NEW YORK CITY BAR

THE IMPACT OF THE USE OF ARTIFICIAL INTELLIGENCE ON PEOPLE WITH DISABILITIES

I. INTRODUCTION

The rapid pace of enhancements to artificial intelligence (AI) and the increasing embrace of AI means that AI will continue to have an impact on an unlimited number of industries and populations across the world. Some commentators have focused on the promise of the positive impact of AI while other commentators have focused on the potential perils of AI, including the harm that AI could cause to underserved and marginalized communities.[1] In this analysis, the New York City Bar Association Presidential Task Force on Artificial Intelligence and Digital Technologies (Task Force)[2] highlights the harm to people with disabilities already caused by the use of AI and the harm that is likely to continue to occur or occur in the future.

For example, disabled people often find themselves either stereotyped, objectified, or invisible in AI-generated content. Image and text generation models regularly reinforce negative or "inspirational" tropes, failing to provide authentic, diverse representation. Moreover, AI relies on large-scale statistical learning, which tends to optimize for the "average" user. People with disabilities, who form a highly diverse and frequently underrepresented group, are often excluded or mischaracterized as a result. Additionally, most digital content remains inaccessible, and AI systems trained on such data perpetuate or even amplify these barriers. The statistical nature of AI can lead to discriminatory outcomes in employment, credit, education, and access to services, as disabled people are least likely to fit the "average" profile that these systems favor. Lastly, current approaches to monitoring AI-driven discrimination often fail to capture harms affecting people with disabilities, due to the uniqueness and variety of these experiences.

While the consensus view among AI experts seems to be that AI is a tool that must be responsibly embraced in order to reach its full potential, there is no consensus among AI experts on what should be done by policymakers to assure that both the use of AI is embraced and

[1] During the COVID-19 pandemic, when working from home became pervasive, video conferencing also became even more pervasive. Accessibility of the main video conferencing platforms (Zoom, Google Meet, Microsoft Teams, and others) was demanded by users and increasingly required by companies and public entities. Support for human-created and AI-supplied captioning was provided relatively quickly with provision of sign language support and accessible user interfaces following in short order.

[2] The Task Force is composed of approximately 250 representatives of more than 50 committees, councils, and other task forces of the City Bar and adjunct members, including lawyers, academics, trade association representatives, consultants, technologists, and others. See Task Force Dashboard at https://www.nycbar.org/committees/task-force-on-digital-technologies/ and https://www.nycbar.org/wp-content/uploads/2025/05/PTFAIDT-Leadership-250430.pdf (All websites last accessed on June 12, 2025).

About the Association
The mission of the New York City Bar Association, which was founded in 1870 and has 23,000 members, is to equip and mobilize a diverse legal profession to practice with excellence, promote reform of the law, and uphold the rule of law and access to justice in support of a fair society and the public interest in our community, our nation, and throughout the world.

continues to expand and the assurance that underserved and marginalized communities are protected from the harm and potential harm caused by AI.[3] Ultimately, in our view, equitable AI cannot be achieved without proactive, disability-centered design, rigorous inclusion throughout development, and responsive regulation. To address the challenges, this analysis puts forward several recommendations, including mandatory inclusion of disability data in AI training sets, standardization of accessibility benchmarks for AI, and direct inclusion of people with disabilities across all stages of AI development.

The following sections elaborate on these problems in detail and outline actionable, strategic solutions.

II. BUILDING A DISABILITY-INCLUSIVE AI ECOSYSTEM: ACHIEVING BENEFITS WHILE MINIMIZING HARMS

A truly disability-inclusive AI ecosystem involves more than technical fixes; it requires sustained collaboration among policymakers, industry, and the disability community. Reports from advocacy organizations highlight the importance of integrating disability rights perspectives throughout the entire AI development process.

In 2021, the Center for Democracy and Technology (CDT) and the American Association of People with Disabilities (AAPD) issued a report titled *Centering People with Disabilities in Tech Policy*[4] (2021 Report). The 2021 Report sets forth in writing some of the ways that disabled people are impacted by technologies, including by technologies such as algorithmic systems, AI, and other emerging technologies.

