# **Human-Centered Machine Learning**

Subjects: Computer Science, Artificial Intelligence Contributor: Tharindu Kaluarachchi

Human-centered Machine Learning (HCML) is about developing adaptable and usable Machine Learning systems for human needs while keeping the human/user at the center of the entire product/service development cycle.

human-centered machine learning HCML HCAI human-centered artificial intelligence

Deep Learning

## **1. Introduction**

Human-Centered Machine Learning, also referred to as Human-Centered Artificial Intelligence (HCAI or HAI), is gaining popularity due to the concerns raised by influential technology firms and research labs about the human context. A workshop in conjunction with the Conference on Human Factors in Computing Systems in 2016 [1] explained that HCML should explicitly recognize the human aspect when developing ML models, re-frame machine learning workflows based on situated human working practices and explore the co-adaptation of humans and systems. In early 2018, Google Design (https://design.google/library/ux-ai/, accessed on 1 April 2021) published an article noting that HCML is the User Experience (UX) of AI. Referring to a real consumer ML product, Google highlighted how ML could focus on human needs while solving them in unique ways that are only possible through ML. Several research projects (https://hcai.mit.edu/, accessed on 1 April 2021) by the Massachusetts Institute of Technology (MIT) on self-driving technologies called their approach Human-Centered Artificial Intelligence. The MIT team recognized both the development of AI systems that are continuously learning from humans and the parallel creation and fulfillment of a human-robot interaction experience. In 2019, the Stanford Institute for Human-Centered Artificial Intelligence (https://hai.stanford.edu/, accessed on 1 April 2021) was initiated with the goal of improving AI research, education, policy, and practice. They recognized the significance of developing AI technologies and applications that are collaborative, augmentative, and enhance human productivity and guality of life. A workshop (https://sites.google.com/view/hcml-2019, accessed on 1 April 2021) held in 2019 with the Conference on Neural Information Processing Systems for Human-Centered Machine Learning focused on the interpretability, fairness, privacy, security, transparency, accountability, and multi-disciplinary approach of AI technologies. Started in 2017, Google People + AI Research initiative (https://pair.withgoogle.com/, accessed on 1 April 2021) published a 2019 book presenting guidelines for building human-centered ML systems. This team is researching the full spectrum of human interactions with machine intelligence to build better AI systems with people.

Considering the scope of the HCML/HCAI prior work and publications by leading industry and academic institutions, we derived a definition for HCML that covers the breadth of this existing work. We validated the definition using feedback from several researchers working in the same domain and further validated with some influential researchers in leading academia and industrial institutions' Human-Centered AI research teams.

Human-centered Machine Learning (HCML): Developing adaptable and usable Machine Learning systems for human needs while keeping the human/user at the center of the entire product/service development cycle.

- "adaptable" includes adding features such as explainability, interpretability, fairness, privacy, security, transparency, and accountability.
- "Usable" refers to the UX of AI, including system usability and user burden.
- "Human needs" implies the significance of the problems we are selecting to solve with AI.
- "Entire development cycle" includes all steps from conceptualization to maintenance, which extends from Human-Centered Design to working systems that are continuously learning.

There is a natural incentive to research all the principles mentioned previously; however, this is seldom achieved in practice. In individual research, the entire development life-cycle is only partially detailed, possibly due to emphasis on the focused technicalities of the research. Therefore, we selected research that demonstrated one or more design elements matching the above definition of HCML research.

As shown in <u>Figure 1</u>, HCML work lies across many aspects of Machine Learning. We define algorithmic work related to HCML as Back-End HCML and work with interactions with humans as Front-End HCML. We excluded algorithm-centric back-end HCML papers as it would divert our focus away from the baseline HCML concepts. For instance, analyzing and classifying explainability algorithms is beyond this paper's scope and may be reviewed in separate works, such as <sup>[2][3]</sup>. However, algorithmic contributions with Front-End HCML practices, such as user evaluations, were included.

**Figure 1.** Human-Centered Machine Learning (marked with dashed lines) research lays over a broad spectrum, as shown here. The intersection of Machine Learning Research and Human-Centered Design is the domain we identify as Human-Centered Machine Learning.



### 2. HCML Related Work: Principles and Guidelines

Among the HCML literature, this category pertains to research that compiles design guidelines and principles for HCML or provides assistance to build HCML products and services. These works stem from different intentions, such as guidelines for developing intelligent user interfaces, visualization, prototyping, and human concerns in general. Amershi et al. <sup>[4]</sup> present a set of guidelines resulting from a comprehensive study conducted with many industry practitioners that worked on 20 popular AI products. Some approaches have been focused on deriving requirements and guidelines for planned sandbox visualization tools <sup>[5]</sup>. One article highlights guidelines related to three areas of HCML, ethically aligned design, tech that reflect human intelligence, and human factors design 6. Browne et al. <sup>[2]</sup> proposed a wizard of oz approach to bridge designers with engineers to build a humanconsiderate machine learning system targeting explainability, usability, and understandability. Some papers attempt to identify what HCML is <sup>[8]</sup>, and discuss how AI systems should understand the human and vice versa. Apart from general perspectives, Chancellor et al. [9] attempted to analyze literature in the mental health-AI domain to understand which humans are focused on such work and compile guidelines to maintain humans as a priority. In a slightly different layout, Ehsan et al. <sup>[10]</sup> attempted to uncover how to classify human-centered explainable AI in terms of prioritizing the human. Wang et al. [11] also tried to design theory driven by a user-centered explainable AI framework and evaluate a tool developed with actual clinicians. Schlesinger et al. [12] explored ways to build chatbots that can handle 'race-talk'. Long et al. <sup>[13]</sup> attempts to define learner-centered AI and figure out design considerations. Yang et al. [14] explore insights for designers and researchers to address challenges in human-AI interaction.

