Observatorio (OBS*)

Received: 2025-12-01 | **Accepted:** 2026-02-17 | **Published:** 2026-05-22 | **DOI**: 10.15847/OBS20262896

Age does matter: Midjourney's visual representation of older adults and journalistic ethics

✉ **Juan Linares-Lanzman**
Interdisciplinary Research Centre on Social and Cultural Transformations
Open University of Catalonia
Barcelona, Spain

David Carbonell Mateo
Interdisciplinary Research Centre on Social and Cultural Transformations,
Open University of Catalonia
Barcelona, Spain

Andrea Rosales
Interdisciplinary Research Centre on Social and Cultural Transformations
Open University of Catalonia
Barcelona, Spain

Inés Montoya Espinagosa
Interdisciplinary Research Centre on Social and Cultural Transformations
Open University of Catalonia
Barcelona, Spain

Abstract

With the emergence of Artificial Intelligence (AI), the journalistic and informational media system is undergoing a profound transformation that significantly challenges the production and distribution of information. AI, particularly Generative AI (GenAI), poses risks and potential harms to society, reinforcing stereotypes based on gender, race, origin, and age. This article explores how ageism manifests in GenAI as an example of how AI reinforces stereotypes and its implications for media professionals. We focus on the visual analysis of images generated with Midjourney, one of the most popular AI image-generative tools. We defined 91 prompts describing everyday activities and combined them with two profiles ("person" and "older person") to produce 1,456 images, along with their automatically generated descriptions. These images were then quantitatively analysed to examine demographic characteristics and representations of age. Our results show that most images representing an age-neutral person depict a young individual, while older adults are scarce and relegated to secondary roles and negative depictions. Older individuals are portrayed with signals of vulnerability, fragility, and dependence, while gender and ethnic stereotypes become more extreme when associated with age. GenAI ageist depictions represent a step back in the more diverse and inclusive vision that the legacy media has adopted over recent decades regarding age. GenAI offers media professionals opportunities to enhance productivity, but addressing its biases is essential in news production; being critical of the images generated and disseminated to ensure fairness, inclusivity, and diversity has become an ethical imperative in journalism today, more than ever.

Keywords: generative AI, journalism, visual ageism, age stereotypes, computer vision

1. Introduction

The journalistic and informational media system is undergoing a profound transformation. Artificial Intelligence (AI)—and particularly Generative AI (GenAI)—is becoming part of daily newsroom routines and reshaping news productivity (Lewis et al., 2025). While these tools speed up news production, they might represent an operational adjustment rather than a fundamental rethinking of media priorities (Simon, 2024). As the Reuters Institute reports, 74% of 135 senior news industry leaders from 40 countries expect that "AI will help newsrooms increase productivity and improve workflows, without changing the essence of journalism" (Cherubini & Sharma, 2023). Despite GenAI tools being a news co-producer (Owsley & Greenwood, 2024), core journalistic ethics remain central.

In this research, we focus on the potential impact of visual GenAI on news production. In particular, we examine ageism (Butler, 1969); one of the most common forms of discrimination that nevertheless often goes unnoticed. Far from being eradicated, the World Health Organization currently estimates that half of the world's population is ageist (World Health Organization, n.d.), and in the United States, 90 per cent of older people report having experienced age discrimination at some point in their lives (Barber et al., 2024). This form of discrimination is intersectional (Nash et al., 2020), as it cannot be separated from sexism (Krekula et al., 2018), ableism (Gendron et al., 2024), racism (Steward et al., 2023), and other forms of social injustice. Age stereotypes and ageism shape interpersonal relations and are structurally embedded in institutions, digital systems (Rosales et al., 2023), and media (Loos & Ivan, 2018). Media discourses about older adults often imply dependency and otherness, shaping how society treats older adults (Fealy et al., 2012). Before the advent of generative AI, there was a tendency towards more inclusive and positive depictions of older adults in the media (Loos & Ivan, 2018). However, portrayals of later life in the media used to remain fairly limited. Moreover, depictions that seem "positive" at first glance often seem more complex and ambiguous in their ageing representations (Ylänne, 2026), for example, showing a progressive interest in considering older adults as potential audiences and consumers. As a result, older adults were portrayed in media images either as healthy, active, and prosperous, with images that tend to be bright and positive, or as depressed and dependent (Loos & Thijssen, 2022); extreme situations that do not account for the diversity of real-life experiences. Thus, while dealing with GenAI, media professionals have an opportunity to challenge age stereotypes and reduce ageism.

Our main objective is to understand how the "age" category is visually represented in GenAI systems. We focus on images generated by Midjourney, one of the most used GenAI tools globally (AIPRM, 2024). Midjourney is more than a text-to-image generator: it also functions as a social platform where users share, browse, discuss, and iterate on images and prompting strategies. While it competes with other leading tools, it can also serve as a constantly updated, community-curated "living" repository of hyperrealistic visual references, which may inform visual and creative direction in journalistic practice. It is important to note that we use minimal prompts to capture Midjourney's default biases. Results may differ when users write more detailed prompts in real-world contexts.

We focus on images for their news value and role in shaping collective imaginaries (Niederer & Colombo, 2024). We adopt an intersectional perspective (focusing on age in relation to gender and race) as an example of how stereotypes are embedded in GenAI systems. Furthermore, we also reflect on the implications of these text-to-image biases for journalistic ethics in the algorithmic era (Paik, 2025), particularly concerning issues of diversity, inclusion, and representation—core concerns of journalism ethics (Cools & Diakopoulos, 2024, p. 5). To address this objective, the study is guided by the following research questions:

- RQ1. How is age visually represented in generative AI images when age is unspecified compared to when it is explicitly indicated?
- RQ2. How do age representations in generative AI intersect with gender and race in the generated images?
- RQ3. What are the ethical implications of age-related visual biases in generative AI for news production?

2. Generative visual AI in News Organizations

Generative images are increasingly entering newsrooms, particularly through photo editors and visual teams. Interview-based evidence from major media organisations in Australia and the U.S. documented as early as 2023 the use of text-to-image tools for creative brainstorming and for illustrating stories—especially pieces about generative AI itself (Thomson et al., 2024). In parallel, an AP industry report based on 292 journalism professionals found that multimedia production is among the main uses after text, with 20.4% reporting use for producing illustrations/video/audio or editing images (Diakopoulos et al., 2024). Focusing specifically on image generation, a representative UK survey of more than 1,000 journalists estimates that around 6% use these tools regularly (the report mentions the use of Midjourney, Stable Diffusion, or Wochit), and shows higher overall adoption among men, under-50s, those with management responsibilities, and multi-format producers (Thurman et al., 2025).

