Youth and Artificial Intelligence: Where We Stand

Citation
Hasse, Alexa, Sandra Cortesi, Andres Lombana-Bermudez, and Urs Gasser. 2019. Youth and Artificial Intelligence: Where We Stand. Berkman Klein Center for Internet & Society publication.

Permanent link
http://nrs.harvard.edu/urn-3:HUL.InstRepos:40268058

Terms of Use

Share Your Story

Youth and Artificial Intelligence:
WHERE WE STAND

Alexa Hasse   Sandra Cortesi   Andres Lombana-Bermudez   Urs Gasser

SUGGESTED CITATIONS

APA
Hasse, A., Cortesi, S., Lombana-Bermudez, A., & Gasser, U. (2019). *Youth and artificial intelligence: Where we stand.* Youth and Media, Berkman Klein Center for Internet & Society. Retrieved from https://cyber.harvard.edu/publication/2019/youth-and-artificial-intelligence/where-we-stand

Chicago (Bibliography)
Hasse, Alexa, Sandra Cortesi, Andres Lombana-Bermudez, and Urs Gasser. "Youth and Artificial Intelligence: Where We Stand," Youth and Media, Berkman Klein Center for Internet & Society (2019), accessed on [Month Day, Year], https://cyber.harvard.edu/publication/2019/youth-and-artificial-intelligence/where-we-stand

Chicago (Footnote)
Alexa Hasse, Sandra Cortesi, Andres Lombana-Bermudez, and Urs Gasser. "Youth and Artificial Intelligence: Where We Stand," Youth and Media, Berkman Klein Center for Internet & Society (2019), accessed on [Month Day, Year], https://cyber.harvard.edu/publication/2019/youth-and-artificial-intelligence/where-we-stand.

MLA
Hasse, Alexa, Sandra Cortesi, Andres Lombana-Bermudez, and Urs Gasser. "Youth and Artificial Intelligence: Where We Stand." Youth and Media, Berkman Klein Center for Internet & Society, 2019. Web. [Day Mon. Year]. <https://cyber.harvard.edu/publication/2019/youth-and-artificial-intelligence/where-we-stand>*.*

Bluebook
Alexa Hasse, Sandra Cortesi, Andres Lombana-Bermudez, and Urs Gasser. YOUTH AND ARTIFICIAL INTELLIGENCE: WHERE WE STAND (2019), *available at* https://cyber.harvard.edu/publication/2019/youth-and-artificial-intelligence/where-we-stand.

KEYWORDS

Youth, young people, children, adolescents, teenagers, high school, middle school, digital media, Internet, online, artificial intelligence, education, health, well-being, privacy, safety, future of work, creativity, entertainment.

PUBLICATION DATE

May 2019

LICENSE

ACKNOWLEDGEMENTS

Youth and Artificial Intelligence:
Where We Stand

By Alexa Hasse, Sandra Cortesi, Andres Lombana-Bermudez, & Urs Gasser.

This spotlight seeks to share Youth and Media's initial learnings and key questions around the intersection between artificial intelligence (AI) and youth (ages 12-18), in the context of domains such as education, health and well-being, and the future of work. It aims to encourage various stakeholders — including policymakers, educators, and parents and caregivers — to consider how we can empower young people to meaningfully interact with AI-based technologies to promote and bolster learning, creative expression, and well-being, while also addressing key challenges and concerns.

TABLE OF CONTENTS

I. INTRODUCTION

While much has been written about the impacts of artificial intelligence (AI) on society more broadly, little work has focused on how youth (12-18-year-olds) interact with and are impacted by AI. This spotlight reviews recent[1] literature and examples[2] on artificial intelligence and its impact on the lives of youth[3] in domains such as education, health and well-being, creativity and entertainment, and the future of work.[4] This piece is not intended to be a comprehensive discussion of all the challenges and opportunities that AI presents for young people; instead, the spotlight aims to highlight some of the Youth and Media team's initial exploratory learnings and questions around the interplay between AI and youth. The main aim is to share preliminary observations as well as highlight questions for a concerned public audience and provide different entry points for further exploration our team is pursuing.

