Examining the technology-mediated cycles of injustice that contribute to digital ageism

Advancing the conceptualization of digital ageism: evidence and implications

Charlene H. Chu*

Lawrence S. Bloomberg Faculty of Nursing; University of Toronto; Cross-appointed, Institute for Life Course and Aging, University of Toronto; KITE-Toronto Rehabilitation Institute, University Health Network, Toronto, charlene.chu@utoronto.ca

Rune Nyrup

Leverhulme Centre for the Future of Intelligence, University of Cambridge, rn330@cam.ac.uk

Simon Donato Woodger

Scarborough Health Network, donatowoodgersimon@gmail.com

Kathleen Leslie

Faculty of Health Disciplines, Athabasca University, kleslie@athabascau.ca

Shehroz Khan

KITE – University Health Network, shehroz.khan@utoronto.ca

Corinne Bernette

University of Toronto, corinne.bernett@mail.utoronto.ca

Amanda Grenier

Norman and Honey Schipper Chair in Gerontological Social Work, Factor Inwentash Faculty of Social Work, University of Toronto, and Baycrest Hospital, amanda.grenier@utoronto.ca

ABSTRACT

Our work draws attention to digital ageism referring to the nexus of ageism (discrimination or bias related to age) that is mediated and perpetuated by artificial intelligent (AI) systems and technologies. Building on the World Health Organization's recently published policy brief entitled "Ageism in AI for Health" and our previous work about digital ageism, this paper aims to advance our current understanding and conceptualization of digital ageism in technology and AI systems broadly and beyond health alone. To do this, we will 1) elaborate on our conceptual model and the ageist technology-mediated cycles of injustice that can produce and reinforce digital ageism; 2) present empirical evidence of our descriptive analysis of seven commonly used facial image datasets to highlight data disparities for older adults which will provide real-world evidence that substantiates one of the elements in our ageist cycles of injustice; and 3) summarize results from our grey literature search of various grey literature databases including the AI ethics guidelines Global Inventory to identify guidance documents that address ageism in AI in research or technology development. This paper uniquely contributes conceptual and empirical evidence of digital ageism which will advance knowledge in the field and deepen our understanding of how ageism in AI is fostered by broader ageist cycles of injustice. Lastly, we will briefly provide future considerations to address digital ageism.

CCS CONCEPTS

• Computing methodologies~Artificial intelligence

KEYWORDS

AI, ageism, bias, discrimination, inequity, algorithmic bias

*corresponding author

1 Introduction

Artificial Intelligence (AI) in the last decade has been followed by increased attention, and at times outcry, towards technological developers and government regulators to combat the presence of bias and discrimination within this field [1]. The varieties and sources of racial- and gender-related discrimination in AI are numerous and have affected opportunities to access healthcare, fair employment, and fair judicial sentencing. The concern about how societal biases are encoded and how they are potentially amplified in AI systems is a threat to equity and inclusivity [2]. Our work draws attention to another type of societal bias — discrimination based on age — or ageism, and how this implicit bias can influence AI systems and technologies, otherwise referred to as *digital ageism* [3]. The World Health Organization (WHO) recently published a brief policy report on ageism in AI in health [4] as part of their global campaign to combat ageism. To expand the current discourse, this paper aims to advance our current understanding and conceptualisation of digital ageism in technology and AI systems *beyond* a health context through four specific objectives. First, we will further elaborate on our conceptual model that outlines ageist technology-mediated cycles of injustice that can produce and reinforce digital ageism in AI systems, and expand on these intersecting cycles that contribute to digital ageism. Second, we will present empirical evidence of our descriptive analysis of commonly used facial image datasets to highlight data disparities for older adults and substantiate one of the elements in the cycles of injustice. Third, we describe the results from our grey literature search of various databases including the AI Ethics Guidelines Global Inventory to identify guidance documents that address ageism in AI in research or technology development. Lastly, we will provide considerations for future directions in digital ageism research. This paper will advance knowledge in the field and deepen our understanding of how ageism in AI is fostered by broader ageist cycles of injustice.

2 Digital ageism: a contemporary form of ageism against older people

Ageism can be conceptualised as having three elements that include prejudicial attitudes towards older people and ageing, discriminatory practices against older people, and institutionalised policies and social practices that foster the attitudes and actions in the first two elements [5]. There is a global campaign to combat ageism (WHO Global Report on Ageism) [4] led by the WHO, which has underscored ageism and addressing ageism as a global priority and part of the public discourse. AI systems have emerged as a major driver of technological development in our 21st century digital world that continues to progress and advance, particularly as networked mobile devices have generated vast quantities of personal data from their users. In essence, the technologies form a new system that can further uphold and amplify discriminatory practices and institutionalised policies and social practices towards older people.