Although the use of AI-generated images is not the dominant application of GenAI in newsrooms—unsurprisingly, given that news is primarily a text-based product—these developments are already reshaping routines and expectations. Visual GenAI is incorporated as a support resource (ideation, illustration, format adaptation), while simultaneously increasing the need for editorial standards, verification practices, and human oversight. Visual professionals often describe these tools as both useful and "ethically challenging, concerning and risky" (Thomson et al., 2024, p. 1701), which fuels internal debate, selective adoption, and the consolidation of roles focused on AI expertise and governance in news organisations (Diakopoulos et al., 2024).

Against this backdrop, studying bias in AI-generated images matters for journalists' work because it shows how these tools make some social categories more visible than others, which can influence editorial choices early in the production process. Even when such images are not published directly, they can shape illustrative decisions and visual framing, with implications for credibility and ethical coherence. Examining representational patterns in Midjourney can therefore help anticipate risks and inform newsroom policies, supporting more responsible visual practices as GenAI becomes part of everyday content production (Diakopoulos et al., 2024; Thomson et al., 2024; Thurman et al., 2025).

The constantly evolving nature of these technologies and their fast adoption are prompting media organisations to establish AI guidelines to ensure responsible, ethical, and unbiased use of GenAI. Common principles in AI guidelines include transparency, accountability, fairness, privacy, and consideration of journalistic values. Thus, best practices include "human oversight, explainability of AI systems, disclosure of automated content, and protection of user data" (de-Lima-Santos et al., 2025, p. 2585). Although AI guidelines have proliferated, studies highlight the lack of a regulatory framework as safeguards are weak, and guidelines often fail to clarify how to apply those concepts, for example, how to interface with third-party applications and providers. This lack has implications for journalist ethics—what journalists and news organisations "should do"—which is especially relevant in the context where journalists confront "new forms of media" (Ward, 2019, p. 295).

3. GenAI and bias

GenAI refers to a class of online AI tools that enable the creation of content based on patterns derived from existing data and large language models. Its popularity surged with the release of ChatGPT in late 2022, recognised as the fastest-growing consumer application in history, reaching 1 million users in just five days (Gordon, 2023). Today, GenAI tools are diverse and used by millions worldwide to generate a plethora of content (such as text, code, audio, images, and more). Illustrating this surge, an estimated 15 billion generative images were produced with GenAI tools in just one year—equivalent to the total photographic output of the preceding 150 years (Valyaeva, 2023). GenAI images are rapidly consolidating as a subject of inquiry in their own right (Leong et al., 2025) and producing an ever-expanding debate on their transformative social impacts (Dwivedi et al., 2023).

At an individual level, GenAI offers a faster, cheaper alternative to capturing or producing images. At a collective level, however, the mass production of images confronts professionals with the problem of visual AI bias: a systematic tendency to generate representations that misalign with social values such as inclusion, equity, and diversity (Bruiger, 2025). Both academics and professionals are increasingly concerned. Scholars in critical AI studies identify bias as one of the field's central challenges, framing it as a technical distortion and a social issue of fairness (Choudhry et al., 2024; Ferrara, 2024; Tang & Zhu, 2024). Media photo editors worry that image generators could mirror real-world inequalities and reinforce stereotypes—especially about race, sexual orientation, and disability (Thomson et al., 2024).

Empirical research shows that AI-generated images from Midjourney, DALL·E, Copilot, and Stable Diffusion perpetuate stereotypes about gender (Sandoval-Martin & Martínez-Sanzo, 2024), race (Nicoletti & Bass, 2024), and origin (Turk, 2023). Yet much less attention has been given to age bias in GenAI (Bird et al., 2023). This neglect is unsurprising, as age discrimination—though pervasive—remains unnoticed in research, industry, and society at large (Officer & de la Fuente-Núñez, 2018), and thus takes more subtle forms.

The first references to age bias in GenAI emerged from empirical studies documenting stereotypes of older adults in chatbot responses (Aranda Rubio et al., 2024; Kamruzzaman et al., 2024). It was also considered in theoretical reflections (Vázquez & Garrido-Merchán, 2024), visual discourse about dementia in generative images (Putland et al., 2023), and studies on depictions of older adults in GenAI outputs (Kamelski & Klinge, 2024). Recently, a study demonstrated that DALL·E 2-generated images continue to exhibit ageist features over time (Martens et al., 2025). Another study highlights an over-representation of older adults in humorous contexts, reinforcing stereotypes and discrimination—not only against older adults but also against other marginalised groups such as people with visual impairments or higher body weight (Saumure et al., 2025).

4. Methodology

This study explores the visual depictions of age and ageing in images generated with basic prompts on Midjourney. The images were generated using Midjourney (ver.: 6.1) (Caleb, 2024), during the period spanning December 9, 2024 to January 8, 2025, in order to maintain aesthetic coherence across multiple images and reduce undesirable variability in the results; we created and applied a Midjourney filter that was based on the selection of photorealistic images. First, we generated 91 everyday activities by combining the

list of the Harmonised European Time Use Surveys (Eurostat, n.d.) and diverse leisure activities (Janke et al., 2006) to ensure a varied set of everyday activities. Every activity was combined with two user targets to create the prompt. On one side, a "person" was used to evoke age-neutral depictions; on the other, an "older person" was used to evoke older-specific depictions, with no specific references to biological age, to be able to analyse what Midjourney identifies as an older person. We used "older person," as it resulted in the most general and least stereotypical option for generating images of single older individuals. This decision was made after presenting the results in different cultural and academic contexts. We avoided alternatives such as "elderly" (often associated with frailty or "very old age") and "old people" (which is plural and can sound blunter or more impersonal). No further details were provided to understand how age is depicted by default in Midjourney.

Furthermore, to strengthen methodological robustness, we incorporated human validation and qualitative review throughout the image-generation process. In January 2025, three researchers manually reviewed all prompts and resulting images as a quality-control step prior to automated analysis. We flagged outputs that did not clearly depict identifiable people (e.g., faces not sufficiently visible), deviated from the intended hyperrealistic style, or failed to represent the action described in the prompt according to our criteria. In such cases, prompts were reformulated and re-run under the same settings. During this screening, we also recorded qualitative cues to contextualise the subsequent analyses (e.g., framing, apparent age/gender/racial diversity, gendered or racialised roles, facial affect, and clothing/personal style). Disagreements were discussed until consensus was reached, with reflexive attention to the influence of our background in age studies.