The writing of this piece was inspired by four main questions:
- What are the domains of young people's lives impacted by AI (e.g., education, work, health and well-being) and, thus far, what does the research tell us about youth and AI in these contexts?
- What takeaways from prior research around youth and other forms of digital technologies can be examined with respect to AI?
- Within each domain of research, what are some relevant readings and initial questions related to youth we should be asking? If you feel that questions are missing in this spotlight, please share them with us, noting the specific domain(s) the questions fall under, through this brief survey (http://brk.mn/yamyouthandaiadditionalquestions).
- How might we empower young people to meaningfully interact with AI-powered technologies to promote and bolster learning, health and well-being, and creative expression, while also addressing key challenges and concerns?

---

1 We largely focused on reviewing reports, articles in scholarly journals, newspaper articles, and other outputs that have been published between 2017-2019.

2 N.B.: For Perma.cc links, for full site functionality (e.g., to be redirected to secondary links and view embedded videos), please click "View the live page" in the top right corner.

3 As our team focuses on youth ages twelve to eighteen, throughout this piece, we primarily excluded research and news articles that concentrate on the ways adults engage with and are impacted by AI-based technologies, as well as younger children (i.e., under age 12) — particularly in cases where it was unclear how and why the content was relevant to the youth population. For additional information on AI and society more broadly, please visit the Berkman Klein Center's website: https://cyber.harvard.edu.

4 If you would like to learn more about the domains discussed in this paper as they relate to youth and digital technologies, a helpful starting point may be *Born Digital: How Children Grow Up in a Digital Age* (2016) by John Palfrey and Urs Gasser.

1. Overview of Initial Themes and Issues

Over the past several years, technologies based on AI have started changing our daily lives. Innovations are rolled out at an accelerated pace, not only in professional working environments but also at home and in schools. Hello Barbie and Cozmo are just two examples of AI-enabled toys that have already made their way into some children's homes in the U.S., with many more in development around the world. AI-powered toys offer playful and creative opportunities for children, with some systems promoting enhanced literacy, social skills, and language development.

Thus far, there is less research on the beneficial impact of AI-based systems specifically on adolescents, compared to young children. However, recent reports and studies indicate that AI is playing an increasingly important role in, for instance, the domains of education and learning, and health and well-being. Within the formal education setting, AI-powered educational technology (ed tech) — such as intelligent tutoring systems, tailored curriculum plans, and intelligent virtual reality — can improve educational outcomes, and offer rich and engaging interactive learning experiences for youth. In informal and connected learning environments, such as MIT Media Lab's Scratch platform, youth have the opportunity to design and program AI-based interactive games, simulations, chatbots, and virtual robots with great benefits for creativity, learning, and self-expression.

Within the domain of health and well-being, AI-driven applications are being designed and deployed to address health care concerns for young people — particularly in the context of mental and behavioral health — in the form of diagnostic tools, therapeutic chatbots, and public health interventions. These technologies open up the potential for earlier intervention for vulnerable youth, increased access to and engagement with therapeutic services, and greater awareness around public health issues.

However, AI-based technologies also create challenges for young people that need to be addressed by a variety of stakeholders (including youth themselves in the conversation) to safeguard and uphold young people's rights (e.g., right to privacy, right to participation). Particularly in relation to privacy, there are serious concerns around how AI-powered toys and ed tech applications can collect and store personal data. There is a risk of undermining youth privacy if the companies that design AI-fueled technologies are not clear and ethical about how they collect user data, where that data is stored, who can access it, and what can be done with it. From the spoken dialogue between a teenager and a digital personal assistant to the record of geographic locations a high school student creates when using an ed tech homework app, a wide range of personal information from young users can be recorded and stored as they interact with AI-powered systems. Thus, it is crucial that AI-based technologies are designed in a responsible way that safeguards youth privacy, especially when, as we have seen in the online platforms ecosystem, personal data on the cloud raises concerns around data privacy, security, and retention.

