

HCAI 2023 Reading Suggestions

- Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267, 1-38.
<https://www.sciencedirect.com/science/article/pii/S0004370218305988>
- Zamfirescu-Pereira, J. D., Wong, R. Y., Hartmann, B., & Yang, Q. (2023, April). Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-21). <https://dl.acm.org/doi/abs/10.1145/3544548.3581388>
- Sarkar, A., Gordon, A. D., Negreanu, C., Poelitz, C., Ragavan, S. S., & Zorn, B. (2022). What is it like to program with artificial intelligence?. *arXiv preprint arXiv:2208.06213*. <https://www.ppig.org/files/2022-PPIG-33rd-sarkar.pdf>
- Danry, V., Pataranutaporn, P., Mao, Y., & Maes, P. (2023, April). Don't Just Tell Me, Ask Me: AI Systems that Intelligently Frame Explanations as Questions Improve Human Logical Discernment Accuracy over Causal AI explanations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-13). <https://dl.acm.org/doi/abs/10.1145/3544548.3580672>
- Mirowski, P., Mathewson, K. W., Pittman, J., & Evans, R. (2023, April). Co-Writing Screenplays and Theatre Scripts with Language Models: Evaluation by Industry Professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-34).
<https://dl.acm.org/doi/abs/10.1145/3544548.3581225>
- Wang, Y., Shen, S., & Lim, B. Y. (2023, April). RePrompt: Automatic Prompt Editing to Refine AI-Generative Art Towards Precise Expressions. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-29).
<https://dl.acm.org/doi/abs/10.1145/3544548.3581402>
- Singh, N., Bernal, G., Savchenko, D., & Glassman, E. L. (2022). Where to hide a stolen elephant: Leaps in creative writing with multimodal machine intelligence. *ACM Transactions on Computer-Human Interaction*.
<https://dl.acm.org/doi/abs/10.1145/3511599>
- Jakesch, M., Bhat, A., Buschek, D., Zalmanson, L., & Naaman, M. (2023, April). Co-writing with opinionated language models affects users' views. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-15).
<https://dl.acm.org/doi/abs/10.1145/3544548.3581196>
- Sarkar, A. (2023, April). Enough With "Human-AI Collaboration". In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-8).
<https://dl.acm.org/doi/abs/10.1145/3544549.3582735>

- Suh, M., Youngblom, E., Terry, M., & Cai, C. J. (2021, May). AI as social glue: uncovering the roles of deep generative AI during social music composition. In *Proceedings of the 2021 CHI conference on human factors in computing systems* (pp. 1-11). <https://dl.acm.org/doi/abs/10.1145/3411764.3445219>
- Advait Sarkar. 2023. Exploring Perspectives on the Impact of Artificial Intelligence on the Creativity of Knowledge Work: Beyond Mechanised Plagiarism and Stochastic Parrots. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work (CHIWORK '23)*. Association for Computing Machinery, New York, NY, USA, Article 13, 1–17. <https://doi.org/10.1145/3596671.3597650>
- Capel, T., & Brereton, M. (2023, April). What is Human-Centered about Human-Centered AI? A Map of the Research Landscape. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-23). <https://dl.acm.org/doi/abs/10.1145/3544548.3580959>
- Kim, S. S., Watkins, E. A., Russakovsky, O., Fong, R., & Monroy-Hernández, A. (2023, April). " Help Me Help the AI": Understanding How Explainability Can Support Human-AI Interaction. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-17). <https://dl.acm.org/doi/abs/10.1145/3544548.3581001>
- Solomon, J. (2016). Heterogeneity in Customization of Recommender Systems By Users with Homogenous Preferences. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16* (pp. 4166–4170). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2858036.2858513>
- Agre, P. E. (2014). Toward a critical technical practice: Lessons learned in trying to reform AI. In *Social science, technical systems, and cooperative work* (pp. 131-157). Psychology Press. <https://www.taylorfrancis.com/chapters/edit/10.4324/9781315805849-8/toward-critical-technical-practice-lessons-learned-trying-reform-ai-philip-agre>
- Liu, M. X., Sarkar, A., Negreanu, C., Zorn, B., Williams, J., Toronto, N., & Gordon, A. D. (2023, April). "What It Wants Me To Say": Bridging the Abstraction Gap Between End-User Programmers and Code-Generating Large Language Models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (pp. 1-31). <https://dl.acm.org/doi/abs/10.1145/3544548.3580817>
- Daee, P., Peltola, T., Vehtari, A., & Kaski, S. (2018, March). User modelling for avoiding overfitting in interactive knowledge elicitation for prediction. In *23rd International Conference on Intelligent User Interfaces* (pp. 305-310). <https://dl.acm.org/doi/abs/10.1145/3172944.3172989>

