

Ageism in AI: new forms of age discrimination in the era of algorithms and artificial intelligence

Justyna Stypińska

{justyna.stypinska@fu-berlin.de}

Free University Berlin, Germany

Abstract. Scholars in fairness and ethics in AI have successfully and critically identified discriminatory outcomes pertaining to the social categories of gender and race. The salient scrutiny of fairness, important for the debate of AI for social good, has nonetheless paid insufficient attention to the critical category of age. The aging population has been largely neglected during the turn to digitality and AI. Ageism in AI can be manifested in five interconnected forms: (1) age biases in algorithms and datasets, (2) age stereotypes, prejudices and ideologies of actors in AI, (3) invisibility of old age in discourses on AI, (4) discriminatory effects of use of AI technology on different age groups, (5) exclusion as users of AI technology, services and products. Furthermore, the paper provides illustrations of these forms of ageism in AI.

Keywords: ageism, population ageing, artificial intelligence, digital inequalities, age discrimination

1. Introduction

Over the last few years, the power of algorithms and artificial intelligence (AI) has become an issue present and discussed in society, media, business, and the social sciences. Artificial intelligence and Machine Learning (ML) systems are frequently pictured as capable of making smarter, faster, better, and presumably more neutral decisions. European Union, along with governments in Europe and around the world see the AI as the biggest promise of the 21st century, making positive contributions to multiple aspects of human life from improving healthcare, to climate change mitigation [1]. At the same time, AI systems entail a number of hazards, such as opaque decision-making, gender-based or other kinds of discrimination, interference in private lives or being used for illegal purposes. The scholars of ethical and fair AI express concern about the way those technologies show hidden biases resulting in exclusion

and discrimination of members of marginalized groups [2], [3] and can pose threat to fundamental human rights and social justice [4], [5].

Particularly, the way sexism and racism operate in AI attracted significant attention of prominent scholars. Studies have shown how face recognition systems work poorly for women with dark skin [6], that word embeddings - a framework used for text analysis in machine learning and neural language processes - exhibit female/male gender stereotypes to a disturbing extent [7]. A seminal study on search engines showed how algorithms systematically retrieved racist and sexist search results, from keywords such as 'black girl', which was termed "algorithmic oppression" [8]. Yet, the category of age has been largely neglected in existing research. The aim of this article is to rectify this gap by pointing to several theoretical aspects and empirical illustrations of manifestations of ageism in AI.

2. Ageing population, ageism, and technology

The need to examine ageism in AI stems not only from the technological acceleration, but primarily from the reality of demographic change, which defines the way societies are ageing around the globe. In Europe, for example, the current median age amounts to 43.9 years and is projected to increase to 48.2 years until 2050. WHO estimates that by the end 2030 the number of people 60 years and older will grow by 56 per cent, from 962 million (2017) to 1.4 billion (2030). This rapid increase created momentum for UN to proclaim the Decade of Healthy Ageing (2021-2030), where combating ageism is seen as a fundamental strategy for securing healthy and dignified future for older adults. Simultaneously the World Health Organization launched a global campaign to fight ageism and issued its very first Global Report on Ageism [9]. Despite decades of research, implementation of anti-discrimination policies and legislation, ageism is still alive and well [10] or even, as recent studies show, has dramatically intensified during the Covid-19 pandemic [11].

The definition of ageism has shifted from its basic understanding as "a process of systematic stereotyping and discrimination against people because they are old" [12] to elaborated and multidimensional conceptualizations of ageism. Theories of ageism are following the digitalization of the phenomenon, that is the presence of age biases, stereotypes, prejudice, and discrimination in their digital form. The concept of "visual ageism" for example, responds to this change [13]. Yet, the existing conceptualizations are still too narrow to address the complexity of various ageism manifestations in AI and algorithmic systems.

