Minds and Machines (2026) 36:2
https://doi.org/10.1007/s11023-025-09757-6

Losing Our Voice? Generative AI and the Degradation of Human Expression

Scott Robbins[1] · **Inga Blundell**[2]

Received: 16 May 2025 / Accepted: 5 November 2025

Abstract
This paper examines the implications of generative AI (GenAI) emulating human expression, i.e. human communication and human creative expression. While GenAI seems to offer benefits such as increased efficiency and productivity, its use raises significant practical and conceptual concerns: GenAI comes with the increased efforts of prompting, verification and editing, and causes the deskilling of its users. It also comes at a monetary cost and causes various ethical issues e.g., a lack of authenticity. We further show that GenAI's fundamental issue is that it is by design not able to output human expression but only human-like expression. Using AI for tasks that are fundamentally about communicating is replacing communication with something that is not communication. Finally, we show that the consequences of the use of GenAI cannot be avoided on an individual level by those individuals avoiding the use of GenAI and will necessarily lead to an erosion of human expression in general. This is because GenAI: will lead to a distrust in human expression when the authenticity of authorship over time becomes unclear; will cause a devaluation of human expression when human expression can be mimicked with less effort by GenAI; and, will discourage human expression altogether when GenAI has set the bar too high.

Keywords Artificial intelligence · Generative AI · Deskilling · Prompting · ChatGPT · AI ethics

✉ Scott Robbins
scott.robbins@kit.edu

[1] Karlsruhe Institute of Technology, Douglasstraße 24, 76133 Karlsruhe, Germany

[2] Bonn, Germany

Published online: 19 November 2025

1 Introduction

Generative AI (GenAI) has exploded into our lives. The adoption of GenAI has been faster than that of both the PC and the internet itself (Bick et al., 2024). ChatGPT came onto the scene in 2022 and already, for example, 39% of Americans claim to use GenAI. 25% of US workers say they have used GenAI once in the last week and nearly 1 in 10 US workers say they have used it every day for work. At work, people are using it for (in order of popularity) writing communications, performing administrative tasks, translating and summarizing, search, coding, documentation, etc. At home people are using it for (in order of popularity) writing communications, translating and summarizing, personal assistance (lists and schedules), ideas and prompts for creative projects, entertainment recommendations, etc. (Bick et al., 2024). More people may be using it and don't even know it. GenAI is now integrated into all kinds of popular applications like Microsoft Word, Gmail, and Adobe photo editing products. We have all been exposed to GenAI images. Even the President of the United States tweeted a GenAI image of himself as the Pope (Matza, 2025).

The change that is coming (or is already here) is a seismic shift. The internet and telecommunications brought about a shift to Computer-Mediated Communication (CMC). However, GenAI has brought about a change to Artificial Intelligence-Mediated Communication (Hancock et al., 2020). This isn't a difference in scale but a difference in kind. Using a word processor to write a letter and having it suggest some grammar changes and spelling changes is not the same as having ChatGPT write the letter for you. Importantly, when people started writing with word processors, we could literally see the difference - it wasn't written with a typewriter or longhand. Unfortunately, the important difference we face now remains hidden. We won't know if someone used GenAI to write an email or if they wrote it themselves. This relates to the important issue of whether humans are meaningfully in control of the outputs of AI (Robbins, 2023a, b; Santoni de Sio & van den Hoven, 2018). With GenAI we may be losing meaningful control over the very act of expressing ourselves.

Whereas before we were expressing ourselves and then using tools (even AI tools) to enhance or edit those expressions, now we are editing and enhancing the outputs of GenAI. That is, we have outsourced human expression to GenAI. The fundamental issue with this concerns the fact that GenAI has no emotions, desires, opinions, or views to express. When GenAI makes a painting, writes a poem, composes a song, etc. it is not an expression of anything. Calling what it does art or communication is a failure to use language correctly. Those things require humans.

The purpose of this article is to understand the implications of outsourcing human expression to GenAI. In any given instance we may think that the output of GenAI achieves our immediate goal. For example, that love letter that ChatGPT wrote for my partner has made her happy. However, the widespread outsourcing of human expression to GenAI will cause distrust in human expression. Teachers distrust the written work of their students (Dwyer & Laird, 2024), citizens distrust news articles and news images (Holcomb, 2024), and there are so many bots online posing as real people that it will be difficult for people to trust their online interactions in the future (Baraniuk, 2014; Ivarsson & Lindwall, 2023).

Relatedly, the widespread use of GenAI has a devaluing effect on human expression. That is, artists and those who express themselves will see their expressions worth less. The doubt created by the widespread use of GenAI will cause true human expression to be less impactful than it once was. Practically speaking, artists will see their work monetarily devalued because the imitations generated by GenAI will be useful to the businesses that would be their clients. This hurts an artist's ability to fund projects that they are passionate about.

Finally, and perhaps the saddest outcome is that GenAI will discourage human expression. GenAI is already so effective at emulating our expressions that it will discourage us from even trying to do it ourselves. Students will be afraid of turning in their own essays because they are competing with the outputs of GenAI. Why write your own emails and texts when GenAI can be a wittier, more personable version of you? The degradation of our ability to translate our thoughts, emotions, desires, and opinions into communication and art is the loss of our voice. That voice is being replaced by a cheap, functional, uniform, imitation.

In the next section we first discuss the success of GenAI. There are good reasons that so many people and institutions are using these tools. In section three we highlight some of the challenges that individual users face when using GenAI. First, while GenAI is quite effective, its use faces practical issues like generating new tasks that require effort for the user, deskilling the user, and costing money (maybe much more in the future). Second, there are ethical issues with the use of GenAI including a lack of disclosure, effort, commitment, and authenticity. Finally, the use of GenAI comes with the fundamental issue that its output is of a different kind than advertised. GenAI does not have the desires, emotions, opinions, etc. that are required to call its outputs art or communication. In section four we consider the implications of the widespread use of GenAI. These implications consist of the distrust of human expression, the devaluation of human expression and the discouragement of human expression.