- Fothergill, S., Mentis, H., Kohli, P., & Nowozin, S. (2012). Instructing people for training gestural interactive systems. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12* (p. 1737). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2207676.2208303>
- Tullio, J., Dey, A. K., Chalecki, J., & Fogarty, J. (2007). How it works: a field study of non-technical users interacting with an intelligent system. In *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07* (p. 31). New York, New York, USA: ACM Press. <http://doi.org/10.1145/1240624.1240630>
- Eiband, M., Völkel, S. T., Buschek, D., Cook, S., & Hussmann, H. (2019). When people and algorithms meet: user-reported problems in intelligent everyday applications. *Proceedings of the 24th International Conference on Intelligent User Interfaces - IUI '19*, 96–106. <https://doi.org/10.1145/3301275.3302262>
- Heer, J. (2019). Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences*, 116(6), 1844–1850. <https://doi.org/10.1073/pnas.1807184115>
- Aoyu Wu, Liwenhan Xie, Bongshin Lee, Yun Wang, Weiwei Cui, and Huamin Qu. 2021. Learning to Automate Chart Layout Configurations Using Crowdsourced Paired Comparison. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 14, 1–13. DOI:<https://doi.org/10.1145/3411764.3445179>
- Litao Yan, Elena L. Glassman, and Tianyi Zhang. 2021. Visualizing Examples of Deep Neural Networks at Scale. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 313, 1–14. DOI:<https://doi.org/10.1145/3411764.3445654>
- Arendt, D., Grace, E., & Volkova, S. (2005). Interactive Machine Learning at Scale with CHISSL. https://www.cs.jhu.edu/~svitlana/papers/AGV_AAAI18_Demo.pdf
- Self, J. Z., Dowling, M., Wenskovitch, J., Crandell, I., Wang, M., House, L., ... North, C. (2018). Observation-Level and Parametric Interaction for High-Dimensional Data Analysis. *ACM Transactions on Interactive Intelligent Systems*, 8(2), 1–36. <https://doi.org/10.1145/3158230>
- Trivedi, G., Pham, P., Hwa, R., Chapman, W., Wiebe, J., & Hochheiser, H. (2015). An Interactive Tool for Natural Language Processing on Clinical Text. *Presented at 4th Workshop on Visual Text Analytics (IUI TextVis 2015)*. Retrieved from <http://vialab.science.uoit.ca/textvis2015/papers/Trivedi-textvis2015.pdf>
- Brooks, M., Amershi, S., Lee, B., Drucker, S. M., Kapoor, A., & Simard, P. (2015). FeatureInsight: Visual Support for Error-Driven Feature Ideation in Text Classification. *Proceedings of IEEE Symposium on Visual Analytics Science and*

Technology (VAST), 105–112. Retrieved from
<http://ieevis.org/year/2015/info/overview-amp-topics/papers-sessions>

Ma, K.-L., & Teoh, S. T. (2003). StarClass: Interactive Visual Classification Using Star Coordinates. *Proceedings of the 3rd SIAM International Conference on Data Mining*, 178–185.

Perer, A., & Wang, F. (2014). Frequency: Interactive Mining and Visualization of Temporal Frequent Event Sequences. In *Proceedings of the 19th international conference on Intelligent User Interfaces - IUI '14* (pp. 153–162). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2557500.2557508>

Ankerst, M., Ester, M., & Kriegel, H.-P. (2000). Towards an effective cooperation of the user and the computer for classification. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 1(212), 179–188. <http://doi.org/http://doi.acm.org/10.1145/347090.347124>

Parisot, O., Didry, Y., Tamisier, T., & Otjacques, B. (2015). Helping predictive analytics interpretation using regression trees and clustering perturbation. *Journal of Decision Systems*, 24(1), 55–72. <http://doi.org/10.1080/12460125.2015.994331>

Amershi, S., & Conati, C. C. (2009). Combining Unsupervised and Supervised Classification to Build User Models for Exploratory. *JEDM-Journal of Educational Data Mining*, 1(1), 1–54.

van den Elzen, S., & van Wijk, J. J. (2011). BaobabView: Interactive construction and analysis of decision trees. In *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)* (pp. 151–160). IEEE.
<http://doi.org/10.1109/VAST.2011.6102453>

de Oliveira, M. C. F., & Levkowitz, H. (2003). From visual data exploration to visual data mining: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 9(3), 378–394. <http://doi.org/10.1109/TVCG.2003.1207445>

Liu, Y., & Salvendy, G. (2007). Interactive Visual Decision Tree Classification. In *Human-Computer Interaction. Interaction Platforms and Techniques* (Vol. 4551, pp. 92–105). Berlin, Heidelberg: Springer Berlin Heidelberg.
http://doi.org/10.1007/978-3-540-73107-8_11

Teoh, S. T., & Ma, K.-L. (2003). PaintingClass: Interactive Construction, Visualization and Exploration of Decision Trees. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '03* (p. 667). New York, New York, USA: ACM Press.
<http://doi.org/10.1145/956750.956837>