Additionally, the complex interplay between data sets and algorithms that power these "black box" AI systems — particularly when these systems are connected to the Internet — lead to pressing questions around bias and discrimination, transparency and accountability, and privacy and safety. There are also important concerns around the impact that AI systems might have in amplifying existing social inequalities among youth of different races, socio-economic statuses, genders, and regions (e.g., Global North and Global South[5]). Obstacles to the adoption of AI-powered technologies in under-resourced schools and underrepresented homes could exacerbate existing gaps within the youth population with respect to access to AI systems and the skills to utilize them.

Last but not least, the deep knowledge gap separating most people affected by AI systems and their designers creates an information asymmetry that complicates and often threatens progress toward a more inclusive future.[6] The youth population is a particularly vulnerable group, as very little has been done to empower young people to critically engage the discourse surrounding the next generation of technologies that have a marked potential to shape their lives for better or for worse. As just one example, innovations in artificial intelligence are helping to reshape the labor market, with important implications for career trajectories, and the requisite skills for youth to thrive in this shifting occupational landscape. Taking into account the massive impact AI is predicted to have on their lives, it is critical that young people have a basic understanding of what AI is and the ethical, societal, and privacy- and safety-related implications of these technologies.

II. EDUCATION AND HEALTH AND WELL-BEING

Artificial intelligence systems are already enhancing the field of education by enabling a more customizable approach to teaching and learning. Additionally, AI-fueled technologies have begun to change the landscape of health and well-being, particularly within the area of mental health. While the interventions discussed in this part II are revolutionary in that they have the potential to promote engagement, learning, and well-being on a larger scale than previously possible, it is important to highlight the fact that, like many AI systems, these technologies are developed in ways that may restrict youth from underrepresented backgrounds from full participation.

1. Education

AI-powered ed tech in the classroom — such as intelligent tutoring systems, AI-based curriculum plans, and intelligent virtual reality — can enhance educational outcomes, and provide engaging learning experiences for young people. AI-fueled curriculum plans, often referred to as "curriculum playlists," and intelligent tutoring systems, offer personalized learning experiences for youth, adapting instruction and feedback to students' capabilities and needs — the former, through daily individualized learning activities based on a curriculum the system curates for each student, and the latter through one-to-one tutoring that aims to mimic human tutoring. Intelligent tutoring systems may also be incorporated into other forms of educational technologies, such as virtual reality systems. A character within the immersive experience, for instance, might offer AI-driven guidance and support, helping the user engage with the virtual world in ways that promote learning. Over time, as reduction in costs and increases in Internet connectivity allow more schools and homes to access AI-fueled technologies, the future of AI in education holds the potential for personalized learning at scale. Further advances in the field of AI-based ed tech may also include the development of "lifelong learning companions" — intelligent systems accessible via mobile devices that accompany youth across all levels of schooling, in and beyond the classroom, offering feedback on tasks, suggested educational resources, and activities that promote 21st century skills, such as leadership and creativity.

Particularly for youth from regions where high-quality education is limited, learning at scale has the potential to help democratize access to such education and provide more diverse learning opportunities. As some early evidence on the impact of massive open online courses (MOOCs) in these regions reveals, students from

low-income backgrounds with low levels of educational attainment have reported career benefits from taking such online courses. Moreover, AI-powered technologies such as automated translation offer the potential to customize and localize educational resources and content into the languages and contexts of these regions, helping to bridge the knowledge access gap.

