

Review

A Review of Recent Deep Learning Approaches in Human-Centered Machine Learning

Tharindu Kaluarachchi * , Andrew Reis and Suranga Nanayakkara 

Augmented Human Lab, Auckland Bioengineering Institute, The University of Auckland, 70 Symonds Street, Grafton, Auckland 1010, New Zealand; andrew@ahlab.org (A.R.); suranga@ahlab.org (S.N.)

* Correspondence: tharindu@ahlab.org

Abstract: After Deep Learning (DL) regained popularity recently, the Artificial Intelligence (AI) or Machine Learning (ML) field is undergoing rapid growth concerning research and real-world application development. Deep Learning has generated complexities in algorithms, and researchers and users have raised concerns regarding the usability and adoptability of Deep Learning systems. These concerns, coupled with the increasing human-AI interactions, have created the emerging field that is Human-Centered Machine Learning (HCML). We present this review paper as an overview and analysis of existing work in HCML related to DL. Firstly, we collaborated with field domain experts to develop a working definition for HCML. Secondly, through a systematic literature review, we analyze and classify 162 publications that fall within HCML. Our classification is based on aspects including contribution type, application area, and focused human categories. Finally, we analyze the topology of the HCML landscape by identifying research gaps, highlighting conflicting interpretations, addressing current challenges, and presenting future HCML research opportunities.

Keywords: human-centered machine learning; HCML; HCAI; human-centered artificial intelligence; Deep Learning



Citation: Kaluarachchi, T.; Reis, A.; Nanayakkara, S. A Review of Recent Deep Learning Approaches in Human-Centered Machine Learning. *Sensors* **2021**, *21*, 2514. <https://doi.org/10.3390/s21072514>

Academic Editor: Biswanath Samanta

Received: 11 February 2021

Accepted: 30 March 2021

Published: 3 April 2021

Publisher's Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2021 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Artificial Intelligence (AI) represents the broad spectrum of automated decision making from conditional logic to Neural Networks. Decisions or predictions made using data-driven techniques fall into Machine Learning (ML), a subset of AI. The subset of Machine Learning techniques that use Deep Neural Networks (DNN) is called Deep Learning (DL) (see Figure 1).




Figure 1. Overview of Artificial Intelligence. In modern-day practice, Deep Learning (DL) and Machine Learning (ML) are commonly referred to as Artificial Intelligence (AI), even though AI also includes rule-based simpler techniques. In this paper, the term AI refers to ML and DL.

Over the last two decades, Artificial Intelligence (AI) research publications have grown to account for 9% of all conference publications and 3% of all journal publications [1]. The

majority of this AI research typically explores developing algorithms and optimizing technologies with an emphasis on performant models benchmarked on accuracy. Apart from academic research, Machine Learning, Artificial Intelligence, and Deep Learning are incredibly pervasive within data-abundant industries. These same industries have built products and services atop an AI back-end with varying success. As the integration of performant models within the industry continues, there is an equally increasing requirement to analyze and improve the translation from the algorithmic model to end-user requirement.

1.1. Rise of Human-Centered Machine Learning (HCML)

HCML terminology has been used in publications over a decade ago. In an assistive tool [2], Balasubramanian et al. claimed they introduced the term “Human-Centered Machine Learning Algorithms” through their investigation on human-in-the-loop ML systems. However, the HCML term started to gain and accelerate popularity in the mid-2010s after the modern Deep Learning era started. The average user does not yet understand the capabilities, limitations, and inner workings of AI. Therefore users have reported concerns [3,4] regarding explainability, user experience (UX) with the interfaces, user-data privacy, security, reliability, and dependability of AI systems. The need to solve these issues has evolved into the field of Human-Centered Machine Learning (HCML). Recognizing that algorithm optimization and innovative neural network architectures alone do not solve usability and adoptability problems, the HCML field seeks to elevate ML systems’ user-centered development. Reference to HCML terminology now exists in formal and informal forums, such as AI publications, workshops, conferences [5], and blog posts and articles from AI-focused corporations (see Figure 2).

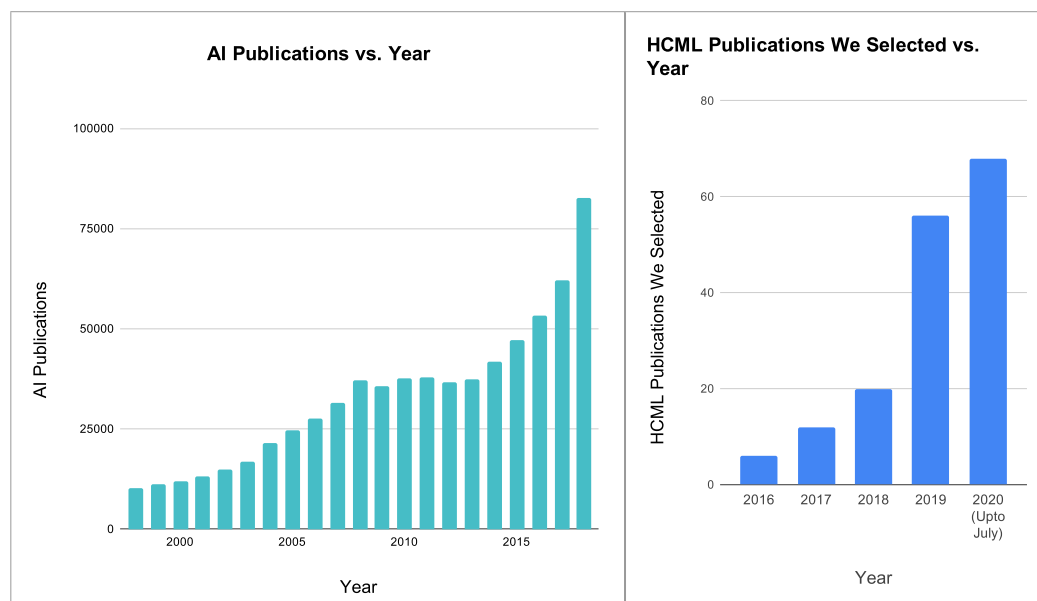


Figure 2. Rise of AI-related publications and Human-Centered Machine Learning (HCML) publications.

HCML emerged as the research area investigating the methods of aligning machine learning systems with human goals, context, concerns, and ways of work [5]. Several different terminology and abbreviations have emerged alongside the many institutions researching the user experience aspect of artificial intelligence. The widespread use of ML and AI interchangeably has led institutions to use either of the terms Human-Centered Machine Learning (HCML) or Human-Centered Artificial Intelligence (HCAI or HAI). Several researchers have also started labeling this field the UX of AI (see Section 2 for detailed description). However, all these terms’ underlying principle remains the same, which is developing usable and adaptable, ‘Human-in-the-Loop’ machine learning systems.

Some of the most common non-expert-user issues are explainability [6–10], interpretability [11–14], privacy and security [15–20], reliability [21–25], and fairness [26–31]. These categories emerged as contributing research areas of HCML, and each plays a role in the broader goal of improving the usability and adoptability of AI systems. It is important to note that these research areas are not new but instead are aligned with the underlying motivations of HCML. The growing demand to adapt ML systems to real-life use cases requires timely and in-depth HCML research; however, it is imperative to survey the landscape of previous works in this emerging field.

1.2. Scope and Contribution

This paper analyzes HCML work related to Deep Learning (as highlighted in Figure 3). HCML emerged in the mid-2010s with increasing popularity, mainly due to the widespread adaptation of Deep Learning. The rapid development of hardware to power Deep Learning algorithms enabled novel neural network architecture research. Deep Learning algorithms' complex nature was the main reason for the rise of some sub-topics in HCML, such as explainability and interpretability. The exponential growth of Deep Learning and the existence of sub research areas within HCML generated diverse research from algorithm development and optimization to user studies research. The research focus varies, focusing on the algorithm, user-studies, concepts, or a mixture of these.




Figure 3. Human-Centered Machine Learning (HCML) is the intersection of Machine Learning and Human-Centered Design (marked with pink stripes). In this paper, our focus lies in HCML approaches related to Deep Learning as marked by the green area.

There has been no prior work done to investigate and detail Deep Learning related prior work within the HCML domain to the best of our knowledge. Therefore, in this work, we analyze recent approaches in HCML related to Deep Learning. In summary, our key contributions of this work are:

- **Working definition of HCML:** Many leading research groups have interpreted HCML through the lens of their respective institutions. Those interpretations present multiple perspectives on the totality of HCML. Therefore, we have analyzed these interpretations and subsequently derived a working definition (see Section 2) that envelopes the prevailing work in the HCML field. We validated the final definition through consultation with leading academics and industry experts who are spearheading HCML research.
- **Systematic Literature Review:** We employ a Systematic Literature Review (SLR) technique to filter and select publications, as explained in Section 3. We categorize and critically analyze 162 publications based on four criteria. Prior work stems from a range of domains such as medicine and software engineering, and the nature of these publications takes several forms. Some focus on specific user-concerns while others

highlight general aspects within HCML. These variations paved the way for us to build a taxonomy of HCML publications (see Section 4).

- Challenges and opportunities for HCML research: Given the emerging nature of the HCML field, research conducted in HCML faces challenges from several directions. Our analysis across the defined scope of HCML reveals opportunities for future research. We discuss these points in detail in the discussion section (see Section 5).

2. Defining HCML

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 [5] 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 quality 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 Figure 4, 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 [32,33]. However, algorithmic contributions with Front-End HCML practices, such as user evaluations, were included.

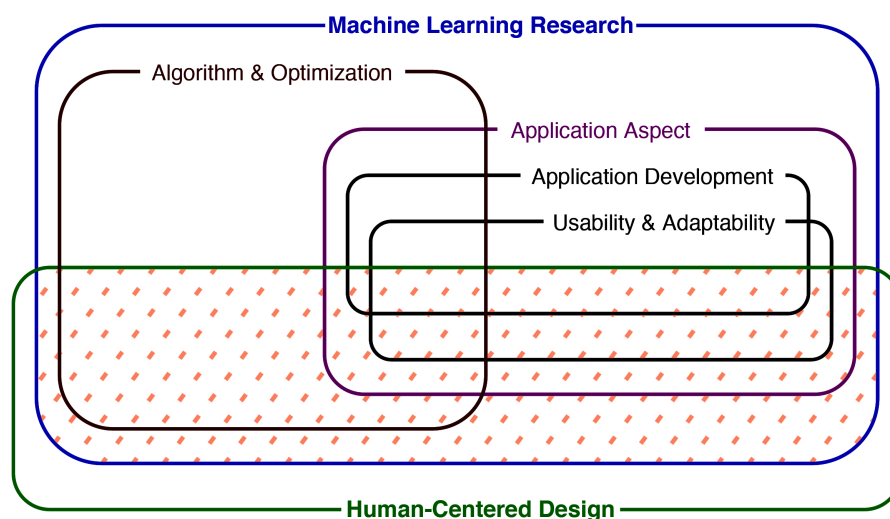


Figure 4. Human-Centered Machine Learning (marked with dashed lines) research lays over a broad spectrum, as shown here. The intersection of Machine Learning Research and Human-Centered Design is the domain we identify as Human-Centered Machine Learning.