- Kapoor, A., Lee, B., Tan, D. S., & Horvitz, E. (2010). Interactive optimization for steering machine classification. In *CHI* (pp. 1343–1352).
<http://doi.org/10.1145/1753326.1753529>
- Talbot, J., Lee, B., Kapoor, A., & Tan, D. S. (2009). EnsembleMatrix: interactive visualization to support machine learning with multiple classifiers. In *Proceedings of the 27th international conference on Human factors in computing systems - CHI 09* (p. 1283). <http://doi.org/10.1145/1518701.1518895>
- Behrisch, M., Korkmaz, F., Shao, L., & Schreck, T. (2014). Feedback-driven interactive exploration of large multidimensional data supported by visual classifier. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)* (pp. 43–52).
<http://doi.org/10.1109/VAST.2014.7042480>
- Mühlbacher, T., Piringer, H., Gratzl, S., Sedlmair, M., & Streit, M. (2014). Opening the Black Box: Strategies for Increased User Involvement in Existing Algorithm Implementations. *IEEE Transactions on Visualization and Computer Graphics*, 20(1), 1643–1652. <http://doi.org/10.1109/TVCG.2014.2346578>
- Mamani, G. M. H., Fatore, F. M., Nonato, L. G., & Paulovich, F. V. (2013). User-driven feature space transformation. *Computer Graphics Forum*, 32(3), 291–299.
<http://doi.org/10.1111/cgf.12116>
- Brown, E. T., Liu, J., Brodley, C. E., & Chang, R. (2012). Dis-function: Learning distance functions interactively. *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)*, 83–92. <http://doi.org/10.1109/VAST.2012.6400486>
- Ankerst, M., Elsen, C., Ester, M., & Kriegel, H.-P. (1999). Visual classification: an interactive approach to decision tree construction. In *Proceedings of the fifth ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '99* (pp. 392–396). New York, New York, USA: ACM Press.
<http://doi.org/10.1145/312129.312298>
- Wu, E., & Madden, S. (2013). Scorpion: Explaining Away Outliers in Aggregate Queries. *Proceedings of the VLDB Endowment*, 6(8), 553–564.
- Bertini, E., & Lalanne, D. (2010). Investigating and reflecting on the integration of automatic data analysis and visualization in knowledge discovery. *ACM SIGKDD Explorations Newsletter*, 11(2), 9. <http://doi.org/10.1145/1809400.1809404>
- Fan, X., Liu, Y., Cao, N., Hong, J., & Wang, J. (2015). MindMiner: Quantifying Entity Similarity via Interactive Distance Metric Learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces Companion - IUI Companion '15* (pp. 93–96). New York, New York, USA: ACM Press.
<http://doi.org/10.1145/2732158.2732173>

- Fan, X., Liu, Y., Cao, N., Hong, J., & Wang, J. (2015). MindMiner: A Mixed-Initiative Interface for Interactive Distance Metric Learning. In J. Abascal, S. Barbosa, M. Fetter, T. Gross, P. Palanque, & M. Winckler (Eds.), (Vol. 9297, pp. 611–628). Cham: Springer International Publishing.
http://doi.org/10.1007/978-3-319-22668-2_47
- Sarkar, A., Blackwell, A. F., Jamnik, M., & Spott, M. (2015). Interaction with uncertainty in visualisations. In *Eurographics/IEEE VGTC Conference on Visualization (EuroVis 2015) Short Papers* (pp. 133–137). <http://doi.org/10.2312/eurovisshort.20151138>
- Sarkar, A., Spott, M., Blackwell, A. F., & Jamnik, M. (2016). Visual discovery and model-driven explanation of time series patterns. In *2016 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)* (Vol. 2016–Novem, pp. 78–86). IEEE. <http://doi.org/10.1109/VLHCC.2016.7739668>
- Rashid Tahir, Brishna Batool, Hira Jamshed, Mahnoor Jameel, Mubashir Anwar, Faizan Ahmed, Muhammad Adeel Zaffar, and Muhammad Fareed Zaffar. 2021. Seeing is Believing: Exploring Perceptual Differences in DeepFake Videos. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 174, 1–16.
DOI:<https://doi.org/10.1145/3411764.3445699>
- van Berkel, Niels, Jorge Goncalves, Daniel Russo, Simo Hosio, and Mikael B. Skov. "Effect of Information Presentation on Fairness Perceptions of Machine Learning Predictors." In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1-13. 2021.
<https://dl.acm.org/doi/10.1145/3411764.3445365>
- Dudley, J. J., & Kristensson, P. O. (2018). A Review of User Interface Design for Interactive Machine Learning. *ACM Transactions on Interactive Intelligent Systems*, 8(2), 1–37. <https://doi.org/10.1145/3185517>
- Gonzalo Mendez, Luis Galárraga, and Katherine Chiluiza. 2021. Showing Academic Performance Predictions during Term Planning: Effects on Students' Decisions, Behaviors, and Preferences. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 22, 1–17. DOI:<https://doi.org/10.1145/3411764.3445718>
- Weisz, J. D., Jain, M., Joshi, N. N., Johnson, J., & Lange, I. (2019). BigBlueBot: teaching strategies for successful human-agent interactions. *Proceedings of the 24th International Conference on Intelligent User Interfaces - IUI '19*, 448–459.
<https://doi.org/10.1145/3301275.3302290>
- Ashktorab, Z., Jain, M., Liao, Q. V., & Weisz, J. D. (2019). Resilient Chatbots: Repair Strategy Preferences for Conversational Breakdowns. *Proceedings of the 2019*

CHI Conference on Human Factors in Computing Systems, 1–12.