Social scientists have studied extensively the way older adults use and interact with digital technology [13]–[17], the way technological advances can create new inequalities in form of digital divide [18], to the way gerontechnology can assist older adults in adapting to ageing processes [19]. Newly emerged theoretical framework of "Socio-gerontechnology" promises to provide a unique understanding of ageing and technology from a social sciences and humanities perspective and contribute to the development of new ontologies, methodologies and theories. However, the current ageing research with regard to intelligent technologies has been limited to several themes: the use of social robots, and other smart technologies to assist and support active ageing and ageing in place [20], use of digital data in smart mobility [21] and the use of smart technologies (e.g. wearable devices) in tracking and quantifying ageing bodies [17]. Yet, the theoretical reflection and empirical analysis on the potential negative impact of the intelligent

technologies and algorithmic systems on older persons has until now not been problematized in sociology of ageing and social gerontology in a systematic way.

3. Going beyond „bias“

In their seminal work, Friedman & Nissenbaum (1996) identified three types of bias in computer systems: pre-existing from social institutions, technically created, and emerging bias from the context of use. Furthermore, the domination of AI ethics and fairness research by computer and data scientists is reflected in the use of language. The concept of “bias”, originating from computer sciences dominates the scholarly discourses about exclusion in AI [22] and poses risk of reducing the complexity of social inequalities to solely technical level. Additional risks are presented by the substantial power asymmetries between those with the resources to design and deploy AI systems, and those who are classified, ranked, and assessed by these systems [23].

Therefore, we need to go beyond the concept of bias and create new concepts and levels of analysis to describe digital inequalities. This paper proposes thus to look at ageism in AI in five forms:

- (1) age biases incorporated in algorithms and digital datasets
- (2) age stereotypes, prejudices, and ideologies of actors in the field of AI
- (3) invisibility of category of age and old age in discourses around AI
- (4) discriminatory effects of use of AI technology on different age groups
- (5) exclusion as users of AI technology, services, and products

4. Ageism in AI applications - illustrations

The identified dimensions of ageism in AI are separated conceptually, however there are significant interrelations and synergies between them. The illustrations of the five forms provided in this section are based on scientific literature review, grey literature and own preliminary research and observations. Hence, they are not exhaustive in the scope of the problem or the depth of investigation. Their fragmentary character reflects the unsystematic and incomplete character of research on ageism in AI and points to the need for further research in this area.

4.1 Age bias in algorithms and digital datasets

Single studies indicate that age bias can occur in machine learning models or big data approaches. Diaz and colleagues (2019) analysed the treatment of age-related terms across 15 sentiment analysis models and 10 widely-used GloVe word embeddings. The study showed evidence that sentiment analysis, disclosed significant age biases – sentences with ‘young’ adjectives were 66% more likely to be scored positively than identical sentences with ‘old’ adjectives [24]. Moreover, it proved that various sentiment analysis methods impact bias in outcomes – particularly that tools validated against social media data exhibit increased bias.

Another study found relevant differences in the outcomes of face recognition models for predicting age and gender from photographs [25]. The researchers used Convolutional Neural Net (CNN) which is an advanced deep learning technique to analyse visual imagery. The model was trained on photos of celebrities from IMDb and Wikipedia, where their picture was matched with their age, as well as data for general public from UTKFace data set. The results showed that age estimation was generally performing poor on older age groups (60+), an effect which was compounded by gender and race - the age estimation worked disappointingly on older women of colour.

The age estimation algorithms working on visual data require large datasets for training. Since an algorithm is only as good as the data it works with [26], this is where ageism in digital datasets becomes apparent. Study of ageism in big data approaches confirms that the most ageist practices in intelligent systems design are related to the data set limitations in terms of the representativeness of the studied population and particularly to recruitment procedures that tend to exclude older people [27]. Most of the datasets include radical age cut offs for their data. For example, The Face and Gesture Recognition Research Network (FG-NET) aging database contains on average 12 pictures for each of its 82 subjects in varying ages between 0 and 69. Other datasets limit the ages at 70 (Tufts-Face-Database), at 77 (MORPH) and the list continues. There are some outstanding exceptions to the rule, such as UTKFace dataset, where the photos depict adults up to age of 116 years old. The cut-offs are also visible in training data used for proprietary algorithms developed in expanding sector of AI industry¹.