3. Systematic Literature Review Approach

Systematic Literature Review (SLR) is a technique used to provide an extensive overview of the literature extracted by a repeatable and objective process. Although SLR was initially used in research within the domain of medicine, it is also gaining popularity in computer science [34–37]. We chose to follow the SLR approach when selecting literature on Human-Centered Machine Learning because it reduces the subjective nature of the process while widening the search area. Following the development and finalization of the working HCML definition, we conducted a three-phase SLR approach.

Phase 1:

We performed a domain-specific keyword and title search for “Human-Centered Machine Learning” and “Human-Centered Artificial Intelligence”. We employed Google Scholar (<https://scholar.google.com/> accessed on 1 April 2021), ACM Digital Library (<https://dl.acm.org/> accessed on 1 April 2021), and IEEE Xplore (<https://ieeexplore.ieee.org/Xplore/home.jsp> accessed on 1 April 2021) to sample a diverse collection of research. We downloaded 200 papers as the initial step. We went through the titles, keywords, and abstracts of all papers and selected 85 publications that matched.

Phase 2:

We used a specific selection process based on a set of specific inclusion and exclusion criteria to select papers. The selection process was to use five independent researchers working in the field of AI and HCML to vote for publications based on the inclusion and exclusion criteria, which is as follows:

- Included—Deep Learning papers published from 2016 (when HCML re-emerged)-May 2020.

- Included—Traits that matched our HCML definition of usability and adoptability.
- Included—Publications surveys and guidelines that are tagged as Human-Centered Machine Learning.
- Excluded—Work that did not contain real human/user inputs in one or more stages of development.

Equally weighted vote distributions from the researchers were used to select papers. From the initial 85 papers of the initial pool, this selection process picked 31 papers. These 31 papers were used as the basis for expanding the initial search base. We used additional HCML keywords, extracted from the selected papers, including “Human-AI interaction”, “user experience”, “user-centered design”, “empirical-study”, “end-user machine learning” and “machine learning”. Using these targeted keywords, we downloaded an additional 50 publications. After analyzing them, we identified a broad and diverse range of keywords used in HCML related work.

Phase 3:

In the third stage, in addition to the keyword search using search engines that we used for the first stage, we expanded our search to proceedings that typically have papers with Human-Centered Design and Machine Learning Research. This includes, CHI proceedings (<https://sigchi.org/conferences/conference-history/chi/> accessed on 1 April 2021), UIST Proceedings (<https://sigchi.org/conferences/conference-history/uist/> accessed on 1 April 2021), IUI proceedings (<https://sigchi.org/conferences/conference-history/iui/> accessed on 1 April 2021), Neural Information Processing Systems proceedings (<https://nips.cc/> accessed on 1 April 2021) and AAAI proceedings (<https://www.aaai.org/> accessed on 1 April 2021). Since our focus is on Deep Learning and HCML, we downloaded all the publications of the conferences and journals mentioned above from 2016. Using python scripts, we filtered out possible HCML work by looking for keywords. In HCI publications, we searched for all Deep Learning related terms and filtered all publications with at least one match. In AI publications, we searched for human-centered design related keywords. As the final step, we manually went through 320 (including the 50 papers from the search in phase 2) filtered publications and selected 131 papers using the same selection process in phase 2. In total, we analyzed 162 publications in this work.

4. Classifying HCML Research

Despite being a nascent field, an attempt to classify HCML related work is beneficial in defining future HCML research boundaries. Categorizing the work can be done along multiple axes, given that HCML consists of many different types of contributions focusing on different domains and different users. While resulting in taxonomy, this also allows us to find desiderata by looking into the missing pieces from the categorization. In this section, we classify the work under several criteria, making it easier for researchers to see the HCML spectrum.

The breadth of our HCML definition also includes work done in areas such as explainable AI (XAI) and data privacy. Sub-research areas such as these have a number of researchers creating different research avenues [11,12,15,17,20,26,28,32,38–41]. For instance, there are works on developing algorithms and novel DL architectures in XAI to add explainability to the models [42–46]. In comparison, there is also work that considers user experience and user requirements for XAI [7–10,47], and evaluates algorithms and models with user studies [48]. However, analyzing and categorizing XAI algorithms is not the focus of this paper. Our focus concerns the usability and adoptability of ML systems per the underlying motivation of HCML.

4.1. Contribution Varieties

4.1.1. User Studies

Per the HCML definition, it is desirable always to keep the human at the center when developing ML systems. This includes understanding the user requirements, the inputs for

the iterative development from a feedback loop, and the evaluation. This can be achieved through many approaches; however, user studies are among the most frequently used approaches to infuse user input into the development. Many existing works in literature have successfully utilized user studies in their development process. Some approaches conduct evaluations and comparisons between existing systems through user studies; however, we will explore work that presented these user studies as the main contribution in this section.

The work of Cai et al. [49] uncovered the requirements of an AI assistant for use in the medical domain. They conducted a thorough study with 21 pathologists to understand how non-AI-experts interact with an AI-tool [49]. They kept the pathologists' involvement active throughout the AI assistant's development and presented the findings in a detailed manner. These efforts can be said to make an excellent contribution to HCML. In contrast, to engage the human throughout the process, Liebling et al. [50] conducted a study only to understand the users' needs for AI language translation applications. Luger et al. [51] conducted a set of interviews to understand people's perception of voice assistants from tech giants. Since they conducted the study in 2016, it may be outdated. Assistants have evolved significantly over the last four years. In the same area, Candello et al. [52] ran a study to identify the effects of typefaces when communicating with chatbots. In a case study, Yang et al. [53] explored ways to sketch Natural Language Processing (NLP) powered user experiences and presented a wizard-of-oz type NLP rapid prototyping method. Yang et al. [54] studied the affective aspect of conversational agents. In a comparison study, Diaz et al. [55] analyzed age-biases in algorithms related to sentiment analysis. Gero et al. [56] conducted a study to understand mental models people build regarding AI using an AI-game. Schaekermann et al. [57] conducted a study to compare two AI assistants that classify medical time series data.

Apart from studies to understand aspects of specific applications, there have been attempts to understand user concerns, such as explainability and fairness, that later translate into ML models' features. One comprehensive study [48] has been conducted to explore current practices of explainable-AI in the industry to understand the desirable techniques for real-world usage. Another study [58] attempts to identify gaps between current XAI algorithmic work and practices towards user-centered XAI. In a similar study using industry personnel [59], Holstein et al. compared AI-fairness practices in literature and the real world. In a user study [60] conducted with people from the Amazon Mechanical Turk (<https://www.mturk.com/> accessed on 1 April 2021), the trust of an AI model based on stated vs. real-world accuracy of the model is explored. Another user study has been conducted while attempting to find general principles for interpretability [61]. They specifically investigate the interpretability of human-simulatable machine learning models. There have been other studies carried out to explore design approaches for interactive ML tools for non-expert ML engineers to develop models quickly [62]. As opposed to a generic interactive ML tool, Krening et al. [63,64] conducted a user study to investigate specific aspects related to interactive reinforcement learning tools for novice ML engineers. In a completely different attempt to mix a UX team with an AI team, Kayacik et al. [65] present a study on how the teams from two distinct domains interacted to create an AI-Music application. Hong et al. [66] investigated how users conceptualize, experience, and reflect on their engagement in machine teaching.

Beede et al. [67] conducted a real-world study to evaluate a Deep Learning system deployed to detect diabetic retinopathy. A study conducted by Santhanam et al. [68] investigated the effects of cognitive biases when evaluating the outputs of conversational agents. Lin et al. [69] investigated the better approach for collaborative ideation from a physical robot and a virtual agent. Pfau et al. [70] investigated the appropriateness of using bots in games when real players drop due to various reasons. Madaio et al. [71] attempted to understand the role checklists play in AI ethics using 48 practitioners. Xu et al. [72] attempted to investigate children's perception of conversational agents available in smart devices. Smith-Renner et al. [73] studied how automatically generated explanations of ML

models shape users' perceptions of ML models. Völkel et al. [74] explored how to mislead chatbots in profiling users. A study conducted by Kaur et al. [75] found out that ML practitioners often over-trust and misuse interpretability tools. Alaqaraawi et al. [76] evaluated saliency maps, a popular explanation technique for CNN-based classification algorithms, with a user study. Ishibashi et al. [77] investigated audio visualization techniques for audio related interactive machine learning applications. Das et al. [78] explored whether humans can improve performance using explainable AI capabilities with computer techniques that surpass humans. Drozdal et al. [79] attempted to discover what is vital to ML engineers' trust in the Auto ML system. Ohn-Bar et al. [80] studied an indoor navigation interface for blind people. Dodge et al. [81] investigated how explanations impact people's fairness judgment. Yang et al. [82] studied how visual explanations affect end-users' appropriate trust in ML.

4.1.2. Applications

This section explores the literature that mainly focuses on AI application development and follows HCML principles at one or more points in the development-life-cycle. For example, this could be in the form of a user needs survey, interviews, or evaluating user studies. 'The Bach Doodle' is an attempt to develop a web-based music-harmonizing app for large scale deployment [83] where user feedback was used to evaluate the system and ensure reliability before public release. Frid et al. [84] developed an AI background music generation app for video creators where they integrated human input before, during, and after the development. In the same domain, Louie et al. [85] investigated how to adjust an AI music generation tool to minimize user burden with user experiments. The HCML approach has been used in developing assistive technologies. For example, Balasubramanian et al. [2] carried out user experiments to understand the user perspective to develop an assistive AI tool for the blind to grasp non-verbal cues. Kacorri et al. [86] also developed an assistive tool for visually impaired people to get information about visual objects using Deep Learning techniques. Similarly, Feiz et al. [87] used AI technologies to develop an application enabling blind people to write on printed forms. Lee et al. [88] attempts to improve object recognition based applications for blind people by leveraging the hand as a point to consider when focusing on an object in a frame. Zhao et al. [89] developed a face recognition tool for visually impaired people to identify their friends.

Fast et al. [90] attempted to create a camera and fiction-text-input based application to detect human activities. Xu et al. [91] developed and evaluated a new chatbot trained on Twitter conversations to be deployed on social media for customer service. EgoScanning [92] is an application to fast forward first-person recorded videos by being content-aware using object detection techniques. Presenting a unique concept, Kimura et al. [93] developed an application (SottoVoce) to decode speech utterances without voice using ultrasound and Deep Learning. Aila [94] is an application that acts as a document labeling assistant and uses attention-based deep neural networks to perform it. Alghofaili et al. [95] present a Deep Learning based application to pull up navigation aids in Virtual Reality (VR) environments by using gaze patterns. Wu et al. [96] developed an AI-powered tool to assist people with dyslexia to write posts on social media. Guo et al. [97] developed an application to visualize event sequence predictions of multiple records. They used inputs from machine learning practitioners before development and used eighteen participants to evaluate the system. Roffo et al. [98] developed a tool to conduct a standard psychiatric test to assess attachment in children for the first time automatically. VizML [99] is an application made to visualize recommendations using neural networks. Laput et al. [100] developed a system to detect hand activities using sensors available in smartwatches by using a convolutional neural network. SmartEye [101] is an application that assists smartphone users in taking good photos using a view proposal network. McCormack et al. [102] developed an AI application that communicated with improvising musicians in collaborative environments. Swire [103] is an application developed to retrieve user

interfaces using sketches. Gamut [104] is a design probe to understand how data scientists understand ML models.

