COMMENT

Knowing you know nothing in the age of generative AI

Sarah V. Bentley [1]✉

Generative AI is a revolutionary new technology whose impact promises to democratise knowledge. And yet, unlike the printing press, which expanded knowledge through the amplification of one voice to many, generative AI reduces many voices to one. Its disruptive nature provides us with a timely reminder of both the power and fallibility of knowledge: its authorship, ownership, and veracity. This Comment situates generative AI within the evolutionary context of human information dissemination and knowledge production. Whilst acknowledging the extraordinary potential of this new tool, it asks the question— given that knowledge is probably our most valuable asset, should we not be applying more of it to better understand the impact of AI-mediated knowledge tools on both our information practices and their associated knowledge outcomes?

[1] Data61, Commonwealth Scientific and Industrial Research Organisation (CSIRO), EcoScience Precinct, Dutton Park, QLD, Australia. ✉email: sarah.bentley@data61.csiro.au

Thanks to the public release of generative AI knowledge tools, such as OpenAI's ChatGPT or Google's Gemini, today anyone with a digital device and online connectivity has access to seemingly endless amounts of personalised knowledge. Rendered in a style, accent, or at a comprehension level of the consumer's choice, the purpose of generative AI is to provide informational responses to any enquiry posed. Since successfully launching in late 2022[1], the enthusiastic uptake of this new technology—evident anywhere from classrooms to workplaces, doctor's surgeries to newsrooms—demonstrates just how much this technology is (and will be) impacting the lifecycle of information production. Accepting that knowledge represents the human internalisation of information (Cress & Kimmerle, 2008; Symons & Johnson, 1997), the pervasiveness of this uptake would suggest that generative AI represents a new form of knowledge tool capable of fundamentally changing not only the production of knowledge, but also our understanding of its value (Birnbaum, 1985).

A brief history of knowledge: from bones to bytes

Early knowledge. The origin story of (human) knowledge began hundreds of thousands of years ago with the transmission of ideas, art, and culture via oral communication, whether in the form of song, dance, prayer, or folklore (Kelly, 2024). Written knowledge started as people began to draw or carve ideographic symbols, whether onto cave walls, bones, or clay tablets (Von Petzinger, 2017). As humans evolved to further value the concept and practice of externalising information—and what it then enabled through communication—shared writing structures began to emerge (Suenaga, 2023). These took the form of phonemic symbols representing different sounds, or other logographic structures, such as cuneiform writing, or indeed Chinese characters in which each symbol represents both sound and meaning (Norman, 2023).

These collective developments qualitatively changed human knowledge practices. They first allowed for the recording of simple transactions, thus enabling trade. But further expansions, such as the insertion of vowels into the Phoenician script by the Ancient Greeks; (considered the first alphabet; Woodard, 1997), allowed for more nuanced written creations. This alphabet spread quickly, its success put down to its simplicity, for unlike other writing structures that included hundreds of symbols (thereby necessitating intensive—and exclusive—training procedures), the Phoenician alphabet contained just 24 characters (Hock & Joseph, 2009). It was, in fact, with the vowel-enriched Phoenician alphabet that Homer's Odyssey was written. A Cumaean adaption of Ancient Greek eventually evolved to become the Latin alphabet known today, and from which, during the 16th century, the 26-symbol English alphabet then emerged.

Written knowledge to printed knowledge. In the century preceding the establishment of the English alphabet, written knowledge production shifted gears significantly with the invention of the printing press. Around 1440, in a small town in Germany, a local craftsman Johannes Gutenberg invented the first fast-moving durable type mechanical printing machine (Briggs & Burke, 2009; Füssel, 2020). The rapidity of its adoption was breath-taking (for the time). By the end of the 15th century the rhythmic thudding of printing presses could be heard in over 270 cities across Europe (Eisenstein, 2005). In fact, it has been estimated that from 1450 to 1500, the same quantity of books was printed as produced in the previous millennia by handwriting scribes (Dewar, 1998).

This machine allowed—for first time ever—the means to circulate knowledge at scale. Previous scribal culture had allowed only for a one-to-one distribution, meaning that knowledge transmission relied on a single person capturing information from another, and conveying it to someone else, either in the form of a manuscript or an oral communication (Dewar, 1998; Eisenstein, 2005). This resulted in a precarious preservation of knowledge, reliant on mnemonics, elitist structures, and the note-taking skills of 'wandering scribes' (Eisenstein, 2005). The printing press now provided the means for one voice to speak to many (Wallace & Green, 1996), and in so doing, this simple device went on to play an unarguably influential role in the significant societal changes that took place over the rest of the millennium.