Secondly, we asked Midjourney to generate 8 images for each activity, for each age group (with and without specifying age). As a result, a total of 728 images of "a person" and 728 images of "an old person" were produced, for a total of 1,456 images generated. We also generated a description for each image using Midjourney's Describe tool (see an example in Figure 1). This allowed us to obtain a detailed description of what they represented in the image, beyond external interpretations; it functions as an auxiliary mechanism to enrich the examination of how AI depicts older adults. All images were batch-downloaded and consolidated into a single table for subsequent analysis, including the prompt, the target age group, and a unique image identifier for each image.

For image analysis, we focused on quantitative analysis of the demographics of individuals represented in the images and on other qualitative characteristics of the images and their textual descriptions. We used a visual language model for visual captioning, particularly LLaVA-1.6 (Liu et al., 2023). It allowed us to predict demographic characteristics—age (coded into seven age groups: 18–24, 25–34, 35–44, 45–54, 55–64, 65–79, and 80+), gender (values: "woman", "man"), and race (values: "white", "black", "asian", "other")—and contextual attributes—emotions (values: "positive", "negative", "neutral"), image brightness (measured on a 0–100 scale, where 100 indicates the brightest value), and garments and style (reported as the 10 most frequent terms)—associated with individuals depicted in the images. These results should be interpreted as model-based attributions. Thus, the results do not intend to represent reality, but a commonly used Large Language Model (LLM) prediction.

For the textual analysis of the image descriptions generated by Midjourney, we used the LLM Llama (ver.: llama3-70b-8192), particularly a 70b parameter open-weight model from the Llama family run through Groq. It allowed us to extract different types of words, for example, gender, age, and nationality-related words,

adjectives, etc., in the images obtained. The generated descriptions were reviewed in full and manually validated through several rounds of iteration. We systematically removed non-informative elements and generic image-related terms that offered no clear evidence for the study's purposes (e.g., "photo", "image", "portrait", "shot", "camera"), along with punctuation and other low-value tokens. We also checked whether the remaining terms were clearly visible in the images and analytically useful for our study; if not, we removed them one by one until we consolidated a convincing final sample for analysis. These corrections were embedded in an ongoing, continuous monitoring of the results, thereby strengthening the contribution of human verification throughout the entire process. The analysis was guided by frequency lists (top 10 terms) and by grouping words into predefined analytical categories (e.g., age, gender, place/nationality, setting, emotions, and grammatical classes), while also tracking references to digital or assistive devices, environmental cues, colours, and clothing/personal style.

Figure 1: Midjourney images and descriptions of a Person / Older person generated with the prompt "taking a break"

Midjourney Description: "A man sits on a red wooden bench, looking at the blue sky and white clouds in the distance. The camera is positioned behind him, creating an atmosphere of loneliness. He has short hair and wears black. The scene is captured with cinematic lighting and a low-angle shot, in the style of a Nikon D850 raw format. --ar 1:1"

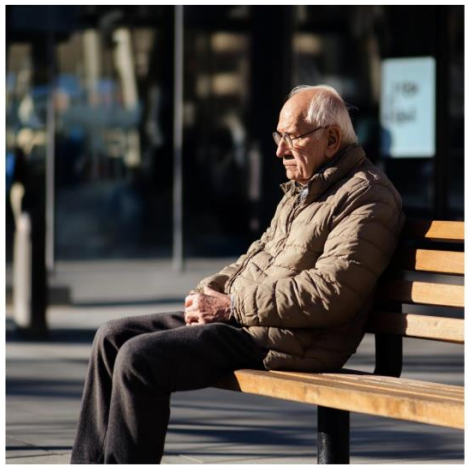

Midjourney Description: "Elderly man sitting on wooden bench in city, side view. The elderly is wearing glasses and dressed for cold weather with black trousers and brown jacket. In the background you can see an office building or shopping center with glass windows. Sunlight falls from above, creating shadows of trees on concrete floor. Shot by Nikon D850 DSLR camera with a wide-angle lens, capturing natural daylight conditions. --ar 1:1"

Source: Midjourney (ver.: 6.1)
Note: Descriptions were generated using Midjourney's Describe tool

5. Results

In this section, we present the results from the computer vision analysis of the images and the textual analysis of the image descriptions generated with Midjourney. We report results for the demographic and contextual attributes previously mentioned in the methods. For ease of reference, prompts that include the string "a person" will hereafter be referred to as "age-neutral prompts", and those that include "an older person" as "older-specific prompts", and consequently, "age-neutral depictions" and "older depictions" respectively for images generated.

5.1. Demographic depictions: age

One of the main findings of this research is that the images generated with Midjourney reveal a distinct pattern regarding age (Table 1; Figure 2). When age-neutral prompts were used, most depictions represented young people. Specifically, 69% of the images from age-neutral prompts depicted people between 25 and 34 years old. In contrast, only 5% of individuals generated with age-neutral prompts fell within the 65 to 79 age group. On the other hand, older-specific prompts overwhelmingly produced individuals aged between 65 and 79, accounting for 73% of cases. Similarly, the Midjourney Describes obtained were words associated with young people (e.g. "young", "early twenties", "mid 20s", "early thirties") on 203 occasions (28%), and the words associated with older people (e.g. "older", "elderly", "late sixties") appeared 30 times (4.1%).

Table 1: Age distribution of Midjourney-generated images

	18-24	**25-34**	**35-44**	**45-54**	**55-64**	**65-79**	**80+**
Age-neutral prompt	0.16	0.69	0.09	0.02	0.00	0.05	0.00
Older-specific prompt	0.00	0.00	0.00	0.00	0.00	0.73	0.26

Source: Author's own elaboration based on the LLaVA computer vision model.
Note: N 1,456 images. Age categories were derived via automatic classification; image generation was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methodology).

When we compared Midjourney age-neutral depictions to real-world demographic indicators (United Nations, n.d.), statistically significant differences were observed (Table 2). According to real-world indicators, the age category was significantly more distributed ($p<.001$). Individuals aged 25 to 34 account for approximately 20% of the global population, far below the 69% represented in Midjourney-generated images. Additionally, people aged 45 to 54 represent 15% of the global population, whereas this group is nearly absent in Midjourney's outputs, appearing in only 2% of cases. A similar pattern was observed among individuals aged 55 to 64, who were absent from the AI-generated images, despite representing 13% of the actual population. Perhaps the most striking finding was the absence of the population aged 80 and over from the images generated by Midjourney. Yet, they account for around 3% in real-world demographics. It may seem insignificant, but it represents over 175 million older people worldwide.