The DeepWriting [105] application used generative neural networks to edit handwritten text, and the paper also published a dataset. A tone-aware chatbot was created by Hu et al. [106] for customer care on social media. Kimura et al. [107] developed a system to augment visual experiences by projecting around the primary display using generative models. Huber et al. [108] developed an image-grounded conversational agent that uses visual sentiment cues. Oh et al. [109] developed an application, DuetDraw, to draw in collaboration with an AI. CheXplain [110] is a tool designed for physicians to understand AI-based analysis on chest X-ray images. Wu et al. [111] developed a Deep Learning system to predict and diagnose user engagement with mobile user interfaces. EarBuddy [112] is an application to control wireless earbuds via interactions on a person's face. PenSight [113] is a tool where a camera is attached to a tablet pen to create seamless hand inputs using both hands. OralCam [114] is a smartphone app for users to self-examine oral health using Deep Learning. EmoG [115] is a generative tool to support storyboarding by incorporating emotional expressions. Wu et al. [116] developed a system for acoustic activity recognition targeting low user burden by using self-supervised learning techniques. ReCog [117] is a mobile app for blind users to recognize personal objects by letting them train a neural network with their photos. Mirror Ritual [118] is an effective mirror that displays a generated poem to engage the user while conceptualizing an emotional state.

Iconate [119] is a tool created to generate complex icons based on text queries. Sun et al. [120] present a system to aid model selection for demand forecasting. Im et al. [121] attempted to create an application to derive social signals for online social platforms computationally. Xiao et al. [122] attempted to develop a prototype of an active-listening chatbot to interview people. WorldGaze [123] tries to leverage the existing cameras of smartphones to include gaze information to enrich voice assistants. Jensen et al. [124] present a system to generate feedback for teachers in classrooms automatically. ARMath [125] is a Deep Learning powered tool for children to discover mathematical concepts with real-life objects with augmented reality. MaraVis [126] is a Deep Learning powered tool for real-time urban marathon visualization and coordinated intervention. Silva [127] is a tool to find out potential sources of unfairness in datasets or ML models. An NLP-powered tool was developed by Wambsganss et al. [128] for students to develop argumentation quality in writing by providing feedback. Sterman et al. [129] developed a tool to visualize and model writing styles using Deep Learning. Opisthenar [130] is a Deep Learning based tool to recognize head poses and finger tappings. Zhang et al. [131] propose several techniques to correct errors more efficiently when typing on mobile phone keyboards. Bylinskii et al. [132] created a tool to predict the relative importance of elements in graphics and visualizations. Cami et al. [133] introduce a new pen input space using hand postures. Sketchforme [134] is a tool to generate intricate sketches based on text descriptions.

Lip-Interact [135] is a tool to provide silent voice commands via lip movement. CodeMend [136] is a tool to support search and integration of code for programmers. HairBrush [137] is an interactive 3D hair modeling system. LabelAR [138] is an augmented-reality-based tool to label objects for computer vision in a novel way. ViZig [139] is an application developed with semi-supervised training to find anchor points in instructional videos. AlterEgo [140] is a wearable silent speech interface that allows users to converse silently. Alcove [141] is an assistive comic reading tool for people with low vision. iSeqL [142] is a tool designed for the rapid construction of customized text mining models. Creative Sketching Partner [143] is a proof-of-concept intelligent interface to inspire designers while sketching. Scones [144] is a tool that generates and modifies sketches iteratively according to text descriptions. CQAVis [145] is an application to filter high-quality comments from online community question answering forums. Zhelezniakov et al. [146] investigated the requirements for pen-centered mathematical expression recognition applications. SaIL [147] is an assistive web navigation tool that injects important ARIA landmarks automatically. Grover et al. [148] designed and evaluated intelligent agents

to improve productivity and focus at work. Cartograph [149] is a visualization system that uses knowledge from Wikipedia to create thematic maps. Kulahcioglu et al. [150] created an affect-aware color palette selection tool for word cloud generation. VASTA [151] is a vision and language-assisted programming-by-demonstration system for smartphone task automation. Davis et al. [152] is an intelligent drawing partner that can improvise and collaborate on abstract sketches. Piano Genie [153] is an intelligent controller for non-musicians to improvise on the piano. SViM [154] is an adaptive screen-magnifier tool for people with low vision to watch videos. Lee et al. [155] developed a tool to assess the quality of rehabilitation exercises. Sun et al. [156] created a tool for developers to learn about the dataset while labeling. These AI applications have incorporated the human aspect in ML systems' development and can thus be classified as HCML research.

4.1.3. Algorithms

Some work [157,158] has been done to improve the interpretability of the Deep Learning algorithms for general and computer vision purposes by employing user evaluations to assess intended outcomes. Arendt et al. [159] developed an efficient image labeling tool with user evaluations to support its intended performance. In a similar tool to efficiently label audio samples, Kim et al. [160] also follow a similar human-centered approach to evaluate success. Although these can be identified as tools, the key contribution, as described in the publications, is developing and evaluating the algorithm. We classify these papers as algorithmic contributions that follow HCML principles. Banovic et al. [161] developed a new weakly supervised algorithm based on inverse reinforcement learning to detect and generate human behaviors. Fridman et al. [162] developed an algorithm to predict the driver's states using videos of driver glances. Crowdverge [163] is an algorithm to figure out whether crowds would agree on one answer for visual question answering tasks. Huang et al. [164] developed an algorithm to augment static panorama images through a realistic audio assignment. Guzdial et al. [165] developed an algorithm to design levels in games like Super Mario and investigated the user aspect with level designers. Kim et al. [166] present a Deep Learning based algorithm to estimate eye-gaze with low latency. Seemo [167] is a framework developed to map emotions into vector representations using representation learning techniques.

Ryolai et al. [168] explored ways to assign laughter to tangible objects. Yuan et al. [169] trained a deep learning model to predict the scannability of webpage content. Cogam [170] tried to generate explanations for machine learning models by incorporating the desired cognitive load. Lai et al. [171] created a technique to annotate visualizations according to text descriptions automatically. Soundr [172] used Deep Learning to figure out the user's spatial location and head orientation using voice. Duan et al. [173] developed a technique to optimize UI interfaces automatically with error correction. Pfau et al. [174] looked at how Deep Learning techniques can improve dynamic difficulty adjustments in games. Bassen et al. [175] developed a reinforcement learning algorithm to optimize educational activities in online courses. Donkers et al. [176] developed a recommendation and explanation method and evaluated the quality with user studies. Arent et al. [177] present a technique that generalizes parallel coordinated visualization techniques to sequences of learned representations. Eshan et al. [178] created a technique to generate rationales automatically. Le et al. [179] developed a model to identify fingers in capacitive touchscreens. CoSummary [180] is a technique to fast-forward surgical videos adaptively. Micallef et al. [181] created a technique that uses interactive visualization to elicit the need of domain experts to improve the accuracy of prediction models. Athukorala et al. [182] created a novel adaptation technique for search engines. Mittal et al. [183] proposed an architecture to generate emoticons using multi-modal input. Weber et al. [184] present a combination of manual and automated Deep-Image-Prior-based image restoration techniques. Schlegel et al. [185] illustrate a model-agnostic visual debugging workflow for multi-target time series classifications.

4.1.4. Non-AI-Expert Tools

Apart from applications that use AI in the back-end where the users are not exposed to the underlying technology or AI model, some tools allow users to interface with the AI algorithms [186–191]. These tools are often referred to as non-AI-expert tools or non-expert tools in general. For instance, a voice assistant is an AI application for the user to engage verbally. The user might not even know that AI algorithms run in the back-end. However, a customized tool designed for a physician to interact with an ML model is a non-AI-expert tool where the non-AI-expert manipulates AI algorithms without having AI-specific knowledge. A non-AI-expert's level of expertise can range from none, such as a lawyer with no AI background, to a moderate level, such as a novice ML engineer. These variations in knowledge require designers to determine the correct level of abstraction for specific user groups. Here, we look at several attempts to develop such non-AI-expert tools following the HCML approach. SMILY [192,193] is a prime example of a non-AI-expert tool. It allows pathologists to change parameters of an ML model to search medical images. The What-If Tool [194] is designed for non-expert ML engineers to carry out interactive probing across different input variations. They have conducted an evaluation study with ML students to validate their tool. To make machine learning models accessible to non-AI-experts, Ramos et al. [195] showed how to leverage intrinsic human capabilities of teaching to teach machines to do machine learning.

4.1.5. 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. [196] 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 [197]. One article highlights guidelines related to three areas of HCML, ethically aligned design, tech that reflect human intelligence, and human factors design [198]. Browne et al. [199] proposed a wizard of oz approach to bridge designers with engineers to build a human-considerate machine learning system targeting explainability, usability, and understandability. Some papers attempt to identify what HCML is [200], and discuss how AI systems should understand the human and vice versa. Apart from general perspectives, Chancellor et al. [201] 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. [202] attempted to uncover how to classify human-centered explainable AI in terms of prioritizing the human. Wang et al. [203] also tried to design theory driven by a user-centered explainable AI framework and evaluate a tool developed with actual clinicians. Schlesinger et al. [204] explored ways to build chatbots that can handle 'race-talk'. Long et al. [205] attempts to define learner-centered AI and figure out design considerations. Yang et al. [206] explore insights for designers and researchers to address challenges in human–AI interaction.

4.1.6. Survey Results

Some work has focused on presenting analyzed results of surveys conducted towards different goals Figure 5, such as understanding human perspectives before developing solutions. For example, Cai et al. [207] present many hurdles software engineers faced working in the machine learning field through a thorough survey and analysis. Other survey studies [208] focused on finding out people's opinions about delegating tasks to AI agents to help ML engineers and product developers. Dove et al. [209] surveyed to understand how design innovation is practiced in the ML domain in terms of user experience.