Emerging Questions

1. How can we equip educators — especially those in low-resource communities — with the skills and support to implement AI-based technologies in the classroom in a thoughtful manner? What kind of training would educators require? How does privilege manifest itself in the means of education chosen by different groups/ stakeholders? What is the prevailing narrative about AI and technology in the classroom, and who shapes it?
2. How can educators and students be included in the design process of AI-powered educational technologies?
3. How can we bring personalized, interest-driven learning to scale while addressing the diversity of pathways and passions that learners from different demographic backgrounds, cultures, and social contexts have?
4. How can we leverage AI-driven personalization at scale to promote access to education in regions of the world where education is challenging for the general population, including youth, to obtain? How can we ensure that these tools are meaningful for the youth who use them — in other words, how can these systems take into account learners' sociocultural context?
5. What kinds of safeguards could best mitigate potential negative issues related to using AI-based technologies in the classroom? How can we balance AI systems' need for data while safeguarding the privacy and security of sensitive student data, both at the individual and cohort levels? How can we ensure that both the AI ed tech itself and the human decision-makers who oversee and engage the AI will approach the acquisition, storage, transmission, and analysis of AI-based student data in an ethical way? Are there educational functions — either inside or outside of the classroom — that AI ed tech is uniquely positioned to fill in a positive manner? Are there functions it should not fill, for privacy, security, or ethical reasons?
6. Now that AI ed tech is gaining more traction in learning ecosystems, how do we implement and iterate upon the earlier wave of ed tech privacy discussions — especially best practice takeaways — that developed around cloud-based ed tech? Many of the opportunities and challenges around AI ed tech and student privacy and security were previewed in this earlier multistakeholder exploration, which included academic research, new industry standards, new state laws, and many other tools. There is a need to take stock: which AI ed tech, student privacy, and security questions might be answered, at least initially, with reference to some of these earlier efforts? Which current and emerging questions around AI ed tech were not covered by these efforts but might find some points of guidance or inspiration from them? Are there privacy and security questions that flow from AI ed tech that are completely outside the realm of prior discussion?

7. Educational institutions may automate basic administrative processes, and in the case of significant or complex decisions, may use algorithms to produce information to inform human decision-making. For instance, schools may employ machine learning and AI-based technology to sort through student applications. How can we make sure that these systems do not (perhaps inadvertently, but systematically) exclude certain types of candidates if left unchecked and unsupervised? How can we ensure that students can understand how decisions are made about them to maintain their confidence in the system and institution?
8. How can we make sure that AI-based technologies used in sorting (e.g., through student applications), selecting, and hiring processes that attempt to classify applicants as prospectively successful (or unsuccessful) based on historical patterns, do not instead identify and reinforce the biases of their designers or available training data?
9. How do we make education around AI, in both formal and informal learning settings, accessible and engaging for young people — starting with an understanding of data[7]/computational fluency? Learning about and working with data is central in understanding the role of big data in AI machine learning systems and the core issues (e.g., biases embedded in data, lack of transparency in the data, etc.) surrounding these systems. Furthermore, how can we engage youth in conversations around the ethics of the design, development, and deployment of AI-based technologies,[8] inside and outside of the classroom? Additionally, how can we democratize access to education around AI and machine learning, particularly in the context of lower-income countries where many youth speak little or no English?

---

[7] In the U.S., the National Council of Teachers of Mathematics recommends data analysis, probability, and statistics as mainstream strands in the k-12 curriculum, which is aligned with the Common Core State Standards in mathematics recently adopted by 42 states, Washington, D.C., and other regions of the U.S. Somewhat similar curricula around data analysis and statics has been implemented in primary and secondary grades in countries such as Singapore, Norway, and New Zealand.

[8] The Digital Citizenship+ (Plus) Resource Platform may be a helpful starting point for resources in this vein.

Recommended Readings

- AI Grand Challenges for Education, Beverly Park Woolf, H. Chad Lane, Vinay K. Chaudhri, Janet L. Kolodner
- The Future of Education: How A.I. and Immersive Tech Will Reshape Learning Forever, Lucas Rizzotto
- Do We Really Want Computerized Systems Controlling the Learning Process?, Mitch Resnick
- Embedded EthiCS: Integrating Ethics Broadly Across Computer Science Education, Barbara J. Grosz, David Gray Grant, Kate Vredenburgh, Jeff Behrends, Lily Hu, Alison Simmons, Jim Waldo
- (Draft) Beijing Consensus on Artificial Intelligence and Education: Outcome Document of the International Conference on Artificial Intelligence and Education 'Plan Education in the AI Era: Lead the Leap,' United Nations Educational, Scientific and Cultural Organization (UNESCO)
- Framing the Law & Policy Picture: A Snapshot of K-12 Cloud-Based Ed Tech & Student Privacy in Early 2014, Leah Plunkett, Alicia Solow-Niederman, Urs Gasser (for background on the cloud-based ed tech discussion)
- Student Privacy and Ed Tech (K-12) Research Briefing, Leah Plunkett, Urs Gasser

