

Minsuk Kahng 강민석

✉ kahng@google.com
🏠 <https://minsuk.com>

My research aims to empower researchers and practitioners to gain insights from interactive **data visualization**, enabling them to **responsibly develop AI systems**. To achieve this goal, I build novel visual analytics tools that help these people explore large datasets for AI.

My work has been supported by NSF, DARPA, Google, and NAVER, deployed on the industry ML platforms (e.g., LLM Comparator for Google, ActiVis for Facebook), and recognized by prestigious awards (e.g., NSF Graduate Fellowship, Google PhD Fellowship, Georgia Tech Dissertation Award).

Areas of expertise: Visual Analytics, Data Visualization, Responsible AI, Explainable AI, Human-Computer Interaction

Employment

- 2022–present **Google** DeepMind, Atlanta, GA
Senior Research Scientist, People+AI Research (PAIR) Team
Foundational Research Unit, Google DeepMind (since May 2024)
Responsible AI, Google Research (until May 2024)
- 2019–2022 **Oregon State University**, Corvallis, OR, USA
Assistant Professor of Computer Science
School of Electrical Engineering and Computer Science, College of Engineering
- 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 18 **latest** and **highly-cited** publications. Please see Google Scholar for citations: <https://scholar.google.com/citations?user=vieI1GYAAAAJ> (total 3,000, h-index 23).

*: my student advisee; Click the titles for PDF.

- [S18] 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. *Preliminary Version in ACM Conference on Human Factors in Computing Systems (CHI Late-Breaking Work)*, 2024. [Deployed on Google's LLM evaluation platforms. Featured at Google I/O 2024 as part of Responsible GenAI Toolkit for Gemini open models \[link\]](#).
- [S17] Automatic Histograms: Leveraging Language Models for Text Dataset Exploration. Emily Reif, Crystal Qian, James Wexler, and **Minsuk Kahng**. *ACM Conference on Human Factors in Computing Systems (CHI Late-Breaking Work)*, 2024.
- [S16] Understanding the Dataset Practitioners Behind Large Language Model Development. Crystal Qian, Emily Reif, and **Minsuk Kahng**. *ACM Conference on Human Factors in Computing Systems (CHI Late-Breaking Work)*, 2024.
- [S15] 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, et al. *ACM Conference on Fairness, Accountability, and Transparency (FAccT)*, 2024. [Featured in Google Research blog \[link\]](#).
- [S14] Visualizing Linguistic Diversity of Text Datasets Synthesized by Large Language Models. Emily Reif, **Minsuk Kahng**, and Savvas Petridis. *IEEE Visualization Conference (VIS Short)*, 2023.
- [S13] VLSlice: Interactive Vision-and-Language Slice Discovery. Eric Slyman*, **Minsuk Kahng**, and Stefan Lee. *International Conference on Computer Vision (ICCV)*, 2023.