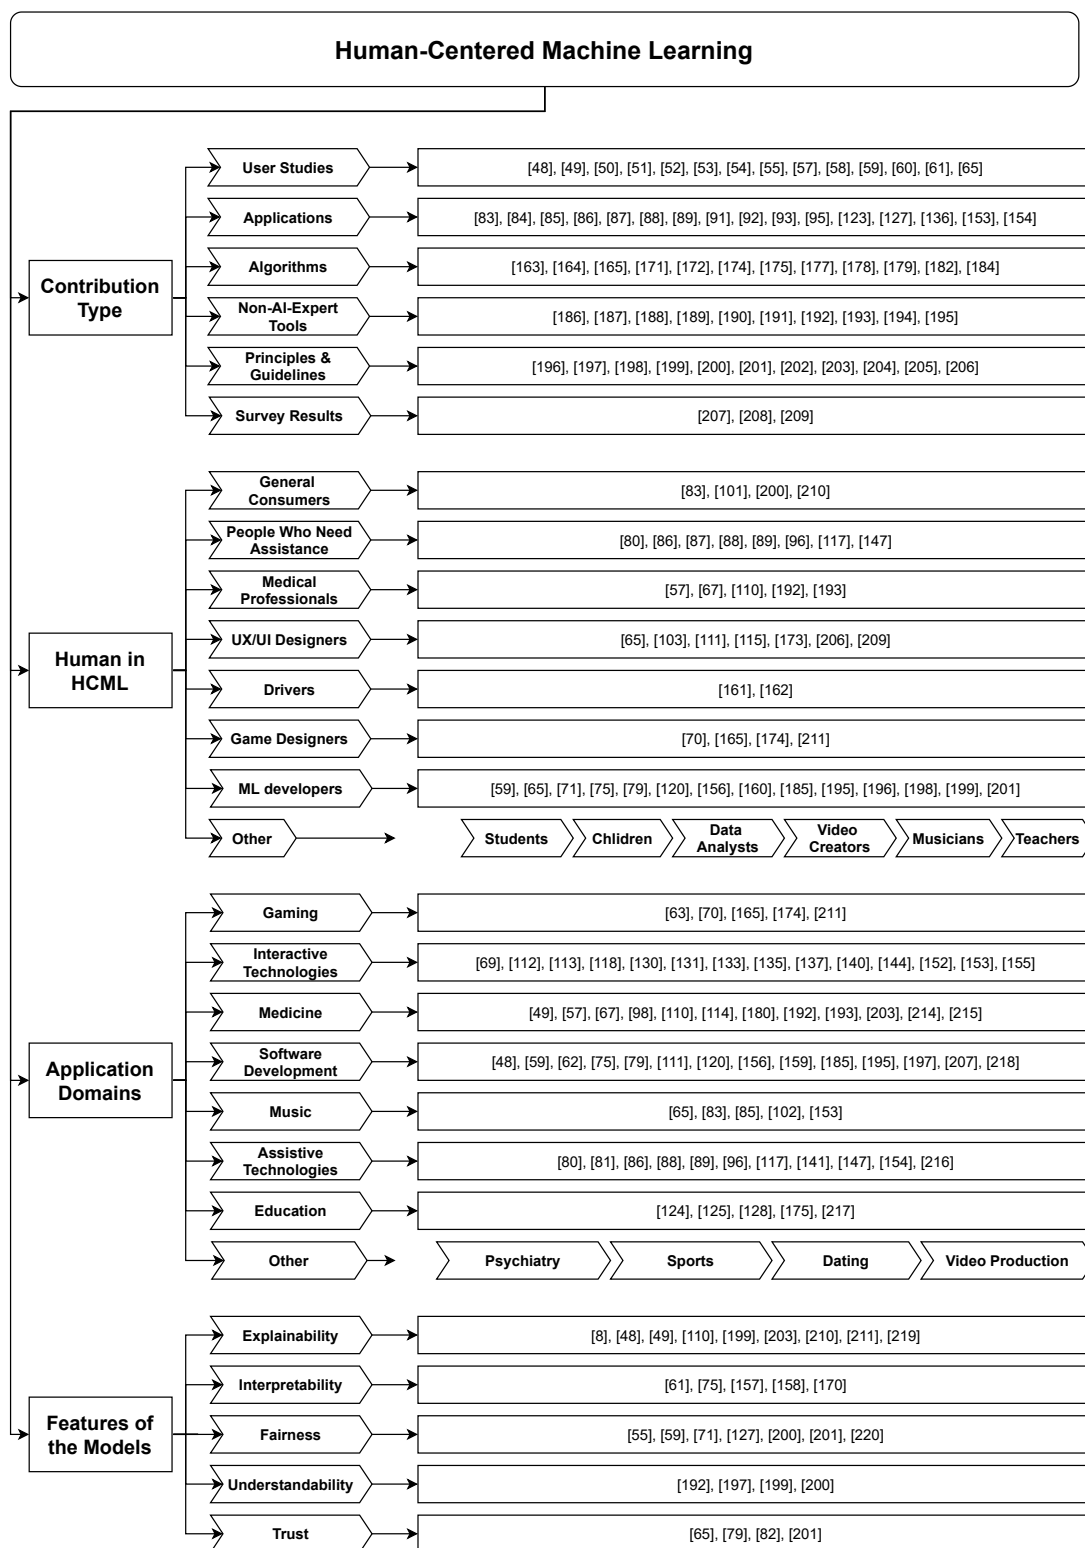


Figure 5. Classification chart. Each category has a few representative examples.

4.2. The ‘Human’ in HCML

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 [83,160]; another could be determining design principles for a particular ML system optimizing usability and adoptability [48,65,196]. 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 [83,101,200,210], while others on specific end-users. Examples for these include people who need assistance [2,80,86–89,96,117,147], medical professionals [57,67,110,192,193], international travelers [50], Amazon Mechanical Turk [60,99], drivers [161,162], musicians [102], teachers [124], students [128], children [72,125], UX designers [65,115,206,209], UI designers [103,111,173], data analysts [97], video creators [84], and game designers [70,165,174,211]. Apart from focusing on a specific user group, some have tried to understand multiple user-perspectives from ML engineers to the end-user [48]. Some of the prior works that target the developer as the human focus on novice ML engineers to help them develop ML systems faster [62,197]. Notably, the majority of works that target the developer side focused on ML engineers [59,65,71,75,79,120,156,160,185,195,196,198,199,201,207].

4.3. 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 [63,70,165,174,211], interactive technologies [69,112,113,118,130,131,133–135,137,140,144,152,153,155,212,213], medicine [49,57,67,110,114,180,192,193,203,214,215], psychiatry [98], music [65,83,85,102,153], sports [126], dating [60], video production [84], assistive technologies [2,80,86,88,89,96,117,141,147,154,216], education [124,125,128,175,217], and mainly software and ML engineering [48,59,62,75,79,111,120,156,159,185,195,197,207,218] based on our selected literature.

4.4. 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 [6,12,20,26,58,73,76,81,104,178,219]. This is not surprising, given the history of XAI research area dates back to 1980s [220,221]. In a comprehensive study, Bhatt et al. [48] 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. [211] discuss how explainable AI should work for designers. A study [210] used 1150 online drawing platform users and compared two explanation approaches to figure out which approach is better. Ashktorab et al. [222] explored explanations of machine learning algorithms concerning chatbots. Although explainability is not the main focus, some research [49,199] investigated the explainability aspect when developing ML systems. Another work [202] 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 [8,203]. Chexplain [110] worked on providing an explainable analysis of chest X-rays to physicians. Das et al. [78] 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 [157,158,170] to develop novel interpretability algorithms using human studies to validate if those algorithms achieved the expected results. Isaac et al. [61] studied what matters to the interpretability of an ML system using a human study. Another study [75] 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 [55,59,71,127,200,201,223], understandability [192,197,199,200], and trust [60,79,82,201]. 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.

5. Discussion

The breadth and depth of the classification presented in the previous section show that HCML grows in many branches despite being a relatively new field of research. However, given the variations in the field, the young age, and the rapid growth of AI, the HCML field has numerous research gaps, challenges, and confusions that are interesting to be discussed and analyzed. For example, a standard definition has yet to be agreed upon, resulting in some liberal interpretations per researchers' prerogative and the level of abstraction. Here we provide a discussion (see Figure 6 for a summary of this section) concerning these gaps, challenges, and confusions under several subtopics. Apart from this, it is also important to discuss the scope and limitations of current research to identify the opportunities for future work.

5.1. Human at the Center

5.1.1. Interpretation of 'Human-Centered'

One of the main confusions we uncovered within HCML literature is identifying who the 'human' is in HCML and how the 'human' is involved. As described in Section 4, 'human' could be any stakeholder of the work, ranging from the owner of the research institute to the end-user. However, merely any kind of human involvement will not make a particular AI research an HCML approach. The HCML approach depends on the phase in which humans interact with the development life cycle, the purpose of the interaction, and how they interact. For example, while many agree that HCML exists to create usable and adaptable systems for users, others [224,225] argue that an ML engineer interfering with the training loop when adjusting parameters is recognizable as HCML. The latter perspective may be interpreted as a valid HCML approach assuming that interferences from ML engineers contribute to improvements in usability and adoptability, apart from accelerating and optimizing the training process. However, creating a tool centered on ML engineers' needs to accelerate the training process differs from claiming HCML is about the human interfering with the training process.

5.1.2. Explainability Moving from Non-Experts to Experts

A thorough study conducted with leading industry research teams stated that, despite being created to cater to user needs, ML engineers mainly use explainability research for debugging purposes in practice [48]. The study investigates how to move the explainability research to inform end-users rather than focus on implementation improvements for engineers. These practical issues may already exist in other HCML research areas or emerge soon if not addressed clearly.

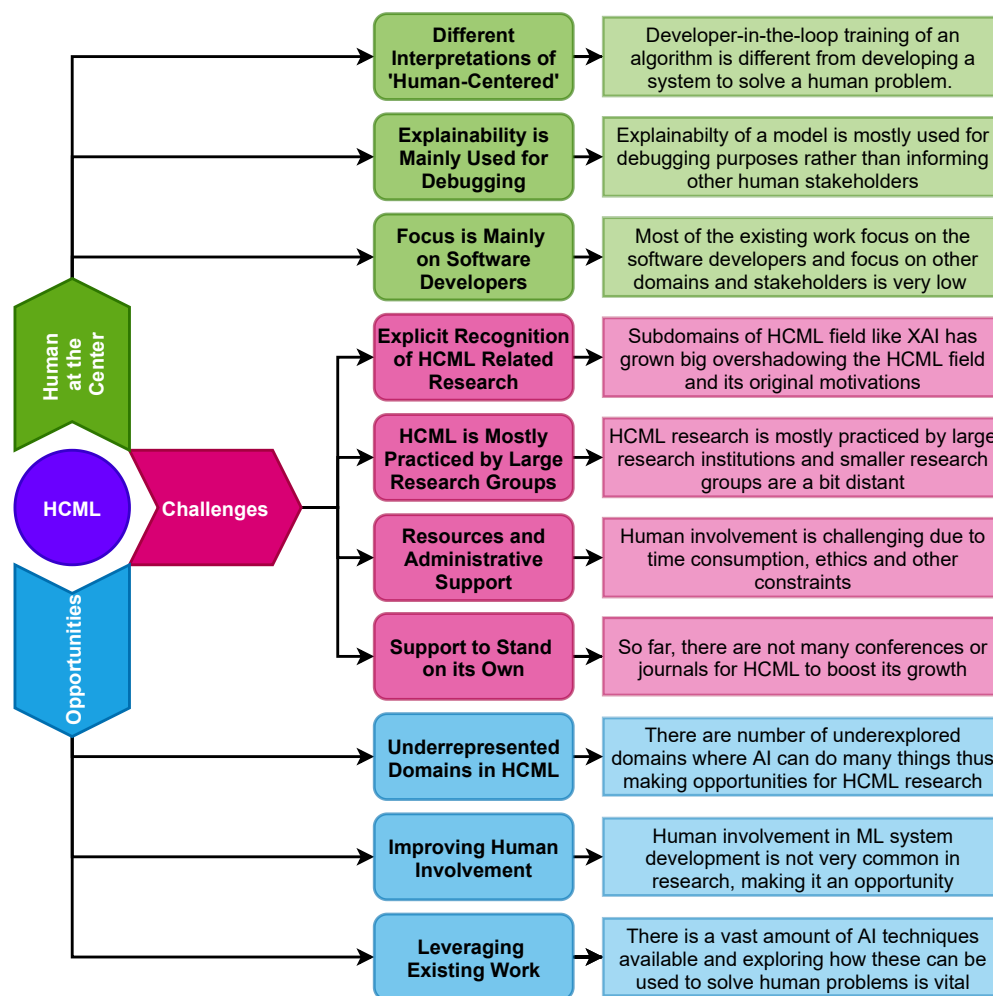


Figure 6. Summary of the Discussion Section. This figure represents the points discussed, and each point is followed by the main idea behind it.