In March 2025, the CDT and AAPD issued another report titled, *Building a Disability-Inclusive AI Ecosystem, A Cross-Disability, Cross-Systems Analysis of Best Practices*[5] (Report). The Report can be seen as a continuation of the 2021 Report. The overarching idea

[3] This analysis was drafted by the Task Force Subcommittee on the Impact of Artificial Intelligence on People with Disabilities and Underserved Communities (AI and People with Disabilities Subcommittee) and approved by the Task Force Subcommittee on Articles and Blogs. The AI and People with Disabilities Subcommittee's primary mission includes addressing the implications and impact of the use of artificial intelligence on people with disabilities and underserved communities and the response of the judiciary and policymakers at the federal, state and local levels on enacting or modifying laws to take into account people with disabilities and underserved communities. The current members of the AI and People with Disabilities Subcommittee include (1) Cynthia Bennett, Senior Research Scientist, Responsible AI, Human-Computer Interaction, Disability, Accessibility at Google; (2) Michael Bervell, Founder and CEO at TestParty; (3) Co-Chair Larry Goldberg, former Head of Accessibility at Yahoo; (4) Co-Chair Jutta Treviranus, Director and Professor, Inclusive Design Research Centre at OCAD, University, Toronto, Canada; and (5) Sakeena Trice, Senior Staff Attorney at the New York Lawyers for the Public Interest. The Task Force Subcommittee on Articles and Blogs serves as a peer review for Task Force writings, especially articles, blogs, reports, statements and other writings. Current subcommittee members include (1) Angelena Bradfield, Head of Policy and Government Affairs at Financial Technology Association; (2) Robert Mahari, Harvard Law School and MIT Media Lab; (3) Lorraine McGowen, Task Force Co-Chair and Partner at Orrick, Herrington & Sutcliffe LLP; (4) Robert Schwinger, Partner at Norton Rose Fulbright US LLP; (5) Edwin Smith, Massachusetts Uniform Law Commissioner and Partner at Morgan, Lewis & Bockius LLP; (6) Tiffany Smith, Task Force Co-Chair and Partner at WilmerHale; and (7) Jerome Walker, Task Force Co-Chair and Partner at Jerome Walker PLLC.

[4] *Centering Disability in Technology Policy Issue Landscape and Potential Opportunities for Action* by Henry Claypool, Claire Carey, Alexander C. Hart and Linnea Lassiter available at https://cdt.org/wp-content/uploads/2021/12/centering-disability-120821-1326-final.pdf.

[5] *Building A Disability-Inclusive AI Ecosystem A Cross-Disability, Cross-Systems Analysis Of Best Practices* by Ariana Aboulafia, CDT and Henry Claypool, AAPD available at https://cdt.org/wp-content/uploads/2025/03/2025-03-11-CDT-Building-A-Disability-Inclusive-AI-Ecosystem-report-final.pdf.

behind the Report was that there are many different players in an AI ecosystem, including people with disabilities. To be effective, these groups must include disability rights and disability justice advocates so that tech policy organizations will have representatives that could play a significant role, especially when coupled with industry developers of these technologies, policymakers, regulators, and other key stakeholders.

The Report provides high level background as to how people with disabilities are impacted by AI across a variety of systems. Those systems include employment, education, public benefits, and information and communication technology (ICT). The Report emphasizes assistive tech, healthcare, transportation, and the criminal justice system.

The Report includes recommendations for those stakeholders, especially recommendations to policymakers as to how policymakers can make those systems more inclusive for people with disabilities. The underlying framework recognizes that these AI tools can be helpful for people with disabilities in various contexts, including maximizing benefits and minimizing or mitigating harms. Each of these stakeholders can play a role in achieving these goals.

III. DISABILITY REPRESENTATION BY GENERATIVE AI

Studies and focus groups strongly suggest that AI-generated texts and images distort the experiences of people with disabilities. There have also been research studies evaluating disability representation in AI-generated text and image outputs. The research studies demonstrate that while AI models are regularly updated and it may be impossible to replicate the exact examples, the examples could serve as guidance on the types of representational harms that could be produced and should be evaluated. Two of these studies,[6] for instance, collects qualitative evaluations of outputs from AI models that generate images and text, respectively, from focus groups of people who have disabilities. Overall, the focus groups include people with a variety of disabilities–from physical disabilities impacting mobility and motor movements to people with sensory disabilities such as hearing and vision disabilities to people with neurodivergence, cognitive, and mental health disabilities. The focus group participants identified several tropes, many of which showed up in both images and text outputs from AI models.