#### 3. The 'Human' in HCML Research

The main component of HCML is the Human and thus elevating the significance of the human. The 'Human' in HCML is defined across varying ML expertise levels, ranging from no ML background to an expert ML scientist. The Human in HCML can also be involved in various stages of the ML system development process in different capacities. For instance, the focus may be on the end-user, the developer, or the investor. One could focus on a certain user-aspect when developing a product or service <sup>[15][16]</sup>; another could be determining design principles for a particular ML system optimizing usability and adoptability <sup>[17][18][4]</sup>. The multidimensionality of what is considered Human within HCML contributes to the complexities within the field.

Considering the works that focused on the user side, some researchers catered to general end-users or consumers <sup>[15][19][8][20]</sup>, while others on specific end-users. Examples for these include people who need assistance <sup>[21][22][23][24][25][26][27][28][29]</sup>, medical professionals <sup>[30][31][32][33][34]</sup>, international travelers <sup>[35]</sup>, Amazon Mechanical Turk <sup>[36][37]</sup>, drivers <sup>[38][39]</sup>, musicians <sup>[40]</sup>, teachers <sup>[41]</sup>, students <sup>[42]</sup>, children <sup>[43][44]</sup>, UX designers <sup>[18][45][14][46]</sup>, UI designers <sup>[47][48][49]</sup>, data analysts <sup>[50]</sup>, video creators <sup>[51]</sup>, and game designers <sup>[52][53][54][55]</sup>. Apart from focusing on a specific user group, some have tried to understand multiple user-perspectives from ML engineers to the end-user <sup>[17]</sup>. Some of the prior works that target the developer as the human focus on novice ML engineers to help them develop ML systems faster <sup>[56][5]</sup>. Notably, the majority of works that target the developer side focused on ML engineers <sup>[57][18][58][59][60][61][62][16][63][64][4][6][7][9][65]</sup>.

#### 4. Application Domains

Machine Learning works well in many scenarios provided that a relationship exists between the task at hand and the availability of data. This power of making decisions or predictions based on data has empowered ML to infiltrate many other domains, such as medicine, pharmacy, law, business, finance, art, agriculture, photography, sports, education, media, military, and politics. Given that the majority in those sectors are not AI experts developing AI systems for them requires us to investigate the human aspect of such systems. Our analysis shows that application domains have specifically targeted gaming [66][52][53][54][55], interactive technologies [67][68][69][70][71][72][73][74][75][76][77] [78][79][80][81][82][83], medicine [84][30][31][32][85][86][33][34][11][87][88], psychiatry [89], music [18][15][90][40][80], sports [91], dating [36], video production [51], assistive technologies [21][22][23][25][26][27][28][92][29][93][94], education [41][44][42][95][96], and mainly software and ML engineering [17][57][56][59][60][48][61][62][97][63][64][51][98] based on our selected literature.

#### 5. Features of the Models

Features of AI models addressing the concerns of users to improve the usability and adoptability of AI systems such as explainability, interpretability, privacy, and fairness have been the focus of many HCML related work <sup>[99][100]</sup> <sup>[101][102][103][104][105][106][107][108][109]</sup>. This is not surprising, given the history of XAI research area dates back to 1980s <sup>[110][111]</sup>. In a comprehensive study, Bhatt et al. <sup>[17]</sup> investigated how explainability is practiced in real-world industrial AI products and presents how to focus explainability research on the end-user. Focusing on game designers, Zhu et al. <sup>[55]</sup> discuss how explainable AI should work for designers. A study <sup>[20]</sup> used 1150 online drawing platform users and compared two explanation approaches to figure out which approach is better.

Ashktorab et al. <sup>[112]</sup> explored explanations of machine learning algorithms concerning chatbots. Although explainability is not the main focus, some research <sup>[84][7]</sup> investigated the explainability aspect when developing ML systems. Another work <sup>[10]</sup> tried to investigate who is the human in the center of human-centered explainable AI. In addition, there exists work that tried to bring a user-centered approach to XAI research <sup>[113][11]</sup>. Chexplain <sup>[32]</sup> worked on providing an explainable analysis of chest X-rays to physicians. Das et al. <sup>[114]</sup> attempted to improve humans' performance by leveraging XAI techniques.

