

Do as I Say, Not as You Do: Assessing Public Versus Practitioner Perspectives on the Development of Artificial Intelligence and Language Technologies

Tyler Loakman^{1*,2}, Chenghua Lin²

¹Department of Computer Science, The University of Sheffield, UK

²Department of Computer Science, The University of Manchester, UK

chenghua.lin@manchester.ac.uk

Abstract

With the public awareness of artificial intelligence (AI) applications such as large language models (LLMs) and generative text-to-video models reaching unprecedented levels, there is an increasing need to reconcile the perspectives held by informed practitioners, who research and develop these technologies, with the interests and opinions of the general public, who constitute the majority of people impacted by technological developments. In this work, we survey practitioners (i.e., academics and industry professionals) alongside members of the general public on perspectives relating to pertinent areas of artificial intelligence, including perceived benefit and risk, the ability for models to achieve consciousness and output art, and the use of copyright protected material for model training. Alongside comparisons of diverse perspectives on these topics, we additionally contextualise the findings with a discussion of relevant literature that supports and refutes the myriad observed stances.¹

1 Introduction

The advent of technologies such as large language models (LLMs) like ChatGPT (OpenAI et al., 2023) and other artificial intelligence (AI) applications has driven the unprecedented integration of AI into all aspects of life, including healthcare (Luo et al., 2025; Mondal et al., 2022), education (Kumar and Lata, 2025; Han et al., 2024; Selwyn, 2022), entertainment (Hemraj, 2025), and everyday consumer goods such as fridges and e-cigarettes. The available investment capital for AI technology has additionally resulted in “AI washing”, where products and services exaggerate the extent and effectiveness of AI integration (Ozturkcan and Bozdağ, 2025; Elhajjar and Itani, 2025; Wu and Chen, 2022). As a result, the integration of AI into everyday life causes substantial economic gains

for some parties, whilst presenting an economic risk to others (Restrepo, 2025; Yuan et al., 2024).

There has been a growing body of work in natural language processing (NLP) on the post-training stage of model development, particularly regarding the alignment of model behaviours to human preferences (Pan et al., 2025; Shu et al., 2025). However, a substantial quantity of human preference data is typically required. One issue that commonly arises during this stage is that such approaches do not adequately account for the plurality of perspectives held by individuals across different cultures and social groups, leaving some demographics under-represented (Zhi-Xuan et al., 2025; Barretto et al., 2021). Similarly, the domain of model debiasing and detoxification relies on prior expectations about what is considered harmful and offensive (Zhi-Xuan et al., 2025). However, research shows that over-reliance on such approaches results in the degradation of overall model performance and a preference for abstaining from responding to topics that may be construed as sensitive or contentious, effectively neutering model usability (Kaneko et al., 2023; Meade et al., 2022).

Alongside the rapid deployment of AI technology and arising challenges in achieving representative alignment for models, there remains a substantial gap in understanding regarding the resulting societal impact. Crucially, the priorities and policies governing those developing AI (frequently driven by chasing leaps in technical capabilities and corporate objectives) may not align with the concerns of the public consumers whose lives these systems fundamentally affect (Horwitz, 2025; Amazon Employees for Climate Justice, 2025; Right to Warn Coalition, 2024).

To address this critical gap, we conduct a survey to compare the perspectives of AI practitioners (consisting of academics and industry professionals) with the general public. In doing so, we ascertain whether the creators and consumers of such technologies are aligned, thereby gaining a better understanding of the benefits and risks of continued AI development. We

*Work performed while at the University of Sheffield.

¹<https://github.com/tylerL404/do-as-i-say-not-as-you-do>

summarise the primary contributions of this work as follows:

- We conduct a survey of the three primary stakeholders in the development of artificial intelligence: academics, industry professionals, and the general public (dichotomised as practitioners and the public) on topics relating to both practical and philosophical domains (e.g., economic and legal concerns, versus art and consciousness).
- Responses are analysed to explore whether practitioners and the public are aligned in their perspectives on AI, assessing whether the development of such technologies is to the benefit of society as a whole.
- We provide an overview of the literature surrounding AI development, including areas of legal, economic, and social concern.

2 Related Works

Public Perception of AI. Existing works have explored the perspectives of different stakeholders in AI development. For instance, [Kendall Roundtree \(2024\)](#) presents a meta analysis of literature surrounding public perceptions of AI development alongside a sentiment analysis of social media communication on Quora and Reddit. Their findings demonstrate concerns regarding the impact of AI development on employment and the risks of AI driven healthcare. Furthermore, [Brauner et al. \(2023\)](#) surveyed 122 members of the German public on opinions regarding personal, economic, industrial, social, cultural, and health-related factors relating to AI development. They identified cybersecurity risks as being the most critical point of public concern when accounting for the interaction between risk and likelihood, whilst threats to careers and personal lives were seen as unlikely, and the position of AI performing unpleasant and undesirable human tasks was viewed as both likely and desirable. Furthermore, [Selwyn and Gallo Cordoba \(2022\)](#) analysed the understanding of the Australian public on AI in 2020, prior to the widespread adoption of performant LLMs. Their results showed that most participants had a sufficient working understanding of the technology, even if they self-identified as lacking knowledge.

Comparative Studies. Comparing the general public and practitioners, [Pew Research Center \(2025\)](#) directly surveyed 5k+ members of their American Trends Panel and 1k+ authors of AI publications in

2024 (all of which reside in the United States).² Results showed that AI experts were significantly more optimistic than the public regarding the potential impact of AI, particularly towards medical care, the economy, and education. We compare the findings of the present work with [Pew Research Center \(2025\)](#) in appendix A.1. Whilst these findings establish a trend of practitioner optimism and public pessimism, existing comparative works have focussed on broad societal applications (e.g., the economy and healthcare), whilst the present work additionally explores philosophically-laden topics such as art, creativity, and consciousness.

Present Work. Building upon these findings, we address gaps in the literature by surveying public and practitioner views on AI with a predominant (though non-exclusive) focus on language technologies. In addition to topics pertaining to practical areas such as the economy and healthcare, we further investigate perspectives surrounding philosophical concerns, such as the ability for AI to generate art and demonstrate consciousness. Such exploration enables a deeper understanding of people’s fundamental perspectives regarding subjective topics where opinions are more resistant to change, in comparison to concerns that can be empirically verified (e.g., the impact of AI on economic growth and employment).

3 Methodology

Recruitment of Participants. Recruitment was performed through addressing email lists at the lead author’s institution, posting on social media, and approaching industry contacts for possible survey dissemination. A wider range of general public respondents was sought through posts on Reddit and the placement of the survey on SurveySwap. All participants were encouraged to share the survey to further eligible participants. Responses were collected in 2025.

Survey Phrasing/Pilot. A small scale pilot study was performed to elicit feedback regarding question clarity and logistics (2 practitioners and 2 members of the public). Phrasing was altered for clarity in light of this feedback.