By today's technological standards, the printing press is a simple, mechanical device through which metal shapes are used to uniformly press ink onto paper. And yet, the industrial-level production of information afforded by this machine democratised society through its role in allowing knowledge to passed to unprecedentedly large numbers of people (Zaret, 2000). As a result, the printing press is now considered pivotal in the disruption of Catholic power during the Middle Ages, the cultural Renaissance that took place as Europe then transitioned towards modernity, and relatedly, the ensuing scientific revolution (Cochrane & Chartier, 2014; Eisenstein, 2005). The shift away from elitist information systems began almost immediately following its invention. In 1476 William Caxton set up a printing press in London to capitalise on the growing demand for books written in the English language, as opposed to the previously dominant Latin texts (Blake, 1991; Morgan Library, 2015). As popular works such as Chaucer's Canterbury Tales spread, the rapid increase in literacy rates contributed to the restructuring of social hierarchies, and with it, the emergence of the 'middle class' (Tarkhnishvili & Tarkhnishvili, 2013). It has been argued that much as the printing press revolutionised society over hundreds of years by allowing for the sharing of human knowledge, the digital or 'information revolution'—which began less than 100 years ago—has equal (if not greater) revolutionary potential (Builder, 1993; Dewar, 1998).

Printed knowledge to digital knowledge. Over the last century, the amplification of knowledge has accelerated exponentially with the arrival of computers, digitalisation, and networking. This next step change in information production, particularly the interconnectivity of computers and users, has evolved the distribution of information from an already revolutionary one-to-many, to a new paradigm-shifting reality of many-to-many. Described as the defining characteristic of the information age, networked computers have radically altered the process, experience, and scope of information sharing (Dewar, 1998), and in so doing, changed not only the nature of knowledge acquisition, but also the production of knowledge itself.

The fundamental and emblematic manifestation of computer networking came in 1960s in the form of the internet. A global system of interconnected computers, the internet emerged thanks to developments in data communication and packet switching (simply speaking, the grouping of data into short, structured forms able to be transmitted across digital networks[2]). The communication protocols enabling functional internetworking came out of R&D collaborations between the US Defence Agency (DARPA) and universities in the US, the UK and France. The transition into the World Wide Web—involving commercial networks and enterprises alongside accelerated public engagement—took place later in the 1990s.

On April 30th 1993, the internet was released into the public domain (Dinmore, 2023; Ring, 2023). All a user had to do was launch a new programme called a 'browser' and then type in an

electronic address, known as a URL (Uniform Resource Locator). More importantly, anyone (leaving aside important issues of digital exclusion; Shaban, 2025) could publish information themselves online. Whether text, photos, music, or videos, this new form of information dispersal could also include links to other people's information, literally accessible at the click of a button. By the end of 1995 more than 24 million people in the US and Canada alone spent on average 5 h per week perusing these new information sources. Today, 64% of the world's population regularly uses the internet (Statista, 2025).

In its early days, the internet was an unregulated, organic, proliferating mass of mostly crowd-sourced information. Within this new and emerging communication culture, an abundance of discussion groups was created. Known as USENETS, these have been described as "an enormous billowing crowd of gossipy, news-hungry people wandering in and through the Internet on their way to various private backyard barbecues" (p. 2, Sterling, 1993). The potential for almost instantaneous mass-scale information sharing was, therefore, enormous. Users could (and would) navigate through thousands of newsgroups, reading hundreds of thousands of articles on a variety of topics, and sharing their own thoughts and reactions.

This new information network changed almost every aspect of society. From hobbyists to political activists, the ability to connect with others across the globe on the basis of shared information and interest was, quite literally, revolutionary (Tudoroiu, 2014). Knowledge was now being crowd-sourced, impacting anything from travel choices to political mobilisations. But equally, the internet began to reshape more formal knowledge processes. Researchers no longer required long visits to libraries but instead could access previously unimaginable quantities of information through browsing online. Scientific collaborations subsequently increased, alongside a substantial acceleration in the dissemination of findings (Gholizadeh et al., 2014; Teasley & Wolinsky, 2001).

This information expansion changed not only the creation of new knowledge, but the validation of existing knowledge. Prior to the printing press, errors written into manuscripts (whether due to a tired scribe or more conscious informational distortions) remained in permanence (Dewar, 1998). The printing of books changed this. Creating a culture in which authorship was associated with ownership (Eisenstein, 2005), the printing revolution meant that facts could be debated, contested, and—when necessary—updated (albeit slowly). Today, the information age allows for facts, ideas, and perspectives to be both created and overturned in real-time (Chan et al., 2018; Flemming et al., 2017). Reflecting Thomas Jefferson's far-sighted description of the nature of ideas, the internet—at least in its early days—held the promise of infinite knowledge expansion: "That ideas should freely spread from one to another over the globe, for the moral and mutual instruction of man, and improvement of his condition, seems to have been peculiarly and benevolently designed by nature, when she made them, like fire, expansible over all space, without lessening their density at any point, and like the air in which we breathe, move, and have our physical being, incapable of confinement or exclusive appropriation" (Jefferson, 1813). Today, as we peruse our AI-generated online search results, the impact of this form of information mediation on the future of human knowledge is unclear.