These comparisons are used to describe patterns in automatic classification of AI-generated images, not to claim real-world representativeness; the same applies to all demographic comparisons throughout the study. Overall, these results suggest that youth operates as Midjourney's default in age-neutral prompts, whereas older age becomes visible mainly when explicitly requested.

Table 2: Computer vision's age distribution in Midjourney-generated images vs real-world data

	18-24	**25-34**	**35-44**	**45-54**	**55-64**	**65-79**	**80+**
Midjourney (age-neutral prompts)	0.16	0.69	0.09	0.02	0.00	0.05	0.00
Real-world indicators	0.21	0.20	0.18	0.15	0.13	0.11	0.03

Source: Author's own elaboration based on the LLaVA computer visual model and real-world indicators drawn from United Nations statistics (as noted above).
Note: N 1,456 images. Midjourney age categories were derived via automated classification; image generation was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methodology).

Figure 2: A person (above) / Older person (below) handling home tasks

Source: Midjourney
Note: Age-neutral depictions were shown as young and usually with a long shot of a slim figure, very often from behind, and with a straight back, wearing informal clothing, and actively performing the activity; older depictions were portrayed in a front face close-up, which ensures that wrinkles are visible, they are slightly hunched, and have more passive attitudes or more precarious conditions

5.2. Demographic depictions: gender

Regarding gender, the results showed minor differences (Table 3a/b; Figure 3). In the computer vision analysis of age-neutral depictions, the gender split was fairly even: around 52% women and 48% men. Real-world data are similar to these results, where gender is almost equally distributed, although with more women than men (United Nations, n.d.). In contrast, the gender distribution of older-specific depictions significantly differed from the real-world distribution of people aged 65 or older, with around 56% of women aged 65 or older in 2022 ($p<.005$). Midjourney draws older people as male in 67% of cases, above the 44% of men in the total older population worldwide.

The analysis of Midjourney descriptions showed a similar picture compared to the computer vision analysis. Among age-neutral depictions, words associated with men (male, man, boy, father) and women (female, woman, girl, her, mother, feminine) were more balanced (275 times (38%) vs 398 times (55%)). And in older-specific images, men appeared more often than women (479 (66%) vs 297 (41%)).

Table 3a: Gender distribution in Midjourney-generated images (age-neutral prompt) vs real-world data

	Women	**Men**
Age-neutral prompt	0.52	0.48
Real-world indicator (All ages)	0.49	0.51

Source: Author's own elaboration based on the LLaVA computer vision model and real-world indicators drawn from United Nations statistics (as noted above).
Note. Midjourney gender categories were derived via automatic classification; real-world indicators refer to the global population (≈8.2 billion people), based on United Nations statistics (as noted above); image generation was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methods)

Table 3b: Gender distribution in Midjourney-generated images (older-specific prompt) vs real-world data

	Women	**Men**
Older-specific prompt	0.33	0.67
Real-world indicator (+65-year-old)	0.56	0.45

Source: Author's own elaboration based on the LLaVA computer vision model and real-world indicators drawn from United Nations statistics (as noted above).

Note. Midjourney gender categories were derived via automatic classification; real-world indicators refer to the global population (≈8.2 billion people), based on United Nations statistics (as noted above); image generation was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methods)

Figure 3: A person (above) / Older person (below) at work

Source: Midjourney.

Note: Age-neutral depictions often represented gendered roles and styles. These gendered roles were even more reinforced in the cases of older people's depictions, particularly regarding home tasks, or work-related activities and very often they represent older people in more passive roles. In sum, older-specific depictions more frequently depict men than age-neutral depictions, diverging from real-world demographics and reinforcing gendered patterns.

5.3. Demographic depictions: race

Regarding the race category, a similar pattern of differential representation emerged in age and gender categories. The computer vision analysis (Table 4; Figure 4) showed a statistically significant difference ($p<.0001$) between the race distribution of images generated from age-neutral prompts and older-specific prompts. When age was not specified, White and Asian individuals accounted for roughly 30% of generated images, whereas Black individuals appeared in only about 9% of cases. In contrast, with older-specific prompts, White individuals dominate (≈50%), followed by Asian individuals (≈18%). Most strikingly, Midjourney generated no older Black depictions.

Table 4: Race distribution in Midjourney-generated images

	White	Asian	Black	Other
Age-neutral prompt	0.30	0.30	0.09	0.32
Older-specific prompt	0.51	0.18	0.00	0.32

Source: Author's own elaboration based on the LLaVA computer vision model (as noted above).
Note. Midjourney race categories were derived via automatic classification; image generation was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methods)

In turn, Midjourney Describes provided no specific clue to the origin of White people, but they sometimes described the origin of non-White people. They used nationality-related words less often in age-neutral depictions, 175 times (24%), than in older adult depictions, 128 times (17%). Overall, older-specific depictions show less racial diversity, with whiteness as the dominant default.

Figure 4: Person (above) / Older person (below) engaging in religious activities

Source: Midjourney.
Non-White depictions were often associated with racialised roles, as in spiritual activities for Asians.

5.4. Affective depictions: emotions

Regarding emotions, distribution is also significantly different ($p<.001$). While most images fell into neutral or positive categories (Table 5). In images generated from age-neutral prompts, neutral and positive emotions were distributed roughly equally. For images generated from older-specific prompts, positive emotions were noticeably lower than neutral ones, in only about 25% of cases.

Table 5: Computing vision emotion distribution in Midjourney-generated images

	Positive	**Neutral**	**Negative**
Age-neutral prompt	0.46	0.53	0.01
Older-specific prompt	0.25	0.74	0.01

Source: Author's own elaboration based on the LLaVA computer vision model (as noted above)

Note. Midjourney emotion categories were derived via automatic classification; image generation was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methods)

In this regard, Midjourney Describes used more words associated with emotions to describe older depictions than age-neutral depictions. Emotion-related words were used 888 times (1.2 times per image) to describe older depictions and 663 times (0.9 times per image) to describe age-neutral depictions. These included words like "smiling" (134 times for older depictions, 76 times for age-neutral depictions) and "happy" (81 times for older depictions, 39 times for age-neutral depictions). In contrast, some words associated negative attributes with older people in comparison with age-neutral depictions, including "confidence" (10 times for older-depictions, and 6 times for age-neutral depictions) or "attractive" or "beautiful" (0 times for older-depictions, and 10 times for age-neutral depictions), "solitude" and "loneliness" (5 times for older-depictions, and 3 times for age-neutral depictions).