4.2 Stereotypes and prejudices in tech industry

Tech culture is homogenous in terms of age, ethnicity, and gender. It is young, predominantly populated by men of Caucasian or Asian origin, which is associated with the structural discrimination embedded in digital technologies [28]. Ageism in IT sector and tech industry is a well-known fact [29]–[31]. These companies show rampant signs of various types of systematic biases and prejudice [29], [32]–[34], ageism being one of them. Large tech companies have phased out older workers over the past few years and they continue to discriminate against anyone old enough to remember the '80s. In 2007, a then 22 year-old, Mark Zuckerberg famously admitted that tech companies should not hire people over 30 because “Young people are just smarter”. A survey among American tech workers shows that 76 percent of respondents say ageism exists in tech globally, whereas 80 percent of those in their late 40’s say they are concerned their age (and ageism attitudes) will affect their careers [35].

Ageism in tech sector is specific as it is targeted at persons at much younger age than in other sectors of economy, where ageism starts to be experienced by someone who reached at least 45 years of age [36]. In fact, online survey among tech workers shows that one fourth of respondents in their early 30’s already regard age as a barrier in getting a new job [35]. Since

¹ For instance, the company YOTI configured its training data set with age brackets of 13-60 and the highest estimated errors in performance of their algorithm for the age group of 50-60. In its White Paper the company admits that : “it seems reasonable to hypothesise that any error will tend to be higher for older people than younger people, because older people will have been exposed to various unpredictable environmental factors for longer” [50].

AI research and development is at the forefront of technological innovation globally, it is plausible to assume these ideologies impact the way AI developers, software engineers, data scientists and other AI practitioners work, think and solve problems. In fact, studies on ageism in digital platforms indicate that the homophily of the community of software developers, who are predominantly young men of high socio-economic status is responsible for prejudices and biases which are coded into algorithms [37].

4.3 Ageism in discourses

Studying discourses about algorithmic systems and processes is crucial to our understanding of the social power of algorithms and AI [38], [39]. Next to the material power, algorithms can exert a discursive power, which can reveal their political entanglements. Authors' preliminary observations of the field of AI ethics and bias in AI and analysis of sample of documents (white papers on inclusivity, brochures from commercial and public organizations) suggest that the mainstream debates around the issues of AI fairness tend to omit the category of age and older persons. The majority of initiatives and/or documents for inclusion and promotion of diversity in AI community are targeted at gender and racial minorities. Following the consciousness on intersectional character of biases, some other inclusion criteria are mentioned in one document, such as class and disability [23]. Age as a category for algorithmic discrimination is not mentioned explicitly in most of the documents. Alternatively, age was mentioned as a category mediated through disability. The absence can be observed in two forms: invisibility of old age as object of discussion and lack of representation of older persons as subjects in those discourses. This might be due to several factors, such as a relatively weak social representation of the rights of older persons in the area of AI, or ideologies and stereotypical beliefs about older persons as users or non-users of AI applications held by software producers. Research suggests that older adults are generally portrayed as frail when described as users of AI assistive technology [40].

4.4 Automatic Decision Making (ADM) systems and their outcomes

This section highlights some of the areas where deployment of AI can result in harm for older persons as a distinct demographic group. Two terms - "algorithmic discrimination" and "discrimination of algorithms" describe the negative outcomes of automated decision-making (ADM) systems or classification systems used in multiple AI applications. Discrimination in context of AI occurs when the outcomes infringe on the rights of persons based on their "protected" characteristic, such as gender, race, age, disability, or nationality. Although not all ADM systems are powered by AI, there is a stable and increasing tendency towards more deployment of AI in those solutions [41].

With regard to ageing population, the risk of discrimination lies in the way the biased algorithms are being used in practice in those realms of social, economic and cultural life where they could infringe on the rights of older persons. Their increased use in the recruitment and hiring practices, despite some drawbacks, can threaten the way age discrimination cases will be possible to detect. Age discrimination in employment is one of the most wide spread types of discrimination in the labour market [10] and the attempts to fight it with anti-discrimination

legislation is challenging. The already famous case of Amazon's hiring algorithm downgrading the resumes of women gives a hint of what it could mean for older workers. For example, if a company had a tendency to hire candidates who graduated from school (or landed their first job) by a certain date, it might introduce a bias toward younger candidates. The company's software developers would need to actively monitor the system to ensure that something like that was not happening. Another example where ADM could negatively impact large groups of older adults, could be in the banking sector. If a mortgage lending model finds that older individuals have a higher likelihood of defaulting it might reduce the lending options based on age leading to excluding older adults from those services [42].