Digital ageism refers to age bias in technology and AI systems, the entrenched the inherent ageist tendencies in our society that creep into the technological infrastructure in our increasingly digital world. Digital ageism is supported and reinforced through the development and application of digital technologies [3]. Our published work has focused on exploring the technology-mediated ageist cycles of injustice as they relate to the development of AI [Figure 1]. These *cycles of injustice* perpetuate ageism through a variety of processes that either exclude, or reinforce ageist stereotypes about older adults in the way technology is designed, how data is gathered from it, how it is allocated, and how older adults are represented within these systems. This work is informed by Whittlestone et al. [29] who conceptualised the framework as a way to understand algorithmic bias within a broader context. Digital ageism and ageism sustain each other through the way ageism is encoded and amplified in AI systems and technology.

An early definition of bias in computer systems refers to a system's ability to systematically and unfairly discriminate against certain individuals or groups of individuals in favour of others [6]. According to Friedman and Nissenbaum [6]: "a system discriminates unfairly if it denies an opportunity or a good or if it assigns an undesirable outcome to an individual or group of individuals on grounds that are unreasonable or inappropriate" (p. 332). There are two distinct types of undesirable outcomes that can result from algorithmic bias. The first is harms of allocation that refer to the distribution of resources and opportunities, for example, deciding who gets access to health care resources or services, or who receives job advertisements or notifications about potential job prospects [7]. Second, harms of representation that pertain to how different groups or identities are represented and perceived by society as a whole [7]. These two types of harm are included in our conceptual understanding of digital ageism. This next section will provide empirical evidence of the design and technical factors that contribute to digital ageism and elaborate on the complex interplay that results in representation and allocation of harms that are then sustained and perpetuated through technology.

2.1 Design

In this section, we will discuss how digital ageism can arise within and from the *design* of digital technologies including, but are not limited to, the hardware, software, interface, and purpose of the technologies. We identified that ageism enters into technology design in three ways that present as ageism stereotypes about older adults; exclusion of older adults in the data being used as well as the process of technology development and testing; and data availability incentives that reinforce ageist stereotypes held by the developers [3].

Human developers and designers of technology are fallible and can be influenced by implicit age-related biases. Ageism is prevalent, deeply ingrained and more socially accepted than other forms of biases [4], and can be viewed as harmless or humorous. This nonchalant attitude is evidence as to by how ageism appears in our vernacular (e.g., having a senior moment, grey tsunami). A survey of 57 countries (N = 83,000) about ageist attitudes found that one in every two people held moderately to highly prejudicial views against older people [8]. Such stereotypes can guide the assumptions they make about older people including their physical and mental capacities, such as being frail or in poor health [8], incompetent, and technologically illiterate or uninterested in developing technological skills [9]. Another study based on the analysis of the historical American English corpus of 400 million words shows that there are increasingly negative stereotypes over the last two centuries associated with the biomedicalisation of ageing and decline [10]. Such over generalisations fail to acknowledge the growing diversity of older adults, particularly in terms of their needs [8]. We find evidence of how this is manifested in the labelling of the datasets, where wide and arbitrary age intervals are assigned to older adults in datasets. For example, we found that in facial image datasets younger age groups may be categorised in arrow age groups like "13–19 years old" and "20–36 years old" whereas older adults are grouped in all-embracing categories like "66+" despite the differences in physical appearances in this group of people spanning four decades. Table 1 indicates how the age groups are labelled.

The exclusion of older adults in the design process of technology is a well-known gap and is discussed in the literature extensively so we will not discuss this established fact in this paper. Rather we want to draw attention to the exclusion of older people in the datasets used to build AI systems. Based on the notion of data disparities defined by Ibrahim as "systematic differences in the quantity and/or quality" [2], the data used in the design of AI systems may exclude older adults to the point of severe under representation [11]. To explore the extent of data disparity, we investigated the presence of data disparities in seven commonly used facial image databases to train facial and age-recognition software: FGNet [12], Academic MORPH & MORPH Longitudinal [13], Labelled Faces in the Wild (LFW) and Labelled Faces in the Wild + (LFW+) [14], Images of Groups (IoG) [15], and Adience [16]. We encountered these databases in our review of the literature about digital ageism. Table 1 presents the data disparity of older adults in each of the datasets with five datasets ranging from 0.001% to 7% and the two datasets with the highest proportion with 17.8% and 21.2%. Commonly, older adults were the least represented age group in these popular datasets. While looking at this table, it is important to bear in mind that in 2019 17.4% of the American population (the nation of origin for all but one of those databases) was over the age of 65 years [17]. Under-representation is a problem in AI training data because AI relies on data to make accurate predictions. A lack of data on a certain demographic reduces the ability of an AI to make predictions about that demographic, compared to better-represented demographics [2] [18]. A lack of data on older adults can spark a feedback loop, producing technology that is not optimised for use by older adults, and in the process discouraging them from using these systems [9]. As a result, less data is gathered from these already under-represented groups for data-driven AI to base future predictions on, disincentivising profit-driven developers from correcting these imbalances [19].