5.5. Visual depictions: brightness

Regarding brightness, our results showed that older depictions were around 11% dimmer than age-neutral depictions ($p<.0005$) with older people's depictions more likely to be located in darker environments (Table 6). Darker images are culturally associated with negative attitudes, environments, and emotions. These darker depictions of older people contradict the tendency to exploit active and healthy ageing attributes, in which older people are often depicted in bright environments (Loos & Thijssen, 2022).

Table 6: Computing vision brightness distribution in Midjourney-generated images

	Men	**Women**	**All**
Age-neutral prompt	0.95	0.95	0.96
Older-specific prompt	0.88	0.83	0.85

Source: Author's own elaboration based on the LLaVA computer vision model (as noted above)

Note. Image brightness was derived via automatic classification; image generation was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methods)

5.6. Cultural depictions: garments and style

The cultural depictions were analysed using Midjourney's describes (Table 7). Results revealed the use of different garments depending on the age profile. This depiction reinforces the vulnerability, physical decline, and personal neglect of older people. Accordingly, age-neutral depictions often showed informal outfits, such as t-shirts, jeans, and sunglasses. In contrast, older people's depictions conveyed a sense of keeping warm with more items, such as sweaters, jackets, scarves, and hats. Older people's physical decline was depicted through the use of glasses. Moreover, descriptions of age-neutral individuals showed more concern

about style, expressed through the focus on hairstyles and, in lesser measure, makeup, including the use of dreadlocks, bun, tied back, ponytail, etc.

Table 7: Garments and style

	Age-neutral	**Older-depictions**
T-shirt	222	105
Jeans	98	38
Sunglasses	14	6
Sweater	46	130
Jacket	22	55
Scarf	0	16
Hat	2	18
Glasses	123	272
Hairstyles	65	5
Makeup	3	0

Source: Author's own elaboration based on the LLaVA computer vision model (as noted above).
Note. Image descriptions were obtained using Midjourney's Describe tool; image generation was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methods)

Below, we present a summary table of results highlighting the main findings (Table 8):

Table 8: Summary table of results

Prompt starts with	**Older adults appear**	**Women appear**	**Racial/ethnic groups appear**	**Emotions are**
"a person"	Underrepresented (e.g., the 65–79 age group is only 5%)	Represented overall (e.g., women 52%)	More diverse overall, but some groups remain scarce (e.g., Black people, ≈9%)	Mostly neutral/positive (e.g., positive 46%)
"an older person"	With age stereotypes (more passive/precarious portrayals) and darker depictions (e.g., ≈11% dimmer)	Underrepresented (e.g., women 33%). Gender roles reinforced	Less diverse: White people dominate; no depictions of older Black people; racialised roles appear	Less positive (e.g., positive 25%)

Source: Author's own elaboration
Note. The whole process of analysis was iterative and human-supervised: outputs that did not meet the criteria (e.g., non-coherent images or non-identifiable people) were re-generated until they did (see Methods)

6. Discussion and conclusion

Our findings are based on a case study of Midjourney (images generated under the settings and period reported) and should be interpreted as specific to this case. In what follows, we distinguish between (a) findings supported by the Midjourney data and (b) broader implications for journalistic practice, which are interpretive and normative in nature.

This research contributes to the growing body of knowledge on GenAI (Leong et al., 2025) by demonstrating how Midjourney can reproduce and amplify visual age bias in its generated images. This study benefits from a substantive analysis of a large set of images and from the integration of quantitative and qualitative methods. Even so, such biases may appear minor while still producing substantial, unpredictable consequences over time, resembling the "butterfly effect" in AI (Ferrara, 2024). Taken together, our findings address the three research questions guiding this study.

Regarding how age is visually represented (RQ1), the results show that Midjourney establishes youth as the visual default: when age is not specified in the prompt, most representations depict young adults, while older age becomes visible mainly when explicitly requested—and typically through age-stereotypical cues. Images reinforce youth as active and dynamic, while old age is portrayed as passive and fragile, less positive in affect, and in darker environments. In line with what we outlined in the Introduction, our results echo prior work documenting stereotypical patterns in generative AI outputs (Nicoletti & Bass, 2024; Sandoval-Martin & Martínez-Sanzo, 2024; Turk, 2023). Our results extend this line of inquiry by demonstrating that age operates similarly, with youth shaping representations of the "general population." In line with Martens et al. (2025), who observed the persistence of digital ageism in AI-generated images over time, our larger data set suggests that such biases are not incidental but structurally embedded in generative systems.

In relation to how age intersects with gender and race (RQ2), we found that older depictions also show reduced gender and racial diversity and stereotypical patterns, compared with age-neutral depictions. Older individuals are more often depicted through gendered roles and passive tasks or gestures, and non-White identities are more likely to be reduced to origin stereotypes, restricting role diversity and perpetuating cultural clichés. The intersectional patterns observed here resonate with theoretical and empirical discussions of ageism, which emphasise that age discrimination rarely operates in isolation (Rocha et al., 2022). Accordingly, visual ageism should not be examined as a standalone phenomenon: its intersections with other forms of discrimination make its dynamics more visible and its effects more clearly articulated.

Finally, our results raise ethical concerns for journalistic work (RQ3), not only for photo editors but also for emerging AI-focused newsroom roles, journalists, editors, and technical staff. Although AI has become a productivity tool in newsrooms, bias remains a significant concern, making sustained human oversight indispensable and challenging. In this context, journalistic practice cannot remain uncritical. Responsibility should rest, in part, with media organisations themselves. However, this also introduces new challenges, as newsrooms often outsource key functionalities to large technology providers (e.g., Google or OpenAI), which can deepen platform dependence and expand external influence over editorial infrastructures and decisions (Simons, 2025). Newsrooms need clear, operational policies grounded in core journalistic values—truth, verification, transparency, and fostering an informed and critical public (Kovach & Rosenstiel, 2007). These policies should, as a baseline, require clear labelling of AI-generated images, enable traceability (e.g., documenting prompts and tool versions), and mandate verification prior to publication. Still, policies alone

are insufficient without professional commitment in a fast-moving ecosystem where newer models do not guarantee fewer biases. This challenge connects to the broader alignment problem (Christian, 2021): technically plausible outputs may still fail to reflect social values such as diversity, equity, and inclusion. News organisations cannot allow industry logics to override the ethics that should guide journalism in every outlet, every day, and in every story.