5.1.3. Too Much Focus on Software Developers

From automated data labeling tools to interactive machine learning tools, many HCML related work made ML engineers the 'human' at the center. The majority of the work we analyzed targets software engineers working in the ML domain. The remainder focuses on an array of different stakeholders and makes software engineers the most focused 'human' in HCML research. These findings also further validate the results of Bhatt et al. [48]. Therefore, it is desirable to expand research to other non-AI-experts.

5.2. Challenges to HCML

After analyzing the confusions, problems, and gaps in HCML research, we scaffolded the underlying challenges that must be addressed in future HCML guided research. These challenges hinder the rapid and coherent development of HCML research. We discuss these challenges under the four categories below.

5.2.1. Explicitly Recognizing as HCML Research

In our search for literature and forming the definition, we understood the confusions that occur when labeling work as HCML. For instance, although XAI (explainable AI) serves the overall goals of HCML, it is an independent field and is not mainly categorized as HCML work. In XAI, research is being conducted to add explainability to models; however, most publications do not follow a human-centered approach. This has reduced the number

of works in studying the end-user or human aspect of explainable AI algorithms. However, there are works that attempt to bring an HCML approach into XAI [8,47,58]. Although XAI is the best example, this problem exists in research on other user-concerns, such as interpretability, privacy, and fairness. We identified that one major reason for this is the lack of extension of the HCML practices to small research teams.

5.2.2. Extending HCML to Small Research Groups

The problem of missing end-user involvement in HCML work seems to be in research interests. Machine Learning researchers are focusing mainly on algorithm development and optimization while the human-centered approach is practiced mainly by Human-Computer Interaction (HCI) researchers. Therefore, collaborations between HCI and ML research produce good HCML outputs. These collaborations and facilitation seem well-established within large industry and academic research groups, such as Google Research (<https://research.google/> accessed on 1 April 2021), Microsoft Research (<https://www.microsoft.com/en-us/research/> accessed on 1 April 2021), MIT (<https://www.mit.edu/> accessed on 1 April 2021), and Stanford (<https://www.stanford.edu/> accessed on 1 April 2021). Therefore, their efforts have to be appreciated. However, this facilitation and collaborations are not well-established within small research groups. As a result, it is not common to see small research groups publish good HCML work. We identified that extending this HCML vision to smaller research groups will accelerate performant outputs with the necessary usability and adoptability.

5.2.3. Resources and Administrative Support

The HCML research we analyzed had several possible improvements that could be observed in terms of the completeness of work. However, achieving this will entail challenges. For instance, some research uses human inputs to identify user requirements, whereas other research used human users for evaluations. Apart from a small selection of works, most of the corpora's human-centered practice is limited to a particular stage of the system development life cycle. While understandable that it is desirable to keep the human at the center throughout the entire development or design cycle, there are many practical hurdles. Access to a group of specific people, institutional ethics approvals, time consumption, meeting internal and publication deadlines are some hurdles to overcome to maximize human involvement throughout the process. These hurdles may result in a lack of human involvement or motivate conducting a user study for superficial purposes only. These reduce the quality of outcomes of HCML research, hence decelerating the development of the field.

5.2.4. Support to Stand on Its Own

Many established research fields have their own high rated conferences and journals with high impact. For example, AI and ML have many conferences including NeurIPS (<https://nips.cc/> accessed on 1 April 2021) and AAAI (<https://www.aaai.org/> accessed on 1 April 2021), and HCI has its own conferences, such as CHI (<https://dl.acm.org/conference/chi> accessed on 1 April 2021) and IMWUT (<https://dl.acm.org/journal/imwut> accessed on 1 April 2021). However, HCML does not have a popular conference or a journal yet. Although venues like IUI (<https://dl.acm.org/conference/iui> accessed on 1 April 2021) and TiiS (<https://dl.acm.org/journal/tiis> accessed on 1 April 2021) include HCML work in their broader scope, those venues are not solely focused on the HCML field. Having dedicated publication venues (examples include ASSETS (<https://dl.acm.org/conference/assets> accessed on 1 April 2021), AHs (<https://augmented-humans.org/> accessed on 1 April 2021) for a field can help streamline the research, as well as assist and solve many challenges mentioned in this section. We believe that HCML has a certain maturity to be represented in its own publication venue, despite being an emerging field.

5.3. Opportunities for HCML and the Future

5.3.1. Underrepresented Domains in HCML

At the inception of HCML, the main focus was on the end-user as in Human-Centered Design. While the 'human' now represents a broader definition, there has been a definite shift in focus to the software or ML engineers. Therefore, it is desirable to see more research targeted to end-users in other general or specific domains. In analyzing industry-specific AI trends and the adoption of HCML driven research, there are apparent shortcomings. AI usage in agriculture is a real need [226]; however, publications related to HCML and agriculture were scarce. Similarly, in medicine, the HCML aspect of medical AI applications is rare, as algorithmic development and optimization dominate the field. Given that medicine is a field filled with non-AI-experts working with critical situations, a human-centered approach will ensure the novel algorithms can be packaged as applications that are useful and adaptable, as Cai et al. [192] demonstrated. Other domains such as law, music, art, chemistry, education, finance, media, politics, photography, and the movie industry face the same under-representation problem in HCML research.

5.3.2. Improving Human Involvement

Exciting work with AI has been produced in law, such as using ML to help judges decide bail [227]. However, such work did not fall within our HCML definition scope due to the lack of investigation on the usability and adoptability of their application. They have either hypothesized the user needs based on assumed user behavior or failed to present how they derived the end-user requirements, namely the judges. This problem of researchers presuming end-user requirements without purposeful investigation is prevalent in many other works [228,229]. We believe that this practice does not work efficiently toward developing usable and adaptable AI systems. Therefore, it is imperative that researchers involve humans in the development process.

5.3.3. Leveraging Existing Work

AI technologies evolve fast and generate usable models and tools in the real-world. Many of these significant works have not yet been studied in terms of real user aspects. An opportunity exists to quickly leverage existing AI technologies to design, develop, and validate real-world AI systems successfully. For instance, it could be investigating models or even useful tools created by other researchers [187,230].

Going beyond small to medium scale, current trends, and innovations in the broader field of high-performance AI suggest a rising significance of HCML. For instance, Deep Learning models are computed at the current technological limitations of available hardware. Models such as GPT-2 [231] have 1.5 billion parameters that give incredibly good results on Natural Language Processing (NLP) tasks. However, it has recently been outperformed by GPT-3 [232], a 175 billion parameter network, which is the largest Deep Learning model to date. These large networks demonstrate that having more parameters improves the model performance and detracts interest from innovation and optimization techniques driven improvements. Few institutions, let alone researchers, have access to the hardware resources to train such monolithic networks. This growing reliance on raw computational power elevates the need to investigate the real-world application aspect of existing models over hardware-dependent improvement in benchmarked accuracy. The arrival of multi-billion parameter networks also brings vast opportunities to HCML, given the rotation of interest from pure performance towards the human aspect of modern AI.

6. Conclusions

Human-Centered Machine Learning (HCML) is an emerging field of research parallel to the exponential growth of Deep Learning (DL) research for several reasons. Some of the reasons are the frequent application of Artificial Intelligence (AI) models in the real world, an increasing number of interactions between humans and AI, professionals raising concerns regarding the black-box nature of DL, and questions about the AI models'

reliability in critical application scenarios. Even though HCML is a nascent field, its origin dates back decades, and the exact term itself is more than a decade old. However, with the AI boom in the early 2010s and the adaptation of Deep Learning techniques, HCML re-emerged with different and modern concerns. Those are mostly usability and adoptability related concerns such as user experience, explainability, understandability, interpretability, privacy, security, fairness, and reliability.

Many research groups started working around HCML with different interpretations of the term. Hence, in this work, we present a working definition for HCML derived from analyzing those interpretations. We followed a systematic literature review approach to make the literature search repeatable and expand its search base. We classified 162 HCML publications we selected in terms of contribution type, application domain, the 'human' at the center, and the models' features. Our search and the analysis of prior work enabled us to understand the spectrum of HCML while identifying the confusions, research gaps, and possible research directions. In conclusion, HCML is an emerging but fast-growing field, given the foundation field—AI—is snowballing. HCML has sub research fields that stand independently, hence reducing the focus on core concepts of HCML. Given the early stage of the field, there are confusions, problems, and opportunities for development, which we discussed in this work. Deep Learning is approaching the limitation of current hardware's abilities while user-aspect investigations are lagging. The field of HCML can bridge the critical divide between model performance and end-user adoption. We believe that this is the prime time to research HCML, given the tremendous opportunities.

Author Contributions: T.K. conceptualised and Investigated as part of his PhD research. A.R. contributed with supporting to the taxonomy and writing. S.N. contributed with overall supervision, review editing. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Tertiary Education Commission Assistive Augmentation research grant under the Entrepreneurial Universities(EU) initiative of New Zealand; grant number 3716751.

Acknowledgments: We are grateful for the careful readings and recommendations provided by three reviewers, which improved the presentation of this paper. We also thank Gimhani Pamarathne for their support to proofread and edit.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Perrault, R.; Shoham, Y.; Brynjolfsson, E.; Clark, J.; Etchemendy, J.; Grosz, B.; Lyons, T.; Manyika, J.; Mishra, S.; Niebles, J.C. *The AI Index 2019 Annual Report*; Technical Report; AI Index Steering Committee, Human-Centered AI Institute, Stanford University: Stanford, CA, USA, 2019.
2. 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, Canada, 8–10 December 2008.
3. Ras, G.; van Gerven, M.; Haselager, P. Explanation methods in deep learning: Users, values, concerns and challenges. In *Explainable and Interpretable Models in Computer Vision and Machine Learning*; Springer: Berlin/Heidelberg, Germany, 2018; pp. 19–36.
4. Manikonda, L.; Deotale, A.; Kambhampati, S. What's up with Privacy? User Preferences and Privacy Concerns in Intelligent Personal Assistants. In Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, New Orleans, LA, USA, 2–4 February, 2018; pp. 229–235.
5. 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; Association for Computing Machinery, New York, NY, USA, 7–12 May 2016; pp. 3558–3565, doi:10.1145/2851581.2856492.
6. 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.
7. Sheh, R.K. Different XAI for different HRI. In Proceedings of the 2017 AAAI Fall Symposium Series, Arlington, VA, USA, 9–11 November 2017.