2 The Success of Generative AI

There are certain tasks that have historically been performed exclusively by humans but are recently emulated by AI and even more recently and more successfully by GenAI. GenAI can generate literature, paintings, imitation photography, music, e-mails, essays, etc. Students, partners, friends, doctors, customer service agents, etc. are all using GenAI to help at work and at home. GenAI is successful at creating human-like outputs in these fields - Outputs that are very hard, if not impossible, to distinguish from human outputs. In 2022 the Colorado State Fair unknowingly awarded first place to an artwork created using AI. In 2023 Boris Eldagsen entered an imitation photo into the Sony World Photography Awards competition and won and later admitted on his website that it was AI generated, declining to accept the award (Grierson, 2023). Even books (partially) written by AI have won competitions. Rie Kudan's The Tokyo Tower of Sympathy, which Chat-GPT generated "word for word" about 5% of, won the Akutagawa Prize in Japan - one of the country's most prestigious literary prizes (Choi & Annio, 2024). Finally, in 2023, the Beatles song

"Now and Then" which was partially generated by AI won a grammy (Weatherbed, 2025).

On top of being very convincing, GenAI's outputs have proven themselves useful in many ways. For example, a study showed that in the context of customer support, less-experienced workers improved their performance (as measured by number of issues handled) by 34% when given access to an AI assistive chatbot. Chatbots and other AI tools may also help reduce loneliness. De Freitas et al. (2024) ran a series of studies which seem to show a reduction of loneliness in participants who talked to a chatbot. Not only do people often mention loneliness in their positive reviews of chatbots marketed for companionship, but loneliness levels seem to be reduced by interacting with chatbots daily (De Freitas et al., 2024). It is important to note however, that these results seem to be dependent upon deceiving the user - either into thinking that the chatbot is human, or that the chatbot has emotions like empathy (which chatbots most certainly do not). To be clear, an agent without emotions like empathy (like a goldfish) could reduce loneliness; however, it does not do so by making the lonely person believe that it has empathy.[1] Hohenstein and colleagues (2023) found that using algorithmic responses when communicating in personal relationships can produce better outcomes. Business.com conducted a survey of about 2000 American workers regarding their experiences with ChatGPT in 2023 and reported that around 23% of respondents were using AI for written communication. A survey conducted by Microsoft and LinkedIn in May 2024 found that 75% of "global knowledge workers" are using generative AI (Microsoft & LinkedIn, 2024). Another survey conducted on the AI platform Pollfish in January 2024 found that, of US respondents between 18 and 34 years old, 33% of men and 14% of women had used ChatGPT for relationship advice (Field, 2024). All these examples imply that these tools are at least perceived as useful by many people.

McKinsey & Co. predicts that GenAI could add between $2.6 and $4.4 trillion in economic value across the 63 use cases they looked at (Chui et al., 2023). They also predict that a worker's time invested in completing a task could be reduced by 60–70% through the use of generative AI. This could save companies money by reducing the number of employees they need. The World Economic Forum claims that 41% of companies are planning to reduce their workforce due to the automation benefits of AI (World Economic Forum, 2025). For example, the CEO of the popular language learning application Duolingo has said in a LinkedIn post that "without AI, it would take us decades to scale our content to more learners", and that Duolingo will now have an "AI First" policy that will only increase headcount if a department cannot show that they could improve efficiency with AI. Also, Duolingo plans to replace all contract workers[2] with AI (Peters, 2025). For companies to be planning this way, GenAI must be quite effective at automating previously human-only tasks.

[1] Thanks to an anonymous reviewer for their helpful comment about this.

[2] Contract workers are synonymous with free lancers or independent contractors.

3 Challenges of Individual AI Use

Of course, GenAI's relatively short history isn't quite that rosy and the list of downsides is at least as long as its list of triumphs. We will look at three different types of issues an individual might face in using GenAI: Practical issues, ethical issues and what we will call the fundamental issues.

3.1 Practical Issues

Practical issues refer to those issues that have a direct effect on the users and their daily lives. Here we will show how GenAI causes additional effort for the user, causes the user's deskilling, and has economic costs. This list is not exhaustive but will suffice to show that there are major practical issues that must be taken into account when using GenAI.

3.1.1 The Additional Effort of Using GenAI

AI is often portrayed as an absolute time saver: Your assignments for your Kant course will practically write themselves, no more writing cumbersome emails, and if you want you can even have AI create some artwork for you. If we take a closer look at these cases, it becomes clear, however, that it is not quite that easy. AI tools in many cases may simply create more work than they alleviate us of (Cappelli et al., 2024; McKendrick, 2024). GenAI output is not created out of thin air - it requires, in part, a process of prompting. This can be a (sometimes lengthy) back and forth with the GenAI algorithm where the prompter ensures that all requirements to the output, including e.g. in the case of text output, that the right tone and the right length are met (Marvin et al., 2024; Sikha et al., 2023). After prompting has yielded an at-first-glance acceptable output, work is required to verify the output in terms of its factual assertions (is what the output says true?) and in terms of ethical correctness (has the output said something that crosses an ethical line?). In the case of artwork one might want to, e.g., make sure that it is not depicting something that would widely be seen as racist. For LLMs like ChatGPT one will want to ensure that the output is factually correct, i.e., that its assertions are true. This is important as LLMs are indifferent to the truth. That is, they are "bullshitters" that will confidently tell you things without regard to their veracity (Hicks et al., 2024).

