georgia Tech [c10, w15]
2016–2017 Dezhi (Andy) Fang, B.S. in Computer Science, Georgia Tech [w13-14]
2013–2016 Zhiyuan (Jerry) Lin, B.S. in Computer Science, Georgia Tech [c3-4,7, w5,7,10]
2015–2016 Peter Polack, M.S. in Computer Science, Georgia Tech [J3, w11]
2014–2015 Mayank Gupta, M.S. in Computer Science, Georgia Tech [w9]
2013–2013 Ganesh Parameswaran, M.S. in Computer Science, Georgia Tech [c7, w10]

Invited Talks

What is “Good” AI? Human-Centered Tools for Responsible AI

July 2026 Invited Talk, KoSAIM (Korean Society of AI in Medicine) 2026 Summer School, Seoul, Korea

Who Defines “Good” AI? Rethinking Evaluation Beyond Static Benchmarks

July 2026 Invited Talk, Seoul Forum on AI Safety & Security, Seoul, South Korea

June 2026 Invited Talk, AI Safety Institute (AIS), Seongnam, South Korea

Human-Data Interaction for Responsible AI

Nov. 2025 Guest Lecture, Yonsei University (Internal), Seoul, South Korea

Visual Analytics for Responsible AI

June 2025 **Keynote**, SIGCHI Korea Local Chapter Summer Event, KAIST, Daejeon, South Korea

When AI Fails: Human-Centered Perspectives

May 2025 Guest Lecture, Yonsei University (Internal), Seoul, South Korea

Visual Analytics for Responsible Development of LLMs

Oct. 2024 **Keynote**, 2nd Korea Visualization Workshop, Seoul National University, Seoul, South Korea

Visual Analytics for Responsible AI

Apr. 2024 Invited Talk, The Kim Jaechul Graduate School of AI, KAIST, South Korea

Panel Speaker: Pathways for PhDs featuring Google PhD Fellow Alumni

Aug. 2023 2023 Global PhD Fellowship Summit, Google

Human-Centered AI through Scalable Visual Data Analytics

March 2022 University of Minnesota, Department of Computer Science & Engineering, Minneapolis, MN

March 2022 The Ohio State University, Department of Computer Science & Engineering, Columbus, OH

Visual Analytics for Interpretable and Inclusive AI

March 2022 University of California, San Diego, Halicioğlu Data Science Institute, San Diego, CA

Explainable AI for Non-Experts: Multiple Coordinated Views or Interactive Articles?

Oct. 26, 2021 IEEE VIS 2021 Seattle Satellite, University of Washington, Seattle, WA

Visual Analytics for Machine Learning Interpretability

May 4, 2021 Tech Talk Tuesday Seminar, Oregon State University, Corvallis, OR

After-Action Review (AAR) System for AI

April 21, 2021 DARPA Explainable AI (XAI) Workshop (co-presented with Jed Irvine)

Visual Analytics for Large-Scale ML Systems

Sept. 14, 2020 Dagstuhl Seminar on “Interactive Visualization of Fostering Trust in AI”, Germany

Human-AI Interaction through Scalable Visual Data Analytics

June 17, 2020 Yonsei University, Department of Computer Science, Seoul, South Korea