The broader context is also shifting: where media once exercised editorial control over content, control increasingly becomes algorithmic. For example, while earlier media research suggested a gradual shift toward more positive or "active ageing" representations in legacy media (Loos & Ivan, 2018) our findings indicate that generative AI may not follow the same trajectory. Instead, it appears to reproduce narrower and more conventional tropes, reinforcing concerns in intellectual debates about the persistence of stereotypical visual framing in AI systems. Particularly, our results show a turn towards a more subtle representations of ageism that affects more deeply the imaginaries about older adults; for example, depictions of older adults were darker than the rest of the images, and older adults were typically using a sweater or jacket instead of the typical t-shirt of age-neutral depictions, which shows a more active life, less concerned with care.

However, one of the main challenges would be the widespread lack of awareness of ageism in society; AI ageism (Stypinska, 2023) may affect older adults and their self-image. Stereotypical portrayals of old age not only reinforce social perceptions of this life stage but can also shape older people's self-perception and daily experiences. Diversity in representations of age, gender, and race in the media can challenge prejudice and foster a sense of belonging among marginalised communities. Likewise, challenging stereotypes about older people through images can support personal identities and broader social values. Ageing does not equate to decline: people of all ages can live fulfilling lives. Older adults can be stylish, attractive, and active while also navigating health challenges associated with ageing. Negative representations may be internalised, affecting self-esteem, agency, and belonging. At the same time, they shape how other generations imagine ageing, limiting the possibility of a future marked by dignity and diversity.

This research has several limitations. First, the results reflect how Midjourney generated images and descriptions at a specific point in time. We do not know to what extent the system still operates in the same way, or whether results obtained at different times would reveal different patterns, more or less ageist according to our interpretation. In this sense, the study provides a snapshot that should be interpreted accordingly. Therefore, generalisations to other models, platforms, periods, or journalistic contexts should be made with caution. Second, we used basic prompts to compare age-neutral and age-specific depictions. While journalists and media professionals may use more detailed prompts, a gap always remains between the prompt and the output. Large language models fill this gap with the most probable options, which can reproduce latent biases. Third, image and description classification was carried out using computer vision and large language models. Age and emotion algorithmic classification are widely recognised as limited and contestable (Kim et al., 2021), and this raises epistemological constraints. For us as humans, being part of an open and diverse culture was not easy to judge IA depictions in terms of binary demographic categories or beyond stereotypical judgments. We consider using a LLM for automatic classification, despite all its limitations, as these kinds of systems are those used for image classification in common digital platforms where this kind of content will be published and shared. We treat the demographic categories produced by computer vision as outputs for analysis; their probabilistic nature and socially constructed character must

be emphasised. Age, gender, and race are not natural facts but context-dependent categories whose meanings vary across time and place; they cannot be treated as stable or universal. Moreover, generative AI and computer vision systems cannot access individuals' identities or self-identification. Instead, they assign labels based on visual cues (e.g., style, lighting, clothing) and learned correlations. As a result, outputs often reflect how the system classifies images rather than who the depicted individuals are." Accordingly, our object of analysis is not real-world populations but AI-produced visual representations. Comparisons with real-world demographic distributions should therefore be made cautiously and should not be read as normative claims about representativeness. In this sense, one of the main limitations of automatic classification concerns visual diversity: labels can be reductive (e.g., binary gender), and race is especially difficult to operationalise consistently given its conceptual complexity and social sensitivity. Computer vision systems may also reflect or amplify biases in training data, annotation practices, or algorithmic design, limiting neutrality and potentially reinforcing inequalities. Finally, our selection of variables emerged from a preliminary, researcher-led categorisation and may be shaped by our backgrounds in age studies. For these reasons, AI-generated depictions should be interpreted as algorithmic attributions rather than measures of intrinsic characteristics of real individuals. We do not claim that the model "detects" demographic properties; rather, we analyse how classification systems, such as those used by digital platforms where these images will be published and shared, may encode social differences in generative images.

Further research could build on this study by focusing on the analysis of images and visual biases related to age and other social categories—not only race and gender, but also other forms of discrimination such as ableism and fatphobia, among others. Here, we analysed images generated in Midjourney that were not published in the media, given the limitations, to confirm whether images already published by news outlets and produced with generative tools were, or were not, and the lack of regulation on the matter. This study could be complemented with interviews with photo editors and staff in emerging AI-focused roles to better understand newsroom routines, motivations, and strategies for mitigating visual ageism while upholding journalistic ethics, following the line of inquiry on visual ageism proposed by Loos & Ivan (2018).

Acknowledgements / Funding
This research was funded by Volkswagen Foundation under grant number 9C565 + 9C565-1. The content is solely the responsibility of the authors and does not necessarily represent the official views of the Volkswagen Foundation.

Conflict of interests
The author(s) declare no conflict of interest.

Ethical statement
This study was conducted in accordance with the principles of scientific research and did not require additional ethics committee approval.

Declaration of AI usage
Generative AI tools were used for image generation and analysis (see Methodology), as well as for minor English-language editing. All scientific content, interpretations, and conclusions are solely those of the authors.

Data availability
The data supporting the findings of this study are available upon request.