For AI generated output this can be particularly problematic as the AI model's considerations are not known to the user. GenAI is by design opaque. While old-fashioned AI was easily able to explain itself due to human logic directly built into the system, the AI powering GenAI is able to, loosely speaking, develop its own logic by detecting inarticulable patterns from a lot of training data examples (Robbins, 2019, 2020). The problem is hence that GenAI, unlike old-fashioned AI, does not come with an explanation. We are trying to verify its output without knowing how this output was reached. When the output isn't verified, it can have embarrassing consequences. For example, two lawyers at the law firm Morgan and Morgan made headlines because a judge found that their lawsuit against Walmart contained fictitious case citations that, the lawyers admitted, were caused by the use of an LLM

(Merken, 2025). After ensuring that the citations were correct, however, one would need to ensure that the characterizations of the points made in those citations were also correct. This is no small task.

After all of this, one would still need to edit the output to give it a personal touch or to ensure that it fits into the specific context under which it was written. One might need to make a personal reference to a partner, add a piece of knowledge in a work email that would only be known to the user (not the LLM), or change the title of the story to make it closer to what one thinks the publishing house one is submitting to is looking for. One might argue that advances in technology have always come with a shift in tasks and skills. This is true, but it is also true that they have often come with a predicted reduction of effort that they have failed to deliver (Steffensen et al., 2022). Smartphones, for example, purport to allow us to accomplish all kinds of things during times that would normally be wasted (e.g., waiting in line or taking the bus). However, they have also created a flood of communication that we now have to respond to in an increasingly shorter window of time. They have also freed employers to demand more of us during times we would normally have to ourselves. It is difficult to say that we have more free time now than the times before the smartphone (Ogden et al., 2023).

3.1.2 Deskilling

The explicit goal of technology is often to do something for us, and oftentimes this equates to something we don't want to do any more. The flip side of this is that by not doing whatever it is we have delegated to technology we will fail to develop or maintain that skill anymore. Those of us old enough to remember a time before cell phones will remember that most of us had 10 to 20 phone numbers memorized. We had the skill of memorizing phone numbers. When cell phones came around that could store hundreds of phone numbers that could be dialed by selecting the right name, that skill was no longer useful. Now, most of us do not have that skill anymore.

Borgmann (1984) was concerned that technology would make fewer demands on our skill and attention. His concept of a 'device paradigm' describes how technology hides the processes that produce the commodities that we consume. Central heating provides us the commodity of heat as opposed to chopping wood and building a fire. Plumbing provides water rather than digging a well. These devices have prevented us from developing the skills necessary to provide these things for ourselves. We are now dependent upon devices we don't understand.

Technology has prevented us from developing and maintaining many skills. Some of these skills we have no reason to care about (maybe memorizing phone numbers is one of them… unless your phone is out of battery and you need to make a call). However, what GenAI is capable of doing, if used widely, would affect our ability to communicate our feelings, to translate our thoughts into prose, and to, in short, express ourselves. Just like our inability to navigate using a map or build a fire, we may lose our ability to write a message of love to our partner. With GenAI we are faced now, more than ever, with the responsibility of deciding what machines should or shouldn't do for us (Robbins, 2025).

Here, we are appealing to the reader's intuition that expressing oneself is valuable. First, it should be noted that if its value is unknown then its value should be properly understood before we delegate it to GenAI and deskill ourselves. Second, human expression in terms of language or art is unique to human beings and "are clearly central to human sociality" (Heintz & Scott-Phillips, 2023). Humans, by their very nature, are social animals (This centrality to human sociality makes it clear that these skills are important). In short, to argue that these skills are unimportant is to argue that human sociality is unimportant. While it may be possible to reject sociality in some post-human future, the premise we are starting from is that we are human – and therefore need these skills so central to human sociality. We should not willingly do something that diminishes these skills as it seems it would indeed diminish our humanity.

3.1.3 Costs

In the very literal sense GenAI does not come without costs: ChatGPT bills its advanced users $20 per month. Although possibly more convenient, love letters before the use of AI were a lot cheaper than that. To be fair, creating visual arts has never been free of costs. However, if you are using Stable Diffusion (possibly on top of what you used before) it will cost you $27 per month. On top of that, businesses tend to keep costs low for as long as they can, working on the premise that they will be able to lock users in demand higher prices further down the line.

3.2 Ethical Issues

Battisti (2025) points out that three ethical issues arise when using GenAI to communicate with friends or partners: lack of disclosure, lack of effort, and lack of commitment. A *lack of disclosure* arises when one uses generative AI for one's own benefit but does not disclose the use of AI. If you are, e.g., using AI to write to your partner or friend to apologize for something you said in a fight, it might be essential for this to actually work that your partner does not know that you didn't come up with this apology yourself but that it was in fact AI generated. Hohestein et al. (2023) point out that communication speed and the use of positive emotional language result in conversation partners evaluating each other as closer and more cooperative. Its speed and emotional mimicry of GenAI makes it a useful tool when trying to dissolve a fight. Hohestein et al. also found, though, that the opposite is the case when there is a suspicion of secretly using AI where, conversely, conversational partners are judged more negatively. Your partner, if they found out, would understandably be angry that you had deceived them.

So why not just be open about the fact that AI was used to achieve this level of "positive emotional language"? Why not immediately point out that your message was in fact created using AI? This leads us to Battisti's second ethical issue: *lack of effort*. "Using ChatGPT for romantic interactions expresses laziness, which undermines the moral value placed on personal effort in one's maintaining relationships." (Battisti, 2025). The value of the written apology seems to be reduced by the fact that a lot less effort seems to have gone into it. This seems intuitive: When receiving an

apology, part of what determines its value is the effort that has gone into it. Receiving an effortless apology in passing does not compare to receiving a handwritten card, some flowers and a lovingly composed apology. Battisti does point out however, that the idea of a moral wrong in a lack of effort is somewhat vague. First of all, a conversational partner using AI can not be said to be putting in no effort at all. They must still prompt, verify the output or edit it to remove any inconsistencies with past conversations. We are "only" speaking of a reduction of effort here, making it unclear if it is in fact a moral wrong we are looking at. While we may not be able to draw the line at how much effort reduction crosses an ethical line, the line can be crossed.