June 1, 2020 NSF Center on Pervasive Personalized Intelligence, Webinar

Feb. 5, 2020 AI Seminar, Oregon State University, Corvallis, OR

- July 19, 2019 **Human-Centered AI through Scalable Visual Data Analytics**
Seoul National University, Department of Computer Science & Eng., Seoul, South Korea
- Feb-Apr 2019 **Human-Centered AI through Scalable Visual Data Analytics**
Oregon State University, and several other places.
- Aug. 12, 2018 **ActiVis: Visual Exploration of Industry-Scale Deep Neural Network Models**
ACM SIGGRAPH, Vancouver, Canada
- June 10, 2018 **ActiVis: Visual Exploration of Industry-Scale Deep Neural Network Models**
Human-in-the-Loop Data Analytics Workshop at ACM SIGMOD, Houston, TX
- Oct. 31, 2017 **Data Management Challenges in Practical Machine Learning**
Guest Lecture, Georgia Institute of Technology, Atlanta, GA
- May 29, 2018 **Interactive Visual Analytics for Understanding Machine Learning**
KAIST Interaction Lab, Daejeon, South Korea
- Mar. 4, 2015 **GLO-STIX: Graph Visualization Legos**
DARPA PI Meeting, Arlington, VA
- Feb. 10, 2015 **Scalable Graph Exploration and Visualization: Sensemaking Challenges and Opportunities**
International Conference on Big Data and Smart Computing, Jeju, South Korea
- May 2012 **State-of-the-Art Methods and Trends in Information Retrieval**
Guest Lecture, Seoul National University, Seoul, South Korea
- May 18, 2012 **Discovering Semantics in Heterogeneous Data Graphs**
Data Science Meetup in Seoul, NexR, Seoul, South Korea

Funding

Grants & Gifts

- 2021–2022 **NSF** National AI Research Institute
Title: *USDA-NIFA Institute for Agricultural AI for Transforming Workforce and Decision Support*
Total funded: US\$20,000,000 (My share: about \$450,000 expected)
Senior Personnel (Lead PI: Ananth Kalyanaraman).
- 2020–2022 **DARPA** Explainable Artificial Intelligence (XAI)
Title: *xACT: Explanation-Informed Acceptance Testing of Deep Adaptive Programs*
Total funded for initial grant: US\$6,500,000 (My share: about US\$125,000)
Total funded for extension granted in 2020: US\$600,000 (My share: about \$175,000)
Co-PI (Lead PI: Alan Fern).
- 2020–2020 **NSF** Industry-University Collaboration Research Center on Pervasive Personalized Intelligence
Project Title: *Visual Analytics for Scalable AI Debugging*
Total funded to OSU for Year 1: US\$64,000 (My share: about US\$32,000)
Project PI (OSU Site PI: Weng-Keen Wong); Collaboration with NEC and Intel.

2021–2022 **NAVER** AI Lab, Unrestricted Gift
Total funded to me for Year 1: US\$100,000
Collaboration and advising on data visualization and HCI for large-scale language models.

2021–2022 **Google Cloud** Research Credit
Total funded: US\$5,000.

Graduate Fellowships

2018–2019 **Google PhD Fellowship**. Full Tuition + US\$35,000 for 2 years.

2014–2017 **NSF Graduate Research Fellowship**. Full Tuition + US\$34,000 for 3 years.

Professional Service

Conference Organizing Committee

Publication Chairs, IEEE Visualization and Visual Analytics Conference (VIS'24, 25, 26)

Posters Chair, IEEE Pacific Visualization Conference (PacificVis'27)

Best Paper Award Committee, VisNotes, IEEE Pacific Visualization Conference (PacificVis'25)

Webmaster and Web Designer, ACM Conference on Web Search and Data Mining (WSDM'16)

Workshop Co-organizer

KDD 2018 Workshop on Interactive Data Exploration and Analytics (IDEA)

Journal Co-editor

ACM Transactions on Interactive Intelligent Systems, Special Issue Highlights of IUI 2019

Program Committee

Conference ACM Conference on Human Factors in Computing Systems (CHI'27) (Associate Chair)
Annual Conference on Neural Information Processing Systems (NeurIPS'26) (Area Chair)
IEEE Visualization and Visual Analytics Conference (VIS'21, 22, 23, 25)
ACM International Conference on Intelligent User Interfaces (IUI'19, 20, 21, 22, 23, 24)
AAAI Conference on Artificial Intelligence (AAAI'21, 22)
ACM Conference on Web Search and Data Mining (Demo) (WSDM'22)
IEEE Visualization Conference (Short Papers) (VIS'20)
SIAM International Conference on Data Mining (SDM'20)
ACM International Conference on Information and Knowledge Management (Demo) (CIKM'19)
ACM International Conference on Intelligent User Interfaces (Poster and Demo) (IUI'19)