<https://doi.org/10.1145/3290605.3300484>

- Chen, N.-C., Suh, J., Verwey, J., Ramos, G., Drucker, S., & Simard, P. (2018). AnchorViz: Facilitating Classifier Error Discovery through Interactive Semantic Data Exploration. In *Proceedings of the 2018 Conference on Human Information Interaction & Retrieval - IUI '18* (pp. 269–280). New York, New York, USA: ACM Press. <https://doi.org/10.1145/3172944.3172950>
- Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, C., Eslami, S. M. A., & Botvinick, M. (2018). Machine Theory of Mind. Retrieved from <http://arxiv.org/abs/1802.07740>
- Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., ... Dennison, D. (2015). Hidden Technical Debt in Machine Learning Systems. *NIPS*, 2494–2502. Retrieved from <http://papers.nips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf>
- Changhoon Oh, Jungwoo Song, Jinhan Choi, Seonghyeon Kim, Sungwoo Lee, and Bongwon Suh. 2018. I Lead, You Help but Only with Enough Details: Understanding User Experience of Co-Creation with Artificial Intelligence. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Paper 649, 13 pages.: <https://doi.org/10.1145/3173574.3174223>
- Dove, G., Halskov, K., Forlizzi, J., & Zimmerman, J. (2017). UX Design Innovation: Challenges for Working with Machine Learning as a Design Material. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems - CHI '17* (pp. 278–288). New York, New York, USA: ACM Press. <https://doi.org/10.1145/3025453.3025739>
- Jasper, R. J., & Blaha, L. M. (2017). Interface Metaphors for Interactive Machine Learning (Vol. 10284, pp. 521–534). http://doi.org/10.1007/978-3-319-58628-1_39
- Horvitz, E. (1999). Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99* (pp. 159–166). New York, New York, USA: ACM Press. <http://doi.org/10.1145/302979.303030>
- Sacha, D., Sedlmair, M., Zhang, L., Lee, J. A., Peltonen, J., Weiskopf, D., ... Keim, D. A. (2017). What you see is what you can change: Human-centered machine learning by interactive visualization. *Neurocomputing*, 0, 1–12. <http://doi.org/10.1016/j.neucom.2017.01.105>
- Shilman, M., Tan, D. S., & Simard, P. (2006). CueTIP: a mixed-initiative interface for correcting handwriting errors. *Proceedings of the ACM Symposium on User*

Interface Software and Technology, 323–332.

<http://doi.org/10.1145/1166253.1166304>

Amershi, S., Fogarty, J., & Weld, D. (2012). ReGroup: interactive machine learning for on-demand group creation in social networks. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12* (p. 21). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2207676.2207680>

Lim, B. Y., & Dey, A. K. (2010). Toolkit to support intelligibility in context-aware applications. In *Proceedings of the 12th ACM international conference on Ubiquitous computing - Ubicomp '10* (p. 13). New York, New York, USA: ACM Press. <http://doi.org/10.1145/1864349.1864353>

Sarkar, A., Jamnik, M., Blackwell, A. F., & Spott, M. (2015). Interactive visual machine learning in spreadsheets. In *2015 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)* (pp. 159–163). IEEE. <http://doi.org/10.1109/VLHCC.2015.7357211>

Gulwani, S., Harris, W. R., & Singh, R. (2012). Spreadsheet data manipulation using examples. *Communications of the ACM*, 55(8), 97. <http://doi.org/10.1145/2240236.2240260>

Valentin, J., Torr, P., Vineet, V., Cheng, M.-M., Kim, D., Shotton, J., ... Izadi, S. (2015). SemanticPaint: Interactive 3D Labeling and Learning at Your Fingertips. *ACM Transactions on Graphics*, 34(5), 1–17. <http://doi.org/10.1145/2751556>

Krause, J., Perer, A., & Bertini, E. (2014). INFUSE: Interactive Feature Selection for Predictive Modeling of High Dimensional Data. *IEEE Transactions on Visualization and Computer Graphics*, 20(12), 1614–1623. <http://doi.org/10.1109/TVCG.2014.2346482>

Bernard, J., Sessler, D., Bannach, A., May, T., & Kohlhammer, J. (2015). A visual active learning system for the assessment of patient well-being in prostate cancer research. In *Proceedings of the 2015 Workshop on Visual Analytics in Healthcare - VAHC '15* (pp. 1–8). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2836034.2836035>

Amershi, S., Chickering, M., Drucker, S. M., Lee, B., Simard, P., & Suh, J. (2015). ModelTracker: Redesigning Performance Analysis Tools for Machine Learning. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15* (pp. 337–346). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2702123.2702509>

Amershi, S., Lee, B., Kapoor, A., Mahajan, R., & Christian, B. (2011). CueT: Human-Guided Fast and Accurate Network Alarm Triage. In *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11* (p.

157). New York, New York, USA: ACM Press.