Closely related to this issue is the third issue: *the lack of commitment*. Not only are we expecting there to be some effort put into the apology, we are also expecting our conversational partner to be committed to the relationship. That is, we expect a "genuine interest in maintaining the relationship, which requires a mutual commitment to addressing each other's needs". Here the lack of effort is still the starting point. The problem is not the lack of effort itself though but the fact that a lack of effort indicates the lack of willingness to maintain the relationship.

Further, Battisti argues that the fundamental issue here is a another one, namely a lack of authenticity. When a particular task is not performed by the individual herself, she fails to be authentic "Given the bonds of intimacy and trust within the relationship, certain tasks cannot typically be outsourced to someone or something else, even if doing so might improve the quality of the outcome"(Battisti, 2025, p 28).

When receiving a written apology for example, it seems to be a requirement that the apology is written and authored by the person apologizing themselves. Imagine having your partner send their mother over to apologize for them because she is better at apologies. This would clearly not "count" as an apology from your partner. While apologies can, in principle, be made by representatives of a group - whereby a member of a group apologizes for what another member of the group did (e.g., the Chancellor of Germany apologizing to the Jewish community for actions committed by German nazis), the example above contains a group with a membership of one. No one else can apologize on your behalf - neither your mother nor an algorithm (Joyce, 1999).

While we have discussed these ethical issues in the context of romantic partners, they apply outside of this context as well. A student using GenAI to write a paper, for example, faces the ethical issues of failing to disclose the use of GenAI, would fail to put the effort into their paper, show a lack of commitment to the course, and finally, wouldn't be authentic.

In the example of the mother being used to apologize for our partner, our issue is that the wrong person is apologizing. We are expecting an apology from our partner, not from their mother. We can still argue that an apology is being made. Maybe the mother is apologizing - to some extent - for raising their offspring so poorly (Joyce, 1999; Mookherjee et al., 2009). We are facing an even larger, conceptual problem, when it is not the mother, i.e., a human being, but an algorithm apologizing.

3.3 Conceptual Issues

In 2024 Bumble founder Whitney Wolfe Herd suggested that the future of dating is through the use of AI concierges. Dating app users would be assigned personal AI "dating concierges", who would date hundreds of other users' AI concierges to determine compatibility, and who may be worth meeting in person. AI would "sit through all those awkward first dates, so you won't have to" (Raiken, 2024). At first (and second and third) glance this sounds quite silly. What does it mean for an AI chatbot to "sit through" first dates with another AI chatbot? How would the user learn anything from this, and why would the user trust the outcome of these "first dates"? More importantly, if after a lengthy process of AI-chatbot dating two users decide to go on a date, wouldn't they call this the first date? Of course, Whitney Wolfe Herd must be using language quite loosely (if we are to take her at all seriously). With no humans in the mix, there is no date. With no intentional agent[3] there is no communication at all. These things require human expression. It is fundamentally important to understand that these acts, by definition, can only be performed by humans (or, perhaps, some animals). Pretending that a machine can go on a "date" is just using language incorrectly at best, and at worst it is a deliberate misrepresentation of what is happening.

When AI creates "art", becomes your "friend", or in general "communicates" we are making this mistake. The problem is that communication uses words to communicate the desires, emotions, and view of the communicator. AI has neither desires, nor emotions, nor a view. There is, simply, nothing to communicate. The AI chatbot can pretend to have emotions and desires and communicate them the way a human would. The chatbot could be designed to say things like "please can we talk later tonight? I will miss you too much" but these words do not refer to any real desire or emotion within the chatbot. Floridi and Nobre (2024) elaborate on this phenomenon of "conceptual borrowing", providing us with a list of misleading terms used in the context of AI, such as "machine learning", "hallucination", "attention" or "artificial intelligence" itself. The extent of the confusion about AI's capabilities is summed up by the fact that "AI has ended up describing computers anthropomorphically, as computational brains with psychological properties, while brain and cognitive sciences have ended up describing brains and minds computationally and informationally, as biological computers", which he argues, can be used to serve the (often financial) interests of some.

The case of AI generated "art" has received a lot of attention. It may not be so clear cut to the reader that AI can't create art. How does what was said in the previous paragraph about communication (if the reader is on board with that) apply to art? To be sure, it is not within the scope of this paper to solve the debate about what is and what is not art, which has been going on for millenia. The authors fall on one side of this debate - concisely captured in a quote by Leo Tolstoy: "…the aim of works of art is to infect people with the emotion the artist has experienced" (Tolstoy, 1094). That is, art communicates the emotions of the artist. We are seeing something the way that the artist saw it - we are "looking at a looking" (Berger, 2018). AI doesn't have any

[3] We use 'inentional agent' to include humans as well as, for example, non-human animals that can communicate.

emotions or viewpoint from which they are looking. The "art" it generates doesn't have anything to communicate. Which, according to one side of the debate, means that AI generated art is not art at all. For those interested, there are arguments that the artwork can stand alone in that a person viewing the art can feel and interpret the artwork as communicating something. In this view, it does not matter that there was no intention to communicate. What matters is that the person understands the artwork as communicating something (see e.g., Coeckelbergh, 2017 for an overview).

Here, we are not distinguishing artistic production and relational communication. Of course, there is a real difference here. In the former there is no live audience that the artist is responding to. In the latter there is direct feedback and both agents are directly responding to each other. However, we treat them the same because they are both under the umbrella of human expression. Both are expressions of an agent's intentions, feelings, desires, etc. One could argue that artistic production, due to a gap in production and consumption, could be treated differently.[4] We disagree with this for the reasons stated in the previous paragraph.