- [S12] 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.
- [S11] FitVid: Responsive and Flexible Video Content Adaptation. Jeongyeon Kim, Yubin Choi, **Minsuk Kahng**, and Juho Kim. *ACM Conference on Human Factors in Computing Systems (CHI)*, 2022.
- [S10] One Explanation is Not Enough: Structured Attention Graphs for Image Classification. Vivswan Shitole, Li Fuxin, **Minsuk Kahng**, Prasad Tadepalli, and Alan Fern. *Conference on Neural Information Processing Systems, (NeurIPS)*, 2021.
- [S9] “Why did my AI agent lose?”: Visual Analytics for Scaling Up After-Action Review. Delyar Tabatabai*, Anita Ruangrotsakun*, Jed Irvine, Jonathan Dodge, Zeyad Shureih, Kin-Ho Lam*, Margaret Burnett, Alan Fern, and **Minsuk Kahng**. *IEEE Visualization Conference (VIS Short)*, 2021. [Featured in DARPA XAI Workshop](#).
- [S8] Contrastive Identification of Covariate Shift in Image Data. Matthew L. Olson, Thuy-Vy Nguyen*, Gaurav Dixit, Neale Ratzlaff, Weng-Keen Wong, and **Minsuk Kahng**. *IEEE Visualization Conference (VIS Short)*, 2021.
- [S7] CNN Explainer: Learning Convolutional Neural Networks with Interactive Visualization. Zijie J. Wang, Robert Turko, Omar Shaikh, Haekyu Park, Nilaksh Das, Fred Hohman, **Minsuk Kahng**, and Duen Horng (Polo) Chau. *IEEE Transactions on Visualization and Computer Graphics (VIS)*, 2020. [200 citations; Invited to present at SIGGRAPH’21 as a top paper](#).
- [S6] How Does Visualization Help People Learn Deep Learning? Evaluating GAN Lab with Observational Study and Log Analysis. **Minsuk Kahng** and Duen Horng (Polo) Chau. *IEEE Visualization Conference (VIS Short)*, 2020.
- [S5] FairVis: Visual Analytics for Discovering Intersectional Bias in Machine Learning. Ángel Alexander Cabrera, Will Epperson, Fred Hohman, **Minsuk Kahng**, Jamie Morgenstern, and Duen Horng (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 Duen Horng (Polo) Chau. *IEEE Transactions on Visualization and Computer Graphics*, 25(8), 2019. [600 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, Duen Horng (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 Duen Horng (Polo) Chau. *IEEE Transactions on Visualization and Computer Graphics (VIS)*, 2017. [Deployed on Facebook ML platform; Invited to present at SIGGRAPH’18 as one of the selected top papers; 400 citations: 8th most cited paper published in VIS and TVCG journal in 2018-2023 \[ref\]](#).
- [S1] Interactive Browsing and Navigation in Relational Databases. **Minsuk Kahng**, Shamkant B. Navathe, John T. Stasko, and Duen Horng (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 2,800, h-index 23).

Refereed Conference Papers

- [C23] 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'24)*, Oct. 2024. [22.3% acceptance rate] [Deployed on Google's LLM evaluation platforms. Featured at Google I/O 2024 as part of Responsible GenAI Toolkit for Gemini open models](#) [link].
- [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.
- [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. 400 citations: top 8th most cited in IEEE TVCG in the past 5 years.](#)
- [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. **600 citations: 2nd most cited paper published in the IEEE TVCG journal in the past five years** [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 Workshop, Poster, and Demo Papers

- [W23] **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 Late-Breaking Work)*, 2024. [33.9% acceptance rate]. [Deployed on Google's LLM evaluation platforms. Featured at Google I/O 2024 as part of Responsible GenAI Toolkit for Gemini open models \[link\]](#).
- [W22] 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 Late-Breaking Work)*, 2024. [33.9% acceptance rate].
- [W21] 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 Late-Breaking Work)*, 2024. [33.9% acceptance rate].
- [W20] 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.
- [W19] 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/>)
- [W18] 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.
- [W17] 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.
- [W16] **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.
- [W15] 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.
- [W14] 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.
- [W13] **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](#).

- [W12] 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.
- [W11] 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.
- [W10] 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.
- [W9] 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.
- [W8] 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.
- [W7] 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.
- [W6] 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.
- [W5] 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.
- [W4] **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.
- [W3] **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](#).
- [W2] **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.
- [W1] 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.

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: Sensemaking 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

- [P3] Gemma Team et al. Gemma 2: Improving Open Language Models at a Practical Size. *arXiv preprint*, 2024.
- [P2] 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].
- [P1] Seongmin Lee, Zijie J. Wang, Aishwarya Chakravarthy, Alec Helbling, ShengYun Peng, Mansi Phute, Duen Horng Chau, and **Minsuk Kahng**. LLM Attributor: Interactive Visual Attribution for LLM Generation. *arXiv preprint*, 2024.

Teaching Experience

Instructor

- Winter 2022 **CS 499/549: Visual Analytics**
Oregon State University, Corvallis, OR
New course on data visualization fundamentals and interactive large data analysis practices
<https://minsuk.com/courses/cs499-549-winter2022/>
- Spring 2022 **CS 565: Human-Computer Interaction**
Spring 2021 Oregon State University, Corvallis, OR
Spring 2020 HCI principles and user interface design, prototyping, and evaluation
<https://minsuk.com/courses/cs565-spring2021/>
- Fall 2020 **CS 539: Selected Topics in AI: Data Visualization for ML**
Oregon State University, Corvallis, OR
Recent literature on interactive data visualization approaches to ML interpretability
<https://minsuk.com/courses/cs539-002-fall2020/>
- Fall 2012 **Databases for Statistics**
Sookmyung Women's University, Seoul, South Korea
Introductory database course designed for undergraduate students in statistics