<http://doi.org/10.1145/1978942.1978966>

Yasha Iravantchi, Karan Ahuja, Mayank Goel, Chris Harrison, and Alanson Sample. 2021. PrivacyMic: Utilizing Inaudible Frequencies for Privacy Preserving Daily Activity Recognition. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 198, 1–13. DOI:<https://doi.org/10.1145/3411764.3445169>

Swati Mishra and Jeffrey M Rzeszotarski. 2021. Designing Interactive Transfer Learning Tools for ML Non-Experts. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 364, 1–15. DOI:<https://doi.org/10.1145/3411764.3445096>

Kai Lukoff, Ulrik Lyngs, Himanshu Zade, J. Vera Liao, James Choi, Kaiyue Fan, Sean A. Munson, and Alexis Hiniker. 2021. How the Design of YouTube Influences User Sense of Agency. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 368, 1–17. DOI:<https://doi.org/10.1145/3411764.3445467>

Dennis Lawo, Thomas Neifer, Margarita Esau, and Gunnar Stevens. 2021. Buying the ‘Right’ Thing: Designing Food Recommender Systems with Critical Consumers. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 85, 1–13. DOI:<https://doi.org/10.1145/3411764.3445264>

Fass, A. M., Bier, E. A., & Adar, E. (2000). PicturePiper: using a re-configurable pipeline to find images on the Web. In *Proceedings of the 13th annual ACM symposium on User interface software and technology - UIST '00* (pp. 51–62). New York, New York, USA: ACM Press. <http://doi.org/10.1145/354401.354411>

Lang, K. J., & Baum, E. B. (1992). Query Learning Can Work Poorly when a Human Oracle is Used. In *IJCNN 1992* (pp. 335–340).

Kulesza, T., Stumpf, S., Burnett, M., Wong, W. K., Riche, Y., Moore, T., ... McIntosh, K. (2010). Explanatory debugging: Supporting end-user debugging of machine-learned programs. In *Proceedings - 2010 IEEE Symposium on Visual Languages and Human-Centric Computing, VL/HCC 2010* (pp. 41–48). <http://doi.org/10.1109/VLHCC.2010.15>

Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. (2014). Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine*, 35(4). <http://doi.org/10.1609/aimag.v35i4.2513>

- Ware, M., Frank, E., Holmes, G., Hall, M., & Witten, I. H. (2001). Interactive machine learning: letting users build classifiers. *International Journal of Human-Computer Studies*, 55(3), 281–292. <http://doi.org/10.1006/ijhc.2001.0499>
- Sanchis-Trilles, G., Ortiz-Martínez, D., Civera, J., Casacuberta, F., Vidal, E., & Hoang, H. (2008). Improving Interactive Machine Translation via Mouse Actions. *EMNLP-2008: Conference on Empirical Methods in Natural Language Processing*, (October), 485–494.
- Ko, A. J., & Myers, B. A. (2008). Debugging reinvented: asking why and why not questions about program behaviour. *2008 ACM/IEEE 30th International Conference on Software Engineering*, 301–310. <http://doi.org/10.1145/1368088.1368130>
- Lim, B. Y., Dey, A. K., & Avrahami, D. (2009). Why and why not explanations improve the intelligibility of context-aware intelligent systems. *Proceedings of the 27th International Conference on Human Factors in Computing Systems - CHI 09*, 2119–2129. <http://doi.org/10.1145/1518701.1519023>
- Kulesza, T., Burnett, M., Wong, W., & Stumpf, S. (2015). Principles of Explanatory Debugging to Personalize Interactive Machine Learning. In *Proceedings of the 20th International Conference on Intelligent User Interfaces - IUI '15* (pp. 126–137). <http://doi.org/10.1145/2678025.2701399>
- Amershi, S., Fogarty, J., Kapoor, A., & Tan, D. S. (2011). Effective End-User Interaction with Machine Learning. *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 1529–1532.
- Kulesza, T., Stumpf, S., Burnett, M., & Kwan, I. (2012). Tell me more? The Effects of Mental Model Soundness on Personalizing an Intelligent Agent. In *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12* (p. 1). <http://doi.org/10.1145/2207676.2207678>
- Groce, A., Kulesza, T., Zhang, C., Shamasunder, S., Burnett, M., Wong, W.-K., ... McIntosh, K. (2014). You Are the Only Possible Oracle: Effective Test Selection for End Users of Interactive Machine Learning Systems. *IEEE Transactions on Software Engineering*, 40(3), 307–323. <http://doi.org/10.1109/TSE.2013.59>
- Kulesza, T., Stumpf, S., Wong, W.-K., Burnett, M. M., Perona, S., Ko, A. J., & Oberst, I. (2011). Why-oriented end-user debugging of naive Bayes text classification. *ACM Transactions on Interactive Intelligent Systems*. <http://doi.org/10.1145/2030365.2030367>