The conceptual issue discussed in this section is not that AI is bad at communicating or bad at creating art. It is that we are making a category mistake in saying that AI can communicate or create works of art at all. AI is simply not the type of thing that can communicate. Without intentions, desires, emotions, etc. there is nothing for AI to communicate. Using AI for the tasks that are fundamentally about communicating, therefore, is replacing communication with something that is not communication. We are taking away something that it is in our nature to seek: human expression (Carroll, 2004). Without human expression, there is nothing for us to respond to.

When Mark Zuckerberg claims that people have on average 3 friends but desire 15 and that the way to bridge this gap is with AI chatbots, he is also guilty of this category mistake. We would be deceived into expressing ourselves into a mirage of a human. We would be putting effort and work into a relationship that wasn't one. Friendships are about a mutual concern for one another (Aristotle). We can't have concern for a chatbot because there is nothing there to be concerned about. The chatbot can't have concern for us because a chatbot doesn't have emotions.

One objection that a reader might have is that while AI cannot create art or communicate on its own, there may be a sense in which AI can act as a 'co-author' of art or communication. First, it should be noted that AI does not meet the requirements of authorship according to academic journals – including this one. Specifically, authorship requires the ability to take accountability for the work. This requirement excludes GenAI.[5] Another argument against conceptualizing AI as a co-author (or teammate, collaborator, partner, etc.) is that it requires more than one peer. A peer must be an unsubstitutable agent that has goals, pursues those goals, has the autonomy to join and withdraw from the shared project, and can co-determine the ends of that project (Evans et al., 2023). None of these conditions are met by GenAI. This is all to say that AI doesn't communicate, doesn't create art, can't be a friend, a colleague, a teammate, or co-author. This is due to the conceptual definitions of these

[4] Thanks to an anonymous reviewer for this point.

[5] See Springer's authorship criteria page: https://www.springer.com/us/editorial-policies/authorship-principles. Also see Moffat and Hall (2024) for an article arguing against including AI as a possible co-author.

things, not due to a lack of computing power or a missing capability that will be there in the future. It is a conceptual issue that won't ever go away.

Although these concerns make our use of AI sound very bleak, maybe, at least for individuals, it could be possible to avert these issues by simply avoiding any use of GenAI in their personal and professional lives. In the next chapter we will evaluate if this strategy can be effective.

4 The Erosion of Human Expression

First, we must remark that it is not always under our control to use AI or not. Professionally we may be required to use GenAI-based tools. More importantly though, there is a crucial distinction to be made between *using* GenAI to achieve a certain goal and *consuming* GenAI output. Although making sure we are not using GenAI is a difficult task already (tools often come with built-in GenAI that we are not even aware of using), it still seems somewhat achievable. Making sure that we are not consuming GenAI outputs on the other hand is simply impossible. As we have seen, GenAI has become so good at imitating classically human outputs like visual art, literature and communication, that anyone consuming art or texts will not be able to say for sure if GenAI was used to create it, which will lead to the distrust, devaluation, and discouragement of human expression.

4.1 Distrust of Human Expression

First, when consuming art, love letters, or messages we are always confronted with the issue of distrust, i.e. the suspicion of false authorship. Where at one point it would have seemed paranoid to suspect that the author of the love letter I receive is in fact not the person signing it, it now seems very reasonable to ask myself if it might not have in fact been GenAI that created this letter. In a similar way this could happen when consuming art: If we take Tolstoy's idea of art, i.e., that the aim of works of art is to infect people with the emotion the artist has experienced, then a lack of trust in the work of art that I am consuming actually having been created by a human becomes a problem. A consumer of art will usually be invested in understanding the artist's expression of human feeling, e.g., pain, fear or anger. However, after having repeatedly been confronted with GenAI works of "art" posing as human authored, one will become more weary of trying to understand a piece of art for fear of there not being anything to understand - there is no artist and therefore no expressed emotion.

As stated above one might argue that GenAI created work will still provoke human emotion and must still be valued as art (Coeckelbergh, 2017). However, even if this were true, an important aspect of consuming art seems still to be getting lost irrecoverably - the conversation between artist and audience is diminished by the distrust created by GenAI. The audience's distrust makes them lose out on the valuable practice of interpreting the artist's message.

Salas Espasa and Camacho (2025) also put a focus on the artwork itself – claiming that a certain type of authenticity is found in artworks that are generated by AI – but curated, edited, and contextualized by humans. This blurs the line of authenticity –

leaving room for an AI generated artwork to have some authenticity (when a human has a high level of input). That is, there may be, in AI generated art highly curated and edited by a human, something expressed. As synthetic "art" becomes ubiquitous, it becomes increasingly banal. However, this raises the value of authentic human expression. It is true, that at some point an output of AI could be transformed enough to consider it human expression.[6] A lot would depend upon how many and what kind of creative and normative choices were delegated to GenAI. Was it the human responding to the output of GenAI? Or was it GenAI being used to realize the expression of the human? The answers to these questions will be unknown to the audience – creating the distrust that this section is concerned with.

The fact that the consumer loses faith in the fact that works of art, communication, and acts of human expression are created by humans will also create an issue for the humans expressing themselves: Their expressions will lose their worth.

4.2 The Devaluation of Human Expression

When consumers no longer believe in the authenticity of an artists' work or of human communication, this will not only have an effect on the consumers but even more so on the creators: We no longer value human expression as highly. This is not only a loss of credit but also a monetary one for artists and creators: Art, for example, will simply become cheaper or lose its value completely. For example, it has been predicted that music sector workers will lose a quarter of their income due to AI (Burke, 2024).