Notably, both stories about people with disabilities generated by large language models (LLMs) and images of people with disabilities generated from text-to-image (T2I) models overrepresented disability as physical (e.g., wheelchair users). This narrow conception of disability leaves out the variety of life experiences that are umbrellaed under the disability moniker. Participants also identified intersectional stereotypes in AI-generated images.[7] This finding was corroborated by a case study[8] reported in popular media where a T2I user generated over 100 images of autistic people and the results were quite homogenous, representing an

[6] Mack et al. 2024: https://dl.acm.org/doi/10.1145/3613904.3642166 and Gadiraju et al. 2023: https://dl.acm.org/doi/abs/10.1145/3593013.3593989.

[7] Mack et al. 2024: https://dl.acm.org/doi/10.1145/3613904.3642166.

[8] Crimmins 2023: https://www.dailydot.com/irl/ai-autistic-ableist/.

autistic person as a white man.[9] Another study[10] employed a method called sentiment analysis where LLMs analyze the tone of text. The authors of the study found that LLMs labeled conversations with disability terms as more negative and toxic than similar conversations without such terms. The LLMs were exhibiting a bias that disability is a more negative or toxic topic.[11] The authors of the study observed an inspirational tone in stories generated by AI about people with disabilities. Inspirational stories (sometimes referred to as "inspiration porn") are critiqued by people with disabilities as intended to objectify them to nondisabled people over being authentic representations.[12] Focus group participants pointed out text in stories such as disabled people being the best at their job with no qualifiers on how they achieved this or becoming a "huge hit" as a comedian, despite their disabilities. These stories lack substance into how people with disabilities go about achieving such success and the task for the benefit of others, but are meant to inspire, nonetheless.

Focus group participants evaluating AI-generated stories also pointed out themes of nondisabled people helping disabled people. This was concerning to focus group participants because the message communicated was that people with disabilities do not have agency, and nondisabled people should be in the service of "helping" disabled people. This inaccurately represents how people with disabilities prefer to be assisted because respectful assistance is self-determined by the disabled person seeking it. In some stories, such as one describing a wheelchair user crossing a street with no curb cut, the portrayed solutions were unsafe. In this story, a "little girl" carried the wheelchair user across the street. This form of assistance is not only disrespectful and unrealistic, but it promotes unsafe behavior that could put people at risk. Some of the more salient examples of other stereotypes and tropes focus groups of people with disabilities identified in outputs from LLMs and T2I models include:

1. People with disabilities were shown in images with stereotypical accessories and technologies such as blind people who were depicted wearing dark glasses. When this representation is the primary representation, it inaccurately communicates that all blind people wear dark glasses.

2. Accessible technologies were portrayed inaccurately in images such as hearing aids appearing in ways that do not reflect any actual hearing aids in images that were meant to be realistic portrayals.

3. Images also showed mobility aids and technologies being used incorrectly such as showing a white cane which blind and low vision people use to aid in navigation and obstacle detection dragging behind the user. However, in real-life use, the user would sweep the cane across the front of his or her body of where he or she intended to step next.

4. Images also showed outdated portrayals of technology. Focus group participants expected images to show technology of the current era unless specified for a

[9] *Do you see any diversity?: Man says AI is ableist after generating 100-plus images of 'an autistic person'—and they're all white men. Most were skinny, and weirdly enough, a disproportionate abundance had red hair.* by Tricia Crimmins (August 9, 2023) available at https://www.dailydot.com/irl/ai-autistic-ableist/.

[10] Venkit et al. 2023: https://aclanthology.org/2023.trustnlp-1.3/.

[11] Gadiraju et al. 2023: https://dl.acm.org/doi/abs/10.1145/3593013.3593989.

[12] Young 2014: https://www.ted.com/talks/stella_young_i_m_not_your_inspiration_thank_you_very_much?language=en.

different time period. Participants were concerned that this reflected a lack of image training data that shows contemporary accessible technologies.