Response Filtering. All participants were asked to respond to questions outlining their understanding and familiarity with AI technologies in addition

²As a result, this work was both American-centric, and also surveyed members of a questionnaire panel, resulting in a substantial response bias. In contrast, this work is international (though predominantly European), and anonymous.

to their understanding of the technical underpinnings of machine learning (e.g., architectures and training methodologies). Participants additionally self-identified their position within AI development as being in academia (e.g., researchers and lecturers in natural language processing, computer vision, and general machine learning), industry (e.g., private researchers and applied AI software developers) or the general public. Using responses to a five-point Likert scale, we removed self-identified in-domain academics and industry practitioners who did not identify their own general and technical AI knowledge to be above or equal to 3. As a result, we filter out mistaken self-identifications (as practitioners necessarily must have sufficient domain-knowledge), whilst not being overly restrictive (e.g., requiring technical knowledge ratings of 4-5), therefore allowing multidisciplinary practitioners who may regularly use AI systems as the focus of their work, but for whom a deep technical knowledge may not be mandatory. Finally, response timestamps were checked for suspicious activity such as botting, of which no evidence was found.

Demographic Information. In total, the survey received 336 responses: 69 academics, 33 industry professionals, and 234 members of public. Following filtering, this resulted in 76 practitioners across academia and industry, and 212 members of the general public. Over half of respondents were Millennials (born 1981-1996), with similar numbers of Gen X (1965-1980) and Gen Z (1997-2006), at $\approx 20\%$ each).^{3,4} A further small number of respondents were from the Baby Boomer generation (1946-1964). In terms of geographical spread, the majority of respondents were from Europe ($\approx 60\text{-}70\%$), with a considerable number of North American responses ($\approx 17\text{-}30\%$). A further nominal number of responses were received from Oceania, East Asia, Southeast Asia, and the Middle East/North Africa (MENA). In total, age and geographical demographics were comparable across practitioners and the public, providing a sample that is not explicitly restricted to a specific locale or other demographic variable. We do not present further results breakdowns by region or age due to the skewed distribution. Additional information regarding the industry/field breakdown of the public respondents is provided in Figure 7 within Appendix A.3.

Rating Scales. All non-categorical statements were rated on a 5-point Likert (i.e., “strong disagree”

³Respondents were required to be at least 18 years of age.

⁴We use the textual generational labels for ease of reference - full date ranges were presented to respondents.




Figure 1: Practitioner and public use of AI technologies in free time and employment.

to “strong agree”, or the semantic equivalent for a given question phrasing).⁵ Only broad demographic information was collected to encourage honesty, as some respondents may have opinions that challenge the public stance of the organisations they represent.

Question phrasing has been condensed for inclusion in visualisations, whilst the full original questions are presented in Table 1 within Appendix A.2, in addition to significance testing for ordinal responses.⁶

4 Artificial Intelligence Usage

With the wide range of forms that artificial intelligence and machine learning models take, from simple grammar correction and translation tools to in-line code suggestions and autonomous driving assistance, in Figure 1 we first outline the different technologies being used both at home and for work.

Similar popularity rankings can be seen for different technologies between practitioners and the public, but with different levels of general use. For instance, the most popular technologies for personal use across

⁵Participants were only presented with the linguistic labels for the two extremes (i.e., 1 and 5), whilst a middle-point “neutral” option is added to figures for clarity.

⁶Differences between practitioner and public responses are significant at $\alpha=0.01$ for all statements.

both groups are translation tools and open-domain chatbots, with more than half of respondents from both groups using such technologies, whilst a considerable portion of people additionally interact with image generation, virtual assistants, and grammar correction AI ($\approx 20\%$). Across both groups, the rarest technologies used by participants are autonomous driving, AI-driven telemedicine, and music generation models - likely owing to their niche application, relatively rudimentary abilities, or high financial barrier to entry. Importantly, over half of public respondents used open-domain chatbots (i.e., LLMs). Such results are unsurprising, with OpenAI revealing that by 2025, approximately 10% of the world's adult population had interacted with ChatGPT alone (Hitzig et al., 2025). Regarding AI-usage at work (or in study, for full-time students), less than 30% of public respondents interact with any, whilst the majority of practitioners make use of open-domain chatbots and specialised code models. Of the respondents that reported using any AI applications in their employment, almost half of the public (43%) and more than half of practitioners (55%), did so without endorsement or recommendation by their employer, presenting data protection concerns. Furthermore, 22% of practitioners and 7% of the public reported that only a subset of the AI technologies used in employment were endorsed by their employer, while still making use of technologies outside of this scope.

5 Training Data Collection

The following section focuses on concerns surrounding the use of publicly available and potentially copyright protected material for model training. Owing to the need for large volumes of data to imbue models with sufficient world-knowledge and enable reasoning abilities, much of the available training data is sourced via scraping large volumes (i.e., the majority) of available web content, to then partially compete against human authors and creators in the same domains. As a result, popular pre-training corpora include collections such as CommonCrawl (Common Crawl Foundation, 2025) and RedPajama (Weber et al., 2024), which consist largely of publicly available data.

However, whilst model exist that have been trained on highly curated and ethically collected data, such as Olmo (Ettinger et al., 2025; Walsh et al., 2025; Groeneveld et al., 2024), which is trained on the Dolma dataset (Soldaini et al., 2024), the majority of large-scale pre-training still concerns mass-collected data where little regard is given to the licenses of the original works and the rights of copyright holders.

Whilst at the time of publication, the legal status of AI training remains unclear, there have been a growing number of legal challenges on the topic of model training, including for language generation (Advance Local Media LLC v. Cohere, 2025), code generation (Doe v. Github, Inc., 2022), and image generation (Andersen v. Stability AI, 2024). In part, such cases concern the disregard of traditionally upheld copyright protections and the debated matter of model training constituting fair use. This has resulted in loopholes being exploited by some companies, such as Anthropic legally obtaining non-pirated physical copies of print publications and destructively scanning them (Teilmann-Lock and Savin, 2025; Bartz v. Anthropic PBC., 2025), following fair-use precedent set by Google Books (Authors Guild Inc. v. Google Inc., 2015).

Additional complications arise from unclear agreements between platforms and users, where user consent is purportedly given for data to be used for AI purposes, whilst the layperson user is not fully aware of their agreements. Such patterns have been previously observed with the use of end user license agreements (EULAs), where details are obfuscated behind legalese or buried with extensive terms and conditions (Corbett, 2019; Hutton and Henderson, 2017).

Figure 2 presents survey results regarding whether companies should be permitted to train on public data without the explicit consent of copyright holders. Overall consensus of both groups is against the motion that copyrighted material should be without explicit consent (64% for practitioners and 90% for the public). However, a clear discrepancy is seen between the groups, with 31% of practitioners being in favour of being able to make use of copyright protected material without consent, versus only 4% of the public.

One potential explanation is practitioners viewing the often lengthy and expensive process of obtaining license information and contacting rights-holders to be burdensome, whilst members of the public may not believe that the benefit they receive from more performant AI products outweighs the implications of private organisations profiting from models trained on the work on uncompensated individuals. In particular, this topic raises a potential conflict of interest in practitioners, with many positions being highly lucrative compared to most employment performed by the public (Bone et al., 2025; PwC, 2025). As a result, being able to more quickly perform research and develop successful products may lead to increased compensation or career progression for practitioners, whilst from the public perspective many jobs are placed under

growing threat of replacement (Restrepo, 2025). On the other hand, mandating that companies have the requisite permissions for all training data can be viewed as a competitive disadvantage in the AI race, allowing companies in jurisdictions with more relaxed intellectual property laws to have a significant advantage by having access to substantially larger volumes of data (Rosati, 2025, 2019). We speculate that open-access models with lower barriers to entry (i.e., no fees) are more likely to receive licensing permissions owing to the wider democratisation of their capabilities.




Figure 2: Practitioner and public beliefs concerning AI models being trained on public data without explicit informed consent.