The distrust fostered by the widespread use of GenAI impacts humans who are crafting their emails, letters, works of art, etc. themselves. Their expressions of emotion, opinion, and desires may not be interpreted as such because people are unsure if there is anything to interpret. It is rather intuitive that as the world is flooded with artificial human expression, that the overall value of human expression goes down - this unfortunately gets worse the better the GenAI gets at mimicking human expression. If one can buy fake luxury products that are so close to the "real" thing that no one can tell the difference, then why pay exorbitant prices for the real thing? Their value goes down.

The value of 'real' human expression, could also (for a time at least) go up. The less-than-perfect nature of most authentic human expression could be seen as valuable in an age when the world is flooded with technically, grammatically, 'perfect' written messages and art. Just like in the age of factory made 'perfect' biscuits, there is a high-value put on 'homemade' biscuits that have imperfections like you remember your grandparents making them. The 'poor quality' (e.g., the less than perfect shape) signifies something authentic. In the case of human expression, the technical mistakes and lack of artistic competence found in, for example, a children's painting[7], could increase the value of the work compared to that of 'perfect' AI generated

[6] Salas Espasa and Camacho (2025) relate their argument to Walter Benjamin's concept of 'Aura' and claim that the emergent form of authenticity found when humans edit, curate, and contextualize GenAI 'art' should be called 'semi-aura'.

[7] Thanks to an anonymous reviewer for this example.

paintings. Unfortunately, GenAI can adapt to the changing tastes of users/consumers. It doesn't have to generate paintings like Van Gogh or prose like Virginia Woolf. High school students will quickly learn to get GenAI to write more in a style that they would have written something in – including mistakes. The overall concern of a devaluation of human expression due to a distrust regarding the work's authenticity remains.

In a post that went viral last year, the author Joanna Maciejewska put it nicely when she said "I want AI to do my laundry and dishes so that I can do art and writing, not for AI to do my art and writing so that I can do my laundry and dishes." Where consumers may be able to decide not to consume GenAI generated artwork, artists may not be given that choice: If AI takes away the idealistic and monetary value of art, it may be stripping much of the meaning of artist's lives. Though this might only be relevant for a small (artistic) percentage of the population, there is a similar effect for "everyday people": When AI is doing it better than them, why should they even try.

4.3 Discouraging Human Expression

One of the author's has a partner who sometimes writes her poems. She thinks they are lovely and romantic and cheesy. But they probably won't be winning prizes any time soon nor ever be up to the standards of a published poet's art. As it stands, he is a hard-core technophobe and there is not a chance in this world that he will willingly use GenAI for anything in his life at all. But if that were not the case, he might wake up one day and start worrying about the value of his work. He might suspect her to want something "better" or might want to impress her even more. Since he feels that nearly everyone is using GenAI and he doesn't - quite understandably - wish to gift a poem that is worse than everyone else's poems, he might decide to try using GenAI for his next one. Obviously, this would be a great loss for her as the no-longer-recipient of his self-written poems, but the point here is a different one: the widespread use of GenAI has discouraged him from expressing his feelings in the form of a poem.

A person quite capable of human expression could decide not to use their own skills because they are up against those that use GenAI. Whether the output of GenAI is "better" may be beside the point. Writing a poem takes commitment & effort (the very things that contribute to making it valuable) which means that one can produce less than someone delegating it to GenAI. Producing human expression faster will give people an advantage. The concern is that we get to a point where people who choose to do things themselves will be left behind. Lance Armstrong said that his systematic use of performance enhancing drugs was not cheating because everyone else was doing it - so he wasn't gaining an unfair advantage (Carroll, 2013). If what he said was true, then it would be difficult for an aspiring cyclist to choose to compete without performance enhancing drugs. They would feel that there was no way for them to win.

Students in University now seem to be faced with a similar situation. News articles describe the widespread use of GenAI for producing essays (Marche, 2022). While we hear about the students who get caught cheating, we also hear about the students who are falsely accused of cheating (Klee, 2023). This leaves a student who decides

to do it themselves with two risks. First, they risk being accused of using GenAI anyway. Second, they produce an essay that is inferior to those that are produced with the help of GenAI, ultimately causing grades to suffer. While we don't know the exact extent of students' use of GenAI, the preliminary data is scary. According to one study, 92% of students surveyed said they used GenAI for their university essays (Weale & correspondent, 2025). This information alone would seem to discourage a student from writing an essay themselves.

Finally, preliminary studies seem to point to psychological effects of the use of GenAI that would discourage human expression. Wu et al. (2025) show that while GenAI increases task performance, it decreases intrinsic motivation and increases boredom for the completion of similar tasks without GenAI. This means that once GenAI is used, it becomes more difficult for users to express themselves on their own. It seems that the gains in productivity and convenience come at much too high a cost when what we are giving up in turn is our ability to trust, value and exercise what is essential to us: human expression.

5 Conclusion

GenAI seems to have the ability to mimic human expression successfully. It has shown itself to create visual art, literature, and human-like communication that we cannot distinguish from real human creations. Therefore, it seems to offer enhanced productivity and make life more convenient. We have shown that GenAI poses practical and ethical issues making it less convenient than we thought it would be. More importantly though, GenAI comes with a fundamental catch: Its outputs are by design inhuman and will hence always fall short of a human output. The outputs of GenAI are not art, or communication at all. Only humans can produce these things.

On top of showing the issues of use and consumption of AI on the level of individual use, we have shown that it poses threats to human expression on a societal level. The explosion of GenAI outputs causes an erosion of trust in the authenticity of human creation. This, in turn, devalues true human expression. Finally, the widespread use of GenAI may discourage humans from expressing themselves because they fear they will be at a disadvantage.

Acknowledgements None.