2. Health and Well-being

Within the domain of health and well-being, AI-driven technologies — delivered, for instance, via mobile and web-based platforms, or on a larger scale, applied through public health interventions — have the potential to facilitate diagnosis, deliver targeted treatment, and increase societal awareness around complex health issues. In the context of the youth population, current diagnostic AI tools allow clinicians to assess suicide risk based on natural language processing, creating opportunities for earlier intervention and pathways of care for vulnerable youth. At the same time, AI-based diagnostic systems raise concerns, particularly around privacy. Outside of the clinical setting, AI-powered technologies available through Internet-enabled devices, such as chatbots, can deliver tailored mental health therapy to youth, allowing for increased user engagement and access to therapeutic services. At the broader community level, AI systems are able to promote awareness around major health problems, such as the human immunodeficiency virus (HIV), within high-risk populations. A recent innovation in the public health field, for example, uses AI technology to select peer leaders within communities of homeless youth to increase awareness around HIV testing and preventative measures.

Emerging Questions

1. As AI holds the potential to predict, assess, provide therapy for, and mitigate diseases/health issues among youth, how do we reduce violations of privacy and ethics in which entities such as credit card companies or employers can use AI to discriminate based on existing conditions or even pre-existing conditions? (For more questions on privacy, see part III, section 2)
2. How can AI-based technologies be designed to reduce stigma around mental health service-seeking behaviors among youth — particularly youth from communities where mental illness and treatment-seeking are highly stigmatized?
3. In developing AI-based mental health diagnostic and therapeutic services, how can we make sure that these systems account for cultural nuances in how young people express mental and emotional distress? How can we also make sure that they consider child development issues?
4. How can AI systems — acting more as community connectors and educators (rather than direct problem solvers) — be applied towards other public health issues (beyond raising awareness around HIV testing and prevention), such as prescription drug overdose, to better reach vulnerable youth?
5. How can youth trust AI without always fully understanding its functionality, and what are the implications of that "blind trust" in relation to psychological and emotional well-being?
6. How will AI voice-operated assistants shape youth's social skills and push boundaries? Will young people become more or less adept at communicating with others?
7. How can we involve youth themselves in developing AI-powered technologies (e.g., chatbots) to deliver accurate health information to young people in an engaging, interactive way?

Recommended Readings

*   Artificial Intelligence-Assisted Online Social Therapy for Youth Mental Health, Simon D'Alfonso, Olga Santesteban-Echarri, Simon Rice, Greg Wadley, Reeva Lederman, Christopher Miles, John Gleeson, Mario Alvarez-Jimenez
*   Teenage Suicide Is Extremely Difficult to Predict. That's Why Some Experts Are Turning to Machines for Help, Peter Holley
*   USC Researcher, and AI, Give Homeless Youth a Helping Hand with HIV Education, Joanna Clay
*   Advancement of Artificial Intelligence Opens Health Data Privacy to Attack, John Hickey
*   Alexa, Should We Trust You?, Judith Shulevitz

III. WORK, PRIVACY AND SAFETY, AND CREATIVITY AND ENTERTAINMENT

As AI-based technologies advance, especially those that can provide automation in the workplace, there is concern about how the labor market will be restructured. Indeed, a 2018 Pew Research Center report indicates that a majority of adults around the world are concerned about robots and computers taking over jobs currently performed by humans during their lifetime. The impact of these changes has the potential to be more acute for youth, as they tend to engage in work (e.g., retail and food services) that has a high probability of automation. These developments could also exacerbate current disparities in youth employment by race and geographies. Youth living in contexts characterized by service economies, low levels of educational attainment, and low levels of consumer demand will likely confront a reduction of current employment opportunities.

AI will not only influence how youth work, but also how they play. Social technologies that collect large amounts of user data and video games that incorporate AI-powered characters are popular with youth (social media platforms, such as Facebook and Twitter, and video games such as *Rocket League* or *No Man's Sky*). These trends necessitate an understanding of how young people's data is being used — in the context of play/leisure, education, and other domains — and how youth can effectively navigate privacy and safety concerns surrounding AI-driven technologies. Such developments also call for further efforts that give youth more agency and knowledge around AI-based systems in the context of playing, learning, and socializing.

