Philosophy & Technology (2026) 39:88
https://doi.org/10.1007/s13347-026-01099-0

COMMENTARY

Inspirational Responsibility, Intellectual Property, and AI Training

Jonas Bozenhard[1]

Received: 20 April 2026 / Accepted: 23 April 2026

Abstract

Generative AI systems, including large language models and image or video generators, are typically trained on data scraped from the internet. Since the creators of the works used for AI training have often neither consented to nor been compensated for such use, this has led to allegations of "theft". At the same time, since generative AI models do not store copies of their training data, there is a widespread view that they no not necessarily infringe copyright. Moreover, in human creative practice, drawing inspiration from and learning from the works of others is common, if not unavoidable. The central challenge, then, is to formulate an intellectual property right that allows creators to restrict AI training without limiting human learning and creativity. McIntyre (2026) proposes new strategies for addressing this challenge, most notably through an argument relying on proprietary rights. Critically engaging with this approach, I argue that it overlooks the power dynamics in the art world and leaves creators in a vulnerable and dependent position. Instead, I outline a different pathway to restricting AI training that centres on what I call "inspirational responsibility". This concept captures the responsibility creators bear for how they use and engage with the work of others. Since human creators can be held inspirationally responsible, whereas AI models cannot, the notion of inspirational responsibility provides a normatively relevant basis for restricting AI training without imposing parallel constraints on human learning and creativity.

Keywords AI ethics · Intellectual property · Training data · Creativity · Responsibility

✉ Jonas Bozenhard
jonas.bozenhard@tuhh.de

[1] Hamburg University of Technology, Hamburg, Germany

Published online: 01 May 2026

1 Introduction

Generative AI systems, such as large language models and image or video generators, are trained on vast amounts of data collected from the internet. Since the creators of the texts, images, and videos used to train AI systems often neither consent to nor receive compensation for the use of their works, AI companies have been accused of "theft". As a recent newspaper article renders it: "Today's large-scale AI systems are founded on what appears to be an extraordinarily brazen criminal enterprise: the wholesale, unauthorised appropriation of every available book, work of art and piece of performance that can be rendered digital" (Funder & Powles, 2025). For the claim that AI training involves theft to be compelling, it must be shown that it violates creators' property rights. Based on the argument that generative AI models encode statistical patterns extracted from the data they were trained on rather than containing actual copies of them, it has been contended that these models do not necessarily infringe the copyright of those whose works were used in their training (see Murray, 2023; Samuelson, 2023; Shoemaker, 2024). In addition, as emphasised by Shoemaker (2024), it is common, if not inevitable, in human creative practice for artists to draw inspiration from, learn from, and imitate stylistic features of other artists' works. Therefore, the central challenge lies in formulating a property right that allows creators to restrict AI training without forestalling human learning and creativity.

Addressing this challenge, James H. McIntyre, in his thought-provoking article "The Right to Restrict AI Training" (2026), advances two arguments intended to provide (moral) restrictions specific to AI training, without extending to human learning. The paper deserves praise for its resourcefulness, for its thoughtful engagement with the problem space, and for bringing fresh perspectives to the debate. However, in what follows, I will point to some limitations of McIntyre's approach. On this basis, I will outline a different pathway to restricting AI training that centres on what I will refer to as "inspirational responsibility".

2 Critique of McIntyre's Arguments

McIntyre's first argument (2026: 8–13) builds on the idea of invoking patent-like intellectual property rights that would allow creators to block others from using their work more extensively than under copyright protection. This would require AI companies to obtain permission to use such works for AI training. To address the central challenge mentioned above, McIntyre argues that, in the context of human learning, such patent-like property rights are typically outweighed by a countervailing right to freedom of thought. Unlike humans, AI models have no such right to freedom of thought since they lack moral status. Thus, in the case of AI training, the right to freedom of thought does not apply so that creators' property rights remain intact. However, due to the nature of authorship, McIntyre is in the end sceptical that creators actually possess such far-ranging, patent-like intellectual property rights.

For this reason, McIntyre advances a second argument, framed as the more compelling of the two, as it avoids relying on the overly extensive property rights on

which the first argument depends. Instead, the second argument draws an analogy between content-hosting websites and art galleries. The latter have the proprietary right to prohibit photography or the use of other technologies within their space. Similarly, website proprietors have the right to restrict copying technologies. More specifically, they have a right to block the use of web crawlers for AI training, while allowing other uses of technologies (for example, those necessary for regular human browsing).

The primary aim of the second argument is to show that "*creators* have the moral right to restrict which copying technologies may be used to capture their work" (McIntyre, 2026: 1, emphasis added). Yet, this right is effectively granted to *proprietors*. However, creators are not always identical to the proprietors of the sites where their creations are displayed. McIntyre acknowledges this, but considers it unproblematic. He argues that "creators are sometimes able to exercise significant leverage over proprietors" and that, "via negotiations with proprietors, creators may acquire protections against their work being copied by particular technologies" (McIntyre, 2026: 15). However, this argument neglects the power dynamics typically at play in the artworld. Artists heavily depend on galleries for exposure to audiences and collectors, artistic legitimacy, and market valuation. So, while a selected few very famous artists might be able to exercise some leverage over proprietors, most artists are not in such a powerful position, especially when confronted with the more prestigious and influential galleries. The same issue applies to the influence creators can exert over the most visible and heavily trafficked online platforms. Thus, the proposed approach leaves creators in a vulnerable and highly dependent position regarding the protection of their work from AI training.