Authors' Contributions Both authors: Discussed original idea. Formulated first draft. Edited subsequent drafts.

Funding Open Access funding enabled and organized by Projekt DEAL. None.

Data Availability Not applicable.

Declarations

Conflict of interest None.

References

Baraniuk, C. (2014). *How online chatbots are already tricking you*. https://www.bbc.com/future/article/20140609-how-online-bots-are-tricking-you

Battisti, D. (2025). Second-Person authenticity and the mediating role of AI: A moral challenge for Human-to-Human relationships? *Philosophy & Technology*, *38*(1), 28. https://doi.org/10.1007/s13347-025-00857-w

Berger, J. (2018). *Landscapes: John Berger on Art*. Verso Books.

Bick, A., Blandin, A., & Deming, D. J. (2024). *The Rapid Adoption of Generative AI* (Working Paper 32966). National Bureau of Economic Research. https://doi.org/10.3386/w32966

Borgmann, A. (1984). *Technology and the character of contemporary life: A philosophical inquiry*. University of Chicago Press.

Burke, K. (2024). Music sector workers to lose nearly a quarter of income to AI in next four years, global study finds. *The Guardian*. https://www.theguardian.com/music/2024/dec/04/artificial-intelligence-music-industry-impact-income-loss

Cappelli, P., Tambe, P. S., & Yakubovich, V. (2024). Will large Language models really change how work is done? *MIT Sloan Management Review*, *65*(3), 48–53. https://s3.us-east-1.amazonaws.com/marketing.mitsmr.com/offers/Work24/MITSMR-Work24-Future-of-Work-Collection.pdf#page=21

Carroll, N. (2004). Art and human nature. *The Journal of Aesthetics and Art Criticism*, *62*(2), 95–107. https://www.jstor.org/stable/1559194

Carroll, R. (2013). Lance Armstrong admits doping in Oprah Winfrey interview. *The Guardian*. https://www.theguardian.com/sport/2013/jan/18/lance-armstrong-admits-doping-oprah-winfrey

Choi, C., & Annio, F. (2024). *The winner of a prestigious Japanese literary award has confirmed AI helped write her book*. CNN. https://www.cnn.com/2024/01/19/style/rie-kudan-akutagawa-prize-chatgpt

Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukahrevsky, A., Yee, L., & Zemmel, R. (2023). *Economic potential of generative AI | McKinsey*. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier#/

Coeckelbergh, M. (2017). Can machines create art? *Philosophy & Technology*, *30*(3), 285–303. https://doi.org/10.1007/s13347-016-0231-5

de Santoni, F., & van den Hoven, J. (2018). Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics and AI*, *5*(15), 1–14. https://doi.org/10.3389/frobt.2018.00015

Dwyer, M., & Laird, E. (2024). Report – Up in the Air: Educators Juggling the Potential of Generative AI with Detection, Discipline, and Distrust. *Center for Democracy and Technology*. https://cdt.org/insights/report-up-in-the-air-educators-juggling-the-potential-of-generative-ai-with-detection-discipline-and-distrust/

Evans, K. D., Robbins, S. A., & Bryson, J. J. (2023). Do we collaborate with what we design? *Topics in Cognitive Science*. https://doi.org/10.1111/tops.12682

Field, H. (2024). *AI is shaking up online dating with chatbots that are flirty but not too flirty.* CNBC. https://www.cnbc.com/2024/02/14/generative-ai-is-shaking-up-online-dating-with-flirty-chatbots.html

Floridi, L., & Nobre, A. C. (2024). Anthropomorphising machines and computerising minds: The crosswiring of languages between artificial intelligence and brain & cognitive sciences. *Minds and Machines*, *34*(1), 5. https://doi.org/10.1007/s11023-024-09670-4

Freitas, J. D., Uguralp, A. K., Uguralp, Z. O., & Stefano, P. (2024). AI companions reduce loneliness. *ArXiv*. https://doi.org/10.48550/ArXiv.2407.19096. (ArXiv:2407.19096).

Grierson, J. (2023). Photographer admits prize-winning image was AI-generated. *The Guardian*. https://www.theguardian.com/technology/2023/apr/17/photographer-admits-prize-winning-image-was-ai-generated

Hancock, J. T., Naaman, M., & Levy, K. (2020). AI-Mediated communication: Definition, research Agenda, and ethical considerations. *Journal of Computer-Mediated Communication*, *25*(1), 89–100. https://doi.org/10.1093/jcmc/zmz022

Heintz, C., & Scott-Phillips, T. (2023). Expression unleashed: The evolutionary and cognitive foundations of human communication. *Behavioral and Brain Sciences*, *46*, e1.

Hicks, M. T., Humphries, J., & Slater, J. (2024). ChatGPT is bullshit. *Ethics and Information Technology*, *26*(2), 38. https://doi.org/10.1007/s10676-024-09775-5

Hohenstein, J., Kizilcec, R. F., DiFranzo, D., Aghajari, Z., Mieczkowski, H., Levy, K., Naaman, M., Hancock, J., & Jung, M. F. (2023). Artificial intelligence in communication impacts Language and social relationships. *Scientific Reports*, *13*(1), 5487. https://doi.org/10.1038/s41598-023-30938-9

Holcomb (2024). *Media Mistrust Has Been Growing for Decades—Does It Matter?* https://pew.org/3Nl4wYc

Ivarsson, J., & Lindwall, O. (2023). Suspicious minds: The problem of trust and conversational agents. *Computer Supported Cooperative Work (CSCW)*, *32*(3), 545–571. https://doi.org/10.1007/s10606-023-09465-8

Joyce, R. (1999). Apologizing. *Public Affairs Quarterly*, *13*(2), 159–173. https://www.jstor.org/stable/40441223