- Fiebrink, R. (2010). Real-time interaction with supervised learning. In *Proceedings of the 28th of the international conference extended abstracts on Human factors in computing systems* (pp. 2935–2938). <http://doi.org/10.1145/1753846.1753889>
- Kulesza, T., Amershi, S., Caruana, R., Fisher, D., & Charles, D. (2014). Structured labeling for facilitating concept evolution in machine learning. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14* (pp. 3075–3084). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2556288.2557238>
- Kulesza, T., Wong, W., Stumpf, S., Perona, S., White, R., Burnett, M. M., ... Ko, A. J. (2009). Fixing the program my computer learned. *Proceedings of the 13th International Conference on Intelligent User Interfaces - IUI '09*, 187. <http://doi.org/10.1145/1502650.1502678>
- Fogarty, J., Tan, D. S., Kapoor, A., & Winder, S. (2008). CueFlik: interactive concept learning in image search. *Proceeding of the Twenty-Sixth Annual CHI Conference on Human Factors in Computing Systems - CHI '08*, 29. <http://doi.org/10.1145/1357054.1357061>
- Fails, J. A., & Olsen, D. R. (2003). Interactive machine learning. *Proceedings of the 8th International Conference on Intelligent User Interfaces - IUI '03*, 39. <http://doi.org/10.1145/604050.604056>
- Mayer, M., Soares, G., Grechkin, M., Le, V., Marron, M., Polozov, O., ... Gulwani, S. (2015). User Interaction Models for Disambiguation in Programming by Example. In *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST '15* (pp. 291–301). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2807442.2807459>
- Weld, D. S., Fogarty, J., & Skinner, M. (2008). Intelligence in Wikipedia. *Artificial Intelligence, pages*(April), 1609–1614. <http://doi.org/10.1.1.141.6970>
- Kapoor, A., Lee, B., Tan, D., & Horvitz, E. (2012). Performance and preferences: Interactive refinement of machine learning procedures. *Proceedings of the National Conference on Artificial Intelligence*, 2, 1578–1584. Retrieved from <http://www.scopus.com/inward/record.url?eid=2-s2.0-84868291810&partnerID=tZOtx3y1>
- Patel, K., Drucker, S. M., Fogarty, J., Kapoor, A., & Tan, D. S. (2011). Using multiple models to understand data. *IJCAI International Joint Conference on Artificial Intelligence*, 1723–1728. <http://doi.org/10.5591/978-1-57735-516-8/IJCAI11-289>
- Patel, K., Bancroft, N., Drucker, S. M., Fogarty, J., Ko, A. J., & Landay, J. (2010). Gestalt: Integrated Support for Implementation and Analysis in Machine Learning.

In *Proceedings of the 23rd annual ACM symposium on User interface software and technology - UIST '10* (p. 37). New York, New York, USA: ACM Press.
<http://doi.org/10.1145/1866029.1866038>

Patel, K., Fogarty, J., Landay, J. a., & Harrison, B. (2008). Investigating statistical machine learning as a tool for software development. In *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08* (p. 667). New York, New York, USA: ACM Press.
<http://doi.org/10.1145/1357054.1357160>

Amershi, S., Fogarty, J., Kapoor, A., & Tan, D. (2010). Examining multiple potential models in end-user interactive concept learning. In *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10* (p. 1357). New York, New York, USA: ACM Press.
<http://doi.org/10.1145/1753326.1753531>

Vidulin, V., Bohanec, M., & Gams, M. (2014). Combining human analysis and machine data mining to obtain credible data relations. *Information Sciences*, 288, 254–278.
<http://doi.org/10.1016/j.ins.2014.08.014>

Sarkar, A. (2016). Constructivist design for interactive machine learning. In *Conference on Human Factors in Computing Systems - Proceedings* (Vol. 07-12-May-).
<http://doi.org/10.1145/2851581.2892547>

Petter Bae Bae Brandtzæg, Marita Skjuve, Kim Kristoffer Kristoffer Dysthe, and Asbjørn Følstad. 2021. When the Social Becomes Non-Human: Young People's Perception of Social Support in Chatbots. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 257, 1–13.
DOI:<https://doi.org/10.1145/3411764.3445318>

Zohar Gilad, Ofra Amir, and Liat Levontin. 2021. The Effects of Warmth and Competence Perceptions on Users' Choice of an AI System. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 583, 1–13.
DOI:<https://doi.org/10.1145/3411764.3446863>

Minha Lee, Peter Ruijten, Lily Frank, Yvonne de Kort, and Wijnand IJsselsteijn. 2021. People May Punish, But Not Blame Robots. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 715, 1–11.
DOI:<https://doi.org/10.1145/3411764.3445284>

Soomi Park, Patrick G. T. Healey, and Antonios Kaniadakis. 2021. Should Robots Blush? *Proceedings of the 2021 CHI Conference on Human Factors in Computing*

Systems. Association for Computing Machinery, New York, NY, USA, Article 717, 1–14. DOI:<https://doi.org/10.1145/3411764.3445561>

Angel Hsing-Chi Hwang and Andrea Stevenson Won. 2021. IdeaBot: Investigating Social Facilitation in Human-Machine Team Creativity. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 712, 1–16. DOI:<https://doi.org/10.1145/3411764.3445270>

Carrie J Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2021. Onboarding Materials as Cross-functional Boundary Objects for Developing AI Assistants. Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 43, 1–7. DOI:<https://doi.org/10.1145/3411763.3443435>

Magdalena Wischniewski, Rebecca Bernemann, Thao Ngo, and Nicole Krämer. 2021. Disagree? You Must be a Bot! How Beliefs Shape Twitter Profile Perceptions. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 160, 1–11. DOI:<https://doi.org/10.1145/3411764.3445109>

Levy, Ariel, Monica Agrawal, Arvind Satyanarayan, and David Sontag. "Assessing the Impact of Automated Suggestions on Decision Making: Domain Experts Mediate Model Errors but Take Less Initiative." In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pp. 1-13. 2021. <https://dl.acm.org/doi/10.1145/3411764.3445522>

Sundararajan, M., Xu, S., Taly, A., Sayres, R., & Najmi, A. (2019). Exploring principled visualizations for deep network attributions. *CEUR Workshop Proceedings*, 2327.