While explainability tries to untangle what is happening inside the Deep Learning black boxes, interpretability investigates how to make AI systems predictable. For instance, if a certain neural network classifies an MRI image as cancer, figuring out how the network makes such a decision falls into explainability research. However, an attempt to build a predictable MRI classification network where a change of network's parameters results in an expected outcome falls into interpretability research. There have been attempts <sup>[115][116][117]</sup> to develop novel interpretability algorithms using human studies to validate if those algorithms achieved the expected results. Isaac et al. <sup>[118]</sup> studied what matters to the interpretability of an ML system using a human study. Another study <sup>[59]</sup> figured out that ML practitioners often over-trust or misuse interpretability tools.

Apart from these two common DL features, some other work considered the aspects of fairness <sup>[119][57][58][120][8][9]</sup> <sup>[121]</sup>, understandability <sup>[33][5][7][8]</sup>, and trust <sup>[36][60][122][9]</sup>. Fairness represents the degree of bias in decisions, such as gender and ethnic skews, that influence the predictive model. For instance, gender and ethnic biases in the models can cause serious impacts on certain tasks. Understandability is a slightly different feature from explainability. While explainability shows how a model makes a certain decision, understandability tries to show how a neural network works to achieve a task. Trust refers to a subjective concern where the user's trust towards the decisions made by a certain model is studied.

#### References

- Gillies, M.; Fiebrink, R.; Tanaka, A.; Garcia, J.; Bevilacqua, F.; Heloir, A.; Nunnari, F.; Mackay, W.; Amershi, S.; Lee, B.; et al. Human-Centred Machine Learning. In Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, CHI EA'16, San Jose, CA, USA, 7–12 May 2016; Association for Computing Machinery: New York, NY, USA, 2016; pp. 3558–3565.
- 2. Adadi, A.; Berrada, M. Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). IEEE Access 2018, 6, 52138–52160.
- 3. Tjoa, E.; Guan, C. A Survey on Explainable Artificial Intelligence (XAI): Towards Medical XAI. arXiv 2019, arXiv:1907.07374.
- 4. Amershi, S.; Weld, D.; Vorvoreanu, M.; Fourney, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P.N.; Inkpen, K.; et al. Guidelines for Human-AI Interaction. In Proceedings of the 2019

CHI Conference on Human Factors in Computing Systems, CHI'19, Glasgow, UK, 4–9 May 2019; pp. 1–13.

- Nodalo, G.; Santiago, J.M.; Valenzuela, J.; Deja, J.A. On Building Design Guidelines for An Interactive Machine Learning Sandbox Application. In Proceedings of the 5th International ACM In-Cooperation HCI and UX Conference, CHIuXiD'19, Jakarta, Surabaya, 1–9 April 2019; pp. 70– 77.
- 6. Xu, W. Toward Human-Centered AI: A Perspective from Human-Computer Interaction. Interactions 2019, 26, 42–46.
- Browne, J.T. Wizard of Oz Prototyping for Machine Learning Experiences. In Proceedings of the Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems, CHI EA'19, Glasgow, UK, 4–9 May 2019; pp. 1–6.
- 8. Riedl, M.O. Human-Centered Artificial Intelligence and Machine Learning. Hum. Behav. Emerg. Technol. 2019.
- Chancellor, S.; Baumer, E.P.S.; De Choudhury, M. Who is the "Human" in Human-Centered Machine Learning: The Case of Predicting Mental Health from Social Media. Proc. ACM Hum. Comput. Interact. 2019, 3.
- 10. Ehsan, U.; Riedl, M.O. Human-centered Explainable AI: Towards a Reflective Sociotechnical Approach. arXiv 2020, arXiv:2002.01092.
- Wang, D.; Yang, Q.; Abdul, A.; Lim, B.Y. Designing theory-driven user-centric explainable AI. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–15.
- 12. Schlesinger, A.; O'Hara, K.P.; Taylor, A.S. Let's talk about race: Identity, chatbots, and AI. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–14.
- Long, D.; Magerko, B. What is AI Literacy? Competencies and Design Considerations. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; pp. 1–16.
- Yang, Q.; Steinfeld, A.; Rosé, C.; Zimmerman, J. Re-Examining Whether, Why, and How Human-Al Interaction Is Uniquely Difficult to Design. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; pp. 1–13.
- 15. Huang, C.Z.A.; Hawthorne, C.; Roberts, A.; Dinculescu, M.; Wexler, J.; Hong, L.; Howcroft, J. The Bach Doodle: Approachable music composition with machine learning at scale. arXiv 2019, arXiv:1907.06637.