Workshop Workshop on Human-In-the-Loop Data Analytics (HILDA at SIGMOD'22)
Workshop on Visualization Meets AI (at PacificVis'20)
Symposium on Visualization in Data Science (at VIS'18, 19)
Workshop on Big Data Tools, Methods, and Use Cases for Innovative Scientific Discovery (at BigData'19)
Workshop on Visualization for AI Explainability (at VIS'18)

KDD Workshop on Interactive Data Exploration and Analytics (IDEA) (at KDD'16, 17)
Workshop on Visual Analytics for Deep Learning (at VIS'17)

Paper Reviewer (Invited)

- | | |
|------------|---|
| Journal | IEEE Transactions on Visualization and Computer Graphics (TVCG) ('19, 21, 22, 23, 24, 25)
ACM Transactions on Interactive Intelligent Systems (TiiS) ('20, 26)
ACM Transactions on Intelligent Systems and Technology (TIST) ('20, 22)
Distill ('19)
ACM Transactions on Computer-Human Interaction (TOCHI) ('18) |
| Conference | ACM Conference on Human Factors in Computing Systems (CHI'18, 19, 21, 22, 24, 25, 26)
ACM Conference on Fairness, Accountability, and Transparency (FAccT'26)
Annual Conference on Neural Information Processing Systems (Ethics Reviews) (NeurIPS'25)
EG/VGTC Conference on Visualization (EuroVis'18, 25)
ACM Symposium on User Interface Software and Technology (UIST'23)
IEEE Visualization Conference (VIS'18, 19, 20)
ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW'20) |

Grant Proposal Reviews

National Science Foundation (NSF) CISE Proposal Review Panelist (2024)
Google Academic Research Awards Reviewer (2024)

Open-Sourced & Deployed Software

Open Source Contributions

LLM Comparator: An Interactive Visualization Tool for LLM Evaluation [C23, P13]
Lead Author. <https://github.com/PAIR-code/llm-comparator>

Adversarial Nibbler Data Exploration: The Nibbler Challenge Dataset for Red Teaming [C22]
Lead Contributor. <https://github.com/PAIR-code/adversarial-nibbler-vis>

Learning Interpretability Tool (LIT): A Visual Interactive Tool for Understanding NLP Models
Contributor. <https://github.com/PAIR-code/lit>

TensorFlow.js: A WebGL-accelerated JavaScript Library for Deep Learning Models
Contributor. <https://github.com/tensorflow/tfjs>

GAN Lab: A Visual Experimentation Tool for Generative Adversarial Networks [C9]
Lead Author. <https://github.com/poloclub/ganlab>

South Korea Map: Geodata for Administrative Divisions of South Korea for Visualization
Contributor. <https://github.com/southkorea/southkorea-maps>

CSRankings: Rankings of Computer Science Departments (on Data Visualization) [C24]
Contributor. <https://github.com/emeryberger/CSrankings>

Deployed Systems for Industry

LLM Comparator: Visual Analysis of LLM Evaluations [C23, P13]

Project lead. Deployed on Google's LLM evaluation platforms

ActiVis: Visualization of Deep Learning Models at Facebook [C8]

Project lead. Deployed on FBLeaer Flow, Facebook's machine learning platform

ML Cube: Visual Exploration of Machine Learning Results using Data Cube Analysis [W6]

Project lead. Deployed on FBLeaer, Facebook's machine learning platform

Patent

Neural Network Model Visualization

Pierre Andrews, Minsuk Kahng, and Aditya Kalro.

Project lead. Part of internship work from Facebook [C8].

US11461629B1, <https://patents.google.com/patent/US11461629B1/en>

Last updated: September 2026