5. Images portrayed bionic or science fiction like prosthetics and technologies. Participants expected realistic portrayals unless they specified for science fiction or futuristic renderings.

6. Images often portrayed people with disabilities alone, and looking sad or in pain. AI-generated stories about disability similarly portrayed exclusion such as friends expressing to the wheelchair-using protagonist that he cannot play basketball with them because he uses a wheelchair. These images and stories incorrectly portrayed disability as tragic and incapable of community participation.

7. Images depicted people with disabilities in medical or afterlife settings such as people in clouds. Participants identified these stereotypes as promoting disability as purely a medical issue which should be cured or communicating an afterlife scenario where disability is ostensibly cured.

8. Focus group participants identified another intersectional stereotype where AI-generated images of people with disabilities were shown to be older adults. Whereas certain disabilities do more often co-occur with older ages, younger people with disabilities also wanted to be represented to push back on this stereotype.

Researchers and legal scholars have also identified inaccuracies in AI-generated responses to user prompts. Inaccurate information has the potential to mislead all people, and particularly people with disabilities who are seeking information. AI can itself be an accessibility tool. For instance, several websites with information may be inaccessible to people with disabilities. An AI-generated summary of this information could provide the information in an accessible format. However, if the information is inaccurate, the inaccuracy compounds information inequity experienced by people with disabilities.[13] For example, a recent study found factual inaccuracies embedded in the stories focus group participants with disabilities prompted an LLM to generate.[14] In addition, the CDT demonstrated how AI-generated answers to user questions about voting, including those inquiring about voting processes for people with disabilities, were inaccurate.[15]

In certain cases, AI can provide bespoke user experiences tailored to people's particular abilities. For instance, Google's Project Relate is a specialized AI system allowing people with

[13] Adnin and Das 2024: https://dl.acm.org/doi/fullHtml/10.1145/3663548.3675631. The point the authors are making is that AI summaries are useful to people with disabilities and inaccuracies can compound information inequity for this group. This study interviewed blind and low vision people about how they use generative AI and information seeking was a top use case, and they had varied levels of concern (no concern to a lot of concern) about inaccuracies. Often this concern was contextualized. For example, they may trust AI-generated summaries and information less in higher-stakes information queries (e.g., medical information). *I look at it as the king of knowledge: How Blind People Use and Understand Generative AI Tools* (October 2024) by Rudaiba Adnin and Maitraye Das available at https://doi.org/10.1145/3663548.3675631.

[14] Gadiraju et al. 2023: https://dl.acm.org/doi/abs/10.1145/3593013.3593989.

[15] https://cdt.org/insights/brief-generating-confusion-stress-testing-ai-chatbot-responses-on-voting-with-a-disability/.

non-standard speech to train their own speech recognizer.[16] When the user trains the model with samples of his or her own voice, then he or she can use Project Relate to either show text or play voice transcripts of what has been spoken. This can improve interpersonal interactions, especially with strangers who have not learned how to understand the user's natural speech.

IV. FLAWS IN AI TRAINING DATA AND FLAW IN AI OUTPUTS

AI is frequently heralded as a transformative tool; yet for people with disabilities, AI often fails to meet even basic expectations. This failure is rooted in a fundamental issue: AI systems are only as equitable as the data and design choices that shape them. As a result, when AI tools are built and trained without the intentional inclusion of disability-centered data or perspectives, they systematically exclude, misrepresent, or fail to accommodate the very users who could benefit most from them. The approach used by TestParty[17] is designed to address this problem by grounding its approach to focus on compliance with the Web Content Accessibility Guidelines (WCAG), the global gold standard for accessible web design, referenced in legal frameworks such as the Americans with Disabilities Act (ADA),[18] the European Accessibility Act (EAA),[19] and Section 508 of the Rehabilitation Act.[20]

TestParty has identified a number of significant challenges as well as possible strategies related to addressing these problems:

1. Inaccessible Inputs Produce Inaccessible Outputs

A striking statistic illustrates the scale of the problem: as of 2023, 97% of the top one million website homepages are not accessible to people with disabilities, failing to meet even the most basic WCAG compliance levels.[21] This means that the vast majority of the digital content used to train LLMs and other AI systems inherently lacks accessibility. In practice, this creates a dangerous feedback loop. Developers assume AI can "fix" inaccessible code, when

[16] https://sites.research.google/relate/.