- 16. Kim, B.; Pardo, B. A human-in-the-loop system for sound event detection and annotation. ACM Trans. Interact. Intell. Syst. (TiiS) 2018, 8, 1–23.
- Bhatt, U.; Xiang, A.; Sharma, S.; Weller, A.; Taly, A.; Jia, Y.; Ghosh, J.; Puri, R.; Moura, J.M.F.; Eckersley, P. Explainable Machine Learning in Deployment. In Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency, FAT'20, Barcelona, Spain, 27–30 January 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 648–657.
- Kayacik, C.; Chen, S.; Noerly, S.; Holbrook, J.; Roberts, A.; Eck, D. Identifying the Intersections: User Experience + Research Scientist Collaboration in a Generative Machine Learning Interface. In Proceedings of the Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–8.
- Ma, S.; Wei, Z.; Tian, F.; Fan, X.; Zhang, J.; Shen, X.; Lin, Z.; Huang, J.; Měch, R.; Samaras, D.; et al. SmartEye: Assisting Instant Photo Taking via Integrating User Preference with Deep View Proposal Network. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
- Cai, C.J.; Jongejan, J.; Holbrook, J. The Effects of Example-Based Explanations in a Machine Learning Interface. In Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI'19, Los Angeles, CA, USA, 7–11 March 2019; pp. 258–262.
- 21. Balasubramanian, V.; Chakraborty, S.; Krishna, S.; Panchanathan, S. Human-centered machine learning in a social interaction assistant for individuals with visual impairments. In Proceedings of the Symposium on Assistive Machine Learning for People with Disabilities at NIPS, Whistler, BC, Canada, 8–10 December 2008.
- 22. Ohn-Bar, E.; Guerreiro, J.A.; Ahmetovic, D.; Kitani, K.M.; Asakawa, C. Modeling Expertise in Assistive Navigation Interfaces for Blind People. In Proceedings of the 23rd International Conference on Intelligent User Interfaces, IUI'18, Tokyo, Japan, 7–11 March 2018; pp. 403–407.
- Kacorri, H.; Kitani, K.M.; Bigham, J.P.; Asakawa, C. People with visual impairment training personal object recognizers: Feasibility and challenges. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, 7–12 May 2017; pp. 5839–5849.
- Feiz, S.; Billah, S.M.; Ashok, V.; Shilkrot, R.; Ramakrishnan, I. Towards Enabling Blind People to Independently Write on Printed Forms. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
- 25. Lee, K.; Kacorri, H. Hands holding clues for object recognition in teachable machines. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.

- 26. Zhao, Y.; Wu, S.; Reynolds, L.; Azenkot, S. A face recognition application for people with visual impairments: Understanding use beyond the lab. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–14.
- 27. Wu, S.; Reynolds, L.; Li, X.; Guzmán, F. Design and Evaluation of a Social Media Writing Support Tool for People with Dyslexia. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–14.
- 28. Ahmetovic, D.; Sato, D.; Oh, U.; Ishihara, T.; Kitani, K.; Asakawa, C. ReCog: Supporting Blind People in Recognizing Personal Objects. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–12.
- 29. Aydin, A.S.; Feiz, S.; Ashok, V.; Ramakrishnan, I. SaIL: Saliency-Driven Injection of ARIA Landmarks. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 111–115.
- 30. Schaekermann, M.; Beaton, G.; Sanoubari, E.; Lim, A.; Larson, K.; Law, E. Ambiguity-Aware Al Assistants for Medical Data Analysis. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–14.
- 31. Beede, E.; Baylor, E.; Hersch, F.; Iurchenko, A.; Wilcox, L.; Ruamviboonsuk, P.; Vardoulakis, L.M. A Human-Centered Evaluation of a Deep Learning System Deployed in Clinics for the Detection of Diabetic Retinopathy. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–12.
- 32. Xie, Y.; Chen, M.; Kao, D.; Gao, G.; Chen, X. CheXplain: Enabling Physicians to Explore and Understand Data-Driven, AI-Enabled Medical Imaging Analysis. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- Cai, C.J.; Reif, E.; Hegde, N.; Hipp, J.; Kim, B.; Smilkov, D.; Wattenberg, M.; Viegas, F.; Corrado, G.S.; Stumpe, M.C.; et al. Human-Centered Tools for Coping with Imperfect Algorithms During Medical Decision-Making. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI'19, Glasgow, UK, 4–9 May 2019; pp. 1–14.
- 34. Hegde, N.; Hipp, J.; Liu, Y.; Emmert-Buck, M.; Reif, E.; Smilkov, D.; Terry, M.; Cai, C.; Amin, M.; Mermel, C.; et al. Similar image search for histopathology: SMILY. NPJ Digit. Med. 2019, 2, 56.
- 35. Liebling, D.J.; Lahav, M.; Evans, A.; Donsbach, A.; Holbrook, J.; Smus, B.; Boran, L. Unmet Needs and Opportunities for Mobile Translation AI. In Proceedings of the 2020 CHI Conference

on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.