Table 1 Data disparity of the oldest age demographics in seven common facial image datasets

Dataset Name	FGNet	Academic MORPH	MORPH Longitudinal	Labelled Faces in the Wild (LFW)	Labelled Faces in the Wild + (LFW+)	Images of Groups (IoG)	Adience
Total number of images in dataset	1,002	55,314	402,055	13,244	15,710	5,080	26,850 (19,847 for which data was available)
Oldest demographic (labelled)	61+	50+	70+	61+	61+	66+	60+
Number of images of oldest demographic	6	3,933	594	2,806	2,806	1,253	875
Overall percentage of images of oldest demographic	0.40%	7%	0.001%	21.20%	17.80%	4%	4.50%

With respect to data availability incentives, one form of this is the way older adults are often targeted for healthcare-related monitoring technologies at a greater rate than other demographics [19]. These monitoring systems often require wearable components which, to them, can be stigmatising and symbolise their frailty such as the call alarm [20]. The commercialisation of ageism-informed technologies underpins and drives powerful financial and economic incentives to developers and industry to continue focusing on this market. Reflective questions that we can ask ourselves include "What technologies are easiest to market to older people?; What technologies will people buy related to older people?; and What types of research and development receive the most funding?" According to Chu et al.: "if most technologies marketed toward older adults are designed to resolve or manage health problems, then this could easily reinforce the impression that older adults are mainly unhealthy, and in need of support" [3]. As part of ageist cycles of injustice, these factors reinforce digital ageism in the design phase, as well as the harms of allocation and representation of older adults [3] [21].

2.2 Technology

Design and development rooted in ageist assumptions produces technology that can perpetuate ageing as a disease process, seeking to hold off growing old for as long as possible [22]. This in turn ties into assumptions about what digital technologies for older people look like such as being mainly focused on healthcare devices and health management systems, companion or social robots, cognitive support tools, medication reminder devices, etc. In essence, the AI systems and technologies are designed to help or support older people to manage vulnerabilities and address problems, rather than be tools for or sources of joy, achieving one's aspirations, pleasure, or simply for having fun.

Another consequence of technology that fails to include or capture healthy ageing, is that the data that is currently being collected from these technologies will only be representative of older people that confirms ageist stereotypes. These assumptions and unrepresentative data create a feedback loop that reinforces negative stereotypes, and the technologies developed. Essentially, not only is there a lack of data for older people, the data we do have is unrepresentative of the ageing process fails to capture healthy ageing.

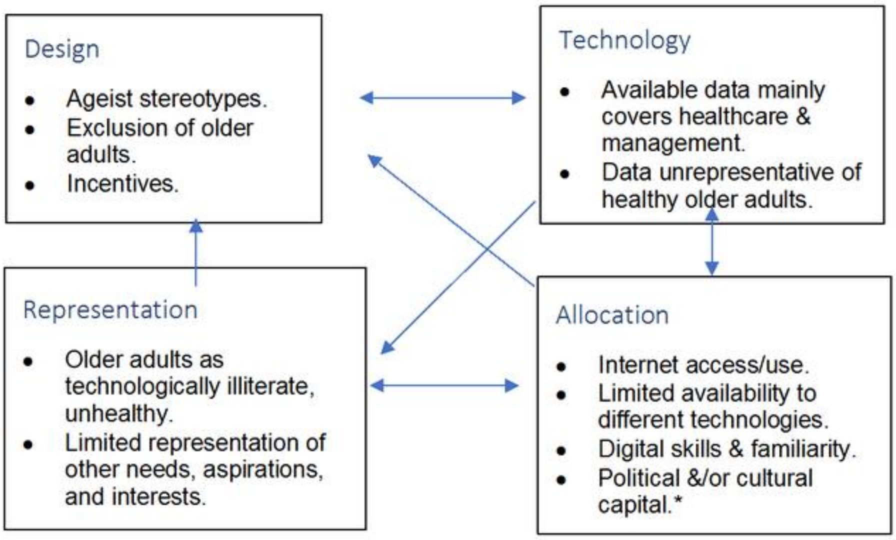

Figure 1: Ageist technology-mediated cycles of injustice that contribute to digital ageism [3]. For this paper, we clarified the allocation component and added political and/or cultural capital*.