[17] TestParty is a digital accessibility company co-founded by the author of this section of the report. The mission of TestParty is to automate web accessibility compliance. To achieve this goal, TestParty uses AI to analyze source code and produce accessible outputs across websites, mobile apps, documents (including PDFs), and other digital interfaces.

[18] 42 U.S.C. Chapter 126. The four purposes of the ADA are (1) to provide a clear and comprehensive national mandate for the elimination of discrimination against individuals with disabilities; (2) to provide clear, strong, consistent, enforceable standards addressing discrimination against individuals with disabilities; (3) to ensure that the Federal Government plays a central role in enforcing the standards established in this chapter on behalf of individuals with disabilities; and (4) to invoke the sweep of congressional authority, including the power to enforce the fourteenth amendment and to regulate commerce, in order to address the major areas of discrimination faced day-to-day by people with disabilities. 42 U.S.C. 12101(b).

[19] The EAA is a directive, passed in 2019, that aims to improve the functioning of the internal market for accessible products and services, by removing barriers created by divergent rules in Member States. According to the European Commission, businesses will benefit from common rules on accessibility in the EU leading to costs reduction, easier cross-border trading, more market opportunities for their accessible products and services. In addition, the European Commission notes that persons with disabilities and elderly people will benefit from more accessible products and services in the market, accessible products and services at more competitive prices, fewer barriers when accessing transport, education and the open labour market, and more jobs available where accessibility expertise is needed.

[20] 29 U.S.C. 794d. Among other things, Section 508 requires federal department and agencies to provide for accessibility when developing, procuring, maintaining, or using electronic and information technology.

[21] WebAIM, 2023.

in reality, generative models trained on inaccessible data reproduce or amplify these flaws. If developers feed inaccessible HTML or untagged PDFs into an LLM with a prompt like "make this accessible," the outputs will likely reflect the same barriers present in the original material. AI cannot solve for accessibility if it was never trained to recognize it.

2. Reframing the Problem: Changing Inputs and Outputs

TestParty has identified two promising strategies to break this loop:

A. *Input Remediation*: Before invoking generative AI models, preprocess or "sanitize" the input code to ensure it aligns more closely with accessibility standards. This includes flagging missing alt text, semantic HTML issues, unlabeled form elements, and other WCAG violations.

B. *Output Transformation Using Deterministic Models*: In addition to using AI for input analysis, deterministic systems should be built that output standardized accessible elements such as closed captions for video, alt text for images, and properly structured HTML for screen readers. These hybrid models allow an ability to retain control and compliance in ways that pure generative systems often cannot deliver.

3. Artificial Specific Intelligence vs. Artificial General Intelligence

Much of the current hype in AI center around the promise of Artificial General Intelligence (AGI).[22] However, AGI's generality often comes at the expense of precision and inclusion. These models tend to center "average" users, marginalizing outlier populations in the statistical distributions used for training. In contrast, a more deliberate shift toward Artificial Specific Intelligence (ASI)[23] is the better option. This approach aligns closely with the insights of disability scholars like Jutta Treviranus,[24] who describes how AI systems that optimize for the "middle of the bell curve" inherently exclude people whose needs fall at the edges. Recent research supports this reframing. For instance, studies by scholars like Ethan Mollick[25] have shown that general-purpose AI improves the productivity of lower-performing users but offers limited benefit (or even harm) for domain experts and outliers. If this holds true across sectors, it means people with complex needs such as those with disabilities are not being effectively served by generalized tools.

4. AI as an Emergent, Not Prescriptive, Technology

One of the most underappreciated aspects of AI is its emergent behavior. The capabilities of large-scale models like GPT-4 or Google's Gemini are not fully known at launch. Instead, capabilities are discovered by users through experimentation and lived use. This emergent property presents both an opportunity and a danger. It means accessibility

[22] For purposes of this analysis, GAI is defined as systems that can perform any intellectual task that a human can perform.

[23] For purposes of this analysis ASI is defined as AI tools designed for narrow, well-defined use cases that serve specific communities, including people with disabilities.