1. Future of Work

Technological innovation is radically transforming the labor market, with the McKinsey Global Institute estimating that between 400 to 800 million jobs worldwide could be automated by 2030. Advances in AI-powered technologies, however, are also opening up opportunities in emerging jobs, creating a demand for new skills (e.g., around coding, data/big data) and "uniquely human" skills (e.g., empathy, creativity, community building) to thrive in this evolving occupational landscape. Against this backdrop, how can youth best prepare for the future labor market?

Emerging Questions

1. How can we most effectively equip youth, both through formal and informal learning spaces, with the skills to participate in the shifting occupational environment — not only technological skills (e.g., data analysis, programming) but also increasingly important higher-order cognitive skills (e.g., creativity, critical thinking) and social-emotional skills (e.g., empathy, adaptability, lifelong learning)? How can we create opportunities to *re*train young people as the skills needed to thrive in today's society shift over time?
2. The lack of diversity in the AI industry begs the question: how do we encourage youth from underrepresented populations to explore and pursue careers in the field of AI, through educational programming and informal learning?
3. How will AI-fueled technologies shape young people's perceptions of potential future career pathways? What aspects of the future of work are they optimistic about and what are their concerns? How do those perceptions impact their psychological and emotional well-being? And how can we include young people in the dialogue around the future of jobs to better support them navigate the challenges and leverage the opportunities related to changes in the occupational landscape?

Recommended Readings

*   Jobs Lost, Jobs Gained: What the Future of Work Will Mean for Jobs, Skills, and Wages, James Manyika, Susan Lund, Michael Chui, Jacques Bughin, Jonathan Woetzel, Parul Batra, Ryan Ko, Saurabh Sanghvi
*   The Future of Jobs and Jobs Training, Lee Rainie, Janna Anderson
*   Developing a More Diverse AI, Stephanie Wykstra

2. Privacy and Safety

As today's youth are sharing large amounts of personal data with others, knowingly (e.g., using social media platforms, posting photos) and unknowingly (e.g., likes and shares; online shopping; and using mobile apps, sensors, and trackers), AI-based technologies — and their reliance on big data — raise serious privacy and safety concerns. The main objective of many AI systems is to make predictions and draw inferences about individuals and groups by algorithmically detecting patterns in large volumes of data. These predictions range from the relatively innocuous, such as recommending YouTube videos based on viewing history, to the deeply personal, and potentially harmful, such as substance abuse from Facebook users' status updates and likes.

In addition, AI-based technologies may lead to other types of privacy and security concerns that go beyond the predictions made by a given AI system. For example, massive amounts of data — big data — collected by an AI system could be accessed by unauthorized third parties and used for illegal, exploitive, and/or dangerous purposes — analogous to the Cambridge Analytica scandal. Even decision-makers with authorized initial access to AI data might then make decisions based on the data that limit youths' current or future opportunities. How can we most effectively equip young people to navigate the privacy and safety issues that surround AI-based applications? Are there safety and privacy concerns for youth that might require action on the part of non-youth stakeholders, such as lawmakers, regulators, or other adult decisionmakers?

Emerging Questions

1. There are currently a variety of frameworks around the concept of digital citizenship addressing online privacy and safety that have helped drive educational reform (e.g., Singapore's Cyber Wellness curriculum; legislative initiatives within the U.S. to require public schools to teach digital citizenship). As AI-powered technologies are changing how we think about issues surrounding privacy and safety in our digital world, how can we involve multiple stakeholders — including policymakers, educators, parents and caregivers, and youth themselves — in developing educational frameworks that address the interplay between AI systems and online safety and privacy?

2. What are the advantages and limitations of using voluntary or compulsory digital citizenship education as a means of empowering youth to understand and protect their own data privacy and security when interacting with AI?

3. How can we empower youth specifically to address disconnects between adult stakeholder decision-making and youth experience around privacy and security concerns in the context of educational technologies? For example, in the U.S., laws and regulations that frame youth digital data privacy typically place all or most authority around consenting to data collection with parents or guardians, rather than with youth themselves. In some circumstances (notably, education) consent to share youth data may be made by another adult party (the school) as a substitute for parental or guardian consent. Thus, with AI ed tech, youth find themselves with limited to no legal rights under federal law to consent or not consent to their private data being collected or used by the AI system. In these instances where concerns about privacy and security impacts on youth may arise, which stakeholders might need to act, through what channels, and how might youth themselves have meaningful input?