Digital knowledge to AI-generated knowledge. Today, generative AI presents us with a powerful new way of accessing information and interacting with knowledge. The term 'generative AI' describes a particular use of Artificial Intelligence, largely based on machine learning and Natural Language

Processing (Zewe, 2023). The product (simply speaking) is a dynamic prediction algorithm, able to respond to questions or prompts with an uncanny human-like demeanour, whether the output is text or voice, visual or musical (Bandi et al., 2023). Generative AI does this by drawing from the vast range of information upon which it is trained, allowing it to compute the statistical probability that informational components (such as words or pixels) cluster together (Salakhutdinov, 2015). For statistical accuracy, this technology must, therefore, train on vast amounts of data, hundreds of billions of words in the case of human language tools (Hughes, 2023).

Therein lies the power of generative AI. Using the fastest processing power (not to mention large amounts of energy; Chien et al., 2023; Luccioni et al., 2024), this technology can—within the parameters of human-coded algorithms—train autonomously and, through recursive self-improvement, adaptively (Madaan et al., 2024). The fuel which drives its training is (at least at this point in time[3]) online sourced human-originated data. This includes books, websites, reports, poems, songs, and articles… ultimately anything digitised and available on the web, including text exuded from community platforms like Reddit, X, or Facebook. Despite the enormity of this data resource, however, generative AI models remain hungry, and current practice—although ethically debatable (Zhang & Yang, 2024)—is to use any and all information input by current human users as an ongoing source of 'live' training material (Fui-Hoon Nah et al., 2023; Lucchi, 2024; Porter, 2023).

This debatable practice encapsulates many of the ethical controversies surrounding generative AI. These include not only issues of Intellectual Property (IP) ownership but also concerns over the quality of the information upon which AI models are trained. Known levels of bias are present in the data, both selection bias (not all segments of society get to put their information online; Giorgi et al., 2022), but there is also substantial evidence of gendered, racial, and ideologically bias in generative AI's outputs, presumably due to these forms of social distortion being present in the original training data (Wach et al., 2023). This uncomfortable reality has required generative AI systems to be further re-trained, adjusted, and controlled—by humans (not algorithms)—to provide an end product that is, loosely speaking, acceptable for human re-consumption (Kirova et al., 2023; Rani & Dhir, 2024).

Despite these controversies, the uptake of generative AI technologies across society has been unprecedented. Marketed as everyone's new best friend, whether co-pilot, research assistant, or life coach, these new knowledge tools are currently free, available 24/7, answer pretty much any question you might care (or dare) to pose, and are (mostly) polite (Quan & Chen, 2024). Further, they are (largely) not domain-specific but can comprehensively answer questions on any subject, from eco-gastronomy to puppetry. Recent surveys are reporting that over one-quarter of American adults say they are using applications like ChatGPT (Motyl et al., 2024; Pew Research Center, 2024), with this figure increasing for younger users as well as people with higher levels of educational attainment (Park & Gelles-Watnick, 2023). 20% say they use it for entertainment, 19% for learning, and 16% to assist them with tasks such as summarising information. In the workplace, 75% of over 31,000 'knowledge workers' (from 31 different countries), reported using generative AI to help them with their jobs (Microsoft & LinkedIn, 2024). In the health domain, a survey of 1,000 British general medical practitioners, reported that 20% were using tools such as ChatGPT for notetaking and diagnostics (Blease et al., 2024). Uptake is also prolific in student populations. A report published by the Digital Education Council showed that 86% of tertiary students use AI tools to assist them with their learning (Digital Education

Council, 2024), and in Australia, a Youth Insight analysis showed that 70% of 14 to 17-year-olds were using generative AI for information sourcing, assignment writing, and general learning (Denejkina, 2023).

It is clear, therefore, that not only is this technology quite literally awe-inspiring, but so too is our uptake of it. Today, its revolutionary potential lies not just in the extraordinary reality of a machine communicating like a human (no small feat), nor the fact that this technology can successfully pass various tests of both intellectual and emotional intelligence (Kosinski, 2024; Sumbal et al., 2024), but more subtly yet potentially more impactfully, the fact that by blithely sidestepping established information-sharing practices, structures, and even laws, this tool has the potential to radically disrupt[4] not just the information we hold dear, but the very landscape in which we, as a global community, currently create, debate, and produce our most valuable commodity, that being knowledge.