- 36. Yin, M.; Wortman Vaughan, J.; Wallach, H. Understanding the Effect of Accuracy on Trust in Machine Learning Models. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI'19, Glasgow, UK, 4–9 May 2019; pp. 1–12.
- Hu, K.; Bakker, M.A.; Li, S.; Kraska, T.; Hidalgo, C. Vizml: A machine learning approach to visualization recommendation. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
- Banovic, N.; Wang, A.; Jin, Y.; Chang, C.; Ramos, J.; Dey, A.; Mankoff, J. Leveraging human routine models to detect and generate human behaviors. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, 6–11 May 2017; pp. 6683–6694.
- 39. Fridman, L.; Toyoda, H.; Seaman, S.; Seppelt, B.; Angell, L.; Lee, J.; Mehler, B.; Reimer, B. What can be predicted from six seconds of driver glances? In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, 7–12 May 2017; pp. 2805–2813.
- 40. McCormack, J.; Gifford, T.; Hutchings, P.; Llano Rodriguez, M.T.; Yee-King, M.; d'Inverno, M. In a silent way: Communication between ai and improvising musicians beyond sound. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Scotland, UK, 4–9 May 2019.
- 41. Jensen, E.; Dale, M.; Donnelly, P.J.; Stone, C.; Kelly, S.; Godley, A.; D'Mello, S.K. Toward Automated Feedback on Teacher Discourse to Enhance Teacher Learning. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- Wambsganss, T.; Niklaus, C.; Cetto, M.; Söllner, M.; Handschuh, S.; Leimeister, J.M. AL: An Adaptive Learning Support System for Argumentation Skills. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–14.
- Xu, Y.; Warschauer, M. What Are You Talking To?: Understanding Children's Perceptions of Conversational Agents. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- 44. Kang, S.; Shokeen, E.; Byrne, V.L.; Norooz, L.; Bonsignore, E.; Williams-Pierce, C.; Froehlich, J.E. ARMath: Augmenting Everyday Life with Math Learning. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–15.

- 45. Shi, Y.; Cao, N.; Ma, X.; Chen, S.; Liu, P. EmoG: Supporting the Sketching of Emotional Expressions for Storyboarding. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–12.
- 46. Dove, G.; Halskov, K.; Forlizzi, J.; Zimmerman, J. 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, Denver, CO, USA, 7–12 May 2017; pp. 278–288.
- Huang, F.; Canny, J.F.; Nichols, J. Swire: Sketch-based user interface retrieval. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–10.
- 48. Wu, Z.; Jiang, Y.; Liu, Y.; Ma, X. Predicting and Diagnosing User Engagement with Mobile UI Animation via a Data-Driven Approach. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- Duan, P.; Wierzynski, C.; Nachman, L. Optimizing User Interface Layouts via Gradient Descent. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–12.
- 50. Guo, S.; Du, F.; Malik, S.; Koh, E.; Kim, S.; Liu, Z.; Kim, D.; Zha, H.; Cao, N. Visualizing uncertainty and alternatives in event sequence predictions. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
- 51. Frid, E.; Gomes, C.; Jin, Z. Music Creation by Example. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- 52. Pfau, J.; Smeddinck, J.D.; Bikas, I.; Malaka, R. Bot or Not? User Perceptions of Player Substitution with Deep Player Behavior Models. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–10.
- 53. Guzdial, M.; Liao, N.; Chen, J.; Chen, S.Y.; Shah, S.; Shah, V.; Reno, J.; Smith, G.; Riedl, M.O. Friend, collaborator, student, manager: How design of an ai-driven game level editor affects creators. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–13.
- 54. Pfau, J.; Smeddinck, J.D.; Malaka, R. Enemy within: Long-Term Motivation Effects of Deep Player Behavior Models for Dynamic Difficulty Adjustment. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; pp. 1–10.

- 55. Zhu, J.; Liapis, A.; Risi, S.; Bidarra, R.; Youngblood, G.M. Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation. In Proceedings of the 2018 IEEE Conference on Computational Intelligence and Games (CIG), Maastricht, The Netherlands, 14–17 August 2018; pp. 1–8.
- 56. Yang, Q.; Suh, J.; Chen, N.C.; Ramos, G. Grounding Interactive Machine Learning Tool Design in How Non-Experts Actually Build Models. In Proceedings of the 2018 Designing Interactive Systems Conference, DIS'18, Hong Kong, China, 9–13 June 2018; pp. 573–584.
- 57. Holstein, K.; Wortman Vaughan, J.; Daumé, H.; Dudik, M.; Wallach, H. Improving Fairness in Machine Learning Systems: What Do Industry Practitioners Need? In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, CHI'19, Glasgow, UK, 4–9 May 2019; Association for Computing Machinery: New York, NY, USA, 2019; pp. 1–16.
- 58. Madaio, M.A.; Stark, L.; Wortman Vaughan, J.; Wallach, H. Co-Designing Checklists to Understand Organizational Challenges and Opportunities around Fairness in AI. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–14.
- Kaur, H.; Nori, H.; Jenkins, S.; Caruana, R.; Wallach, H.; Wortman Vaughan, J. Interpreting Interpretability: Understanding Data Scientists' Use of Interpretability Tools for Machine Learning. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–14.
- Drozdal, J.; Weisz, J.; Wang, D.; Dass, G.; Yao, B.; Zhao, C.; Muller, M.; Ju, L.; Su, H. Trust in AutoML: Exploring Information Needs for Establishing Trust in Automated Machine Learning Systems. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 297–307.
- 61. Sun, D.; Feng, Z.; Chen, Y.; Wang, Y.; Zeng, J.; Yuan, M.; Pong, T.C.; Qu, H. DFSeer: A Visual Analytics Approach to Facilitate Model Selection for Demand Forecasting. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- Sun, Y.; Lank, E.; Terry, M. Label-and-Learn: Visualizing the Likelihood of Machine Learning Classifier's Success During Data Labeling. In Proceedings of the 22nd International Conference on Intelligent User Interfaces, IUI'17, Limassol, Cyprus, 13–16 March 2017; pp. 523–534.
- 63. Schlegel, U.; Cakmak, E.; Arnout, H.; El-Assady, M.; Oelke, D.; Keim, D.A. Towards Visual Debugging for Multi-Target Time Series Classification. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 202–206.