4.5 Marginalization of older persons as users

The last form of age discrimination in AI discussed shortly is the exclusion as users. The compounded effects of ever increasing complexity of the digital technology and the low algorithmic awareness among older adults [43] creates structures which marginalize or exclude older persons as end users of AI technology. Ageism in technology design is not a new phenomenon. Studies showed different patterns in use of digital technology by older adults [44]–[46]. However, the question arises if the data-driven technologies using AI and ML technologies pose any additional risks of harm and ageism?

An example of AI driven technology where older adults might experience ageism as users is the group of products called “conversational AI” which includes virtual assistants and chatbots. Conversational AI agents are increasingly used by companies for customer services and by consumers as personal assistants. Virtual assistants are trained on massive datasets and deploy various algorithms. The challenge is that unless corrected for, the virtual assistant also learns and replicates human biases in the dataset [47]. Recent studies showed that the algorithms used by the conversational AI agents do not always recognise user's gender or understand context-bound and culture-bound language. As a result, the conversations may be biased. Similar problems might occur when testing virtual assistants and chatbots for their sensitivity to issues of age and ageism. Claims are being made that chatbots and virtual assistant are already ageist and sexist in the way they are profiled (usually as young women), but the question is whether they could also be ageist in their conduct (e.g. treat older customers differently than younger, answer questions differently or exhibit ageist stereotypes – jokes, etc). Furthermore, if assistant implemented in a health care application perform more poorly with seniors it could impact the quality of care provision and ultimately the health of the user.

Further examples of areas where older adults might experience ageism as users of AI-driven technology include diverse smartphone applications where incomplete data for older age cohorts might result in inaccurate results [37]. Further research is needed to reveal other aspects of this form of AI ageism.

5. Conclusions

The aim of this paper was to turn the scholarly attention in the direction of ageing population as one of the socio-demographic groups, which could be defined as vulnerable in relation to the

date-driven social transformations happening as the result of increasing use of AI technology in all realms of modern life. By showcasing different forms ageism can take in AI applications this article contributes to the scholarly efforts to advance our knowledge of the harmful ways AI can impact the vulnerable group of older adults. Furthermore, I argue that it essential to go beyond the understanding of inequalities in AI dictated by the narrow use of the term “bias”. As social scientists, we are aware of the structural, institutional, and otherwise “non-quantifiable” forms of injustice and oppression in social world [22], [48].

The current debates on ethical AI happening globally at all levels of stakeholders, from public entities, through large and small companies, AI practitioners and scholars from various disciplines show the urgency of the need to regulate AI with regard to its ethical standing. The landscape of currently drafted regulations, recommendations and guidelines for ethical AI is voluminous and diverse [49], but is still largely missing the representation of the interests of ageing population. Moreover, a large portion of the older population is unfamiliar with complexity of AI, algorithms, big data [43] and also do not want to engage with this new technology due to a lack of trust in these developments. To enable informed decisions on their part, communicative efforts must be made to explain various aspects of AI in formats that older age groups can respond to.