AI-generated knowledge to disruptive knowledge. Despite the unprecedented uptake of generative AI knowledge tools, as of 2024, there is little peer-reviewed, validated, or replicated research informing us as to the effect that outsourcing our thinking to AI knowledge tools and systems may have on our (human) capacity to reflect, understand, and learn (not to mention feel, experience, or empathise; Smith et al., 2024; Yan et al., 2024). The closest findings we have relate to what we know about how human memory processes changed as a result of internet use. Referred to as 'the Google effect' (or digital amnesia), research has demonstrated that when faced with difficult questions, today people's minds jump first to a digital device for solutions (Sparrow et al., 2011), and that further, memory is reduced—both quantitively and qualitatively—the more time a person spends on their digital devices (Çinar et al., 2020; Robert et al., 2024). Alarming as this sounds, this *is* the nature of disruptive innovation.

AI is a disruptive technology, and *generative* AI in particular— a product promising to curate every human's knowledge needs— takes disruption to potentially unprecedented levels. In the light of the socio-technical dilemmas presented by the processes of disruption, there are those who advocate for the application of a 'precautionary principle' (rather like the old adage, better safe than sorry). However, there are also those who argue that such a principle suppresses innovation (Cross, 1996; Hellström, 2003; Jonas, 2009). Regardless of one's position on this debate, the question of what generative AI might be doing to our most valued commodity presents us with an interesting (and possibly long overdue) opportunity to reflect—at this timely moment in history —on the nature of human knowledge, and the need (or not) to apply caution.

A case study: Miscalibrated algorithms. Think back to one of 2024's most interesting generative AI misadventures. The diversity algorithms in a prominent image generator over-zealously re-calibrated the ethnicity ratio of its historical depictions. When asked to visualise a British medieval king, users were presented with images of a fine-looking African man dressed in fur robes and sporting a crown (Field, 2024). Further historical faux pas came in the form of female popes, black Vikings, and ethnically diverse American 'founding fathers'.

What generative AI inadvertently did was to expose the highly 'constructed' nature of knowledge, both its origin story (who, where and why) but also the subsequent representation of that knowledge as it is composited onto our digital screens. This particular form of AI, along with its limitations, biases, and hallucinations (Sun et al., 2024), serves us with a timely reminder of the fallible nature of knowledge (Emsley, 2023). For social

scientists, this may not be news. But for the world at large—the voting electorate, the summoned juror, or even the doomscroller —this is big news. And when it comes to the need to understand and tackle the reality of mass-scale mis- or disinformation (which significantly predates the arrival of generative AI), this is much-needed news (Posetti, 2018).

Despite the concerning implications of a privately funded, opaque, and certainly not peer-reviewed technology quite literally re-writing history, this socio-technical 'incident' presents us with an unprecedented and 'live' insight into the nature of knowledge fabrication. More specifically, knowledge artefacts are—thanks to generative AI— being revealed as potentially subjective, idiosyncratic, and financially tainted. This is not something for us to shy away from. For anyone remotely interested in knowledge practices, this technological misadventure provides the raw material for one of the most exciting, engaging, and illuminating exposés in recent times, and seems like a fitting finale for the already dramatic social trajectory of human knowledge production.

In fact, in an ironic twist on this futuristic technology, the arrival of generative AI in our homes, classrooms, and offices transports us right back to some of the most fundamental of debates waged 3000 years ago in ancient Greece. The Sophists were ancient Greek 'knowledge-makers' whose job it was to deliver expertise, for a fee (Säfström, 2023). Little evidence of their work exists today, and there remains debate as to their ultimate purpose, with some considering them experimental thinkers using speculative reason (Crick, 2010), and others portraying them as intellectually dishonest charlatans[5] (Duke, 2025). Like generative AI today, the Sophists claimed to democratise knowledge, but it was their commodification of information that ultimately became their downfall, revealing as it did, the malleable, serviceable, and fallacious potential of knowledge.

Disruptive knowledge to responsible knowledge. Today, with the dramatic entrance of generative AI into our midst, we find ourselves back at the intersection of knowledge, society, and economics. In order to ensure the beneficent impact of this novel and disruptive technology, we would do well to use it to spotlight important epistemological debates—where knowledge comes from, how it impacts thoughts, attitudes, and behaviours, and how we can trust it. Given the advanced sophistry involved in each and every individually AI-enhanced digitised production of our daily dose of news, information, or educative material, these questions are more important than ever (Datta et al., 2021; Rodilosso, 2024). Answering them will prove fundamental to our ability to use generative AI responsibly, equitably, and for the betterment of future generations.