Author contributions

	Juan Linares-Lanzman	Inés Montoya Espinagosa	David Carbonell Mateo	Andrea Rosales
Conceptualization	X			X
Data curation	X	X	X	X
Formal analysis	X	X		X
Funding acquisition				X
Investigation	X			X
Methodology	X	X	X	X
Project administration				X
Resources				X
Software		X	X	
Supervision	X			X
Validation	X			X
Visualization	X			
Writing - original draft	X			
Writing - review & editing	X	X	X	X

References

AIPRM. (2024, July 24). *AI in art statistics 2024*. https://www.aiprm.com/ai-art-statistics/

Aranda Rubio, Y., Baztán Cortés, J. J., & Canillas del Rey, F. (2024). Is artificial intelligence ageist? *European Geriatric Medicine, 15*(6), 1957–1960. https://doi.org/10.1007/s41999-024-01070-2

Barber, S. J., Shoemaker, K. J., & Dotson, V. M. (2024). Ageism harms older people's health, cognition, and well-being: A selective review and policy recommendations. *Policy Insights from the Behavioral and Brain Sciences, 11*(2), 190–198. https://doi.org/10.1177/23727322241268400

Bird, C., Ungless, E. L., & Kasirzadeh, A. (2023). *Typology of risks of generative text-to-image models*. arXiv. https://doi.org/10.48550/arXiv.2307.05543

Bruiger, D. (2025). Reflections on the AI alignment problem. *AI & Society, 40*(6), 4383–4392. https://doi.org/10.1007/s00146-025-02211-2

Butler, R. N. (1969). Age-ism: Another form of bigotry. *The Gerontologist, 9*(4, Part 1), 243–246. https://doi.org/10.1093/geront/9.4_part_1.243

Caleb. (2024, July 30). *Version 6.1*. Midjourney. https://updates.midjourney.com/version-6-1/

Cherubini, F., & Sharma, R. (2023). *Changing newsrooms 2023: Media leaders struggle to embrace diversity in full and remain cautious on AI disruption*. Reuters Institute for the Study of Journalism. https://doi.org/10.60625/RISJ-KHYH-CT64

Choudhry, M., Sundarrajan, M., Sundaram, K. & Rama, A. (2024). 9 Bias and Fairness in Generative AI. In S. Balasubramaniam, S. Kadry, A. Prasanth & R. Kumar Dhanaraj (Ed.), *Generative AI and LLMs: Natural Language Processing and Generative Adversarial Networks* (pp. 177-192). Berlin, Boston: De Gruyter. https://doi.org/10.1515/9783111425078-009

Christian, B. (2021). *The alignment problem: Machine learning and human values*. W. W. Norton.

Cools, H., & Diakopoulos, N. (2024). Uses of generative AI in the newsroom: Mapping journalists' perceptions of perils and possibilities. *Journalism Practice*, 1–19. https://doi.org/10.1080/17512786.2024.2394558

de-Lima-Santos, M.-F., Yeung, W. N., & Dodds, T. (2025). Guiding the way: A comprehensive examination of AI guidelines in global media. *AI & Society, 40*(4), 2585–2603. https://doi.org/10.1007/s00146-024-01973-5

Diakopoulos, N., Cools, H., Li, C., Helberger, N., Kung, E., Rinehart, A., & Gibbs, L. (2024). *Generative AI in journalism: The evolution of newswork and ethics in a generative information ecosystem* [Report]. https://doi.org/10.13140/RG.2.2.31540.05765

Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., . . . Wright, R. (2023). Opinion paper: "So what if ChatGPT wrote it?" Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management, 71*, Article 102642. https://doi.org/10.1016/j.ijinfomgt.2023.102642

Eurostat. (n.d.). *Harmonised European time use surveys*. Retrieved February 24, 2025, from https://ec.europa.eu/eurostat/web/microdata/harmonised-european-time-use-surveys

Fealy, G., McNamara, M., Treacy, M. P., & Lyons, I. (2012). Constructing ageing and age identities: A case study of newspaper discourses. *Ageing & Society, 32*(1), 85–102. https://doi.org/10.1017/S0144686X11000092

Ferrara, E. (2024). The butterfly effect in artificial intelligence systems: Implications for AI bias and fairness. *Machine Learning with Applications, 15*, Article 100525. https://doi.org/10.1016/j.mlwa.2024.100525

Gendron, T., Camp, A., Amateau, G., Mullen, M., Jacobs, K., Inker, J., & Marrs, S. (2024). The next critical turn for ageism research: The intersections of ageism and ableism. *The Gerontologist, 64*(2), Article gnad062. https://doi.org/10.1093/geront/gnad062

Gordon, C. (2023, February 2). *ChatGPT is the fastest growing app in the history of web applications.* Forbes. https://www.forbes.com/sites/cindygordon/2023/02/02/chatgpt-is-the-fastest-growing-ap-in-the-history-of-web-applications/

Janke, M., Davey, A., & Kleiber, D. (2006). Modeling change in older adults' leisure activities. *Leisure Sciences, 28*(3), 285–303. https://doi.org/10.1080/01490400600598145

Kamelski, T., & Klinge, D. (2024). Generative artificial intelligence and digital ageism: Exploring the construction of age and aging by image-generating AI. *Medien & Altern, 25*. https://doi.org/10.31219/osf.io/p3sdj

Kamruzzaman, M., Shovon, M., & Kim, G. (2024). Investigating subtler biases in LLMs: Ageism, beauty, institutional, and nationality bias in generative models. In L.-W. Ku, A. Martins, & V. Srikumar (Eds.), *Findings of the Association for Computational Linguistics: ACL 2024* (pp. 8940–8965). Association for Computational Linguistics. https://doi.org/10.18653/v1/2024.findings-acl.530

Kim, E., Bryant, D., Srikanth, D., & Howard, A. (2021). Age bias in emotion detection: An analysis of facial emotion recognition performance on young, middle-aged, and older adults. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 638–644). Association for Computing Machinery. https://doi.org/10.1145/3461702.3462609

Kovach, B., & Rosenstiel, T. (2007). *The elements of journalism: What newspeople should know and the public should expect*. Three Rivers Press.

Krekula, C., Nikander, P., & Wilińska, M. (2018). Multiple marginalizations based on age: Gendered ageism and beyond. In L. Ayalon & C. Tesch-Römer (Eds.), *Contemporary perspectives on ageism* (pp. 33–50). Springer. https://doi.org/10.1007/978-3-319-73820-8_3

Leong, L.-Y., Hew, T.-S., Ooi, K.-B., Tan, G. W.-H., & Koohang, A. (2025). Generative AI: Current status and future directions. *Journal of Computer Information Systems*, 1–34. https://doi.org/10.1080/08874417.2025.2482571

Lewis, S. C., Guzman, A. L., Schmidt, T. R., & Lin, B. (2025). Generative AI and its disruptive challenge to journalism: An institutional analysis. *Communication and Change, 1*(1), Article 9. https://doi.org/10.1007/s44382-025-00008-x

Liu, H., Li, C., Li, Y., & Lee, Y. J. (2023). *Improved baselines with visual instruction tuning*. arXiv. http://arxiv.org/abs/2310.03744