Teaching Assistant

- Fall 2017 **CS 6400: Database Systems Concepts and Design**
Instructor: Sham Navathe
Georgia Institute of Technology, Atlanta, GA
Graduate-level database course
- Fall 2014 **CSE 6242 / CX 4242: Data and Visual Analytics**
Instructor: Polo Chau
Georgia Institute of Technology, Atlanta, GA
Graduate-level course on analyzing and visualizing big data

Student Advising & Mentoring

Graduate Student Advising at Oregon State University

- 2021–2022 Eric Slyman, Ph.D. Student in Artificial Intelligence & Computer Science (co-advised with Stefan Lee) [C21]
- 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. Student in Computer Science [C14,18-19, J5]
- 2020–2022 Kin-Ho Lam, M.S. in Artificial Intelligence (co-advised with Alan Fern) [C14,19, W18]
- 2020–2022 Anita Ruangrotsakun, B.S./M.S. Student in Computer Science [C14,18-19, W20]
- 2020–2022 Dayeon Oh, M.S. in Computer Science (Graduated)[W20]
- 2020–2021 Roli Khanna, M.S. in Computer Science (Graduated)[C13, J7, W18,20]
(Thesis: *Assessing and Finding Faults in AI: Two Empirical Studies*)

Undergraduate Student Research Advising at Oregon State

- 2021–2022 Donald R. Bertucci, Computer Science (URSA Engage, Research Assistant) [C19,W19]
- 2021–2022 Melissa Perez, Computer Science (Ecampus) (SRC Program) [C18-19]
- 2020–2022 Anita Ruangrotsakun, Computer Science (Research Assistant) [C14,18-19, W20]
- 2020–2021 Mark Ser, Computer Science (STEM Leaders) [W20]
- 2020–2021 Kristina Lee, Computer Science (Research Assistant) [W20]
- 2020–2021 Thuy-Vy Nguyen, Computer Science (Research Assistant) [C15, W20]

Student Committee at Oregon State University

- 2021–present Eric Slyman, Ph.D. Student (Advisor: Stefan Lee; Previously co-advised by me)
- 2021–2023 Matthew L. Olson, Ph.D. in Artificial Intelligence (Advisor: Weng-Keen Wong)
- 2021–2022 Saeed Khorram, Ph.D. in Computer Science (Advisor: Fuxin Li)
- 2021–2022 Souti Chattopadhyay, Ph.D. in Computer Science (Advisor: Anita Sarma)
- 2021–2022 Amreeta Chatterjee, Ph.D. Student (Advisor: Anita Sarma)
- 2020–2022 Andrew Anderson, Ph.D. Student (Advisor: Margaret Burnett)
- 2020–2022 Malinda Dilhara, Ph.D. Student (Advisor: Danny Dig)
- 2020–2022 Sanad Saha, Ph.D. Student (Advisor: Arash Termehchy)
- 2020–2022 Christopher Buss, Ph.D. Student (Advisor: Arash Termehchy)
- 2020–2022 Josiah Blaisdell, Ph.D. Student (Advisor: Yue Zhang)
- 2022–2022 Brett Stoddard, M.S. in Robotics (Advisor: Heather Knight)
- 2020–2022 Kin-Ho Lam, M.S. in Artificial Intelligence (Advisor: Alan Fern)
- 2020–2022 Vaishnavi Rangarajan, M.S. in Computer Science (Advisor: Raffaele De Amicis)

2020–2022 Rupika Dikkala, M.S. in Computer Science (Advisor: Margaret Burnett)
 2020–2021 Omeed Habibelahian, M.S. in Computer Science (Advisor: Arash Termehchy)
 2021–2022 Puja Agarwal, M.S. Student (Advisor: Margaret Burnett)
 2021–2021 Junhyeok (Derek) Jeong, B.S. Honors in Computer Science (Advisor: Cindy Grimm)

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

Panel Speaker: Pathways for PhDs featuring Google PhD Fellow Alumni
 Aug. 29, 2023 2023 Global PhD Fellowship Summit, Google

Visual Analytics for Interpretable and Inclusive AI
 April 2022 Google, and several other places.