2.3 Representation

The relevant social meanings include stereotypical beliefs about older people and their relation to technology, as well as the affective dimensions of ageism. Relevant stereotypes include that older people struggle with digital technology, are resistant to change and slow to adopt innovations, and are simply less interested in new technology. The social and cultural values impact the design and technology for older adults. Older adults may be unable to harness the benefits of these technologies to achieve their goals [Figure 1]. With the ongoing ageist cycles of injustice, older adults will continue to be stereotyped, viewed as homogeneous, and possessing a limit ranged of needs or uses for digital technologies other than health related needs.

Design decisions also directly feed into the technology that is produced. By reinforcing these stereotypes, these technologies are biased and influence the design of future technologies that continue to embed age-related bias into the fabric of our technological society. These biased technologies impact resource allocation, for example older users are not presented with advertisements for job opportunities, or do not have access to valuable resources such as healthcare. The accuracy of diagnostic algorithms, such as those studied by Schrouff et al., drops for older adults which means that older adults lose out on receiving timely medical interventions [18]. Technologies that are not designed with the needs of older adults may be less suitable and satisfying to use, which may deprive older people from accessing information. Within the context of the COVID-19 pandemic, this would have been the difference between life and death. For example, using technology and the Internet became a lifeline during lockdowns, but older adults unable to use technologies potentially miss benefits of digital information technology including accessible health information [23]; opportunities for social connectivity over distances [24]; business and employment [25]; and greater civic participation [26]. This has the potential of both reducing the data available to developers of future technologies, and reinforcing ageist stereotypes [3] [21].

These technologies also impact how older adults are represented in society. For instance, there is evidence of significant age bias in text data as demonstrated by Díaz et al. and sentences with the word "young" were more likely to be scored positively than the same sentences containing the word "old" [27]. The needs of older adults are either not represented or are misrepresented, fueling ageist fears about a globally ageing population and the impending collapse of the social safety net reinforcing the assumption that older people are unable to care for themselves [28]

[4]. Social representations and meanings affect resource allocation and vice versa [29]. Any negative associations with older adults that biased technologies reinforce will in turn impact material resources and how they are distributed.

2.4 Allocation

Allocation refers to the broader distribution of resources within society, which includes material resources, opportunities, access, and prioritisation within/in social institutions (political and/or cultural capital) like laws, regulations, policies, and guidelines [29]. Within the element of allocation, we encounter inequity to resources fortified by a digital divide, referring to a socio-economic gap that exists in our society between those with the capital and education needed to harness the benefits of digital technology and those who do not [30]. Older people are disadvantaged by the digital divide for a number of reasons including that older individuals have less access to the Internet compared to younger people, new technology is rarely developed with their needs in mind, and they may, therefore, be less familiar with the latest technologies [22]. Not having Internet access means being unable to work, access vital health information, access additional resources, and engage in the freedom of expression. It also prevents AI from gaining much-needed training data from this demographic which feeds into the other elements of the technology-mediated ageist cycles of injustice. The impact of the digital divide is demonstrated by the fact that age remains a significant predictor of information technology usage [31]. Additionally, with the COVID-19 pandemic, recent years have seen restrictions put on in-person care, forcing both healthcare providers and patients to depend on digital communication technologies in lieu of face-to-face discussions. Both the developers of new technology and government regulators have largely overlooked the unique and diverse needs of older adults in their design and implementation processes. For example, the vaccination booking website was not user-friendly for older adults in Ontario, Canada, leaving many older adults struggling to get an appointment. Without an honest look at how ageist design practices may produce ageist results, we risk perpetuating inequity.