References

- [1] European Commission, “On Artificial Intelligence - A European approach to excellence and trust. White Paper,” Brussels, 2020.
- [2] V. Eubanks, *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. St. Martin’s Press, 2018.
- [3] B. D. Mittelstadt, P. Allo, M. Taddeo, S. Wachter, and L. Floridi, “The ethics of algorithms: Mapping the debate,” *Big Data Soc*, vol. 3, no. 2, pp. 1–21, 2016.
- [4] E. Aizenberg and J. van den Hoven, “Designing for human rights in AI,” *Big Data Soc*, vol. 7, no. 2, 2020.
- [5] T. M. Cruz, “Perils of data-driven equity: Safety-net care and big data’s elusive grasp on health inequality,” *Big Data Soc*, vol. 7, no. 1, 2020.
- [6] J. Buolamwini and T. Gebru, “Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification,” *Proc Mach Learn Res*, vol. 81, pp. 1–15, 2018.
- [7] T. Bolukbasi, K.-W. W. Chang, J. Zou, V. Saligrama, and A. Kalai, “Man is to computer programmer as woman is to homemaker? Debiasing word embeddings,” *Adv Neural Inf Process Syst*, pp. 285–287, 2016.
- [8] S. U. Noble, *Algorithms of oppression: How Search Engines Reinforce Racism*. New York: NYU Press, 2018.
- [9] WHO, “Global report on ageism,” Geneva, 2021.
- [10] J. Stypinska and K. Turek, “Hard and soft age discrimination: the dual nature of workplace discrimination,” *Eur J Ageing*, vol. 14, no. 1, pp. 49–61, 2017.
- [11] L. et al. Ayalon, “Aging in Times of the COVID-19 Pandemic: Avoiding Ageism and Fostering Intergenerational Solidarity,” *Journals Gerontol - Ser B Psychol Sci Soc Sci*, 2020.

- [12] R. Butler, *Why Survive? Being old in America*. New York: Harper & Row, 1975.
- [13] L. Ivan, E. Loos, and G. Tudorie, "Mitigating Visual Ageism in Digital Media: Designing for Dynamic Diversity to Enhance Communication Rights for Senior Citizens," *Societies*, vol. 10, no. 4, p. 76, 2020.
- [14] E. Loos, M. Sourbati, and F. Behrendt, "The role of mobility digital ecosystems for age-friendly urban public transport: A narrative literature review," *Int J Environ Res Public Health*, vol. 17, no. 20, pp. 1–16, 2020.
- [15] S. J. Czaja, "The Potential Role of Technology in Supporting Older Adults," *Public Policy Aging Rep*, vol. 27, no. 2, pp. 44–48, 2017.
- [16] A. Wanka and V. Gallistl, "Doing Age in a Digitized World—A Material Praxeology of Aging With Technology," *Front Sociol*, vol. 3, no. April, pp. 1–16, 2018.
- [17] S. Katz and B. L. Marshall, "Tracked and fit: FitBits, brain games, and the quantified aging body," *J Aging Stud*, vol. 45, pp. 63–68, 2018.
- [18] E. Y. Choi, Y. Kim, E. Chipalo, H. Y. Lee, and S. Meeks, "Does Perceived Ageism Widen the Digital Divide? And Does It Vary by Gender?," *Gerontologist*, vol. 60, no. 7, pp. 1213–1223, 2020.
- [19] A. Klimczuk, "Supporting the Development of Gerontechnology as Part of Silver Economy Building," *Ad Alta J Interdiscip Res*, vol. 2, no. 2, pp. 52–56, 2012.
- [20] I. Pedersen, S. Reid, and K. Aspevig, "Developing social robots for aging populations: A literature review of recent academic sources," *Sociol Compass*, vol. 12, no. 6, pp. 1–10, 2018.
- [21] M. Sourbati and F. Behrendt, "Smart mobility, age and data justice," *New Media Soc*, no. February, 2020.
- [22] M. Zajko, "Conservative AI and social inequality: conceptualizing alternatives to bias through social theory," *AI Soc*, no. 0123456789, 2021.
- [23] M. Whittaker, M. Alper, and C. L. Bennett, *Disability, Bias, and AI*. New York, 2019.
- [24] M. Díaz, I. Johnson, A. Lazar, A. M. Piper, and D. Gergle, "Addressing age-related bias in sentiment analysis," *IJCAI Int Jt Conf Artif Intell*, vol. 2019-Augus, no. August, pp. 6146–6150, 2019.
- [25] R. Meade, A. Camilleri, R. Geoghegan, S. Osorio, and Q. Zou, "Bias in Machine Learning: How Facial Recognition Models Show Signs of Racism, Sexism and Ageism," p. 6, 2021.
- [26] S. Barocas and A. D. Selbst, "Big Data's Disparate Impact," *SSRN Electron J*, vol. 671, pp. 671–732, 2018.
- [27] A. Rosales and M. Fernández-Ardèvol, "Structural ageism in big data approaches," *Nord Rev*, vol. 40, no. s1, pp. 51–64, 2019.
- [28] S. Wachter-Boettcher, *Technically wrong: Sexist apps, biased algorithms, and other threats of toxic tech*. New York: W. W. Norton & Company, 2017.
- [29] K. Cook, *Psychology of Silicon Valley*. Cham, Switzerland: Palgrave Macmillan, 2020.
- [30] A. Rosales and J. Svensson, "Perceptions of age in contemporary tech," *Nord Rev*, vol. 42, no. 1, pp. 79–91, 2021.
- [31] V. W. Marshall, "A life course perspective on information technology work," *J Appl Gerontol*, vol. 30, no. 2, pp. 185–198, 2011.
- [32] L. S. H. Park and D. N. Pellow, "Racial formation, environmental racism, and the emergence of Silicon Valley," *Ethnicities*, vol. 4, no. 3, pp. 403–424, 2004.
- [33] J. Shih, "Circumventing discrimination: Gender and ethnic strategies in silicon valley," *Gend Soc*, vol. 20, no. 2, pp. 177–206, 2006.
- [34] A. T. Wynn, "Pathways toward Change: Ideologies and Gender Equality in a Silicon Valley