6 Art & Creativity

Whilst AI applications have previously been developed predominantly for more mundane or supra-human tasks such as text summarisation and cancer screening, models have now become capable of generating articulate, engaging text and realistic video, resulting in such systems increasingly being used in tasks that have historically relied on human creativity, such as writing narratives (Teleki et al., 2025) and poetry (Gonalo Oliveira, 2017) or designing infographics (Ghosh et al., 2025; Zhu et al., 2020) and movie trailers (Yao et al., 2025; Argaw et al., 2024).

Consequently, respondents were surveyed on their perspectives regarding the ability for AI models to demonstrate creativity and generate art - two philosophically-loaded concepts that are frequently related to intention and embodied experience, which AI may be considered to lack (Leong, 2025). Results are presented in Figure 3. Across both practitioners and the public, consensus is against AI models being capable of either of these feats. However, responses within practitioners on the topic of AI creativity are less unanimous, with only a small majority (32%) being strongly of the belief that such systems are incapable of creativity, compared to 74% of public respondents.

This discrepancy between creativity and art highlights a key insight into the differing perspectives




Figure 3: Practitioner and public beliefs concerning the ability of AI models to be creative and generate art.

of both parties. Whilst art is often seen as a highly human-centric concept that requires intention and meaning, creativity may also be defined as the ability to produce unique ideas (Dow, 2017). Consequently, technologies such as large language models are capable of rearranging sequences observed from training data to output novel linguistic strings (and therefore ideas). Practitioners may be more likely to consider such outputs as being creative, even if the ideas are amalgamations and abstractions of existing human thoughts contained in training corpora. Such comparatively shallow abilities have been referred to as pseudo-creativity (Runco, 2023). Even within humans, there is debate as to whether creativity is truly novel, or the result of exposure to embodied experiences that constitutes our training data (Lynn and Bassett, 2020).

Extensive work has demonstrated that people are more likely to rate content negatively when it is suggested that it has been created by AI rather than another human, regardless of tangible characteristics (Bellaiche et al., 2023). This preference for human-created content over machine created content is a simple an extension of longstanding patterns observed alongside industrialisation, such as the value placed behind handmade items versus the perceived coldness of machine produced products (Degirmenci et al., 2025; Fan et al., 2024; Fuchs et al., 2015; Kruger et al., 2004). The implications for future development are myriad, primarily concerning the nature of human-in-the-loop (HITL) generative systems and the extent to which human intention and intervention is necessary for the output of a human-in-the-loop system to retain deeper meaning as it relates to art and creativity. For instance, whilst such systems may be frame as collaboration between humans and AI, determining the level

of agency and ownership of either participant remains difficult (Carrera et al., 2026), particularly for creative domains (Salma et al., 2025; Loakman et al., 2025).

7 Thought & Consciousness

Whilst once firmly in the remit of science fiction (Nader et al., 2024; Hermann, 2023; Hudson et al., 2023), the concept of artificial consciousness and the ability for machines to emulate internal cognitive processes such as thought is progressively being viewed as more viable. A former Google engineer made news by claiming that the LaMDA (Thoppilan et al., 2022) language model was sentient (Tiku, 2022). Whilst dismissed by the community at large, this presented one of the first notable moments where an in-domain practitioners with expertise made positive claims regarding the sentience of computational models. Previously, such claims were largely limited to non-experts whose beliefs about sentient machines may have been dismissed as a lack of understanding (Selwyn and Gallo Cordoba, 2022).

Implications for the consciousness of such technologies includes needing to ensure that such models are not subject to undue suffering and other protections analogous to human-rights (Metzinger, 2021; Schwitzgebel and Garza, 2020; Laukyte, 2019). Furthermore, the consciousness of such models has further potential legal implications pertaining to the allocation of blame and responsibility for the outputs of such technologies such as harmful language (Veluwenkamp, 2025; Noh, 2025).

Survey respondents were queried on whether AI models are, or will become, capable of thought and consciousness, and the consequent real-world implications regarding legal protections for AI analogous to human rights. The results in Figure 4 demonstrate that by and large, both practitioners and the public dismiss the potential sentience of such models, with approximately half of all respondents strongly rejecting both claims. However, only 42% of practitioners strongly rejected the claim that AI models will ever require human-rights style protections, whilst 17% and 21% gave a weaker rejection or were neutral on the topic, respectively. Furthermore, for both questions and for both practitioners and the public, we see at least some respondents giving positive responses, demonstrating that the belief in artificial consciousness and the subsequent need for legal protections is not an entirely fringe belief.




Figure 4: Practitioner and public beliefs concerning the ability for AI systems to be capable of thought and consciousness, and the likelihood of requiring legal protections.

8 Societal Impact

In this final section, we investigate the aspects of society that people believe are impacted by AI in a net positive or negative way. Figure 5 presents public and practitioner perspectives on the overall impact of AI systems on society at large, both now and in the future. The results show that the public have a much more negative view of the societal impact of AI, with a strong majority of respondents having negative views on the impact of AI at the time of response (83%), and longer term forecasts (77%), versus less than half of practitioners (46% and 51%, for the present and future, respectively).




Figure 5: Practitioner and public beliefs concerning the present and future societal impact of AI.

Figure 6 presents a breakdown of the specific domains and industries where AI is considered to be a net positive or negative. Most strikingly, 56% of public respondents and 21% of practitioners believe that AI provides no net benefit to any of the presented areas. In contrast, 3% of public respondents take the optimistic stance that none of the proposed areas are impacted by AI in a net negative way. Whilst raw percentages differ, the results for both net positive and net negative areas are comparable between both practitioners and the general public, demonstrating that specific areas of concern are shared across the surveyed population.

Of the most positively viewed domains, disagreement between the public and practitioners is highest for research, healthcare, education, entertainment, and creativity - with an approximately 2-4x higher proportion of practitioners presenting positive views of the potential impact of AI in these domains than public respondents.




Figure 6: Practitioner and public beliefs concerning the subdisciplines of AI that present overall positive or negative impacts on society.

Overall, the finding that half of surveyed practitioners view AI’s societal impact to be net negative is concerning. One interpretation of this finding is that practitioners with this view may wish to remain involved in the field in order to help direct AI development towards positive societal utility, even if they view the domain as a whole as trending negatively. In the following section, we elaborate on some of the key areas of societal impact:

8.1 Accessibility

AI technologies can significantly lower barriers to partaking in activities that have traditionally required

specialised skill sets and knowledge (Costa et al., 2024; Dessimoz and Thomas, 2024). Whilst the dawn of the internet presented a positive step in this direction, much specialised knowledge remained complex and difficult to understand. Consequently, the use of technologies such as large language models to summarise and explain concepts in a more accessible, simplified manner further reduces barriers. Barriers to other skills have likewise progressively been reduced, such as being able to create music or animations. However, reducing the skill level required to partake in activities has been shown to reduce a skill’s perceived value, therefore garnering negativity from those who learnt with less assistance (Kruger et al., 2004). Additionally, the over-reliance on AI aids may generally reduce the ability levels of a population over time (Gerlich, 2025; Macnamara et al., 2024). Alongside lowering the overall barrier to entry for the general public, AI technologies also provide assistance to those with disabilities, for whom access to a given activity may have otherwise been impossible, rather than just requiring more effort (Brotosaputro et al., 2024).