So, how to do that? Strangely, turning back once more to Ancient Greece can provide us with some ideas. Socrates was the only Sophist who refused to charge for his services (Cooper, 1998; Taylor, 2006). The original advocate of open-sourced information, this Ancient Greek philosopher developed the 'Socratic method,' a technique designed to scrutinise commonly held beliefs and assumptions through rigorous questioning to determine their validity, veracity, and applicability (Benson, 2011). Applying Socratic scrutiny to AI-generated knowledge would most likely conclude with an unsettling awareness of what is either distorted, politely overlooked, or categorically silenced (Worrell & Johns, 2024). At the very least, it should point out that the power of yesterday's printing press to expand one voice to many, or indeed the power of digitisation to provide a 'many to many' information-sharing platform, is today being replaced by the power of generative AI to reduce 'many voices to one'. Given the urgent need for society to become its own discerning audience (think pandemics, wars, climate crises, and ever-widening global inequities) this in itself would represent a profoundly useful piece

of contemporary learning. Socrates is most famous for saying "The only true wisdom is in knowing you know nothing" (Cooper, 1998). If, as it seems, we may have 'unknowingly' moved into a period in which we are fast becoming information rich but thinking poor, acknowledging that we both know everything and nothing may be the one piece of knowledge we need right now.

A brief conclusion: from knowing to not knowing

Knowledge is a complex social construct, the very nature of which means that it will always be debated… is it truth or merely a belief, contextual or universal, transitory or finite (Harley, 2008; Hetherington, 2011; IEP (2025))? Regardless of these debates, no one can deny that humans hold the concept of knowledge in high regard, and for this reason, it commands infinite worth. Looking back through history, we can see how the advent of new types of knowledge shaped the trajectory of entire nations—from boat building to nuclear fission, from psychoanalysis to the discovery of penicillin. In response to its power, ever since the 4th century BCE, civilisations, societies, and nations have constructed institutions alongside entire systems devoted to the creation, assessment, and containment of knowledge.

Today, the launch of generative AI marks an entirely new phase in the trajectory of human knowledge. One in which the very concept of knowledge—its authorship, ownership, and ultimately its veracity—is being challenged. Given the significant role that knowledge plays in human life, the navigation of such a step change would, one might imagine, require a certain level of reflection, deliberation, and consideration (Owen et al., 2013). The speed with which this technology is permeating society, however, has (so far) left little room for such reflexivity (Singla et al., 2024). In fact, today, we are seeing humans reaching for their generative AI chatbots to help them grapple with anything from scientific hypothesising to their next holiday destination (Wang et al., 2023; Wong et al., 2023).

Not wishing to negate the enormous potential of generative AI, nor dampen its enthusiastic uptake, it would seem wise to deploy a little more circumspection when it comes to the trending tendency to farm out our knowledge practices to this latest wave of technological product. Given the value we, as a species, place on knowledge, as well as the influential role that it plays in learning, innovation, and—more generally—societal progress, an understanding of the impact of these new tools on both our knowledge practices and knowledge outcomes—both quantitative and qualitative—would seem in order. In fact, without wishing to over egg the point, admitting that we currently know little (nothing?) about the impact of AI-mediated knowledge tools on our ability to think, reason, debate, and learn, might be a good place to start.

Data availability

No data was collected, analysed, or reported on for this manuscript.

Notes

1 Generative AI in other forms significantly predates the launch in November 2022 of OpenAI's ChatGPT. To illustrate; an earlier and much simpler AI model was the 'Markov chain', created in 1906 to model the behaviour of random processes, and subsequently used in machine learning algorithms to model next word prediction tasks, such as autocompletion (Ching & Ng, 2006; Zewe, 2023).

2 'Packet switching' or 'message blocks' as they were originally referred to, were designed as a communication methodology able to withstand extreme societal breakdowns (such as a nuclear war), the idea being that even in the event of a system collapse the separate message blocks could be re-routed and subsequently rejoined to deliver the intended information (RAND, 2018).

3 Investigations are underway to assess the role that synthetically created data can have in the pursuit of training ever more powerful AI models. The role of synthetic data in research is a contentious topic, not only due to the unquantifiable limits of its artificiality (in terms of ethics, data bias, representation, transparency, and accountability) but also due to the technical problem known as 'model collapse' in which the quality of generated output diminishes rapidly when models are trained on their own synthetic data (Abrams, 2023; Messeri & Crockett, 2024; Neuman & Cohen, 2024).

4 'Disruptive innovation' or 'disruptive technology' are terms used to describe the impact of a new product, service, or idea that significantly challenges society, whether in the form of economies, values, or behaviours. Some categorise disruptive innovation to include smart phones, telemedicine, or Airbnb, whereas others underscore the need for truly disruptive innovation to come from niche market needs challenging established market leaders (Christensen et al., 2013; Guttentag, 2015).