This also becomes apparent when considering that proprietors can change how they regulate their website or, in the case of some powerful platforms, even use the content posted on them to train their own AI models. For example, Meta uses public posts and photos from Instagram and Facebook dating back to 2007 for AI training, and in many countries users have no practical way to opt out (Weatherbed, 2024). Social media platforms like Instagram have become a crucial platform for creators to showcase their works online (Bishop, 2025). Given the extensive reach these widely used platforms provide, it is extremely difficult, if not practically impossible, for many creators to abandon them in favour of a less ubiquitous alternative. In such cases, creators are confronted with an unfortunate trade-off: allowing their content to be used for AI training or sacrificing the exposure and visibility the platform provides. Hence, if the goal is to establish moral restrictions that protect *creators* and give them a more robust and widely applicable right to prevent their works from being used for AI training, an approach based on proprietary rights seems unsuitable.

3 Inspirational Responsibility

In the remainder, I would like to sketch an alternative pathway. This approach highlights an important difference between how AI systems and humans draw on existing works. McIntyre (2026: 4, 7) argues that the relevant ways of drawing on the works

of other creators, for example, through imitation and inspiration, all rely on "learning", which he defines as the extraction of patterns from information. He briefly considers the idea of identifying a critical difference in how humans learn compared to machines that could justify restricting AI training exclusively, but he maintains that "it's not remotely clear what the qualitative difference would amount to or why it should be normatively relevant" (McIntyre, 2026: 7).

However, even if one follows McIntyre's characterisation of learning, one can still point to a crucial distinction that implies a fundamental, morally significant difference in the way humans and AI systems learn and draw on antecedent work: Humans are persons, whereas AI models are not. McIntyre's first argument, which relies on the idea that AI lacks moral status, implicitly reflects this view. Human learning (or human pattern extraction) is typically guided by a person's intentions, desires, and motivations, and a significant portion of what is learned is accessible to their consciousness and thought. None of this applies to machine learning, since AI systems lack intent, desires, motivation, consciousness, and – as McIntyre (2026: 10) himself notes – probably even thought.

As persons, humans can be held morally responsible. Since AI systems are not persons, they are widely denied the capacity to be bearers of responsibility (see e.g. Kiener 2025b, Chaps. 6 and 7). (This circumstance has given rise to the current debate over AI responsibility gaps (see e.g. Matthias, 2004; Da Silva, 2024)). I would like to introduce a specific type of responsibility that may prove productive in addressing the central challenge of restricting AI training without limiting human learning and creativity. The responsibility under consideration is one that creators have for how they use and engage with the works of other creators in their creative process. One may refer to it as "inspirational responsibility".

Humans can be held inspirationally responsible when they use and engage with the works of other creators, for example, when they imitate existing styles or produce what McIntyre calls "works of inspiration". Thus, in scope, inspirational responsibility concerns how creators relate to one another and to their works. It combines aspects of moral responsibility and aesthetic responsibility which is typically treated as non-moral (Wolf, 2016). In other words, inspirational responsibility encompasses both aesthetic and moral dimensions: The aesthetic dimension considers whether a creator's engagement with antecedent work is innovative and meaningful. By contrast, the moral dimension concerns whether such engagement respects values such as attribution, recognition, and respect for the contributions and interests of other creators. In cases of plagiarism, for example, concerns about the moral dimension typically take precedence, whereas, in cases of artistically unoriginal work, concerns about the aesthetic dimension are primary.

Responsibility is often framed in terms of blameworthiness and praiseworthiness (Fischer & Ravizza, 1998; Pereboom, 2001; Strawson, 2003; Shoemaker, 2015). Plagiarists, for example, are typically subject to blame (moral dimension). By contrast, James Joyce can be praised for his creative engagement with Homer's *Odyssey* in *Ulysses*, while certain fan fiction authors might be blamed for adding nothing interesting or inventive to the original work (aesthetic dimension). Since AI systems are not persons, they cannot be praised or blamed in the same way as

human creators.[1] On a different account, responsibility is understood in terms of answerability which requires from a person to "give her (justificatory) reasons for thinking, feeling, or acting in the way she has" (Smith, 2015: 103). Human creators can give justificatory reasons for how they engage with the works of others. Depending on the context, these reasons can relate to aesthetic or moral concerns. AI systems are widely denied the status of morally answerable agents (Kiener, 2025b). Hence, on both accounts of responsibility, AI systems do not qualify as bearers of responsibility.