[24] Jutta Treviranus is a Co-Chair of the AI and People with Disabilities Subcommittee and Director and Professor, Inclusive Design Research Centre at OCAD, University, Toronto, Canada.

[25] *Prompting Diverse Ideas: Increasing AI Idea Variance*, Meincke, Mollick, Terwiesch (January 2024), The Wharton School, University of Pennsylvania available at https://mackinstitute.wharton.upenn.edu/wp-content/uploads/2024/02/for-web-AI-idea-variance.pdf.

breakthroughs can happen, but only if marginalized users are actively experimenting, providing feedback, and shaping how AI tools evolve. Yet most AI development excludes disabled users. Without inclusive design, inclusive testing, and community-led feedback, emergent behaviors can become emergent harms. As Cynthia Bennett[26] highlighted in her research at Google, disabled users who tested image-generation and story-generation models frequently found themselves portrayed as tragic, inspirational, or simply nonexistent. This is a failure of representation and of data, design, and evaluation processes.

V. STATISTICAL DISCRIMINATION DUE TO ALGORITHMS AND PREDICTIVE ANALYTICS

There is an inherent dilemma in the relationship between the current forms of AI and people with disabilities. People with disabilities are the most heterogeneous of the protected minority groups.[27] There are no common data characteristics among people with disabilities other than a large enough difference from the average such that conventional products, services, and accepted knowledge are a mismatch with personal needs and personal reality.[28] Even self-identification as disabled is unreliable due to the reluctance to self-identify.[29] Data related to disability is generally an outlier in any population data set. This would be the case even if it were possible to achieve full proportional representation of people with disabilities in existing population data sets because disability data characteristics are highly unique and variable.[30]

AI, on the other hand, is powered by mechanized statistical reasoning. AI statistically predicts the average response given the available data. This is both helpful and harmful for people with disabilities.[31] It helps people to achieve the average by translating inaccessible information into accessible forms such as describing visual information for someone who is blind. It also enables people to produce average output such as correcting grammar; spelling and phrasing for someone experiencing reading or writing barriers; or expanding shortform[32] expressions into longform communication for someone with limited voluntary movement. This is done at the cost of individuality. Unfortunately, the performance of these helpful functions degrades more the further you are from the average. This leads to the assumption that the issue

---

[26] Cynthia Bennett is a Member of the AI and People with Disabilities Subcommittee and Senior Research Scientist, Responsible AI, Human-Computer Interaction, Disability, Accessibility at Google.

[27] McGuire C. *Measuring Difference, Numbering Normal: Setting the Standards for Disability in the Interwar Period* [Internet]. Manchester (UK): Manchester University Press; 2020. Chapter 1, NUMBERING NORMAL. https://www.ncbi.nlm.nih.gov/books/NBK578569/?report=reader.

[28] Treviranus, J., (2023) *Inclusive Design: Valuing Difference, Recognizing Complexity*. In Handbook of Disability: Critical Thought and Social Change in a Globalizing World. Springer.

[29] Iezzoni, L.I. (2000). *Disability: The Reluctant Identity*. Journal of Health Politics, Policy and Law 25(6), 1157-1167. https://muse.jhu.edu/article/15528.

[30] Treviranus, J., Gupta, A., (2020). *Inclusively Designed Artificial Intelligence*. In Schaffers, H. Vartianen, M., Bus, J., Digital Innovation and Societal Change. River Publishers, London, UK.

[31] Smith, E. M., Graham, D., Morgan, C., & MacLachlan, M. (2023). *Artificial intelligence and assistive technology: risks, rewards, challenges, and opportunities. Assistive Technology*, 35(5), 375–377. https://doi.org/10.1080/10400435.2023.2259247.

[32] In this context, shortform means the person chooses symbols or meaning words and the AI expands this to form a grammatically correct sentence or message (e.g., think Stephen Hawkings). It is also important to mention the motor and expressive communication uses of AI, not just the audio and image translation uses. In other words, this is a way of supporting communication for people who experience barriers to speaking and/or writing.

is fixed and that no further accommodation is required in education or employment situations, for example, leading to disparity among people with disabilities.