8. 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.
9. Andres, J.; Wolf, C.T.; Cabrero Barros, S.; Oduor, E.; Nair, R.; Kjærsum, A.; Tharsgaard, A.B.; Madsen, B.S. Scenario-based XAI for Humanitarian Aid Forecasting. In Proceedings of the Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, 25 April 2020; pp. 1–8.
10. Ferreira, J.J.; Monteiro, M.S. What Are People Doing About XAI User Experience? A Survey on AI Explainability Research and Practice. In *Design, User Experience, and Usability. Design for Contemporary Interactive Environments*; Marcus, A., Rosenzweig, E., Eds.; Springer International Publishing: Cham, Switzerland, 2020; pp. 56–73.
11. Rudin, C. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* **2019**, *1*, 206–215.
12. Hooker, S.; Erhan, D.; Kindermans, P.J.; Kim, B. A Benchmark for Interpretability Methods in Deep Neural Networks. *arXiv* 2018, arXiv:cs.LG/1806.10758.
13. Yang, M.; Kim, B. Benchmarking Attribution Methods with Relative Feature Importance. *arXiv* 2019, arXiv:cs.LG/1907.09701.
14. Khanna, R.; Kim, B.; Ghosh, J.; Koyejo, O. Interpreting Black Box Predictions Using Fisher Kernels; In Proceedings of Machine Learning Research, Long Beach, CA, USA, 11 April 2019; pp. 3382–3390.
15. Papernot, N.; McDaniel, P.; Sinha, A.; Wellman, M. Towards the science of security and privacy in machine learning. *arXiv* **2016**, arXiv:1611.03814.
16. Stahl, B.C.; Wright, D. Ethics and privacy in AI and big data: Implementing responsible research and innovation. *IEEE Secur. Priv.* **2018**, *16*, 26–33.
17. Bartoletti, I. AI in healthcare: Ethical and privacy challenges. In Proceedings of the Conference on Artificial Intelligence in Medicine in Europe; Poznan, Poland, 26–29 June 2019; pp. 7–10.
18. Tadapaneni, N.R. Artificial Intelligence Security and Its Countermeasures. *Int. J. Adv. Res. Comput. Sci. Technol.* **2020**, *8*.
19. Hossain, E.; Khan, I.; Un-Noor, F.; Sikander, S.S.; Sunny, M.S.H. Application of big data and machine learning in smart grid, and associated security concerns: A review. *IEEE Access* **2019**, *7*, 13960–13988.
20. 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.
21. Chan, S.C.; Fishman, S.; Canny, J.; Korattikara, A.; Guadarrama, S. Measuring the Reliability of Reinforcement Learning Algorithms. *arXiv* **2019**, arXiv:1912.05663.
22. Hicks, A.L.; Wheeler, N.; Sánchez-Busó, L.; Rakeman, J.L.; Harris, S.R.; Grad, Y.H. Evaluation of parameters affecting performance and reliability of machine learning-based antibiotic susceptibility testing from whole genome sequencing data. *PLoS Comput. Biol.* **2019**, *15*, e1007349.
23. Hanif, M.A.; Khalid, F.; Putra, R.V.W.; Rehman, S.; Shafique, M. Robust machine learning systems: Reliability and security for deep neural networks. In Proceedings of the 2018 IEEE 24th International Symposium on On-Line Testing And Robust System Design (IOLTS), Platja d’Aro, Spain, 2–4 July 2018; pp. 257–260.
24. Corchado Rodríguez, J.M. *Efficiency and Reliability in Bringing AI into Transport and Smart Cities Solutions*; Gredos: Berlin, Germany, 21 October 2020.
25. Mittal, S. A survey on modeling and improving reliability of DNN algorithms and accelerators. *J. Syst. Archit.* **2020**, *104*, 101689.
26. 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.
27. Prost, F.; Thain, N.; Bolukbasi, T. Debiasing Embeddings for Reduced Gender Bias in Text Classification. In Proceedings of the First Workshop on Gender Bias in Natural Language Processing; Association for Computational Linguistics, Florence, Italy, 28 July–2 August 2019; pp. 69–75, doi:10.18653/v1/W19-3810.
28. Bellamy, R.K.; Dey, K.; Hind, M.; Hoffman, S.C.; Houde, S.; Kannan, K.; Lohia, P.; Martino, J.; Mehta, S.; Mojsilovic, A.; et al. AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias. *arXiv* **2018**, arXiv:1810.01943.
29. Rahmattalabi, A.; Vayanos, P.; Fulginiti, A.; Rice, E.; Wilder, B.; Yadav, A.; Tambe, M. Exploring algorithmic fairness in robust graph covering problems. In Proceedings of the Advances in Neural Information Processing Systems, Vancouver, BC, Canada, 8–14 December 2019; pp. 15776–15787.
30. Verma, S.; Rubin, J. Fairness definitions explained. In Proceedings of the 2018 IEEE/ACM International Workshop on Software Fairness (FairWare), Gothenburg, Sweden, 29 May 2018; pp. 1–7.
31. Chen, M.; Shahverdi, A.; Anderson, S.; Park, S.Y.; Zhang, J.; Dachman-Soled, D.; Lauter, K.; Wu, M. Transparency Tools for Fairness in AI (Luskin). *arXiv* **2020**, arXiv:2007.04484.
32. Adadi, A.; Berrada, M. Peeking inside the black-box: A survey on Explainable Artificial Intelligence (XAI). *IEEE Access* **2018**, *6*, 52138–52160, doi:10.1109/ACCESS.2018.2870052.
33. Tjoa, E.; Guan, C. A Survey on Explainable Artificial Intelligence (XAI): Towards Medical XAI. *CoRR* **2019**, abs/1907.07374.

34. Bargas-Avila, J.A.; Hornbæk, K. Old wine in new bottles or novel challenges: a critical analysis of empirical studies of user experience. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems; Vancouver BC, Canada, 7–12 May 2011; pp. 2689–2698.
35. Farias, C.M.D.; Li, W.; Delicato, F.C.; Pirmez, L.; Zomaya, A.Y.; Pires, P.F.; Souza, J.N.D. A systematic review of shared sensor networks. *ACM Comput. Surv. (CSUR)* **2016**, *48*, 1–50.
36. Kitchenham, B. *Procedures for Performing Systematic Reviews*; Keele University: Keele, UK, 2004; Volume 33, pp. 1–26.
37. Kitchenham, B.; Brereton, O.P.; Budgen, D.; Turner, M.; Bailey, J.; Linkman, S. Systematic literature reviews in software engineering—a systematic literature review. *Inf. Softw. Technol.* **2009**, *51*, 7–15.
38. Gunning, D. Explainable artificial intelligence (xai). *Def. Adv. Res. Proj. Agency (DARPA) Web* **2017**, *2*, 2.
39. Arrieta, A.B.; Díaz-Rodríguez, N.; Del Ser, J.; Bennetot, A.; Tabik, S.; Barbado, A.; García, S.; Gil-López, S.; Molina, D.; Benjamins, R.; et al. Explainable Artificial Intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* **2020**, *58*, 82–115.
40. Byrne, R.M. *Counterfactuals in Explainable Artificial Intelligence (XAI): Evidence from Human Reasoning*; IJCAI: Macao, China, 10–16 August 2019; pp. 6276–6282.
41. Das, A.; Rad, P. Opportunities and Challenges in Explainable Artificial Intelligence (XAI): A Survey. *arXiv* **2020**, arXiv:2006.11371.
42. Islam, M.A.; Anderson, D.T.; Pinar, A.; Havens, T.C.; Scott, G.; Keller, J.M. Enabling explainable fusion in deep learning with fuzzy integral neural networks. *IEEE Trans. Fuzzy Syst.* **2019**, *28*, 1291–1300.
43. Futia, G.; Vetrò, A. On the Integration of Knowledge Graphs into Deep Learning Models for a More Comprehensible AI—Three Challenges for Future Research. *Information* **2020**, *11*, 122.
44. Vijay, A.; Umadevi, K. Explainable AI controlled architecture of D2D system for massive MIMO based 5G networks. In Proceedings of the Second International Conference on Nexgen Technologies, Mumbai, India, 8–9 March 2019.
45. Papanastopoulos, Z.; Samala, R.K.; Chan, H.P.; Hadjiiski, L.; Paramagul, C.; Helvie, M.A.; Neal, C.H. Explainable AI for medical imaging: deep-learning CNN ensemble for classification of estrogen receptor status from breast MRI. In *Medical Imaging 2020: Computer-Aided Diagnosis*; International Society for Optics and Photonics: Houston, TX, USA, 2020; Volume 11314, p. 113140Z. doi.org/10.1117/12.2549298.
46. Oyekanlu, E. Distributed osmotic computing approach to implementation of explainable predictive deep learning at industrial iot network edges with real-time adaptive wavelet graphs. In Proceedings of the 2018 IEEE First International Conference on Artificial Intelligence and Knowledge Engineering (AIKE), Laguna Hills, CA, USA, 26–28 September 2018; pp. 179–188.
47. Jin, W.; Carpendale, S.; Hamarneh, G.; Gromala, D. *Bridging AI Developers and End Users: an End-User-Centred Explainable AI Taxonomy and Visual Vocabularies*; In Proceedings of the IEEE Visualization, Vancouver, BC, Canada, 20–25 October 2019.
48. 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, doi:10.1145/3351095.3375624.
49. 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*, doi:10.1145/3359206.
50. 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, doi:10.1145/3313831.3376261.
51. Luger, E.; Sellen, A. “Like Having a Really Bad PA” The Gulf between User Expectation and Experience of Conversational Agents. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, 7–12 May 2016; pp. 5286–5297.
52. Candello, H.; Pinhanez, C.; Figueiredo, F. Typefaces and the perception of humanness in natural language chatbots. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, 7–12 May 2017; pp. 3476–3487.
53. Yang, Q.; Cranshaw, J.; Amershi, S.; Iqbal, S.T.; Teevan, J. Sketching NLP: A case study of exploring the right things to design with language intelligence. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
54. Yang, X.; Aurisicchio, M.; Baxter, W. Understanding affective experiences with conversational agents. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
55. 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.
56. Gero, K.I.; Ashktorab, Z.; Dugan, C.; Pan, Q.; Johnson, J.; Geyer, W.; Ruiz, M.; Miller, S.; Millen, D.R.; Campbell, M.; et al. Mental Models of AI Agents in a Cooperative Game Setting. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems; Association for Computing Machinery, CHI’20, Honolulu, HI, USA, 25 April 2020; Association for Computing Machinery: New York, NY, USA, 2020; pp. 1–12, doi:10.1145/3313831.3376316.
57. Schaekermann, M.; Beaton, G.; Sanoubari, E.; Lim, A.; Larson, K.; Law, E. Ambiguity-Aware AI 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, doi:10.1145/3313831.3376506.