- Ramos, G.; Meek, C.; Simard, P.; Suh, J.; Ghorashi, S. Interactive machine teaching: A humancentered approach to building machine-learned models. Hum. Comput. Interact. 2020, 35, 413– 451.
- 65. Cai, C.J.; Guo, P.J. Software Developers Learning Machine Learning: Motivations, Hurdles, and Desires. In Proceedings of the 2019 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC), Memphis, TN, USA, 14–18 October 2019; pp. 25–34.
- 66. Krening, S.; Feigh, K.M. Effect of Interaction Design on the Human Experience with Interactive Reinforcement Learning. In Proceedings of the 2019 on Designing Interactive Systems Conference, DIS'19, San Diego, CA, USA, 23–28 June 2019; pp. 1089–1100.
- 67. Lin, Y.; Guo, J.; Chen, Y.; Yao, C.; Ying, F. It Is Your Turn: Collaborative Ideation with a Co-Creative Robot through Sketch. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–14.
- 68. Xu, X.; Shi, H.; Yi, X.; Liu, W.; Yan, Y.; Shi, Y.; Mariakakis, A.; Mankoff, J.; Dey, A.K. EarBuddy: Enabling On-Face Interaction via Wireless Earbuds. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–14.
- 69. Matulic, F.; Arakawa, R.; Vogel, B.; Vogel, D. PenSight: Enhanced Interaction with a Pen-Top Camera. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–14.
- Rajcic, N.; McCormack, J. Mirror Ritual: An Affective Interface for Emotional Self-Reflection. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- Yeo, H.S.; Wu, E.; Lee, J.; Quigley, A.; Koike, H. Opisthenar: Hand Poses and Finger Tapping Recognition by Observing Back of Hand Using Embedded Wrist Camera. In Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology, UIST'19, New Orleans, LA, USA, 20–23 October 2019; pp. 963–971.
- Zhang, M.R.; Wen, H.; Wobbrock, J.O. Type, Then Correct: Intelligent Text Correction Techniques for Mobile Text Entry Using Neural Networks. In Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology, UIST'19, New Orleans, LA, USA, 20–23 October 2019; pp. 843–855.
- 73. Cami, D.; Matulic, F.; Calland, R.G.; Vogel, B.; Vogel, D. Unimanual Pen+Touch Input Using Variations of Precision Grip Postures. In Proceedings of the 31st Annual ACM Symposium on

User Interface Software and Technology, UIST'18, Berlin, Germany, 14–17 October 2018; pp. 825–837.

- 74. Huang, F.; Canny, J.F. Sketchforme: Composing Sketched Scenes from Text Descriptions for Interactive Applications. In Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology, UIST'19, New Orleans, LA, USA, 20–23 October 2019; pp. 209–220.
- 75. Sun, K.; Yu, C.; Shi, W.; Liu, L.; Shi, Y. Lip-Interact: Improving Mobile Device Interaction with Silent Speech Commands. In Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology, UIST'18, Berlin, Germany, 14–17 October 2018; pp. 581–593.
- 76. Xing, J.; Nagano, K.; Chen, W.; Xu, H.; Wei, L.Y.; Zhao, Y.; Lu, J.; Kim, B.; Li, H. HairBrush for Immersive Data-Driven Hair Modeling. In Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology, UIST'19, New Orleans, LA, USA, 20–23 October 2019; pp. 263–279.
- 77. Kapur, A.; Kapur, S.; Maes, P. AlterEgo: A Personalized Wearable Silent Speech Interface. In Proceedings of the 23rd International Conference on Intelligent User Interfaces, IUI'18, Tokyo, Japan, 7–11 March 2018; pp. 43–53.
- Huang, F.; Schoop, E.; Ha, D.; Canny, J. Scones: Towards Conversational Authoring of Sketches. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 313–323.
- Davis, N.; Hsiao, C.P.; Yashraj Singh, K.; Li, L.; Magerko, B. Empirically Studying Participatory Sense-Making in Abstract Drawing with a Co-Creative Cognitive Agent. In Proceedings of the 21st International Conference on Intelligent User Interfaces, IUI'16, Sonoma, CA, USA, 7–10 March 2016; pp. 196–207.
- Bonahue, C.; Simon, I.; Dieleman, S. Piano Genie. In Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI'19, Los Angeles, CA, USA, 7–11 March 2019; Association for Computing Machinery: New York, NY, USA; pp. 160–164.
- Lee, M.H.; Siewiorek, D.P.; Smailagic, A.; Bernardino, A.; Badia, S.B.I. Learning to Assess the Quality of Stroke Rehabilitation Exercises. In Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI'19, Los Angeles, CA, USA, 7–11 March 2019; pp. 218–228.
- Gallego Cascón, P.; Matthies, D.J.; Muthukumarana, S.; Nanayakkara, S. Chewlt. An Intraoral Interface for Discreet Interactions. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–13.
- 83. Matthies, D.J.; Perrault, S.T.; Urban, B.; Zhao, S. Botential: Localizing on-body gestures by measuring electrical signatures on the human skin. In Proceedings of the 17th International

Conference on Human-Computer Interaction with Mobile Devices and Services, Copenhagen, Denmark, 24–27 August 2015; pp. 207–216.