Ariful Islam Anik and Andrea Bunt. 2021. Data-Centric Explanations: Explaining Training Data of Machine Learning Systems to Promote Transparency. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 75, 1–13. DOI:<https://doi.org/10.1145/3411764.3445736>

Andrew Ross, Nina Chen, Elisa Zhao Hang, Elena L. Glassman, and Finale Doshi-Velez. 2021. Evaluating the Interpretability of Generative Models by Interactive Reconstruction. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 80, 1–15. DOI:<https://doi.org/10.1145/3411764.3445296>

Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 2021. Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance. Proceedings of the 2021 CHI Conference on Human Factors in Computing

Systems. Association for Computing Machinery, New York, NY, USA, Article 81, 1–16. DOI:<https://doi.org/10.1145/3411764.3445717>

Harini Suresh, Steven R. Gomez, Kevin K. Nam, and Arvind Satyanarayan. 2021. Beyond Expertise and Roles: A Framework to Characterize the Stakeholders of Interpretable Machine Learning and their Needs. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 74, 1–16. DOI:<https://doi.org/10.1145/3411764.3445088>

Upol Ehsan, Q. Vera Liao, Michael Muller, Mark O. Riedl, and Justin D. Weisz. 2021. Expanding Explainability: Towards Social Transparency in AI systems. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 82, 1–19. DOI:<https://doi.org/10.1145/3411764.3445188>

Min Kyung Lee and Katherine Rich. 2021. Who Is Included in Human Perceptions of AI?: Trust and Perceived Fairness around Healthcare AI and Cultural Mistrust. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 138, 1–14. DOI:<https://doi.org/10.1145/3411764.3445570>

Koh, P. W., Nguyen, T., Tang, Y. S., Mussmann, S., Pierson, E., Kim, B., & Liang, P. (2020). Concept Bottleneck Models. *Proceedings of the 37th International Conference on Machine Learning*. <http://arxiv.org/abs/2007.04612>

Eldon Schoop, Forrest Huang, and Bjoern Hartmann. 2021. UMLAUT: Debugging Deep Learning Programs using Program Structure and Model Behavior. Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Article 310, 1–16. DOI:<https://doi.org/10.1145/3411764.3445538>

Sokol, K., & Flach, P. (2019). Desiderata for Interpretability: Explaining Decision Tree Predictions with Counterfactuals. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 10035–10036. <https://doi.org/10.1609/aaai.v33i01.330110035>

Chang, J., Gerrish, S., Wang, C., & Blei, D. M. (2009). Reading Tea Leaves: How Humans Interpret Topic Models. *Advances in Neural Information Processing Systems* 22, 288–296. <https://doi.org/10.1.1.100.1089>

Kaminskas, M., & Bridge, D. (2016). Diversity, Serendipity, Novelty, and Coverage: A Survey and Empirical Analysis of Beyond-Accuracy Objectives in Recommender Systems. *ACM Transactions on Interactive Intelligent Systems*, 7(1), 1–42. <https://doi.org/10.1145/2926720>

- Su, J., Vargas, D. V., & Kouichi, S. (2017). One pixel attack for fooling deep neural networks. Retrieved from <http://arxiv.org/abs/1710.08864>
- Baron, B., & Musolesi, M. (2017). Interpretable Machine Learning for Privacy-Preserving IoT and Pervasive Systems, 1–10. Retrieved from <http://arxiv.org/abs/1710.08464>
- Tan, S., Caruana, R., Hooker, G., & Lou, Y. (2017). Detecting Bias in Black-Box Models Using Transparent Model Distillation, (1). Retrieved from <http://arxiv.org/abs/1710.06169>
- Miller, T. (2017). Explanation in Artificial Intelligence: Insights from the Social Sciences. Retrieved from <http://arxiv.org/abs/1706.07269>
- Hendricks, L. A., Akata, Z., Rohrbach, M., Donahue, J., Schiele, B., & Darrell, T. (2016). Generating visual explanations. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 9908 LNCS, 3–19. http://doi.org/10.1007/978-3-319-46493-0_1
- Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why Should I Trust You?”: Explaining the Predictions of Any Classifier. In M. Gillies & R. Fiebrink (Eds.), *ACM SIGCHI Workshop on Human-Centered Machine Learning*. San Jose: ACM. Retrieved from <http://hcml2016.goldsmithsdigital.com/program/>
- Tamagnini, P., & Bertini, E. (2017). Interpreting Black-Box Classifiers Using Instance-Level Visual Explanations. <http://doi.org/10.1145/nnnnnnnn.nnnnnnnn>
- Kahng, M., Andrews, P., Kalro, A., & Chau, D. H. (2017). ActiVis: Visual Exploration of Industry-Scale Deep Neural Network Models. Retrieved from <http://arxiv.org/abs/1704.01942>
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 8689 LNCS(PART 1), 818–833. http://doi.org/10.1007/978-3-319-10590-1_53
- Krause, J., Perer, A., & Ng, K. (2016). Interacting with Predictions: Visual Inspection of Black-box Machine Learning Models. <http://doi.org/10.1145/2858036.2858529>
- Mahendran, A., & Vedaldi, A. (2015). Understanding deep image representations by inverting them. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 5188–5196). IEEE. <http://doi.org/10.1109/CVPR.2015.7299155>
- Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., & Elhadad, N. (2015). Intelligible Models for HealthCare. In *Proceedings of the 21th ACM SIGKDD International*