4. Facial recognition technology holds the potential to promote young people's safety (e.g., helping to reunite missing children with their loved ones) but poses serious privacy concerns through increases in surveillance, and, by extension, collection of personal data (e.g., school-wide surveillance systems). How can we reconcile our aim of collecting data with protecting the data of young users?
5. As youth are often at the forefront of adopting the latest digital technologies, will they be able to opt out of using some (though likely not all) AI systems? If not, how will they be informed about the short- and long-term implications these technologies may have on their lives?
6. How can AI be used to create richer social-emotional learning experiences for young people (e.g., through AI-powered characters in virtual environments), and the adults who teach them (e.g., educators, parents and caregivers), around society's most charged inclusion issues — such as immigration, police violence against vulnerable individuals, and hate speech? How can such experiences be created in a way that permits meaningful learning without youth having to fear negative consequences from engaging AI applications honestly? For instance, if a young person who lacks legal immigration status discloses this feature of her identity through an interaction with an AI designed to teach youth about immigrants' rights, how can her privacy be protected such that she does not need to fear governmental or other institutional action being taken against her as a result of this disclosure?
7. How can AI devices be used outside of the educational realm to promote equity for vulnerable youth while protecting their privacy? With so many youth living in poverty, facing violence in their neighborhoods, and experiencing many other difficult and dangerous circumstances, there is no end to the challenges that AI devices conceivably might help tackle. For example, could there be an AI system that provides safe and effective child care solutions such that single parents working multiple jobs to make ends meet would not need to leave their children home alone or in a substandard child care facility while safeguarding children's privacy? As innovators consider how to use AI in these and other equity-promoting ways, how can they avoid creating a false dichotomy between equity and privacy?

Recommended Readings

*   Children's Privacy in the Big Data Era: Research Opportunities, Kathryn C. Montgomery, Jeff Chester, Tijana Milosevic
*   AI Now Report 2018, Meredith Whittaker, Kate Crawford, Roel Dobbe, Genevieve Fried, Elizabeth Kaziunas, Varoon Mathur, Sarah Myers West, Rashida Richardson, Jason Schultz, Oscar Schwartz
*   Hey Alexa, What Can You Hear? and What Will You Do with It?, Sapna Maheshwari

3. Creativity and Entertainment

A variety of video games (e.g., Overwatch, Uncharted, and Mario Kart) currently feature AI-powered virtual characters that players can interact with as they navigate the gaming environment. Advances in gaming and other types of interactive digital platforms, however, are now moving youth from being consumers to producers of AI-driven content. In the online learning community Scratch, for example, young people have the opportunity to design AI-based games, simulations, and chatbots. The creators of the video game Minecraft are developing a youth-friendly platform that allows young people to train AI agents in the game's virtual world, helping youth better understand the inner workings of machine learning algorithms. In the future, AI-based technologies may allow young people to work together with AI systems to produce a variety of creative content, such as music and art.

Emerging Questions

1. How can we extend and increase access to opportunities for collaboration with AI systems beyond gaming, appealing to the backgrounds and interests of a wide array of youth, in ways that promote young people's conceptual understanding of AI?
2. If a trained human cannot tell the difference between the work of a great composer and the work of a machine, how will this affect young people's motivation to create their own artistic work? How will AI's capacity to enhance human creativity inspire young people to engage in creative work that leverages AI?
3. How will the increasingly interactive nature of AI-powered video game characters — which can lead users to feel as if they are part of a "functioning social world" — impact youth's interpersonal skills and social-emotional development?

Recommended Readings

- How Do We Define Creativity?, Ian Sample, Jordan Erica Webber
- There Will Always Be Limits to How Creative a Computer Can Be, Tony McCaffrey
- Deep Learning Is Not a Replacement for Human Creativity, Period, Ben Dickson
- Drones & Dreams: A Speculative Sprint Story Collection, Edited by Amy Johnson, and produced by Malavika Jayaram (Digital Asia Hub)