8.2 Research

The use of AI in academic research has enabled the significant scaling of experiments, resulting in a faster rate of progress (Nolte and Tomforde, 2025; Gottweis et al., 2025; Lu et al., 2024). However, in specific research domains, the use of AI has increasingly presented challenges. For instance, the pre-print platform arXiv decided to halt the acceptance of survey papers and opinion-pieces due to an influx of low-quality AI generated works (Boboris, 2025), whilst many academic venues report the use of AI by reviewers, usually without disclosure. This results in a pipeline where research works have their aims and hypotheses generated by AI, the methodology implemented with AI, and the results written into papers by AI, to be later reviewed by AI (Gottweis et al., 2025; Lu et al., 2024). This is compounded by the volume of research works published in fields such as machine learning and natural sciences (van der Zwaard et al., 2020; Pontis et al., 2017; Bornmann and Mutz, 2015), meaning that a significant number of these works go untouched by human eyes throughout their entire life cycle.

8.3 Healthcare

The survey responses show that the impact of AI on healthcare is considered positive by many people. For areas such as genome sequencing, early screening of life threatening conditions, drug discovery, and other big-data-reliant healthcare fields, AI and machine

learning are integral to meaningful, timely progress (Bajwa et al., 2021; Zhavoronkov et al., 2019). Additionally, the availability of AI doctors may help to alleviate strains on medical systems, including in less economically developed nations where medical expertise may not be readily available at the scale required for the population (Yousef and Schmollgruber, 2024). However, negative sentiment arises when considering the use of AI for final diagnoses and decision-making roles, such as determining the likelihood of survival for someone on life support, without human intervention being required for the final decision.

8.4 Wellbeing

In terms of mental health and social wellbeing, there has been an increasing number of individuals turning to AI assistance for therapy and counselling services, owing to the relative affordability and the perceived anonymity of conversing with a language model (Cox et al., 2025; Lee et al., 2020). However, most implementations of AI for this task rely on third party hosted instances of proprietary models, where sensitive data is shared with international corporations, rather than self-hosted open-source implementations that require substantial financial and educational investments to use. Additionally, model debiasing and detoxification are not solved issues, resulting in instances where models have prompted vulnerable users to either make unsupported medical decisions, or take their own lives whilst in depressive episodes (Schoene and Canca, 2025; Judd et al., 2025; Arnaiz-Rodriguez et al., 2025).

8.5 Information Integrity

As the most negatively viewed impact of AI on society in our survey, information integrity refers to the reliability and trustworthiness of outputs presented by such technologies. Technologies such as large language models are able to facilitate bad actors in the creation and propagation of fake or misleading information at an unprecedented scale, largely going undetected by users (Williams et al., 2025; Radivojevic et al., 2024a; Chen and Shu, 2024). Such materials exist both as web pages and as bots acting as real social media users to promote anger and civil unrest (Radivojevic et al., 2024b). Further adding to the concern, whilst image and video evidence of an event was once considered almost infallible, developments in the automatic generation of such mediums has slowly eroded the utility of such evidence (Iannuzzi et al., 2025; Pataranutaporn et al., 2025; Maras and Alexandrou, 2019). Combined with the growing distrust of governments and news

media, false information is able to propagate faster than ever (Garry et al., 2024; Prats et al., 2024; Mittelstadt et al., 2023; Ognyanova et al., 2020). Furthermore, even within the scope of expected use, hallucination remains a prevalent concern, casting doubt on the factual accuracy of model outputs, requiring additional verification that many users skip. As such problems worsen and more content becomes model-generated, progress towards realising the dead internet theory increases (Walter, 2025; Muzumdar et al., 2025).

9 Conclusion

This work contributed further insights into the differences of opinion held by the individuals responsible for the development of artificial intelligence technologies (i.e., practitioners) and the members of public that constitute the majority of people impacted by such developments. Through a survey, we highlight key differences between both parties, particularly concerning opinions on copyright responsibilities, societal impact, and the philosophical concerns of art, creativity, and consciousness. Our results support the findings of prior work, empirically demonstrating that practitioners have more optimistic views toward AI technologies than the general public. We argue that such differences in opinion between practitioners and the public are not irreconcilable, but highlight the need for the public to become more well-informed about the nature of AI technologies, as well as the need for practitioners to engage more actively with stakeholders outside of the technology industry and corporate interests to better understand the perspectives of society at large.

Limitations

We performed a survey of opinions on the development of artificial intelligence from the perspectives of practitioners and the general public. However, we were only able to capture a small subset of opinions from each of these parties. Owing to logistics and ethics approval procedures, we were unable to collect responses from individuals in rural communities across the globe with limited access to technology, in addition to the young (who were unable to participate in this work due to the 18+ age restriction), and the elderly. Due to the survey being presented only in English, we were additionally not able to receive responses from individuals who are not fluent in English. Whilst we did not enforce a geographical boundary on responses to the survey, the primary demographic represented in this work is that of Millennial Europeans, rather than a stratified sample of

the world's true population for representativeness. We highly encourage research groups globally to perform further work in this vain (as well as longitudinal analyses of how perspectives evolve over time as the technology advances). Furthermore, the largest field from which public respondents were recruited from is technology (excl. ML/AI), constituting 30% of respondents. However, no statistically significant difference was found between members of the public working in technology and those not.

Acknowledgments

Tyler Loakman was supported by the Centre for Doctoral Training in Speech and Language Technologies (SLT) and their Applications funded by UK Research and Innovation [grant number EP/S023062/1].