- Cai, C.J.; Winter, S.; Steiner, D.; Wilcox, L.; Terry, M. "Hello AI": Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making. Proc. ACM Hum.-Comput. Interact. 2019, 3.
- 85. Liang, Y.; Fan, H.W.; Fang, Z.; Miao, L.; Li, W.; Zhang, X.; Sun, W.; Wang, K.; He, L.; Chen, X. OralCam: Enabling Self-Examination and Awareness of Oral Health Using a Smartphone Camera. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- 86. Abibouraguimane, I.; Hagihara, K.; Higuchi, K.; Itoh, Y.; Sato, Y.; Hayashida, T.; Sugimoto, M. CoSummary: Adaptive Fast-Forwarding for Surgical Videos by Detecting Collaborative Scenes Using Hand Regions and Gaze Positions. In Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI'19, Los Angeles, CA, USA, 7–11 March 2019; pp. 580–590.
- Haescher, M.; Matthies, D.J.; Trimpop, J.; Urban, B. SeismoTracker: Upgrade any smart wearable to enable a sensing of heart rate, respiration rate, and microvibrations. In Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems, San Jose, CA, USA, 7–12 May 2016; pp. 2209–2216.
- 88. Elvitigala, D.S.; Matthies, D.J.; Nanayakkara, S. StressFoot: Uncovering the Potential of the Foot for Acute Stress Sensing in Sitting Posture. Sensors 2020, 20, 2882.
- Roffo, G.; Vo, D.B.; Tayarani, M.; Rooksby, M.; Sorrentino, A.; Di Folco, S.; Minnis, H.; Brewster, S.; Vinciarelli, A. Automating the Administration and Analysis of Psychiatric Tests: The Case of Attachment in School Age Children. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
- Louie, R.; Coenen, A.; Huang, C.Z.; Terry, M.; Cai, C.J. Novice-Al Music Co-Creation via Al-Steering Tools for Deep Generative Models. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- 91. Li, Q.; Lin, H.; Wei, X.; Huang, Y.; Fan, L.; Du, J.; Ma, X.; Chen, T. MaraVis: Representation and Coordinated Intervention of Medical Encounters in Urban Marathon. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–12.
- Rayar, F.; Oriola, B.; Jouffrais, C. ALCOVE: An Accessible Comic Reader for People with Low Vision. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 410–418.

- 93. Aydin, A.S.; Feiz, S.; Ashok, V.; Ramakrishnan, I. Towards Making Videos Accessible for Low Vision Screen Magnifier Users. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 10–21.
- 94. Matthies, D.J.; Roumen, T.; Kuijper, A.; Urban, B. CapSoles: Who is walking on what kind of floor? In Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services, Vienna, Austria, 4–7 September 2017; pp. 1–14.
- 95. Bassen, J.; Balaji, B.; Schaarschmidt, M.; Thille, C.; Painter, J.; Zimmaro, D.; Games, A.; Fast, E.; Mitchell, J.C. Reinforcement Learning for the Adaptive Scheduling of Educational Activities. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; pp. 1–12.
- 96. Matthies, D.J.; Bieber, G.; Kaulbars, U. AGIS: Automated tool detection & hand-arm vibration estimation using an unmodified smartwatch. In Proceedings of the 3rd International Workshop on Sensor-Based Activity Recognition and Interaction, Rostock, Germany, 23–24 June 2016; pp. 1–4.
- 97. Arendt, D.; Saldanha, E.; Wesslen, R.; Volkova, S.; Dou, W. Towards rapid interactive machine learning: Evaluating tradeoffs of classification without representation. In Proceedings of the 24th International Conference on Intelligent User Interfaces, Los Angeles, CA, USA, 7–11 March 2019; pp. 591–602.
- Cheng, G.; Xu, D.; Qu, Y. Summarizing Entity Descriptions for Effective and Efficient Human-Centered Entity Linking. In Proceedings of the 24th International Conference on World Wide Web, WWW'15, Florence, Italy, 18–22 May 2015; International World Wide Web Conferences Steering Committee: Geneva, Switzerland, 2015; pp. 184–194.
- 99. Holzinger, A.; Biemann, C.; Pattichis, C.S.; Kell, D.B. What do we need to build explainable AI systems for the medical domain? arXiv 2017, arXiv:1712.09923.
- 100. Hooker, S.; Erhan, D.; Kindermans, P.J.; Kim, B. A Benchmark for Interpretability Methods in Deep Neural Networks. arXiv 2018, arXiv:1806.10758.
- 101. Lee, S.; Suk, J.; Ha, H.R.; Song, X.X.; Deng, Y. Consumer's Information Privacy and Security Concerns and Use of Intelligent Technology. In Proceedings of the International Conference on Intelligent Human Systems Integration, Modena, Italy, 19–21 February 2020; pp. 1184–1189.
- 102. Liu, H.; Dacon, J.; Fan, W.; Liu, H.; Liu, Z.; Tang, J. Does gender matter? towards fairness in dialogue systems. arXiv 2019, arXiv:1910.10486.
- 103. Liao, Q.V.; Gruen, D.; Miller, S. Questioning the AI: Informing Design Practices for Explainable AI User Experiences. In Proceedings of the 2020 CHI Conference on Human Factors in Computing

Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–15.