5 The term sophistry, meaning the use of clever but false arguments, comes from a historical interpretation of the Sophists.

References

Abrams Z (2023) AI is changing every aspect of psychology. Here's what to watch for. Monit Psychol 54(5):46

Bandi A, Adapa PVSR, Kuchi YEVPK (2023) The power of generative AI: a review of requirements, models, input–output formats, evaluation metrics, and challenges. Future Internet 15(8):260

Benson HH (2011) Socratic method. The Cambridge Companion to Socrates. Cambridge University Press, pp. 179–200

Birnbaum JS (1985) Toward the domestication of microelectronics. Commun ACM 28(11):1225–1235

Blake NF (1991) William Caxton and English literary culture. A&C Black

Blease CR, Locher C, Gaab J, Hägglund M, Mandl KD (2024) Generative artificial intelligence in primary care: an online survey of UK general practitioners. BMJ Health Care Inform 31(1):e101102

Briggs A, Burke P (2009) A social history of the media: from Gutenberg to the Internet. Polity

Builder CH (1993) Is it a transition or a revolution? Futures 25(2):155–168

Chan T, Trueger NS, Roland D, Thoma B (2018) Evidence-based medicine in the era of social media: scholarly engagement through participation and online interaction. Can J Emerg Med 20(1):3–8

Chien AA, Lin L, Nguyen H, Rao V, Sharma T, Wijayawardana R (2023) Reducing the Carbon Impact of Generative AI Inference (today and in 2035). Proceedings of the 2nd Workshop on Sustainable Computer Systems, pp. 1–7

Ching W-K, Ng MK (2006) Markov chains. Models, algorithms and applications. Springer

Christensen C, Raynor ME, McDonald R (2013) Disruptive innovation. Harvard Business Review, Brighton, MA, USA. pp. 20151–20111

Çinar M, Bahçeci F, Dikmen S (2020) Examining internet addiction levels of high school last-grade students. Educ Res Rev 15(5):233–241

Cochrane LG, Chartier R (2014) The culture of print: power and the uses of print in early modern Europe (vol. 1005). Princeton University Press

Cooper JM (1998) Socrates (469–399 BC). The Routledge Encyclopedia of Philosophy

Cress U, Kimmerle J (2008) A systemic and cognitive view on collaborative knowledge building with wikis. Int J Comput Support Collab Learn 3:105–122

Crick N (2010) The sophistical attitude and the invention of rhetoric. Q J Speech 96(1):25–45

Cross FB (1996) Paradoxical perils of the precautionary principle. Wash Lee L Rev 53:851

Datta P, Whitmore M, Nwankpa JK (2021) A perfect storm: social media news, psychological biases, and AI. Digit Threats Res Pract 2(2):1–21

Denejkina A (2023) Young people's perception and use of generative AI. http://studentedge.org

Dewar JA (1998) The information age and the printing press: looking backward to see ahead. Rand Santa Monica, CA

Digital Education Council (2024) Global AI student survey 2024. Digital Education Council

Dinmore N (2023) 30 years of a free and open web. CERN. https://home.cern/news/news/computing/30-years-free-and-open-web

Duke G (2025) The Sophists (Ancient Greek). Internet Encyclopedia of Philosophy. IEP

Eisenstein EL (2005) The printing revolution in early modern Europe. Cambridge University Press

Emsley R (2023) ChatGPT: these are not hallucinations–they're fabrications and falsifications. Schizophrenia 9(1):52

Field H (2024) Google to relaunch Gemini AI picture generator in a 'few weeks' following mounting criticism of inaccurate images. CNBC

Flemming D, Cress U, Kimmerle J (2017) Processing the scientific tentativeness of medical research: an experimental study on the effects of research news and user comments in online media. Sci Commun 39(6):745–770

Fui-Hoon Nah F, Zheng R, Cai J, Siau K, Chen L (2023) Generative AI and ChatGPT: Applications, challenges, and AI-human collaboration. J Inf Technol Case Appl Res 25(Issue 3):277–304

Füssel S (2020) Gutenberg and the impact of printing. Routledge

Gholizadeh H, Salehi H, Embi MA, Ale Ebrahim N, Danaee M, Motahar SM, Tanha FH, Osman NAA (2014) Relationship among economic growth, internet usage and publication productivity: comparison among ASEAN and world's best countries. Mod Appl Sci 8(2):160–170

Giorgi S, Lynn VE, Gupta K, Ahmed F, Matz S, Ungar LH, Schwartz HA (2022) Correcting sociodemographic selection biases for population prediction from social media. Proc Int AAAI Conf Web Soc Media 16:228–240