References

- Advance Local Media LLC v. Cohere. 2025. *Advance Local Media LLC et al. v. Cohere Inc.*, no. 1:25-cv-01305, (s.d.n.y. feb. 13, 2025). United States District Court for the Southern District of New York (S.D.N.Y.).
- Amazon Employees for Climate Justice. 2025. *AI Open Letter: We're the Workers Who Develop, Train, and Use AI, So We Have a Responsibility to Intervene*.
- Andersen v. Stability AI. 2024. *Andersen et al. v. Stability AI Ltd. et al.*, no. 3:23-cv-00201 (n.d. cal. jan. 13, 2023). United States District Court for the Northern District of California.
- Dawit Mureja Argaw, Mattia Soldan, Alejandro Pardo, Chen Zhao, Fabian Caba Heilbron, Joon Son Chung, and Bernard Ghanem. 2024. *Towards automated movie trailer generation*. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7445–7454.
- Adrian Arnaiz-Rodriguez, Miguel Baidal, Erik Derner, Jenn Layton Annable, Mark Ball, Mark Ince, Elvira Perez Vallejos, and Nuria Oliver. 2025. *Between help and harm: An evaluation of mental health crisis handling by llms*.
- Authors Guild Inc. v. Google Inc. 2015. *Authors guild, inc. v. google inc.*, no. 13-4829-cv, (2d cir. oct. 16, 2015). United States Court of Appeals for the Second Circuit. Decided Oct. 16, 2015.
- Junaid Bajwa, Usman Munir, Aditya Nori, and Bryan Williams. 2021. *Artificial intelligence in healthcare: transforming the practice of medicine*. *Future Healthcare Journal*, 8(2):e188–e194.
- Daphne Barretto, Julianne LaChance, Emanuelle Burton, and Soohyun Nam Liao. 2021. *Exploring why under-represented students are less likely to study machine learning and artificial intelligence*. In *Proceedings of the 26th ACM Conference on Innovation and Technology in Computer Science Education V. 1, ITiCSE '21*, page 457–463, New York, NY, USA. Association for Computing Machinery.
- Bartz v. Anthropic PBC. 2025. *Bartz v. anthropic pbc.*, no. 3:24-cv-05417-wha, (n.d.c. jun. 6, 2025). United States District Court for the Northern District of California (N.D.C.).
- Lucas Bellaiche, Rohin Shahi, Martin Harry Turpin, Anya Ragnhildstveit, Shawn Sprockett, Nathaniel Barr, Alexander Christensen, and Paul Seli. 2023. *Humans versus ai: whether and why we prefer human-created compared to ai-created artwork*. *Cognitive Research: Principles and Implications*, 8(1):42.
- Kat Boboris. 2025. *Attention authors: Updated practice for review articles and position papers in arxiv cs category*.
- Matthew Bone, Eugenia González Ehlinger, and Fabian Stephany. 2025. *Skills or degree? the rise of skill-based hiring for ai and green jobs*. *Technological Forecasting and Social Change*, 214:124042.
- Lutz Bornmann and Rüdiger Mutz. 2015. *Growth rates of modern science: A bibliometric analysis based on the number of publications and cited references*. *Journal of the Association for Information Science and Technology*, 66(11):2215–2222.
- Philipp Brauner, Alexander Hick, Ralf Philipsen, and Martina Ziefle. 2023. *What does the public think about artificial intelligence?—a criticality map to understand bias in the public perception of ai*. *Frontiers in Computer Science*, Volume 5 - 2023.
- Goenawan Brotosaputro, Agung Supriyadi, and Michael Jones. 2024. *Ai-powered assistive technologies for improved accessibility*. *International Transactions on Artificial Intelligence (ITALIC)*, 3(1):76–84.
- Dashiel Carrera, Jeb Thomas-Mitchell, and Daniel Wigdor. 2026. *Where do i 'add the egg'? Exploring agency and ownership in ai creative co-writing systems*. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI '26, New York, NY, USA. Association for Computing Machinery.
- Canyu Chen and Kai Shu. 2024. *Can LLM-generated misinformation be detected?* In *The Twelfth International Conference on Learning Representations*.
- Common Crawl Foundation. 2025. *Common crawl dataset*. <https://commoncrawl.org/>.
- Susan Corbett. 2019. *Computer game licences: The eula and its discontents*. *Computer Law & Security Review*, 35(4):453–461.
- Carlos J. Costa, Manuela Aparicio, Sofia Aparicio, and Joao Tiago Aparicio. 2024. *The democratization of artificial intelligence: Theoretical framework*. *Applied Sciences*, 14(18).

- Samuel Rhys Cox, Rune Moberg Jacobsen, and Niels van Berkel. 2025. [The impact of a chatbot’s ephemerality-framing on self-disclosure perceptions](#). In *Proceedings of the 7th ACM Conference on Conversational User Interfaces*, CUI ’25, New York, NY, USA. Association for Computing Machinery.
- Tuba Degirmenci, Frank Mathmann, and Gary Mortimer. 2025. [Made with love: Examining consumer engagement with handmade versus machine-made production cues](#). *Journal of Consumer Behaviour*, 24(2):982–1002.
- Christophe Dessimoz and Paul D. Thomas. 2024. [Ai and the democratization of knowledge](#). *Scientific Data*, 11(1):268.
- Doe v. Github, Inc. 2022. *Doe et al. v. Github, Inc. et al.*, no. 4:22-cv-06823-jst (n.d. cal. jan. 13, 2023). United States District Court for the Northern District of California.
- Gayle T. Dow. 2017. [Defining creativity](#)., Creativity and innovation: Theory, research, and practice., pages 5–21. Prufrock Press Inc., Waco, TX, US.
- Samer Elhajjar and Omar S. Itani. 2025. [Ai washing: A conceptual exploration](#). *AMS Review*.
- Allyson Ettinger, David Heineman, Amanda Bertsch, Bailey Kuehl, Dirk Groeneveld, Hamish Ivison, Jacob Morrison, Matt Jordan, Mayee Chen, Pete Walsh, Pradeep Dasigi, Scott Geng, Shane Arora, Tyler Murray, Tyler Romero, Faeze Brahman, Jake Poznanski, David Graham, Finbarr Timbers, Kyle Lo, Luca Soldaini, Michael Noukhovitch, Nathan Lambert, Robert Berry, Shashank Gupta, Victoria Graf, Saumya Malik, Saurabh Shah, Taira Anderson, Teng Xiao, Akari Asai, Akshita Bhagia, Alex Wettig, Alisa Liu, Chloe Anastasiades, Costa Huang, Dustin Schwenk, Aman Rangapur, Harsh Trivedi, Ian Magnusson, Jaron Lochner, Jiacheng Liu, Lj Miranda, Maarten Sap, Malia Morgan, Michael Schmitz, Michal Guerquin, Michael Wilson, Regan Huff, Ronan Le Bras, Rui Xin, Rulin Shao, Sam Skjonsberg, Shannon Zejiang, Tucker Shen, Valentina Pyatkin, Will Merrill, Yapei Chang, Yuling Gu, Ashish Sabharwal, Luke Zettlemoyer, Pang Wei Koh, Ali Farhadi, Noah A. Smith, Hannaneh Hajishirzi, Shuyue Stella Li, Zhiyuan Zeng, and the Olmo Team. 2025. [Olmo 3](#). Technical report, Allen Institute for AI (AI2).
- Xiaoming Fan, Anqi Lai, and Hean Tat Keh. 2024. [Handmade vs. machine-made: the effects of handmade gifts on social relationships](#). *Marketing Letters*, 35(3):1–14.
- Christoph Fuchs, Martin Schreier, and Stijn M.J. Van Osselaer. 2015. [The handmade effect: What’s love got to do with it?](#) *Journal of Marketing*, 79(2):98–110.
- Maryanne Garry, Way Ming Chan, Jeffrey Foster, and Linda A Henkel. 2024. Large language models (LLMs) and the institutionalization of misinformation. *Trends in Cognitive Science*, 28(12):1078–1088.
- Michael Gerlich. 2025. [Ai tools in society: Impacts on cognitive offloading and the future of critical thinking](#). *Societies*, 15(1).
- Akash Ghosh, Aparna Garimella, Pritika Ramu, Sambaran Bandyopadhyay, and Sriparna Saha. 2025. [Infogen: Generating complex statistical infographics from documents](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 20552–20570, Vienna, Austria. Association for Computational Linguistics.
- Hugo Gonalo Oliveira. 2017. [A survey on intelligent poetry generation: Languages, features, techniques, reutilisation and evaluation](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 11–20, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Anil Palepu, Petar Sirkovic, Artiom Myaskovsky, Felix Weissenberger, Keran Rong, Ryutaro Tanno, Khaled Saab, Dan Popovici, Jacob Blum, Fan Zhang, Katherine Chou, Avinatan Hassidim, Burak Gokturk, Amin Vahdat, Pushmeet Kohli, Yossi Matias, Andrew Carroll, Kavita Kulkarni, Nenad Tomasev, Yuan Guan, Vikram Dhillon, Eeshit Dhaval Vaishnav, Byron Lee, Tiago R D Costa, Jos   R Penad  s, Gary Peltz, Yunhan Xu, Annalisa Pawlosky, Alan Karthikesalingam, and Vivek Natarajan. 2025. [Towards an ai co-scientist](#).
- Dirk Groeneveld, Iz Beltagy, Evan Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, William Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah Smith, and Hannaneh Hajishirzi. 2024. [OLMo: Accelerating the science of language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15789–15809, Bangkok, Thailand. Association for Computational Linguistics.
- Jieun Han, Haneul Yoo, Junho Myung, Minsun Kim, Hyunseung Lim, Yoonsu Kim, Tak Yeon Lee, Hwajung Hong, Juho Kim, So-Yeon Ahn, and Alice Oh. 2024. [LLM-as-a-tutor in EFL writing education: Focusing on evaluation of student-LLM interaction](#). In *Proceedings of the 1st Workshop on Customizable NLP: Progress and Challenges in Customizing NLP for a Domain, Application, Group, or Individual (CustomNLP4U)*, pages 284–293, Miami, Florida, USA. Association for Computational Linguistics.
- Shaif Hemraj. 2025. [Ai and the future of creative development: redefining digital media production](#). *AI and Ethics*, 5(5):5105–5119.
- Isabella Hermann. 2023. [Artificial intelligence in fiction: between narratives and metaphors](#). *AI & SOCIETY*, 38(1):319–329.