There is increasing focus on bias in AI within the context of fairness, accountability, and transparency. With respect to political or social capital, we looked beyond research papers tackling specifically these issues, and towards governance and policy papers from public, private, and civil organisations to examine the extent to which ageism or bias against older people were discussed in these documents. We conducted a review of the grey literature following standard grey literature review methods [32] [33] on the topic of digital ageism. Grey literature is an excellent source for recent research across disciplines and sectors [32], and can be used to supplement academic research. This search was done as part of our larger scoping review about the topic. With the help of an information specialist, grey literature was retrieved by searching databases with grey literature (Scopus, Web of Science, and Proquest); grey literature databases (OpenGrey, Canadian Public Documents Collection, U.S. Declassified Documents Online, Cochrane Library, and Trip Pro); and a search of the AI Ethics Guidelines Global Inventory a repository that compiles documents about how AI systems can conduct ethical automated decision making. We also conducted an Advanced Search on Google for any text on websites related to “ageism,” “ageing,” or “older people” and “AI” published in English and published after the year 2010. The grey literature databases provided a holistic approach to our search strategy as each database specialises in specific topics or document types. The search resulted in 2,639 documents from the databases and Google. These documents were scanned, and obviously irrelevant documents were excluded (e.g., “this is the age of AI…”) that did not speak to bias in AI or age-related bias in AI. Of these, 233 full-text documents were downloaded in PDF form and a full-text screening was completed. These documents were scored from 0 to 4 with a rank of one being the least relevant and four being the most relevant (0 = did not speak to age and bias in AI; 1 = mentioned “age” in a list of types of biases; 2 = a sentence of text related to the age-related bias; 3 = 2 or 3 sentences related to age-related bias; 4 = more than 3 sentences about the topic). Of the 23 citations that were relevant, the vast majority of these (n = 17 out of 23) were ranked one, as age was mentioned in a list of types of bias an AI system could have, without labeling this bias as “Ageism” or “Age-bias” directly.

From the search of the databases and Google, only two websites received a rank of four: both were guiding documents for HR from private industry entitled “How AI & mindfulness can tackle age bias in the modern workplace” [34] and “Ageism is hurting the hiring at your tech company more than you realize” [1]. Both of these sources discussed ageism in the workforce related to fair access for employment opportunities and during the hiring process (allocation). The 2016 PayScale study of Silicon Valley’s top 18 tech companies found a median age of 30 or younger [1]. Both discussed the potential of AI to disrupt ageism by identifying filters for older adults in the hiring process, such as salary expectations, years of experience, or the age of an applicant’s credentials [1] [32]. One of the sources described a series of gamified cognitive tasks for potential employment candidates. An AI system would then analyse the results of the tasks in an effort to find the ideal candidate for a given position. One of the touted benefits of this system was its ability to detect desirable traits among players, regardless of race, gender or age [35]. In our search results of the AI Ethics Guidelines Global Inventory that contain 146 documents available in English created by government, private, civil society, and international organisations. Only 34 (23.3%) of these documents mention ageism as a bias for a total of 53 unique mentions with some documents mentioning the word “age” multiple times in the document. Of these, 12 (8.2%) of the examined documents provided slightly more context — often no more than one or two sentences — about bias against older adults [3]. Taken together the grey literature and documents about AI ethics rarely discussed ageism in AI or technology systems. The key white report, which was released after our grey literature review in February 2022, is the WHO’s report entitled “Ageism in artificial intelligence for health” [36].

3 Moving forward and future research

The prevalence of digital ageism and our growing population of older adults would appear to be on a collision course: on the one hand technology-mediated cycles of injustice that perpetuate digital ageism, and on the other you have a growing number of older adults who will have to depend on that technology as it becomes ever more pervasive. It is important to note that the underlying causes are complex and require a thoughtful approach to elucidate and address the contributing factors that are both societal and technical in nature. This paper adds to an emerging body of knowledge focused on digital ageism that is still nascent. In the time being, there are actions on multiple fronts that can counter digital ageism.

While the cycles of injustice that perpetuate digital ageism are unfortunately deeply engrained in our society, we optimistically want to highlight that the interconnected nature of the cycle also means that breaking it at one point may have a reciprocal impact. Firstly, researchers and stakeholders need to bring attention to the notion of digital ageism. They need to collectively challenge this implicit bias through open dialogue and inclusion of older adults throughout the technology design and development pipeline. We recognise digital ageism as a distinctive problem that requires attention and public discourse. This increased attention may translate into policies and protocols for the private sector/industry or the government regulators can become involved in setting and upholding standards for new technologies, including representation standards in databases that AI systems rely on [2] [3] [21], training and algorithm development. For example, the United Kingdom will perform a racial bias review of pulsimeters after it came to light that the medical equipment was not as accurate for patients with darker skin [37]. Funding agencies also need to prioritise exploring this issue to support researchers in identifying the impact of this bias and create financial incentives. Further, more participatory approaches like user-centred design that involves older adults in the design process but can also increase the likelihood that the end result will be technology that is suitable and relevant to them. Doing so is not only supports older adults better, but may be financially profitable. Jointly, these could begin to help target digital ageism at the design and representative elements of the framework