In cases of stylistic imitation or works of inspiration, the presence of inspirational responsibility is intuitively plausible. However, it is even present in more passive forms of learning that do not directly result in a specific creative output. For example, under certain circumstances, someone might be blamed or asked to provide justificatory reasons for reading a book with morally problematic content, for example, a book containing misogynistic jokes. By contrast, an AI model would not be blamed or asked to give justificatory reasons for extracting objectionable patterns from data. (One might reply that the developers of the relevant AI system can be held responsible for including, or failing to exclude, objectionable content in the training data. Then, however, responsibility may be so widely distributed that we face the problem of an abundance of responsibility (Kiener, 2025a)). So, even in basic instances of learning, the identified difference between human and machine learning becomes evident: only humans can bear responsibility for what they learn.

In summary, the key difference in how humans and AI systems use, engage with, and learn from the works of other creators can thus be framed as follows: as persons, human creators can be held inspirationally responsible, whereas AI systems cannot. This difference is not arbitrary, but normatively significant since inspirational responsibility is crucial for sustaining healthy and effective creative practices. Wolf (1981: 400) argues that abandoning the notion of responsibility would undermine our practice of moral thought. Likewise, abandoning the notion of inspirational responsibility would erode concerns for attribution, originality, creative accountability, collegial trust, and authorship, thereby drastically destabilising and threatening creative practices. For that reason, inspirational responsibility merits protection.

4 Outlook

The notion of inspirational responsibility offers a promising basis for restricting AI training, while leaving human learning and creativity intact, as human creators can be held inspirationally responsible, whereas AI models cannot. Of course, this approach requires further elaboration. Crucially, it remains to be explored what kind of concrete and plausible property right could be formulated based on inspirational responsibility, particularly in the face of the challenges highlighted by McIntyre and others.

[1] For a dissenting perspective, see Altehenger et al. (2024).

Authors' Contributions Listed author is the sole author.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability Not applicable.

Declarations

Ethical Approval and Consent to Participate Not applicable.

Consent for Publication Not applicable.

Competing Interests Not applicable.

References

Altehenger, H., Menges, L., & Schulte, P. (2024). How AI systems can be blameworthy. *Philosophia*, *52*(4), 1083–1106. https://doi.org/10.1007/s11406-024-00779-5

Bishop, S. (2025). Influencer creep: How artists strategically navigate the platformisation of art worlds. *New Media & Society*, *27*(4), 2109–2126. https://doi.org/10.1177/14614448231206090

Da Silva, M. (2024). Responsibility gaps. *Philosophy Compass*, *19*, 9–10. https://doi.org/10.1111/phc3.70002.e70002

Fischer, J. M., & Ravizza, M. (1998). *Responsibility and control: A theory of moral responsibility*. Cambridge University Press.

Funder, A., & Powles, J. (2025). Tech companies are stealing our books, music and films for AI. It's brazen theft and must be stopped. The Guardian. https://www.theguardian.com/commentisfree/2025/sep/10/tech-companies-are-stealing-our-books-music-and-films-for-ai-its-brazen-theft-and-must-be-stopped. Accessed 4 Mar 2026

Kiener, M. (2025a). AI and responsibility: No gap, but abundance. *Journal of Applied Philosophy*, *42*(1), 357–374. https://doi.org/10.1111/japp.12765

Kiener, M. (2025b). *Moral responsibility and artificial intelligence*. Hart Publishing.

Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and Information Technology*, *6*(3), 175–183. https://doi.org/10.1007/s10676-004-3422-1

McIntyre, J. H. (2026). The Right to restrict AI training. *Philosophy & Technology*, *39*(1), 25. https://doi.org/10.1007/s13347-025-01032-x

Murray, M. D. (2023). Generative AI art: Copyright infringement and fair use. *SMU Science and Technology Law Review*, *26*(2), 259. https://doi.org/10.25172/smustlr.26.2.4

Pereboom, D. (2001). *Living without free will*. Cambridge University Press.

Samuelson, P. (2023). Generative AI meets copyright. *Science*, *381*(6654), 158–161. https://doi.org/10.1126/science.adi0656

Shoemaker, D. (2015). *Responsibility from the margins*. Oxford University Press.

Shoemaker, E. (2024). Is AI art theft? The moral foundations of copyright law in the context of AI image generation. *Philosophy & Technology*, *37*(3), 114. https://doi.org/10.1007/s13347-024-00797-x

Smith, A. M. (2015). Responsibility as answerability. *Inquiry : A Journal Of Medical Care Organization, Provision And Financing*, *58*(2), 99–126. https://doi.org/10.1080/0020174X.2015.986851

Strawson, P. F. (2003). Freedom and resentment. In G. Watson (Ed.), *Free will,* 72–93). Oxford University Press.

Weatherbed, J. (2024). Meta fed its AI on almost everything you've posted publicly since 2007. The Verge. https://www.theverge.com/2024/9/12/24242789/meta-training-ai-models-facebook-instagram-photo-post-data. Accessed 4 Mar 2026.

Wolf, S. (1981). The importance of free will. *Mind*, *90*(359), 386–405.

Wolf, S. (2016). Aesthetic responsibility. *The Amherst Lecture in Philosophy*, *11*, 1–25.