58. 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, doi:10.1145/3313831.3376590.
59. 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, doi:10.1145/3290605.3300830.
60. 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, doi:10.1145/3290605.3300509.
61. 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.
62. 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, doi:10.1145/3196709.3196729.
63. 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, doi:10.1145/3322276.3322379.
64. Krening, S.; Feigh, K.M. Interaction Algorithm Effect on Human Experience with Reinforcement Learning. *J. Hum.-Robot Interact.* **2018**, *7*, doi:10.1145/3277904.
65. Kayacik, C.; Chen, S.; Noerly, S.; Hollbrook, 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, doi:10.1145/3290607.3299059.
66. Hong, J.; Lee, K.; Xu, J.; Kacorri, H. Crowdsourcing the Perception of Machine Teaching. 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, doi:10.1145/3313831.3376428.
67. 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, doi:10.1145/3313831.3376718.
68. Santhanam, S.; Karduni, A.; Shaikh, S. Studying the Effects of Cognitive Biases in Evaluation 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, doi:10.1145/3313831.3376318.
69. 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, doi:10.1145/3313831.3376258.
70. 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, doi:10.1145/3313831.3376223.
71. 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, doi:10.1145/3313831.3376445.
72. 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, doi:10.1145/3313831.3376416.
73. 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, doi:10.1145/3313831.3376624.
74. Völkel, S.T.; Haeuslschmid, R.; Werner, A.; Hussmann, H.; Butz, A. How to Trick AI: Users' Strategies for Protecting Themselves from Automatic Personality Assessment. 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, doi:10.1145/3313831.3376877.
75. 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, doi:10.1145/3313831.3376219.

76. 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, doi:10.1145/3377325.3377519.
77. Ishibashi, T.; Nakao, Y.; Sugano, Y. Investigating Audio Data Visualization for Interactive Sound Recognition. 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. 67–77, doi:10.1145/3377325.3377483.
78. 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, doi:10.1145/3377325.3377512.
79. 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, doi:10.1145/3377325.3377501.
80. 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, doi:10.1145/3172944.3173008.
81. 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, doi:10.1145/3301275.3302310.
82. 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, doi:10.1145/3377325.3377480.
83. 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.
84. 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, doi:10.1145/3313831.3376514.
85. Louie, R.; Coenen, A.; Huang, C.Z.; Terry, M.; Cai, C.J. Novice-AI Music Co-Creation via AI-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.
86. 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.
87. 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.
88. 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.
89. 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.
90. Fast, E.; McGrath, W.; Rajpurkar, P.; Bernstein, M.S. Augur: Mining human behaviors from fiction to power interactive systems. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, San Jose, CA, USA, 7–12 May 2016; pp. 237–247.
91. Xu, A.; Liu, Z.; Guo, Y.; Sinha, V.; Akkiraju, R. A new chatbot for customer service on social media. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, 7–12 May 2017; pp. 3506–3510.
92. Higuchi, K.; Yonetani, R.; Sato, Y. Egoscanning: Quickly scanning first-person videos with egocentric elastic timelines. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, 7–12 May 2017; pp. 6536–6546.
93. Kimura, N.; Kono, M.; Rekimoto, J. SottoVoce: an ultrasound imaging-based silent speech interaction using deep neural networks. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–11.
94. Choi, M.; Park, C.; Yang, S.; Kim, Y.; Choo, J.; Hong, S.R. Aila: Attentive interactive labeling assistant for document classification through attention-based deep neural networks. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
95. Alghofaili, R.; Sawahata, Y.; Huang, H.; Wang, H.C.; Shiratori, T.; Yu, L.F. Lost in style: Gaze-driven adaptive aid for vr navigation. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
96. 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.

97. 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.
98. 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.
99. 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.
100. Laput, G.; Harrison, C. Sensing fine-grained hand activity with smartwatches. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–13.
101. 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.
102. 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.
103. 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.
104. 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.
105. Aksan, E.; Pece, F.; Hilliges, O. Deepwriting: Making digital ink editable via deep generative modeling. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–14.
106. Hu, T.; Xu, A.; Liu, Z.; You, Q.; Guo, Y.; Sinha, V.; Luo, J.; Akkiraju, R. Touch your heart: A tone-aware chatbot for customer care on social media. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–12.
107. Kimura, N.; Rekimoto, J. ExtVision: augmentation of visual experiences with generation of context images for a peripheral vision using deep neural network. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–10.
108. Huber, B.; McDuff, D.; Brockett, C.; Galley, M.; Dolan, B. Emotional dialogue generation using image-grounded language models. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–12.
109. Oh, C.; Song, J.; Choi, J.; Kim, S.; Lee, S.; Suh, B. 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, Montreal, QC, Canada, 22–27 April 2018; pp. 1–13.
110. 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, doi:10.1145/3313831.3376807.
111. 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, doi:10.1145/3313831.3376324.
112. 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, doi:10.1145/3313831.3376836.
113. 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, doi:10.1145/3313831.3376147.
114. 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, doi:10.1145/3313831.3376238.
115. 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, doi:10.1145/3313831.3376520.
116. Wu, J.; Harrison, C.; Bigham, J.P.; Laput, G. Automated Class Discovery and One-Shot Interactions for Acoustic Activity Recognition. 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, doi:10.1145/3313831.3376875.

117. 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, doi:10.1145/3313831.3376143.
118. 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, doi:10.1145/3313831.3376625.
119. Zhao, N.; Kim, N.W.; Herman, L.M.; Pfister, H.; Lau, R.W.; Echevarria, J.; Bylinskii, Z. ICONATE: Automatic Compound Icon Generation and Ideation. 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, doi:10.1145/3313831.3376618.
120. 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, doi:10.1145/3313831.3376866.
121. Im, J.; Tandon, S.; Chandrasekharan, E.; Denby, T.; Gilbert, E. Synthesized Social Signals: Computationally-Derived Social Signals from Account Histories. 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, doi:10.1145/3313831.3376383.
122. Xiao, Z.; Zhou, M.X.; Chen, W.; Yang, H.; Chi, C. If I Hear You Correctly: Building and Evaluating Interview Chatbots with Active Listening 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, doi:10.1145/3313831.3376131.
123. Mayer, S.; Laptu, G.; Harrison, C. Enhancing Mobile Voice Assistants with WorldGaze. 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, doi:10.1145/3313831.3376479.
124. 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, doi:10.1145/3313831.3376418.
125. 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, doi:10.1145/3313831.3376252.
126. 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, doi:10.1145/3313831.3376281.
127. 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, doi:10.1145/3313831.3376447.
128. 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, doi:10.1145/3313831.3376732.
129. Sterman, S.; Huang, E.; Liu, V.; Paulos, E. Interacting with Literary Style through Computational Tools. 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, doi:10.1145/3313831.3376730.
130. 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, doi:10.1145/3332165.3347867.
131. 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, doi:10.1145/3332165.3347924.
132. Bylinskii, Z.; Kim, N.W.; O'Donovan, P.; Alsheikh, S.; Madan, S.; Pfister, H.; Durand, F.; Russell, B.; Hertzmann, A. Learning Visual Importance for Graphic Designs and Data Visualizations. In Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology, UIST'17, Quebec, QC, Canada, 22–25 October 2017; pp. 57–69, doi:10.1145/3126594.3126653.
133. 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, doi:10.1145/3242587.3242652.
134. 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, doi:10.1145/3332165.3347878.

135. 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, doi:10.1145/3242587.3242599.
136. Rong, X.; Yan, S.; Oney, S.; Dontcheva, M.; Adar, E. CodeMend: Assisting Interactive Programming with Bimodal Embedding. In Proceedings of the 29th Annual Symposium on User Interface Software and Technology, UIST'16, Tokyo, Japan, 16–19 October 2016; pp. 247–258, doi:10.1145/2984511.2984544.
137. 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, doi:10.1145/3332165.3347876.
138. Laielli, M.; Smith, J.; Biamby, G.; Darrell, T.; Hartmann, B. LabelAR: A Spatial Guidance Interface for Fast Computer Vision Image Collection. 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. 987–998, doi:10.1145/3332165.3347927.
139. Yadav, K.; Gandhi, A.; Biswas, A.; Shrivastava, K.; Srivastava, S.; Deshmukh, O. ViZig: Anchor Points Based Non-Linear Navigation and Summarization in Educational Videos. In Proceedings of the 21st International Conference on Intelligent User Interfaces, IUI'16, Sonoma, CA, USA, 7–10 March 2016; pp. 407–418, doi:10.1145/2856767.2856788.
140. 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, doi:10.1145/3172944.3172977.
141. 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, doi:10.1145/3377325.3377510.
142. Shrivastava, A.; Heer, J. ISeqL: Interactive Sequence 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. 43–54, doi:10.1145/3377325.3377503.
143. Karimi, P.; Rezwana, J.; Siddiqui, S.; Maher, M.L.; Dehbozorgi, N. Creative Sketching Partner: An Analysis of Human-AI Co-Creativity. 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. 221–230, doi:10.1145/3377325.3377522.
144. 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, doi:10.1145/3377325.3377485.
145. Hoque, E.; Joty, S.; Marquez, L.; Carenini, G. CQAVIS: Visual Text Analytics for Community Question Answering. In Proceedings of the 22nd International Conference on Intelligent User Interfaces, Limassol, Cyprus, 13–16 March 2017; IUI'17, pp. 161–172, doi:10.1145/3025171.3025210.
146. Zhelezniakov, D.; Cherneha, A.; Zaytsev, V.; Ignatova, T.; Radyvonenko, O.; Yakovchuk, O. Evaluating New Requirements to Pen-Centric Intelligent User Interface Based on End-to-End Mathematical Expressions Recognition. 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. 212–220, doi:10.1145/3377325.3377482.
147. Aydin, A.S.; Feiz, S.; Ashok, V.; Ramakrishnan, I. SaLL: 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, doi:10.1145/3377325.3377540.
148. Grover, T.; Rowan, K.; Suh, J.; McDuff, D.; Czerwinski, M. Design and Evaluation of Intelligent Agent Prototypes for Assistance with Focus and Productivity at Work. 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. 390–400, doi:10.1145/3377325.3377507.
149. Sen, S.; Swoap, A.B.; Li, Q.; Boatman, B.; Dippenaar, I.; Gold, R.; Ngo, M.; Pujol, S.; Jackson, B.; Hecht, B. Cartograph: Unlocking Spatial Visualization Through Semantic Enhancement. In Proceedings of the 22nd International Conference on Intelligent User Interfaces, IUI'17, Limassol, Cyprus, 13–16 March 2017; pp. 179–190, doi:10.1145/3025171.3025233.
150. Kulahcioglu, T.; de Melo, G. Paralinguistic Recommendations for Affective Word Clouds. In Proceedings of the 24th International Conference on Intelligent User Interfaces, IUI'19, Los Angeles, CA, USA, 7–11 March 2019; pp. 132–143, doi:10.1145/3301275.3302327.
151. Sereshkeh, A.R.; Leung, G.; Perumal, K.; Phillips, C.; Zhang, M.; Fazly, A.; Mohamed, I. VASTA: A Vision and Language-Assisted Smartphone Task Automation System. 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. 22–32, doi:10.1145/3377325.3377515.
152. 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, doi:10.1145/2856767.2856795.
153. Donahue, 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, doi:10.1145/3301275.3302288.