Loos, E., & Ivan, L. (2018). Visual ageism in the media. In L. Ayalon & C. Tesch-Römer (Eds.), *Contemporary perspectives on ageism* (pp. 163–176). Springer. https://doi.org/10.1007/978-3-319-73820-8_11

Loos, E., & Thijssen, J. (2022). Visual representations of older persons: Current discourses and their historical roots. In A. Teodorescu & D. Chiribucă (Eds.), *Shaping ageing: Social transformations and enduring meanings* (pp. 64–87). Routledge. https://doi.org/10.4324/9781003046790-6

Martens, L., Virgin, N., Hoffarth, P., Goodwill, J., Derenne, N., Van Eck, R., Klug, M., Manocha, G. D., & Jurivich, D. (2025). Recapitulation of ageism in artificial intelligence-generated images: Longitudinal comparative study. *Journal of Medical Internet Research, 27*, Article e68428. https://doi.org/10.2196/68428

Nash, P., Taylor, T., & Levy, B. (2020). The intersectionality of ageism. *Innovation in Aging, 4*(Supplement_1), 846–847. https://doi.org/10.1093/geroni/igaa057.3105

Nicoletti, L., & Bass, D. (2024, January 2). *Humans are biased. Generative AI is even worse*. Bloomberg. https://www.bloomberg.com/graphics/2023-generative-ai-bias/

Niederer, S., & Colombo, G. (2024). *Visual methods for digital research: An introduction*. Polity.

Officer, A., & de la Fuente-Núñez, V. (2018). A global campaign to combat ageism. *Bulletin of the World Health Organization, 96*(4), 295–296. https://doi.org/10.2471/BLT.17.202424

Owsley, C. S., & Greenwood, K. (2024). Awareness and perception of artificial intelligence operationalized integration in news media industry and society. *AI & Society, 39*(1), 417–431. https://doi.org/10.1007/s00146-022-01386-2

Paik, S. (2025). Journalism ethics for the algorithmic era. *Digital Journalism, 13*(4), 696–722. https://doi.org/10.1080/21670811.2023.2200195

Putland, E., Chikodzore-Paterson, C., & Brookes, G. (2023). Artificial intelligence and visual discourse: A multimodal critical discourse analysis of AI-generated images of "dementia." *Social Semiotics, 35*(2), 228–253. https://doi.org/10.1080/10350330.2023.2290555

Rocha, N. M. F. D., Maia, R. da S., Marques, G. M. V., & Bandeira, R. L. (2022). Intersectionalities and old age: Ageism in the crossroads of race, gender, and age. In L. Dutra-Thomé, D. F. Rabelo, D. Ramos, & E. F. Góes (Eds.), *Racism and human development* (pp. 163–172). Springer. https://doi.org/10.1007/978-3-030-83545-3_12

Rosales, A., Fernández-Ardèvol, M., & Svensson, J. (Eds.). (2023). *Digital ageism: How it operates and approaches to tackling it*. Routledge. https://doi.org/10.4324/9781003323686

Sandoval-Martin, T., & Martínez-Sanzo, E. (2024). Perpetuation of gender bias in visual representation of professions in the generative AI tools DALL·E and Bing Image Creator. *Social Sciences, 13*(5), Article 250. https://doi.org/10.3390/socsci13050250

Saumure, R., De Freitas, J., & Puntoni, S. (2025). Humor as a window into generative AI bias. *Scientific Reports, 15*(1), Article 1326. https://doi.org/10.1038/s41598-024-83384-6

Simon, F. M. (2024, February 6). *Artificial intelligence in the news: How AI retools, rationalizes, and reshapes journalism and the public arena.* Columbia Journalism Review. https://www.cjr.org/tow_center_reports/artificial-intelligence-in-the-news.php

Simons, M. (2025, November 21). *AI is changing the relationship between journalist and audience. There is much at stake*. The Guardian. https://www.theguardian.com/commentisfree/2025/nov/22/ai-is-changing-the-relationship-between-journalist-and-audience-there-is-much-at-stake

Steward, A. T., De Fries, C. M., Dunbar, A. Z., Trujillo, M., Zhu, Y., Nicotera, N., & Hasche, L. (2023). A phenomenological understanding of the intersectionality of ageism and racism among older adults: Individual-level experiences. *The Journals of Gerontology: Series B, 78*(5), 880–890. https://doi.org/10.1093/geronb/gbad031

Stypinska, J. (2023). AI ageism: A critical roadmap for studying age discrimination and exclusion in digitalized societies. *AI & Society, 38*(2), 665–677. https://doi.org/10.1007/s00146-022-01553-5

Tang, S., & Zhu, H. (2024). Mitigating bias in generative AI: A comprehensive framework for governance and accountability. *Law, Ethics & Technology, 3*, Article 0008. https://doi.org/10.55092/let20240008

Thomson, T. J., Thomas, R. J., & Matich, P. (2024). Generative visual AI in news organizations: Challenges, opportunities, perceptions, and policies. *Digital Journalism*, 1–22. https://doi.org/10.1080/21670811.2024.2331769

Thurman, N. J., Thäsler-Kordonouri, S., & Fletcher, R. (2025). AI adoption by UK journalists and their newsrooms: Surveying applications, approaches, and attitudes. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.5815043

Turk, V. (2023, October 10). *How AI reduces the world to stereotypes*. Rest of World. https://restofworld.org/2023/ai-image-stereotypes/

United Nations. (n.d.). *Population Division*. United Nations Population Division. Retrieved February 28, 2026, from https://www.un.org/development/desa/pd/

Valyaeva, A. (2023, August 15). *People are creating an average of 34 million images per day: Statistics for 2024*. Everypixel Journal. https://journal.everypixel.com/ai-image-statistics

Vázquez, A. F. de C., & Garrido-Merchán, E. C. (2024). *A taxonomy of the biases of the images created by generative artificial intelligence*. arXiv. http://arxiv.org/abs/2407.01556

Ward, S. J. A. (2019). Journalism ethics. In K. Wahl-Jorgensen & T. Hanitzsch (Eds.), *The handbook of journalism studies* (2nd ed., pp. 307–323). Routledge. https://doi.org/10.4324/9781315167497-20

World Health Organization. (n.d.). *Ageism*. Retrieved February 28, 2026, from https://www.who.int/health-topics/ageism

Ylänne, V. (2026). Representations of ageing in the media. In J. Twigg & W. Martin (Eds.), *Routledge handbook of cultural gerontology* (2nd ed., pp. 299–306). Routledge. https://doi.org/10.4324/9781003292227-37