Guttentag D (2015) Airbnb: disruptive innovation and the rise of an informal tourism accommodation sector. Curr Issues Tour 18(12):1192–1217

Harley JB (2008) Maps, knowledge, and power. In: Geographic thought. Routledge. pp. 129–148

Hellström T (2003) Systemic innovation and risk: technology assessment and the challenge of responsible innovation. Technol Soc 25(3):369–384

Hetherington S (2011) The Gettier problem. In: The Routledge companion to epistemology. Routledge. pp. 119–129

Hock HH, Joseph BD (2009) Language history, language change, and language relationship: an introduction to historical and comparative linguistics. Mouton de Gruyter

Hughes A (2023) ChatGPT: Everything you need to know about OpenAI's GPT-4 tool. BBC Science Focus

IEP (2025) Epistemology. Internet Encyclopedia of Philosophy. IEP

Jefferson T (1813) Thomas Jefferson to Isaac McPherson, 13 August 1813. Founders Online. https://founders.archives.gov/documents/Jefferson/03-06-02-0322

Jonas H (2009) Technology and responsibility. Readings in the philosophy of technology. Rowman & Littlefield, pp. 173–184

Kelly L (2024) The knowledge gene. Allen & Unwin

Kirova VD, Ku CS, Laracy JR, Marlowe TJ (2023) The ethics of artificial intelligence in the era of generative AI. J Syst Cybern Inform 21(4):42–50

Kosinski M (2024) Evaluating large language models in theory of mind tasks. Proc Natl Acad Sci 121(45):e2405460121

Lucchi N (2024) ChatGPT: a case study on copyright challenges for generative artificial intelligence systems. Eur J Risk Regul 15(3):602–624

Luccioni S, Jernite Y, Strubell E (2024) Power hungry processing: Watts driving the cost of AI deployment? The 2024 ACM Conference on Fairness, Accountability, and Transparency. pp. 85–99

Madaan A, Tandon N, Gupta P, Hallinan S, Gao L, Wiegreffe S, Alon U, Dziri N, Prabhumoye S, Yang Y (2024) Self-refine: Iterative refinement with self-feedback. Advances in Neural Information Processing Systems. 36. https://arxiv.org/abs/2303.17651

Messeri L, Crockett MJ (2024) Artificial intelligence and illusions of understanding in scientific research. Nature 627(8002):49–58

Microsoft & LinkedIn (2024) 2024 Annual Work Trend Index. https://news.microsoft.com/annual-wti-2024/

Morgan Library (2015) William Caxton and the Birth of English Printing. The Morgan Library and Museum

Motyl M, Narang J, Fast N (2024) Tracking Chat-Based AI Tool Adoption, Uses, and Experiences. Designing Tomorrow. https://psychoftech.substack.com/p/tracking-chat-based-ai-tool-adoption

Neuman Y, Cohen Y (2024) A data set of synthetic utterances for computational personality analysis. Sci Data 11(1):623

Norman J (2023) Chinese writing: traditions and transformations. Asia Society. https://asiasociety.org/education/chinese-writing

Owen R, Stilgoe J, Macnaghten P, Gorman M, Fisher E, Guston D (2013) A framework for responsible innovation. Responsible innovation: managing the responsible emergence of science and innovation in society. Wiley, pp. 27–50

Park E, Gelles-Watnick R (2023) Most Americans haven't used ChatGPT; few think it will have a major impact on their job

Pew Research Center (2024) ChatGPT use has ticked up since July, particularly among younger adults. https://www.pewresearch.org/short-reads/2024/03/26/americans-use-of-chatgpt-is-ticking-up-but-few-trust-its-election-information/sr_24-03-26_chat-bot_1-png/

Porter J (2023) ChatGPT bug temporarily exposes AI chat histories to other users. The Verge

Posetti J (2018) A short guide to the history of 'Fake News' and disinformation. International Center for Journalist

Quan Z, Chen Z (2024) Human–computer pragmatics trialled: some (im) polite interactions with ChatGPT 4.0 and the ensuing implications. Interact Learn Environ 1:1–20

RAND (2018) Paul Baran and the Origins of the Internet. RAND Research

Rani U, Dhir RK (2024) AI-enabled business model and human-in-the-loop (deceptive AI): implications for labor. In: Handbook of artificial intelligence at work. Edward Elgar Publishing. pp. 47–75

Ring J (2023) 30 years ago, one decision altered the course of our connected world. NPR. https://www.npr.org/2023/04/30/1172276538/world-wide-web-internet-anniversary

Robert SJ, Kadhiravan S, McKay D (2024) The development and validation of digital amnesia scale. Curr Psychol 43(22):1–10