Next, tackling the issues related to the technology element will also require a multifaceted approach. Ibrahim et al. offers several suggestions to tackle data issues affecting AI health systems at a variety of levels [2]. By raising awareness of the potential benefits of AI, more older adults may volunteer their data for these systems, expanding the data available for AI to train on. Effective communication to older adults about the potential benefits of these systems could bolster awareness efforts about AI. Finally, Ibrahim et al. acknowledge that building representative training datasets is the most important part of digital health equity [2]. These datasets should be representative of the populations they will be deployed to. Remedying these issues will help give older adults the representation in the digital domain that they deserve along with every other group. Related to allocation, national policies supporting access to the Internet as a human right will improve access to education, information, employment, and freedom of assembly which may promote autonomy and help older people with aspects unrelated to health. When older adults have access to the Internet and technology that is suited for their needs and goals, the digital divide can begin to close [2] [21]. These suggestions are preliminary ideas that can begin to address the problem of digital ageism. Certainly, new developments in technology capabilities and knowledge about the area will change what can be done and the actions required to respond to digital ageism.

References

<bib id="bib1"><number>[1]</number>Kelly Kinnard. 2018. Ageism Is hurting your tech company's hiring more than you realize. August 29. https://www.entrepreneur.com/article/318816</bib>
<bib id="bib2"><number>[2]</number>Hussein Ibrahim, Xiaoxuan Liu, Nevine Zariffa, Andrew D. Morris, and Alastair K. Denniston. 2021. Health data poverty: An assailable barrier to equitable digital health care. Lancet Digit 3, 4, e260–e265. DOI: https://doi.org/10.1016/S2589-7500(20)30317-4</bib>
<bib id="bib3"><number>[3]</number>Charlene H. Chu, Rune Nyrup, Kathleen Leslie, Jiamin Shi, Andria Bianchi, Alexandra Lyn, Molly McNicholl, Shehroz Khan, Samira Rahimi, Amanda Grenier, and Lawrence S. Bloomberg. 2022. Digital ageism: Challenges and opportunities in artificial intelligence for older adults. Gerontologist gnab167. DOI: https://doi.org/10.1093/GERONT/GNAB167</bib>
<bib id="bib4"><number>[4]</number>World Health Organization (WHO). 2021. Global report on ageism. Geneva. World Health Organization. Licence: CC BY-NC-SA 3.0 IGO. https://www.who.int/teams/social-determinants-of-health/demographic-change-and-healthy-ageing/combatting-ageism/global-report-on-ageism</bib>
<bib id="bib5"><number>[5]</number>Jody A. Wilkinson and Kenneth F. Ferrar. 2022. Thirty years of ageism research. Cambridge, MA. MIT Press.</bib>
<bib id="bib6"><number>[6]</number>Batya Friedman and Helen Nissenbaum. 1996. "Bias in computer systems." ACM Transactions on Information Systems 14.3, 1996, 330-347. https://nissenbaum.tech.cornell.edu/papers/Bias%20in%20Computer%20Systems.pdf</bib>
<bib id="bib7"><number>[7]</number>Kate Crawford. 2017. The trouble with bias. NIPS 2017 Keynote. Kate Crawford #NIPS2017. YouTube. https://www.youtube.com/watch?v=fMym_BKWQzk&ab_channel=TheArtificialIntelligenceChannel</bib>
<bib id="bib8"><number>[8]</number>Alana Officer, Jotheeswaran A. Thiyagarajan, Mira L. Schneiders, Paul Nash, Vânia de la Fuente-Núñez. 2020. Ageism, healthy life expectancy and population ageing: How are they related? Int J Environ Res Public Health 17, 9, 3159. https://doi.org/10.3390/ijerph17093159</bib>
<bib id="bib9"><number>[9]</number>Andrea Rosales and Mireia Fernández-Ardèvol. 2019. Structural ageism in big data approaches. Nord. Rev 40, s1. DOI: https://doi.org/10.2478/nor-2019-0013</bib>
<bib id="bib10"><number>[10]</number>Reuben Ng, Heather G. Allore, Mark Trentalange, Joan K. Monin, Becca R. Levy. 2015. Increasing negativity of age stereotypes across 200 years: Evidence from a database of 400 million words. PLoS One 10, 2, e0117086. https://doi.org/10.1371/journal.pone.0117086</bib>
<bib id="bib11"><number>[11]</number>Andrea Rosales and Mireia Fernández-Ardèvol. 2020. Ageism in the era of digital platforms. Convergence 26, 5–6. DOI: https://doi.org/10.1177/1354856520930905</bib>
<bib id="bib12"><number>[12]</number>Gabriel Panis, Andreas Lanitis, Nicholas Tsapatsoulis, and Timothy F. Cootes. 2016. Overview of research on facial ageing using the FG-NET ageing database. IET Biom 5, 37–46. DOI: https://doi.org/10.1049/iet-bmt.2014.0053</bib>