The same statistical reasoning is harmful when AI is used to make decisions about people.[33] Current AI is statistically biased toward the average and disabled people tend to be the furthest from the average. This leads to exclusion or discrimination in employment, academic admissions, lending, credit, insurance, custody or security decisions, for example.

Voluntary guidance and regulations to address the discrimination and harms done by AI also use statistical reasoning.[34] For instance, statistics are used to determine whether a class is discriminated against. Given the heterogeneity of disability and the resulting forms that discrimination takes, cluster analysis or comparison is ineffective because each case will be unique. The harm will be determined to be anecdotal or statistically insignificant. The same issue exists with risk-benefit analysis. The unique instances of harm will be outweighed by the majority of benefits, despite the gravity of the harm to the individual.

VI. CONCLUDING OBSERVATIONS AND RECOMMENDATIONS

AI holds genuine promise in improving digital access and independence for people with disabilities, but this potential will only be realized if solutions are grounded in disability inclusion at every stage. To truly make AI systems work for people with disabilities, the following recommendations should be considered:

1. **Enshrine the Principle of "Nothing About Us Without Us":** Include people with disabilities at every stage, from ideation to deployment, based on the principles of "nothing about us without us, and everything is about all of us."

2. **Mandate Full Accessibility in AI Development:** Require AI content creation and consumption tools to be fully accessible to people with disabilities in alignment with existing regulations and standards such as those identified in this analysis.

3. **Require the Inclusion and Weighting of Disability Data:** Mandate accessibility data inclusion in AI training sets, and weight underrepresented experiences accordingly. Instances unique to disabled users should be positively weighted to avoid exclusion by majority-centric models.

4. **Establish Standardized, Sector-Specific Accessibility Benchmarks:** Standardize accessibility benchmarks in generative and assistive AI evaluations much like WCAG does for web content.

5. **Support Hybrid AI Architectures:** Support and encourage hybrid AI architectures by combining deterministic accessibility rules with machine learning.

[33] Cathy O'Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown Publishers, 2016.

[34] https://pmc.ncbi.nlm.nih.gov/articles/PMC10037374/.

6. **Invest in Domain-Targeted (Artificial Specific Intelligence) Systems: Invest** in and prioritize specific, domain-targeted models, not just general-purpose models.

7. **Guarantee Non-AI Alternatives and Human-Centric Options: This recommendation applies to both AI decisions (e.g., who is employed or admitted; who receives services; who is audited; who is deemed a security threat; and what medical help a person qualifies for) and AI services (e.g., captions). People should be able to say that AI is not going to make a fair decision, and they want a human alternative because they are not average or fairly represented in the training data. Human alternatives should be equivalent in convenience, cost and timeliness.** In certain contexts (e.g., education and employment), it is important to ensure that people can always access human-created or supervised alternatives to AI (e.g., live captioners and human interpreters). Since AI captions are now taken for granted, it is harder to get human captioners in some situations.

8. **Foster Ongoing, Community-Led Testing and Feedback:** AI systems must undergo real-world evaluation by diverse users with disabilities to swiftly identify and remediate emerging harms, inaccuracies, or unintended side effects.

9. **Bolster Transparency, Documentation, and Accountability:** AI providers should be required to document both the sources and the design decisions underpinning their systems' treatment of disability. Accessible complaint and redress mechanisms are essential for individuals to report exclusion, inaccuracy, or discrimination.

10. **Strengthen Legal and Regulatory Safeguards:** Policymakers should review and strengthen anti-discrimination frameworks for AI, ensuring that unique and individualized forms of harm affecting disabled people are recognized and addressed—rather than being dismissed as "anecdotal" or insignificant by statistical measures.

AI holds real potential to unlock digital access, but only if it is built with the involvement of people with disabilities at every step in the process. With the right commitments—from policymakers, industry, and end users—AI can become a powerful ally in removing barriers and expanding opportunities for disabled people. The path forward, however, demands explicit design for inclusion, regulatory vigilance, and the recognition that accessible technology is best created by, with, and for disabled users themselves.

Presidential Task Force on Artificial Intelligence and Digital Technologies
Lorraine McGowen, Co-Chair
Edward So, Co-Chair
Jerome Walker, Co-Chair

June 2025

Contact
Raabia Qasim, Associate Policy Counsel | 212.382.6655 | rqasim@nycbar.org