- Technology Company,” *Gend Soc*, vol. 34, no. 1, pp. 106–130, 2020.
- [35] Dice, “Dice Diversity and Inclusion Report,” 2018.
- [36] K. Harris, S. Krygsman, J. Waschenko, and D. Laliberte Rudman, “Ageism and the Older Worker: A Scoping Review,” *Gerontologist*, vol. 58, no. 2, pp. e1–e14, 2018.
- [37] A. Rosales and M. Fernández-Ardèvol, “Ageism in the era of digital platforms,” *Convergence*, vol. 26, no. 5–6, pp. 1074–1087, 2020.
- [38] S. Natale, “If software is narrative: Joseph Weizenbaum, artificial intelligence and the biographies of ELIZA,” *New Media Soc*, vol. 21, no. 3, pp. 712–728, 2019.
- [39] The Royal Society, “Portrayals and perceptions of AI and why they matter,” London, 2018.
- [40] D. Burema, “A critical analysis of the representations of older adults in the field of human–robot interaction,” *AI Soc*, no. 0123456789, 2021.
- [41] F. Chiusi, S. Fischer, and M. Spielkamp, “Automated Decision-Making Systems in the COVID-19 Pandemic : A European Perspective,” Berlin, 2020.
- [42] J. Silberg and J. Manyika, “Tackling bias in artificial intelligence (and in humans) McKinsey,” *Notes from AI Front Tackling bias AI (and humans)*, pp. 1–8, 2019.
- [43] A. B. Gran, P. Booth, and T. Bucher, “To be or not to be algorithm aware: a question of a new digital divide?,” *Inf Commun Soc*, 2020.
- [44] L. Ivan and F. Mocanu, *Older people and technology use: The importance of using video stimuli in group discussions*, vol. 12207 LNCS. Springer International Publishing, 2020.
- [45] V. Gallistl, R. Rohner, A. Seifert, and A. Wanka, “Configuring the older non-user: Between research, policy and practice of digital exclusion,” *Soc Incl*, vol. 8, no. 2, pp. 233–243, 2020.
- [46] B. Barbosa Neves and F. Vetere, *Ageing and Emerging Digital Technologies. Designing and Evaluating Emerging Technologies for Older Adults*. Springer, Singapore, 2019.
- [47] A. Schlesinger, K. P. O’Hara, and A. S. Taylor, “Let’s talk about race: Identity, chatbots, and AI,” *Conf Hum Factors Comput Syst - Proc*, vol. 2018-April, pp. 1–14, 2018.
- [48] S. Wachter, B. Mittelstadt, and C. Russell, “Why fairness cannot be automated: Bridging the gap between eu nondiscrimination law and ai,” *arXiv*, pp. 1–72, 2020.
- [49] T. Hagedorff, “The Ethics of AI Ethics: An Evaluation of Guidelines,” *Minds Mach*, vol. 30, no. 1, pp. 99–120, 2020.
- [50] YOTI, “White Paper Yoti Age Scan - Public version,” no. October, 2020.