Rodilosso E (2024) Filter bubbles and the unfeeling: how ai for social media can foster extremism and polarization. Philos Technol 37(2):71

Säfström CA (2023) Socratic Versus Sophist Teaching. In: Education for everyday life: a sophistical practice of teaching. Springer. pp. 51–63

Salakhutdinov R (2015) Learning deep generative models. Annu Rev Stat Appl 2(1):361–385

Shaban A (2025) Digitalization and exclusion—digital divides and development. In: Digital geographies—theory, space, and communities: a machine-generated literature review. Springer. pp. 255–496

Singla A, Sukharevsky A, Yee L, Chui M, Hall B (2024) The state of AI in early 2024: Gen AI adoption spikes and starts to generate value. McKinsey & Company

Smith LGE, Owen R, Cork A, Brown O (2024) How and why psychologists should respond to the harms associated with generative AI. Commun Psychol 2(1):60

Sparrow B, Liu J, Wegner DM (2011) Google effects on memory: Cognitive consequences of having information at our fingertips. Science 333(6043):776–778

Statista (2025) Number of internet and social media users worldwide as of October 2024. Statista

Sterling B (1993) Short history of the Internet. Internet Society

Suenaga K (2023) An evolutionary theory of human knowledge progress: The new combinations of anthropology, archaeology and economics. Cogent Arts Humanit 10(1):2188780

Sumbal A, Sumbal R, Amir A (2024) Can ChatGPT-3.5 pass a medical exam? A systematic review of ChatGPT's performance in academic testing. J Med Educ Curric Dev 11:23821205241238640

Sun Y, Sheng D, Zhou Z, Wu Y (2024) AI hallucination: towards a comprehensive classification of distorted information in artificial intelligence-generated content. Humanit Soc Sci Commun 11(1):1–14

Symons CS, Johnson BT (1997) The self-reference effect in memory: a meta-analysis. Psychol Bull 121(3):371

Tarkhnishvili A, Tarkhnishvili L (2013) Middle class: Definition, role and development. Glob J Hum Soc Sci Sociol Cult 13(7):21–31

Taylor CCW (2006) Socrates the sophist. In: Remembering Socrates. Oxford University Press, pp. 157–168

Teasley S, Wolinsky S (2001) Scientific collaborations at a distance. In: Science (vol. 292, Issue 5525). American Association for the Advancement of Science. pp. 2254–2255

Tudoroiu T (2014) Social media and revolutionary waves: the case of the Arab spring. N Polit Sci 36(3):346–365

Von Petzinger G (2017) The first signs: unlocking the mysteries of the world's oldest symbols. Simon and Schuster

Wach K, Duong CD, Ejdys J, Kazlauskaitė R, Korzynski P, Mazurek G, Paliszkiewicz J, Ziemba E (2023) The dark side of generative artificial intelligence: a critical analysis of controversies and risks of ChatGPT. Entrep Bus Econ Rev 11(2):7–30

Wallace J, Green M (1996) Bridging the analogy gap: The Internet, the printing press and freedom of speech. Seattle UL Rev 20:711

Wang H, Fu T, Du Y, Gao W, Huang K, Liu Z, Chandak P, Liu S, Van Katwyk P, Deac A (2023) Scientific discovery in the age of artificial intelligence. Nature 620(7972):47–60

Wong IA, Lian QL, Sun D (2023) Autonomous travel decision-making: An early glimpse into ChatGPT and generative AI. J Hosp Tour Manag 56:253–263

Woodard RD (1997) Greek Writing from Knossos to Homer: a linguistic interpretation of the origin of the Greek alphabet and the continuity of ancient Greek literacy. Oxford University Press

Worrell T, Johns D (2024) Indigenous considerations of the potential harms of generative AI. Agora 59(2):33–36

Yan L, Greiff S, Teuber Z, Gašević D (2024) Promises and challenges of generative artificial intelligence for human learning. Nat Hum Behav 8(10):1839–1850

Zaret D (2000) Origins of democratic culture: Printing, petitions, and the public sphere in early-modern England. Princeton University Press

Zewe A (2023) Explained: Generative AI. MIT News

Zhang A, Yang SA (2024) The high cost of GPT-4.o. Australian Strategic Policy Institute

Acknowledgements

Author contributions

Sarah Bentley was to sole creator of this manuscript.

Competing interests

The author certifies that they have no affiliations with or involvement in any organisation or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

Ethical approval

No human participant data was collected for this manuscript; therefore, no informed consent was required.

Informed consent

No human participant data was collected for this manuscript; therefore, no informed consent was required.

Additional information

Correspondence and requests for materials should be addressed to Sarah V. Bentley.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.