- 104. Smith-Renner, A.; Fan, R.; Birchfield, M.; Wu, T.; Boyd-Graber, J.; Weld, D.S.; Findlater, L. No Explainability without Accountability: An Empirical Study of Explanations and Feedback in Interactive ML. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- 105. Alqaraawi, A.; Schuessler, M.; Weiß, P.; Costanza, E.; Berthouze, N. Evaluating Saliency Map Explanations for Convolutional Neural Networks: A User Study. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 275–285.
- 106. Dodge, J.; Liao, Q.V.; Zhang, Y.; Bellamy, R.K.E.; Dugan, C. Explaining Models: An Empirical Study of How Explanations Impact Fairness Judgment. In Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI'19, Los Angeles, CA, USA, 7–11 March 2019; pp. 275–285.
- 107. Hohman, F.; Head, A.; Caruana, R.; DeLine, R.; Drucker, S.M. Gamut: A design probe to understand how data scientists understand machine learning models. In Proceedings of the 2019 CHI conference on human factors in computing systems, Glasgow, UK, 4–9 May 2019; pp. 1–13.
- 108. Ehsan, U.; Tambwekar, P.; Chan, L.; Harrison, B.; Riedl, M.O. Automated Rationale Generation: A Technique for Explainable AI and Its Effects on Human Perceptions. In Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI'19, Los Angeles, CA, USA, 7–11 March 2019; pp. 263–274.
- 109. Dominguez, V.; Messina, P.; Donoso-Guzmán, I.; Parra, D. The Effect of Explanations and Algorithmic Accuracy on Visual Recommender Systems of Artistic Images. In Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI'19, Los Angeles, CA, USA, 7–11 March 2019; pp. 408–416.
- 110. Fagan, L.; Shortliffe, E.; Buchanan, B. Computer-Based Medical Decision Making: From MYCIN to VM; Automedica, Gordon and Breach Science Publishers: Zurich, Switzerland, 1984.
- 111. Clancey, W.J. Knowledge-Based Tutoring: The GUIDON Program; MIT Press: Cambridge, MA, USA, 1987.
- 112. Ashktorab, Z.; Jain, M.; Liao, Q.V.; Weisz, J.D. Resilient chatbots: Repair strategy preferences for conversational breakdowns. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
- 113. Hall, M.; Harborne, D.; Tomsett, R.; Galetic, V.; Quintana-Amate, S.; Nottle, A.; Preece, A. A Systematic Method to Understand Requirements for Explainable AI (XAI) Systems. In

Proceedings of the IJCAI Workshop on eXplainable Artificial Intelligence (XAI 2019), Macau, China, 11 August 2019.

- 114. Das, D.; Chernova, S. Leveraging Rationales to Improve Human Task Performance. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 510–518.
- 115. Ghorbani, A.; Wexler, J.; Zou, J.Y.; Kim, B. Towards automatic concept-based explanations. In Proceedings of the Advances in Neural Information Processing Systems 32: 32nd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, BC, Canada, 8–14 December 2019; pp. 9273–9282.
- 116. Lage, I.; Ross, A.; Gershman, S.J.; Kim, B.; Doshi-Velez, F. Human-in-the-loop interpretability prior. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2018; pp. 10159–10168.
- 117. Abdul, A.; von der Weth, C.; Kankanhalli, M.; Lim, B.Y. COGAM: Measuring and Moderating Cognitive Load in Machine Learning Model Explanations. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–14.
- 118. Lage, I.; Chen, E.; He, J.; Narayanan, M.; Kim, B.; Gershman, S.J.; Doshi-Velez, F. Human Evaluation of Models Built for Interpretability. In Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, Washington, USA, 28–30 October 2019; Volume 7, pp. 59–67.
- 119. Díaz, M.; Johnson, I.; Lazar, A.; Piper, A.M.; Gergle, D. Addressing age-related bias in sentiment analysis. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–14.
- 120. Yan, J.N.; Gu, Z.; Lin, H.; Rzeszotarski, J.M. Silva: Interactively Assessing Machine Learning Fairness Using Causality. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–13.
- 121. Barbosa, N.M.; Chen, M. Rehumanized crowdsourcing: A labeling framework addressing bias and ethics in machine learning. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
- 122. Yang, F.; Huang, Z.; Scholtz, J.; Arendt, D.L. How Do Visual Explanations Foster End Users' Appropriate Trust in Machine Learning? In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 189–201.

Retrieved from https://encyclopedia.pub/entry/history/show/34882