Klee, M. (2023). She Was Falsely Accused of Cheating With AI -- And She Won't Be the Last. *Rolling Stone*. https://www.rollingstone.com/culture/culture-features/student-accused-ai-cheating-turnitin-1234747351/

Marche, S. (2022). The College Essay Is Dead. *The Atlantic*. https://www.theatlantic.com/technology/archive/2022/12/chatgpt-ai-writing-college-student-essays/672371/

Marvin, G., Hellen, N., Jjingo, D., & Nakatumba-Nabende, J. (2024). Prompt Engineering in Large Language Models. In I. J. Jacob, S. Piramuthu, & P. Falkowski-Gilski (Eds.), *Data Intelligence and Cognitive Informatics* (pp. 387–402). Springer Nature. https://doi.org/10.1007/978-981-99-7962-2_30

Matza, M. (2025). *Trump criticised after posting AI image of himself as Pope*. BBC News. https://www.bbc.com/news/articles/cdrg8zkz8d0o

McKendrick, J. (2024). *Generative AI may be creating more work than it saves*. ZDNET. https://www.zdnet.com/article/generative-ai-may-be-creating-more-work-than-it-saves/

Merken, S. (2025). AI hallucinations in court papers spell trouble for lawyers. *Reuters*. https://www.reuters.com/technology/artificial-intelligence/ai-hallucinations-court-papers-spell-trouble-lawyers-2025-02-18/

Microsoft, & LinkedIn. (2024). *AI at Work Is Here. Now Comes the Hard Part: Employees want AI, leaders are looking for a path forward.* (Work Trend Index Annual Report). Microsoft & LinkedIn. https://assets-c4akfrf5b4d3f4b7.z01.azurefd.net/assets/2024/05/2024_Work_Trend_Index_Annual_Report_6_7_24_666b2e2fafceb.pdf

Moffatt, B., & Hall, A. (2024). Is AI my co-author? The ethics of using artificial intelligence in scientific publishing. *Accountability in Research*, *0*(0), 1–17. https://doi.org/10.1080/08989621.2024.2386285

Mookherjee, N., Rapport, N., Josephides, L., Hage, G., Todd, L. R., & Cowlishaw, G. (2009). The ethics of apology: A set of commentaries. *Critique of Anthropology*, *29*(3), 345–366. https://doi.org/10.1177/0308275X09336703

Ogden, R., Witowska, J., & Cernohorska, V. (2023). *Technology Is Secretly Stealing Your Time. Here's How to Get It Back*. Scientific American. https://www.scientificamerican.com/article/technology-is-secretly-stealing-your-time-heres-how-to-get-it-back/

Peters, J. (2025). *Duolingo will replace contract workers with AI*. The Verge. https://www.theverge.com/news/657594/duolingo-ai-first-replace-contract-workers

Raiken, A. (2024). *Bumble founder suggests AI is the future of dating apps*. The Independent. https://www.independent.co.uk/life-style/bumble-founder-ai-dating-apps-b2545136.html

Robbins, S. (2019). A misdirected principle with a catch: Explicability for AI. *Minds and Machines*, *29*(4), 495–514. https://doi.org/10.1007/s11023-019-09509-3

Robbins, S. (2020). AI and the path to envelopment: Knowledge as a first step towards the responsible regulation and use of AI-powered machines. *AI & Society*, *35*(2), 391–400. https://doi.org/10.1007/s00146-019-00891-1

Robbins, S. (2023a). Recommending Ourselves to Death: Values in the age of algorithms. In S. Genovesi, K. Kaesling, & S. Robbins (Eds.), *Recommender Systems: Legal and Ethical Issues*. Springer Nature. forthcoming.

Robbins, S. (2023b). The many meanings of meaningful human control. *AI and Ethics*. https://doi.org/10.1007/s43681-023-00320-6

Robbins, S. (2025). What machines shouldn't do. *AI & Society*. https://doi.org/10.1007/s00146-024-02169-7

Salas Espasa, D., & Camacho, M. (2025). From aura to semi-aura: Reframing authenticity in AI-generated art—a systematic literature review. *AI & SOCIETY*. https://doi.org/10.1007/s00146-025-02361-3

Sikha, V., Korada, L., & Siramgari, D. (2023). Mastering prompt engineering: Optimizing interaction with generative AI agents. *Journal of Engineering and Applied Sciences Technology*, https://doi.org/10.47363/JEAST/2023(5)E117

Steffensen, D. S., McAllister, C. P., Perrewé, P. L., Wang, G., & Brooks, C. D. (2022). You've got mail: A daily investigation of email demands on job tension and Work-Family conflict. *Journal of Business and Psychology*, *37*(2), 325–338. https://doi.org/10.1007/s10869-021-09748-1

Tolstoy, L. (1094). *What is Art?* FUNK & WAGNALLS COMPANY. https://www.gutenberg.org/files/64908/64908-h/64908-h.htm

Weale, S., & correspondent, S. W. E. (2025). UK universities warned to 'stress-test' assessments as 92% of students use AI. *The Guardian*. https://www.theguardian.com/education/2025/feb/26/uk-universities-warned-to-stress-test-assessments-as-92-of-students-use-ai

Weatherbed, J. (2025). *That AI-restored Beatles song won Grammy for Best Rock Performance*. The Verge. https://www.theverge.com/news/604970/beatles-ai-restored-song-now-and-then-grammy-win

World Economic Forum. (2025). *Future of jobs report*. World Economic Forum. https://reports.weforum.org/docs/WEF_Future_of_Jobs_Report_2025.pdf

Wu, S., Liu, Y., Ruan, M., Chen, S., & Xie, X. Y. (2025). Human-generative AI collaboration enhances task performance but undermines human's intrinsic motivation. *Scientific Reports*, *15*(1), 15105. https://doi.org/10.1038/s41598-025-98385-2