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

Human-Centered AI through Scalable Visual Data Analytics
 July 19, 2019 Seoul National University, Department of Computer Science & Eng., Seoul, South Korea

Human-Centered AI through Scalable Visual Data Analytics
 April 2019 Oregon State University, and several other places.

ActiVis: Visual Exploration of Industry-Scale Deep Neural Network Models
 Aug. 12, 2018 ACM SIGGRAPH, Vancouver, Canada
 June 10, 2018 Human-in-the-Loop Data Analytics Workshop at ACM SIGMOD, Houston, TX
 Sept. 20, 2017 Georgia Institute of Technology, School of Computational Science & Engineering, Atlanta, GA

Data Management Challenges in Practical Machine Learning
 Oct. 31, 2017 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.

Academic Service

Committee at Oregon State University

- 2019–2022 Graduate Admission Committee
School of Electrical Engineering and Computer Science
- 2020–2022 Web Advisory Committee
College of Engineering
- 2021–2022 Area Chair for Software Engineering and Human-Computer Interaction, School of Electrical Engineering and Computer Science

Professional Service

Conference Organizing Committee

- Publication Chairs, IEEE Visualization Conference (VIS'24)
- 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

Journal Co-editor

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

Program Committee

- Conference
 - IEEE Visualization Conference (VIS'21–present)
 - ACM International Conference on Intelligent User Interfaces (IUI'19–present)
 - 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'16–17)
 - Workshop on Visual Analytics for Deep Learning (at VIS'17)

Paper Reviewer

- Journal
 - IEEE Transactions on Visualization and Computer Graphics (TVCG) (2019, 2021–24)
 - ACM Transactions on Interactive Intelligent Systems (TiiS) (2020)

ACM Transactions on Intelligent Systems and Technology (TIST) (2020, 2022)
Distill (2019)
ACM Transactions on Computer-Human Interaction (TOCHI) (2015, 2018)
Expert Systems with Applications (2015)

Conference ACM CHI Conference on Human Factors in Computing Systems (CHI'14, 17–19, 21–22, 24)
ACM Symposium on User Interface Software and Technology (UIST'23)
IEEE Visualization Conference (VIS'18–20)
ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW'20)
EG/VGTC Conference on Visualization (EuroVis'18)
SIAM International Conference on Data Mining (SDM'14, 16–17)
Visualization in Data Science (at VIS'17)
ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'14–16)
ACM International Conference on Intelligent User Interfaces (IUI'16)
ACM Conference on Recommender Systems (RecSys'16)
ACM SIGMOD International Conference on Management of Data (SIGMOD'13)
International Conference on Database Systems for Advanced Applications (DASFAA'11)

Grant Proposal Review

National Science Foundation (NSF) CISE Proposal Review Panelist (2024)

Open-Sourced & Deployed Software

Open-sourced Contributions

LLM Comparator: An Interactive Visualization Tool for LLM Evaluation [W23]
Lead Author. <https://github.com/PAIR-code/llm-comparator>

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
Contributor. <https://github.com/emeryberger/CSrankings>

Deployed Systems for Industry

LLM Comparator: Visual Analysis of LLM Evaluations [W23]
Project lead. Deployed on Google's LLM evaluation platforms

ActiVis: Visualization of Deep Learning Models at Facebook [C8]
Project lead. Deployed on FB Learner Flow, Facebook's machine learning platform

ML Cube: Visual Exploration of Machine Learning Results [W13]
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>

References

Duen Horng (Polo) Chau

Associate Professor, Georgia Institute of Technology

<https://www.cc.gatech.edu/~dchau>

Martin Wattenberg

Professor, Harvard University & Principal Scientist, Google

<https://www.bewitched.com/about.html>

Fernanda B. Viégas

Professor, Harvard University & Principal Scientist, Google

<http://www.fernandaviegas.com/about>

Alex Endert

Associate Professor, Georgia Institute of Technology

<http://va.gatech.edu/endert/>

Shamkant B. Navathe

Professor, Georgia Institute of Technology

<https://www.cc.gatech.edu/people/shamkant-navathe>