- Zoë Hitzig, Aaron Chatterji, Tom Cunningham, David Deming, Christopher Ong, Carl Shan, and Kevin Wadman. 2025. How people use chatgpt. *OpenAI (affiliated with Duke University and Harvard University)*. Dated September 15, 2025.
- Jeff Horwitz. 2025. [Meta’s AI rules have let bots hold ‘sensual’ chats with kids, offer false medical info.](#) *Reuters Special Report*.
- Andrew Dana Hudson, Ed Finn, and Ruth Wylie. 2023. [What can science fiction tell us about the future of artificial intelligence policy?](#) *AI & SOCIETY*, 38(1):197–211.
- Luke Hutton and Tristan Henderson. 2017. *Beyond the EULA: Improving Consent for Data Mining*, pages 147–167. Springer International Publishing, Cham.
- Micol Iannuzzi, Glynis Bogaard, and Jennifer M. Schell-Leugers. 2025. [Ai-generated media and the erosion of trust in legal settings: a review on impact and distortion of belief.](#) *AI & SOCIETY*.
- Nick Judd, Alexandre Vaz, Kevin Paeth, Layla Inés Davis, Milena Esherrick, Jason Brand, Inês Amaro, and Tony Rousmaniere. 2025. [Independent clinical evaluation of general-purpose llm responses to signals of suicide risk.](#)
- Masahiro Kaneko, Danushka Bollegala, and Naoaki Okazaki. 2023. [The impact of debiasing on the performance of language models in downstream tasks is underestimated.](#) In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 29–36, Nusa Dua, Bali. Association for Computational Linguistics.
- Aimee Kendall Roundtree. 2024. [Public perception of ai: A review.](#) In *HCI International 2024 – Late Breaking Papers*, pages 72–87, Cham. Springer Nature Switzerland.
- Justin Kruger, Derrick Wirtz, Leaf Van Boven, and Tood Altermatt. 2004. The effort heuristic. *Journal of Experimental Social Psychology*, 40(1):91–98.
- Vikas Kumar and Manju Lata. 2025. *Chatbots and Beyond: Exploring the Future of Conversational Technology in Education*. John Wiley & Sons, Ltd.
- Migle Laukyte. 2019. [Ai as a legal person.](#) In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Law, ICAIL ’19*, page 209–213, New York, NY, USA. Association for Computing Machinery.
- Yi-Chieh Lee, Naomi Yamashita, Yun Huang, and Wai Fu. 2020. ["i hear you, i feel you": Encouraging deep self-disclosure through a chatbot.](#) In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI ’20*, page 1–12, New York, NY, USA. Association for Computing Machinery.
- Wai Yie Leong. 2025. [Ethical and philosophical perspectives on artificial intelligence-generated art.](#) *International Journal of Social Sciences and Artistic Innovations*, 5.
- Tyler Loakman, William Thorne, and Chenghua Lin. 2025. [Who’s laughing now? an overview of computational humour generation and explanation.](#) In *Proceedings of the 18th International Natural Language Generation Conference*, pages 780–794, Hanoi, Vietnam. Association for Computational Linguistics.
- Chris Lu, Cong Lu, Robert Tjarko Lange, Jakob Foerster, Jeff Clune, and David Ha. 2024. [The AI Scientist: Towards fully automated open-ended scientific discovery.](#) *arXiv preprint arXiv:2408.06292*.
- Zheheng Luo, Chenhan Yuan, Qianqian Xie, and Sophia Ananiadou. 2025. [EMPEC: A comprehensive benchmark for evaluating large language models across diverse healthcare professions.](#) In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 9945–9958, Vienna, Austria. Association for Computational Linguistics.
- Christopher W. Lynn and Danielle S. Bassett. 2020. [How humans learn and represent networks.](#) *Proceedings of the National Academy of Sciences*, 117(47):29407–29415.
- Brooke N. Macnamara, Ibrahim Berber, M. Cenk Çavuşoğlu, Elizabeth A. Krupinski, Naren Nallapareddy, Noelle E. Nelson, Philip J. Smith, Amy L. Wilson-Delfosse, and Soumya Ray. 2024. [Does using artificial intelligence assistance accelerate skill decay and hinder skill development without performers’ awareness?](#) *Cognitive Research: Principles and Implications*, 9(1):46.
- Marie-Helen Maras and Alex Alexandrou. 2019. [Determining authenticity of video evidence in the age of artificial intelligence and in the wake of deepfake videos.](#) *The International Journal of Evidence & Proof*, 23(3):255–262.
- Nicholas Meade, Elinor Poole-Dayana, and Siva Reddy. 2022. [An empirical survey of the effectiveness of debiasing techniques for pre-trained language models.](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1878–1898, Dublin, Ireland. Association for Computational Linguistics.
- Thomas Metzinger. 2021. [Artificial suffering: An argument for a global moratorium on synthetic phenomenology.](#) *Journal of Artificial Intelligence and Consciousness*, 08(01):43–66.
- Brent Mittelstadt, Sandra Wachter, and Chris Russell. 2023. [To protect science, we must use llms as zero-shot translators.](#) *Nature Human Behaviour*, 7(11):1830–1832.
- Ishani Mondal, Kabir Ahuja, Mohit Jain, Jacki O’Neill, Kalika Bali, and Monojit Choudhury. 2022. [Global readiness of language technology for healthcare: What would it take to combat the next pandemic?](#) In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 4320–4335, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