<bib id="bib13"><number>[13]</number>Morph Facial Recognition Database: Office of Innovation and Commercialization. University of North Carolina Wilmington. https://uncw.edu/oic/tech/morph.html</bib>

<bib id="bib14"><number>[14]</number>Hu Han and Anil K. Jain. 2014. Age, gender and race estimation from unconstrained gace images. MSU technical report (2014). MSU-CSE-14-5 (2014). http://biometrics.cse.msu.edu/Publications/Face/HanJain_UnconstrainedAgeGenderRaceEstimation_MSUTechReport2014.pdf</bib>

<bib id="bib15"><number>[15]</number>Andrew C. Gallagher and Tsuhan Chen. 2009. Understanding images of groups of people. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, June 20–25. IEEE 256-263. DOI: https://doi.org/10.1109/CVPR.2009.5206828</bib>

<bib id="bib16"><number>[16]</number>Eran Eidinger, Roee Enbar, and Tal Hassner. 2014. Age and gender estimation of unfiltered faces. IEEE Trans. Inf. Forensics and Secur 9, 12, 2170–2179. DOI: https://doi.org/10.1109/TIFS.2014.2359646</bib>

<bib id="bib17"><number>[17]</number>Age and Sex Tables. United States Census Bureau. 2019. https://www.census.gov/topics/population/age-and-sex/data/tables.2019.List_897222059.html</bib>

<bib id="bib18"><number>[18]</number>Jessica Schrouff, Natalie Harris, Oluwasanmi Koyejo, Ibrahim Alabdulmohsin, Eva Schnider, Krista Opsahl-Ong, Alex Brown, Subhrajit Roy, Diana Mincu, Christina Chen, Awa Dieng, Yuan Liu, Vivek Natarajan, Alan Karthikesalingam, Katherine Heller, Silvia Chiappa, and Alexander D'Amour. 2022. Maintaining fairness across distribution shift: Do we have viable solutions for real-world applications? February. https://doi.org/10.48550/arXiv.2202.01034</bib>

<bib id="bib19"><number>[19]</number>Ittay Mannheim, Ella Schwartz, Wanyu Xi, Sandra C. Buttigieg, Mary McDonnell-Naughton, Eveline J. M. Wouters, and Yvonne van Zaalen. 2019. Inclusion of older adults in the research and design of digital technology. Int. J. Environ. Res. P 16, 19, 3718. DOI: https://doi.org/10.3390/ijerph16193718</bib>

<bib id="bib20"><number>[20]</number>Sandra A. Zwijsen, Alistair R. Niemeijer, and Cees M. P. M. Hertogh. 2011. Ethics of using assistive technology in the care for community-dwelling elderly people: An overview of the literature. Aging Ment. Health 15, 4, 419–427. DOI: https://doi.org/10.1080/13607863.2010.543662</bib>

<bib id="bib21"><number>[21]</number>Charlene Chu, Kathleen Leslie, Rune Nyrup, and Shehroz Khan. 2022. Artificial intelligence can discriminate on the basis of race and gender, and also age. Conversation January 18. https://theconversation.com/artificial-intelligence-can-discriminate-on-the-basis-of-race-and-gender-and-also-age-173617</bib>

<bib id="bib22"><number>[22]</number>Alexander Peine and Louis Neven. 2019. From intervention to co-constitution: New directions in theorizing about aging and technology. Gerontologist 59, 1, January. DOI: https://doi.org/10.1093/geront/gny050</bib>

<bib id="bib23"><number>[23]</number>Amanda Hunsaker and Eszter Hargittai. 2018. A review of Internet use among older adults. New Media Soc 20, 10, October, 3937–3954. DOI: https://doi.org/10.1177/1461444818787348</bib>