154. 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, doi:10.1145/3377325.3377494.
155. 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, doi:10.1145/3301275.3302273.
156. 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, doi:10.1145/3025171.3025208.
157. 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.
158. 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.
159. 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.
160. 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.
161. 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.
162. 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.
163. Gurari, D.; Grauman, K. CrowdVerge: Predicting if people will agree on the answer to a visual question. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, CO, USA, 7–12 May 2017; pp. 3511–3522.
164. Huang, H.; Solah, M.; Li, D.; Yu, L.F. Audible panorama: Automatic spatial audio generation for panorama imagery. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–11.
165. 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.
166. Kim, J.; Stengel, M.; Majercik, A.; De Mello, S.; Dunn, D.; Laine, S.; McGuire, M.; Luebke, D. Nvgaze: An anatomically-informed dataset for low-latency, near-eye gaze estimation. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, Glasgow, UK, 4–9 May 2019; pp. 1–12.
167. Liu, Z.; Xu, A.; Guo, Y.; Mahmud, J.U.; Liu, H.; Akkiraju, R. Seemo: A computational approach to see emotions. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–12.
168. Ryokai, K.; Durán López, E.; Howell, N.; Gillick, J.; Bamman, D. Capturing, representing, and interacting with laughter. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, Montreal, QC, Canada, 22–27 April 2018; pp. 1–12.
169. Yuan, A.; Li, Y. Modeling Human Visual Search Performance on Realistic Webpages Using Analytical and Deep Learning Methods. 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, doi:10.1145/3313831.3376870.
170. 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, doi:10.1145/3313831.3376615.
171. Lai, C.; Lin, Z.; Jiang, R.; Han, Y.; Liu, C.; Yuan, X. Automatic Annotation Synchronizing with Textual Description for Visualization. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; pp. 1–13, doi:10.1145/3313831.3376443.
172. Yang, J.J.; Banerjee, G.; Gupta, V.; Lam, M.S.; Landay, J.A. Soundr: Head Position and Orientation Prediction Using a Microphone Array. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI'20, Honolulu, HI, USA, 25 April 2020; pp. 1–12, doi:10.1145/3313831.3376427.
173. 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, doi:10.1145/3313831.3376589.
174. 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, doi:10.1145/3313831.3376423.

175. 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, doi:10.1145/3313831.3376518.
176. Donkers, T.; Kleemann, T.; Ziegler, J. Explaining Recommendations by Means of Aspect-Based Transparent Memories. In Proceedings of the 25th International Conference on Intelligent User Interfaces; Association for Computing Machinery: New York, NY, USA; 17–20 March 2020; IUI'20, pp. 166–176, doi:10.1145/3377325.3377520.
177. Arendt, D.L.; Nur, N.; Huang, Z.; Fair, G.; Dou, W. Parallel Embeddings: A Visualization Technique for Contrasting Learned Representations. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; pp. 259–274, doi:10.1145/3377325.3377514.
178. 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, Los Angeles, CA, USA, 7–11 March 2019; IUI'19, pp. 263–274, doi:10.1145/3301275.3302316.
179. Le, H.V.; Mayer, S.; Henze, N. Investigating the Feasibility of Finger Identification on Capacitive Touchscreens Using Deep Learning. In Proceedings of the 24th International Conference on Intelligent User Interfaces, Marina del Ray, CA, USA, 7–11 March 2019; IUI'19, pp. 637–649, doi:10.1145/3301275.3302295.
180. 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, doi:10.1145/3301275.3302284.
181. Micallef, L.; Sundin, I.; Marttinen, P.; Ammad-ud din, M.; Peltola, T.; Soare, M.; Jacucci, G.; Kaski, S. Interactive Elicitation of Knowledge on Feature Relevance Improves Predictions in Small Data Sets. In Proceedings of the 22nd International Conference on Intelligent User Interfaces, IUI'17, Limassol, Cyprus, 13–16 March 2017; pp. 547–552, doi:10.1145/3025171.3025181.
182. Athukorala, K.; Medlar, A.; Oulasvirta, A.; Jacucci, G.; Glowacka, D. Beyond Relevance: Adapting Exploration/Exploitation in Information Retrieval. In Proceedings of the 21st International Conference on Intelligent User Interfaces, Sonoma, CA, USA, 7–10 March 2016; IUI'16, pp. 359–369, doi:10.1145/2856767.2856786.
183. Mittal, P.; Aggarwal, K.; Sahu, P.P.; Vatsalya, V.; Mitra, S.; Singh, V.; Veera, V.; Venkatesan, S.M. Photo-Realistic Emoticon Generation Using Multi-Modal Input. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; pp. 254–258, doi:10.1145/3377325.3377533.
184. Weber, T.; Hußmann, H.; Han, Z.; Matthes, S.; Liu, Y. Draw with Me: Human-in-the-Loop for Image Restoration. In Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI'20, Cagliari, Italy, 17–20 March 2020; pp. 243–253, doi:10.1145/3377325.3377509.
185. 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, doi:10.1145/3377325.3377528.
186. Berg, S.; Kutra, D.; Kroeger, T.; Straehle, C.N.; Kausler, B.X.; Haubold, C.; Schiegg, M.; Ales, J.; Beier, T.; Rudy, M.; et al. Ilastik: interactive machine learning for (bio) image analysis. *Nat. Methods* **2019**, *16*, 1–7.
187. Kahng, M.; Thorat, N.; Chau, D.H.; Viégas, F.B.; Wattenberg, M. GAN Lab: Understanding Complex Deep Generative Models using Interactive Visual Experimentation. *IEEE Trans. Vis. Comput. Graph.* **2019**, *25*, 310–320.
188. Fillmore, N.; Do, N.V.; Brophy, M.; Zimolzak, A.J. *Interactive Machine Learning for Laboratory Data Integration*; MedInfo: Lyon, France, 2019; pp. 133–137.
189. Bernardo, F.; Grierson, M.; Fiebrink, R. User-centred design actions for lightweight evaluation of an interactive machine learning toolkit. *J. Sci. Technol. Arts* **2018**, *10*, 2–25.
190. Diaz, C.G.; Perry, P.; Fiebrink, R. Interactive Machine Learning for More Expressive Game Interactions. In Proceedings of the 2019 IEEE Conference on Games (CoG), London, UK, 20–23 August 2019, pp. 1–2.
191. Gogia, S. Insight: Interactive Machine Learning for Complex Graphics Selection. Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, 2016.
192. 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; p. 1–14, doi:10.1145/3290605.3300234.
193. 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, doi:10.1038/s41746-019-0131-z.
194. Wexler, J.; Pushkarna, M.; Bolukbasi, T.; Wattenberg, M.; Viégas, F.; Wilson, J. The What-If Tool: Interactive Probing of Machine Learning Models. *IEEE Trans. Vis. Comput. Graph.* **2020**, *26*, 56–65.
195. Ramos, G.; Meek, C.; Simard, P.; Suh, J.; Ghorashi, S. Interactive machine teaching: a human-centered approach to building machine-learned models. *Hum. Comput. Interact.* **2020**, *35*, 413–451, doi:10.1080/07370024.2020.1734931.
196. Amershi, S.; Weld, D.; Vorvoreanu, M.; Fournery, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S.; Bennett, P.N.; 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, doi:10.1145/3290605.3300233.

197. 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, doi:10.1145/3328243.3328253.
198. Xu, W. Toward Human-Centered AI: A Perspective from Human-Computer Interaction. *Interactions* **2019**, *26*, 42–46, doi:10.1145/3328485.
199. 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, doi:10.1145/3290607.3312877.
200. Riedl, M.O. Human-Centered Artificial Intelligence and Machine Learning. *CoRR* **2019**, abs/1901.11184.
201. 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*, doi:10.1145/3359249.
202. Ehsan, U.; Riedl, M.O. Human-centered Explainable AI: Towards a Reflective Sociotechnical Approach. *arXiv* **2020**, arXiv:2002.01092.
203. 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.
204. 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, Canada, 22–27 April 2018; pp. 1–14.
205. 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, doi:10.1145/3313831.3376727.
206. Yang, Q.; Steinfeld, A.; Rosé, C.; Zimmerman, J. Re-Examining Whether, Why, and How Human-AI 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, doi:10.1145/3313831.3376301.
207. 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.
208. Lubars, B.; Tan, C. Ask not what AI can do, but what AI should do: Towards a framework of task delegability. In Proceedings of the Advances in Neural Information Processing Systems; Vancouver, BC, Canada, 8–14 December 2019; pp. 57–67.
209. 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.
210. 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, Los Angeles, CA, USA, 7–11 March 2019; IUI'19, pp. 258–262, doi:10.1145/3301275.3302289.
211. 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.
212. Gallego Cascón, P.; Matthies, D.J.; Muthukumarana, S.; Nanayakkara, S. ChewIt. 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.
213. 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.
214. 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.
215. 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.
216. 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.
217. 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.
218. 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, doi:10.1145/2736277.2741094.
219. 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, doi:10.1145/3301275.3302274.

220. Fagan, L.; Shortliffe, E.; Buchanan, B. *Computer-Based Medical Decision Making: from MYCIN to VM*; Automedica, Gordon and Breach Science Publishers: Zurich, Switzerland, 1984.
221. Clancey, W.J. *Knowledge-Based Tutoring: The GUIDON Program*; MIT Press: Cambridge, MA, USA, 1987.
222. 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.
223. 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.
224. Sacha, D.; Sedlmair, M.; Zhang, L.; Lee, J.; Peltonen, J.; Weiskopf, D.; North, S.; Keim, D. What You See Is What You Can Change: Human-Centered Machine Learning By Interactive Visualization. *Neurocomputing* **2017**, *268*, doi:10.1016/j.neucom.2017.01.105.
225. Sacha, D.; Sedlmair, M.; Zhang, L.; Lee, J.A.; Weiskopf, D.; North, S.C.; Keim, D.A. *Human-Centered Machine Learning through Interactive Visualization: Review and Open Challenges*; ESANN: Bruges, Belgium, 2016.
226. Magomadov, V. *Deep Learning and Its Role in Smart Agriculture*; Journal of Physics: Conference Series; IOP Publishing: Madiun, Indonesia, 2019; Volume 1399, p. 044109.
227. Kleinberg, J.; Lakkaraju, H.; Leskovec, J.; Ludwig, J.; Mullainathan, S. Human Decisions and Machine Predictions. *Q. J. Econ.* **2017**, *133*, 237–293. doi:10.3386/w23180.
228. Pacaux-Lemoine, M.P.; Trentesaux, D.; Rey, G.Z.; Millot, P. Designing intelligent manufacturing systems through Human-Machine Cooperation principles: A human-centered approach. *Comput. Ind. Eng.* **2017**, *111*, 581–595.
229. Liu, J.; Wilson, A.; Gunning, D. Workflow-Based Human-in-the-Loop Data Analytics. In Proceedings of the 2014 Workshop on Human Centered Big Data Research, HCBDR'14, Raleigh, NC, USA, 1–3 April 2014; pp. 49–52, doi:10.1145/2609876.2609888.
230. Smilkov, D.; Thorat, N.; Assogba, Y.; Yuan, A.; Kreeger, N.; Yu, P.; Zhang, K.; Cai, S.; Nielsen, E.; et al. TensorFlow.js: Machine Learning for the Web and Beyond. *arXiv* **2019**, arXiv:1901.05350v2.
231. Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; Sutskever, I. Language models are unsupervised multitask learners. *OpenAI Blog* **2019**, *1*, 9.
232. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models are Few-Shot Learners. *arXiv* **2020**, arXiv:cs.CL/2005.14165.