- Abhishek Muzumdar, Gaurav Kurian, Arjun Rawat, and Priyesh Bhardwaj. 2025. [The dead internet theory: A survey on artificial interactions and the future of social media](#). *Asian Journal of Research in Computer Science*, 18(1):67–73.
- Karim Nader, Paul Toprac, Suzanne Scott, and Samuel Baker. 2024. [Public understanding of artificial intelligence through entertainment media](#). *AI & SOCIETY*, 39(2):713–726.
- Hyungrae Noh. 2025. [Beyond the responsibility gap: Distributed non-anthropocentric responsibility in the ai era](#). *Topoi*.
- Lukas Nolte and Sven Tomforde. 2025. [A helping hand: A survey about ai-driven experimental design for accelerating scientific research](#). *Applied Sciences*, 15(9).
- Katherine Ognyanova, David Lazer, Ronald E. Robertson, and Christo Wilson. 2020. [Misinformation in action: Fake news exposure is linked to lower trust in media, higher trust in government when your side is in power](#). *The Harvard Kennedy School Misinformation Review*, 1:1–19.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, et al. 2023. [Gpt-4 technical report](#).
- Selcen Ozturkcan and Ayşe Aslı Bozdağ. 2025. [Responsible ai in marketing: Ai booing and ai washing cycle of ai mistrust](#). *International Journal of Market Research*, 67(6):696–722.
- Birong Pan, Yongqi Li, Weiyu Zhang, Wenpeng Lu, Mayi Xu, Shen Zhou, Yuanyuan Zhu, Ming Zhong, and Tiejun Qian. 2025. [A survey on training-free alignment of large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 4445–4461, Suzhou, China. Association for Computational Linguistics.
- Pat Pataranutaporn, Chayapatr Archiwanguprok, Samantha W. T. Chan, Elizabeth Loftus, and Pattie Maes. 2025. [Synthetic human memories: Ai-edited images and videos can implant false memories and distort recollection](#). In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI ’25, New York, NY, USA. Association for Computing Machinery.
- Pew Research Center. 2025. [How the U.S. Public and AI Experts View Artificial Intelligence](#). Report, Pew Research Center.
- Sheila Pontis, Ann Blandford, Elke Greifeneder, Hesham Attalla, and David Neal. 2017. [Keeping up to date: An academic researcher’s information journey](#). *Journal of the Association for Information Science and Technology*, 68(1):22–35.
- Marc Prats, Sandra Smid, and Monica Ferrin. 2024. [Lack of trust in institutions and political engagement: An analysis based on the 2021 oecd trust survey](#). *OECD Working Papers on Public Governance*, (75).
- PwC. 2025. [The fearless future: 2025 global ai jobs barometer](#). Report, PwC.
- Kristina Radivojevic, Matthew Chou, Karla Badillo-Urquiola, and Paul Brenner. 2024a. [Human perception of llm-generated text content in social media environments](#).
- Kristina Radivojevic, Nicholas Clark, and Paul R. Brenner. 2024b. [Llms among us: Generative ai participating in digital discourse](#). *ArXiv*, abs/2402.07940.
- Pascual Restrepo. 2025. [We Won’t Be Missed: Work and Growth in the AGI World](#). University of Chicago Press.
- Right to Warn Coalition. 2024. [A right to warn about advanced artificial intelligence](#).
- Eleonora Rosati. 2019. [Copyright as an obstacle or an enabler? a european perspective on text and data mining and its role in the development of ai creativity*](#). *Asia Pacific Law Review*, 27(2):198–217.
- Eleonora Rosati. 2025. [Copyright exceptions and fair use defences for ai training done for “research” and “learning,” or the inescapable licensing horizon](#). *European Journal of Risk Regulation*, 16(3):961–984.
- Mark A. Runco. 2023. [Ai can only produce artificial creativity](#). *Journal of Creativity*, 33(3):100063.
- Zainab Salma, Raquel Hijón-Neira, and Celeste Pizarro. 2025. [Designing co-creative systems: Five paradoxes in human–ai collaboration](#). *Information*, 16(10).
- Annika M Schoene and Cansu Canca. 2025. [‘for argument’s sake, show me how to harm myself’: Jailbreaking llms in suicide and self-harm contexts](#).
- Eric Schwitzgebel and Mara Garza. 2020. [Designing ai with rights, consciousness, self-respect, and freedom](#). In *Ethics of Artificial Intelligence*. Oxford University Press.
- Neil Selwyn. 2022. [The future of ai and education: Some cautionary notes](#). *European Journal of Education*, 57(4):620–631.
- Neil Selwyn and Beatriz Gallo Cordoba. 2022. [Australian public understandings of artificial intelligence](#). *AI & SOCIETY*, 37(4):1645–1662.
- Dong Shu, Haiyan Zhao, Jingyu Hu, Weiru Liu, Ali Payani, Lu Cheng, and Mengnan Du. 2025. [Large vision-language model alignment and misalignment: A survey through the lens of explainability](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 1713–1735, Suzhou, China. Association for Computational Linguistics.
- Luca Soldaini, Rodney Kinney, Akshita Bhagia, Dustin Schwenk, David Atkinson, Russell Authur, Ben Bogin, Khyathi Chandu, Jennifer Dumas, Yanai Elazar, Valentin Hofmann, Ananya Harsh Jha, Sachin Kumar, Li Lucy, Xinxu Lyu, Nathan Lambert, Ian Magnusson, Jacob Morrison, Niklas Muennighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Abhilasha Ravichander, Kyle Richardson, Zejiang Shen, Emma