Conference on Knowledge Discovery and Data Mining - KDD '15 (pp. 1721–1730). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2783258.2788613>

Herlocker, J. L., Konstan, J. a, & Riedl, J. (2000). Explaining collaborative filtering recommendations. *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work*, 241–250. <http://doi.org/10.1145/358916.358995>

Lim, B., & Dey, A. (2009). Assessing demand for intelligibility in context-aware applications. *Proceedings of the 11th International Conference on Ubiquitous Computing*, 195. <http://doi.org/10.1145/1620545.1620576>

Kulesza, T., Stumpf, S., Burnett, M., Yang, S., Kwan, I., & Wong, W.-K. (2013). Too much, too little, or just right? Ways explanations impact end users' mental models. In *Proceedings of IEEE Symposium on Visual Languages and Human-Centric Computing, VL/HCC* (pp. 3–10). <http://doi.org/10.1109/VLHCC.2013.6645235>

Gabriel Lima, Nina Grgić-Hlača, and Meeyoung Cha. 2021. Human Perceptions on Moral Responsibility of AI: A Case Study in AI-Assisted Bail Decision-Making. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 235, 1–17. DOI:<https://doi.org/10.1145/3411764.3445260>

John Zoshak and Kristin Dew. 2021. Beyond Kant and Bentham: How Ethical Theories are being used in Artificial Moral Agents. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 590, 1–15. DOI:<https://doi.org/10.1145/3411764.3445102>

Ali Alkhatib. 2021. To Live in Their Utopia: Why Algorithmic Systems Create Absurd Outcomes. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 95, 1–9. DOI:<https://doi.org/10.1145/3411764.3445740>

Liang, L., & Grauman, K. (2014). Beyond comparing image pairs: Setwise active learning for relative attributes. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (pp. 208–215). <https://doi.org/10.1109/CVPR.2014.34>

Krishna, R. A., Hata, K., Chen, S., Kravitz, J., Shamma, D. A., Fei-Fei, L., & Bernstein, M. S. (2016). Embracing Error to Enable Rapid Crowdsourcing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16* (pp. 3167–3179). New York, New York, USA: ACM Press. <https://doi.org/10.1145/2858036.2858115>

- Chang, J. C., Amershi, S., & Kamar, E. (2017). Revolt: Collaborative Crowdsourcing for Labeling Machine Learning Datasets. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*, 2334–2346.
<https://doi.org/10.1145/3025453.3026044>
- Chang, J. C., Kittur, A., & Hahn, N. (2016). Alloy: Clustering with Crowds and Computation. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, 3180–3191.
<https://doi.org/10.1145/2858036.2858411>
- Sarkar, A., Morrison, C., Dorn, J. F., Bedi, R., Steinheimer, S., Boisvert, J., ... Lindley, S. (2016). Setwise Comparison: Consistent, Scalable, Continuum Labels for Computer Vision. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16* (pp. 261–271). New York, New York, USA: ACM Press. <http://doi.org/10.1145/2858036.2858199>
- Chang, J. C., Amershi, S., & Kamar, E. (2017). Revolt: Collaborative Crowdsourcing for Labeling Machine Learning Datasets. *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*, 2334–2346.
<http://doi.org/10.1145/3025453.3026044>
- Williams, J., Negreanu, C., Gordon, A. D., & Sarkar, A. (2020). Understanding and Inferring Units in Spreadsheets. *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC), 2020-Augus*, 1–9.
<https://doi.org/10.1109/VL/HCC50065.2020.9127254>
- Chia-Ming Chang, Chia-Hsien Lee, and Takeo Igarashi. 2021. Spatial Labeling: Leveraging Spatial Layout for Improving Label Quality in Non-Expert Image Annotation. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 306, 1–12. DOI:<https://doi.org/10.1145/3411764.3445165>
- Evgenia Christoforou, Pinar Barlas, and Jahna Otterbacher. 2021. It's About Time: A View of Crowdsourced Data Before and During the Pandemic. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. Association for Computing Machinery, New York, NY, USA, Article 649, 1–14.
DOI:<https://doi.org/10.1145/3411764.3445317>