<bib id="bib24"><number>[24]</number>Daniela G. Praničević, Judita Peterlin, and Mariette J. Bućan. 2017. Do older people benefit from digital services? In DIEM: Dubrovnik International Economic Meeting, 3, 1, 145–160. https://hrcak.srce.hr/187373</bib>

<bib id="bib25"><number>[25]</number>Daniel Nunan and Maria Laura di Domenico. 2019. Older consumers, digital marketing, and public policy: A review and research agenda. J. Public Policy Mark. 38, 4, October, 469–483. DOI: https://doi.org/10.1177/0743915619858939</bib>

<bib id="bib26"><number>[26]</number>Ruth Abbey and Sarah Hyde. 2009. No country for older people? Age and the digital divide. J Inf. Commun. Eth Soc 7, 4, 225–242. DOI: https://doi.org/10.1108/14779960911004480</bib>

<bib id="bib27"><number>[27]</number>Mark Diaz, Isaac Johnson, Amanda Lazar, Anne Marie Piper, and Darren Gergle. 2018. Addressing age-related bias in sentiment analysis. In Proceedings of the 2018 CHI conference on human factors in computing systems. April 2018, Paper No. 412, 1–14.</bib>

<bib id="bib28"><number>[28]</number>Louis Neven. 2010. 'But obviously not for me': Robots, laboratories and the defiant identity of elder test users. Sociol. Health Illn 32, 2, 335–347. DOI: https://doi.org/10.1111/j.1467-9566.2009.01218.x</bib>

<bib id="bib29"><number>[29]</number>Jess Whittlestone, Rune Nyrup, Anna Alexandrova, Kanta Dihal, and Stephen Cave. 2019. Ethical and societal implications of algorithms, data, and artificial intelligence: A roadmap for research. London. Nuffield Foundation. http://www. nuffieldfoundation.org/sites/default/files/files/Ethical-and-Societal-Implications-of-Data-and-AI-report-Nuffield-Foundat.pdf</bib>

<bib id="bib30"><number>[30]</number>Jyoti Choudrie, Sutee Pheeraphuttranghkoon, and Soheil Davari. 2020. The digital divide and older adult population adoption, use and diffusion of mobile phones: A quantitative study. Inf. Syst. Front. 22, 3, June. DOI: https://doi.org/10.1007/s10796-018-9875-2</bib>

<bib id="bib31"><number>[31]</number>Alicia Y. Hong and Jinmyoung Cho. 2016. Has the digital health divide widened? Trends of health-related Internet use among older adults from 2003 to 2011. J. Gerontol. Series B: Psychol. Sci. Soc. Sci 72, 5, 856-863. DOI: https://doi.org/10.1093/geronb/gbw100</bib>

<bib id="bib32"><number>[32]</number>Karen M. Benzies, Shahirose Premji, K. A. Hayden, Karen Serrett. State-of-the-evidence reviews: advantages and challenges of including grey literature. Worldviews Evid Based Nurs 3, 2, 55–61. DOI: 10.1111/j.1741-6787.2006.00051.x</bib>

<bib id="bib33"><number>[33]</number>Quenby Mahood, Dwayne V. Eerd, Emma Irvin. 2014. Searching for grey literature for systematic reviews: Challenges and benefits. Res Synthesis Methods. 5, 3, 221–s34.</bib>

<bib id="bib34"><number>[34]</number>Andres Blank. 2021. How AI & mindfulness can tackle age bias in the modern workplace. Toolbox Hr. toolbox.com. December 16. https://www.toolbox.com/hr/hr-strategy/guest-article/how-ai-mindfulness-can-tackle-age-bias-in-the-modern-workplace/</bib>

<bib id="bib35"><number>[35]</number>Josh Constine. 2017. Pymetrics attacks discrimination in hiring with AI and recruiting games. TechCrunch. September 20, 2017. https://techcrunch.com/2017/09/20/unbiased-iring/?ncid=rss&utm_source=feedburner&utm_medium=feed&utm_campaign=Feed%3A+Techcrunch+%28TechCrunch%29</bib>

<bib id="bib36"><number>[36]</number>World Health Organization. 2022. Ageism in artificial intelligence for health: WHO policy brief. Geneva. World Health Organization. Licence: CC BY-NC-SA 3.0 IGO. https://apps.who.int/iris/rest/bitstreams/1408281/retrieve</bib>

<bib id="bib37"><number>[37]</number>BBC. 2021. Sajid Javid orders racial bias review after Covid deaths. November 21 2021. https://www.bbc.com/news/uk-59363544</bib>