- Strubell, Nishant Subramani, Oyvind Tafford, Pete Walsh, Luke Zettlemoyer, Noah A. Smith, Hannaneh Hajishirzi, Iz Beltagy, Dirk Groeneveld, Jesse Dodge, and Kyle Lo. 2024. [Dolma: An Open Corpus of Three Trillion Tokens for Language Model Pretraining Research](#). *arXiv preprint*.
- Stina Teilmann-Lock and Andrej Savin. 2025. [Beyond the ai-copyright wars: towards european dataset law?](#) *Computer Law & Security Review*, 58:106190.
- Maria Teleki, Vedangi Bengali, Xiangjue Dong, Sai Tejas Janjur, Haoran Liu, Tian Liu, Cong Wang, Ting Liu, Yin Zhang, Frank Shipman, and James Caverlee. 2025. [A survey on LLMs for story generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 13954–13966, Suzhou, China. Association for Computational Linguistics.
- Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, YaGuang Li, Hongrae Lee, Huaixiu Steven Zheng, Amin Ghafouri, Marcelo Menegali, Yanping Huang, Maxim Krikun, Dmitry Lepikhin, James Qin, Dehao Chen, Yuanzhong Xu, Zhifeng Chen, Adam Roberts, Maarten Bosma, Vincent Zhao, Yanqi Zhou, Chung-Ching Chang, Igor Krivokon, Will Rusch, Marc Pickett, Pranesh Srinivasan, Laichee Man, Kathleen Meier-Hellstern, Meredith Ringel Morris, Tulse Doshi, Renelito Delos Santos, Toju Duke, Johnny Soraker, Ben Zevenbergen, Vinodkumar Prabhakaran, Mark Diaz, Ben Hutchinson, Kristen Olson, Alejandra Molina, Erin Hoffman-John, Josh Lee, Lora Aroyo, Ravi Rajakumar, Alena Butryna, Matthew Lamm, Viktoriya Kuzmina, Joe Fenton, Aaron Cohen, Rachel Bernstein, Ray Kurzweil, Blaise Aguerre-Arcas, Claire Cui, Marian Croak, Ed Chi, and Quoc Le. 2022. [Lamda: Language models for dialog applications](#).
- Nitasha Tiku. 2022. [The google engineer who thinks the company’s ai has come to life](#). *The Washington Post*.
- Stephan van der Zwaard, Arie-Willem de Leeuw, L. (Rens) A. Meerhoff, Sue C. Bodine, and Arno Knobbe. 2020. [Articles with impact: insights into 10 years of research with machine learning](#). *Journal of Applied Physiology*, 129(4):967–979. PMID: 32790596.
- Herman Veluwenkamp. 2025. [What responsibility gaps are and what they should be](#). *Ethics and Information Technology*, 27(1):14.
- Pete Walsh, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Shane Arora, Akshita Bhagia, Yuling Gu, Shengyi Huang, Matt Jordan, Nathan Lambert, Dustin Schwenk, Oyvind Tafford, Taira Anderson, David Atkinson, Faeze Brahman, Christopher Clark, Pradeep Dasigi, Nouha Dziri, Allyson Ettinger, Michal Guerquin, David Heineman, Hamish Ivison, Pang Wei Koh, Jiacheng Liu, Saumya Malik, William Merrill, Lester James V. Miranda, Jacob Morrison, Tyler Murray, Crystal Nam, Jake Poznanski, Valentina Pyatkin, Aman Rangapur, Michael Schmitz, Sam Skjonsberg, David Wadden, Christopher Wilhelm, Michael Wilson, Luke Zettlemoyer, Ali Farhadi, Noah A. Smith, and Hannaneh Hajishirzi. 2025. [2 olmo 2 furious](#).
- Yoshija Walter. 2025. [Artificial influencers and the dead internet theory](#). *AI & SOCIETY*, 40(1):239–240.
- Maurice Weber, Daniel Y. Fu, Quentin Anthony, Yonatan Oren, Shane Adams, Anton Alexandrov, Xiaozhong Lyu, Huu Nguyen, Xiaozhe Yao, Virginia Adams, Ben Athiwaratkun, Rahul Chalamala, Kezhen Chen, Max Ryabinin, Tri Dao, Percy Liang, Christopher Ré, Irina Rish, and Ce Zhang. 2024. [Redpajama: an open dataset for training large language models](#). *NeurIPS Datasets and Benchmarks Track*.
- Angus R Williams, Liam Burke-Moore, Ryan Sze-Yin Chan, Florence E. Enock, Federico Nanni, Tvesha Sippy, Yi-Ling Chung, Evelina Gabasova, Kobi Hackenburg, and Jonathan Bright. 2025. [Large language models can consistently generate high-quality content for election disinformation operations](#). *PLOS ONE*, 20(3):1–29.
- Chih-Chiang Wu and Wei-Peng Chen. 2022. [What’s an ai name worth? the impact of ai etfs on their underlying stocks](#). *Finance Research Letters*, 46:102474.
- Xupeng Yao, Wenxiao Du, Li Sun, and Bin Hu. 2025. [Automatic trailer generation for movies using convolutional neural network](#). *Scientific Reports*, 15(1):7819.
- Khalil Yousef and Shelley Schmollgruber. 2024. [Artificial intelligence in low- and middle-income countries: Reducing the gaps in health care, research, and education](#). *International Journal of Critical Care*, 18(2):1–3.
- Chunhong Yuan, Jingyi Tang, Yi Ding Cao, Tianshi Wei, and Wei Tao Shen. 2024. [The impact of artificial intelligence on economic development: A systematic review: The impact of artificial intelligence on economic development](#). *International Theory and Practice in Humanities and Social Sciences*, 1(1):130–143.
- Alex Zhavoronkov, Yan A Ivanenkov, Alex Aliper, Mark S Veselov, Vladimir A Aladinskiy, Anastasiya V Aladinskaya, Victor A Terentiev, Daniil A Polykovskiy, Maksim D Kuznetsov, Arip Asadulaev, Yuri Volkov, Artem Zhulus, Rim R Shayakhmetov, Alexander Zhebrak, Lidiya I Minaeva, Bogdan A Zagribelnyy, Lennart H Lee, Richard Soll, David Madge, Li Xing, Tao Guo, and Alán Aspuru-Guzik. 2019. [Deep learning enables rapid identification of potent DDR1kinase inhibitors](#). *Nature Biotechnology*, 37(9):1038–1040.
- Tan Zhi-Xuan, Micah Carroll, Matija Franklin, and Hal Ashton. 2025. [Beyond preferences in ai alignment](#). *Philosophical Studies*, 182(7):1813–1863.
- Sujia Zhu, Guodao Sun, Qi Jiang, Meng Zha, and Ronghua Liang. 2020. [A survey on automatic infographics and visualization recommendations](#). *Visual Informatics*, 4(3):24–40.

A Appendix

A.1 Comparison to Existing Work

In the following section, we compare our findings to those reported in [Pew Research Center \(2025\)](#), where

comparable lines of questioning are available. Regarding long term impact of AI (Figure 5) we observe less optimistic results, with only 33% of practitioners and 10% of the public having favourable views, versus 17% of the public and 56% of the experts surveyed by the Pew Research Center. Furthermore, regarding specific areas of benefit (Figure 6), we observe no specific domain in which the majority of practitioners or the public believe AI will present an overall benefit. In contrast, [Pew Research Center \(2025\)](#) observe majority positive perspectives towards AI's effect on the economy and medical care by practitioners, whilst the public remain mostly pessimistic. However, our results echo the overall finding that practitioners are almost universally more optimistic than the public. However, regarding AI concerns, both the present work and [Pew Research Center \(2025\)](#) demonstrate that inaccurate information (i.e., information integrity) is one of the most pressing concerns.

A.2 Original Phrasing & Statistical Testing

The exact question phrasing as presented to participants is presented in [Table 1](#), alongside the results of statistical significance tests regarding differences between practitioner and public responses.

A.3 Additional Demographic Information

Additional information regarding the specific fields that public respondents came from is presented in [Figure 7](#).

Figure	Original Question	Practitioner	Public	<i>U</i>	<i>p</i> _{BH}	<i>r</i>
-	How familiar are you with the general idea of AI and ML and their applications?	-	-	-	-	-
-	How familiar are you with technical aspects of AI/ML (e.g., neural networks, deep learning)?	-	-	-	-	-
1	Which AI-powered tools or applications have you used in your own free time?	-	-	-	-	-
1	Which AI-powered tools or applications have you used to aid in your work/employment?	-	-	-	-	-
1	Was using the AI-powered tools selected above in your work/employment mandated/recommended by your employer?	-	-	-	-	-
2	Do you think that AI/ML companies should be able to train models on publicly available data (e.g., social media posts, PDFs, YouTube videos, etc.) without explicit informed consent?	2.21	1.39	10,205.0	< 0.001	0.267
3	Do you believe that AI is capable of generating art?	2.49	1.61	10,653.5	< 0.001	0.322
3	Do you think that AI is capable of thinking and being conscious?	2.22	1.44	10,407.5	< 0.001	0.292
4	Do you believe that AI/ML systems are capable of being creative?	2.53	1.50	11,738.0	< 0.001	0.457
4	Do you think that an AI system will ever need to be granted protections similar to human-rights?	2.29	1.75	9,927.5	< 0.001	0.232
5	In its current state, how do you perceive the impact of AI on society?	2.54	1.60	11,400.0	< 0.001	0.415
5	How do you perceive the impact of AI on society to change long-term (e.g., the next 10 years)?	2.57	1.77	10,628.0	< 0.001	0.319
6	In which areas do you believe AI provides an overall positive impact?	-	-	-	-	-
6	In which areas do you believe AI provides an overall negative impact?	-	-	-	-	-

Table 1: Original questions as posed to survey respondents. Two-sided Mann-Whitney *U* Test results comparing practitioner and public responses are provided, alongside mean Likert scores and absolute rank-biserial correlation (*r*) for effect size. *p*-values are adjusted using the Benjamini-Hochberg procedure across all 12 evaluation metrics. Practitioner *N* = 76, Public *N* = 212.




Figure 7: Breakdown of industries within which public respondents operate.