



The work of art in the age of artificial intelligibility

John McLoughlin¹

Received: 2 February 2023 / Accepted: 14 December 2023

© The Author(s) 2024

Abstract

The emergence of complex deep-learning models capable of producing novel images on a practically innumerable number of subjects and in an equally wide variety of artistic styles is beginning to highlight serious inadequacies in the ethical, aesthetic, epistemological and legal frameworks we have so far used to categorise art. To begin tackling these issues and identifying a role for AI in the production and protection of human artwork, it is necessary to take a multidisciplinary approach which considers current legal precedents, the practice of software engineering, historical attitudes towards technological innovation and a sustained technical analysis of the models themselves. This paper queries the location and nature of substantive artistic work in the developmental stages of an AI-generated image, offering critiques of existing assumptions and posing questions for future research. The emergence of convincing AI creative output, artistic or literary, has significant long-term implications for the humanities, including the need for re-appraisal of foundational ideas about authorship and creativity in general. The effects of artificial intelligence, whether generalised or task-specific, cannot be ignored or displaced now that easy-access, scalable image and text production is a reality.

Keywords Photography · AI · GAN · DALL-E · Gérôme · Copyright · Programming · Intentionality

1 Introduction

Photography is an art. It forces artists to discard their old routine and forget their old formulas. It has opened our eyes and forced us to see that which previously we have not seen: a great and inexpressible service for Art. It is thanks to photography that Truth has finally come out of her well. She will never go back (Gérôme 1902).

Is there an ontological difference between early computer-generated art, net art and the more recent forms of AI-driven art? Or is it just a difference of degree, i.e. of the mode and intensity of technological entanglement? Should the recent applications of AI to image making and image curating encourage us to (re)turn to bigger questions concerning the very purpose of artistic production? What are art, photography and other forms of image making for? Who are they for? Does art exist outside the clearly designated realm of human cultural practice? Will AI create new conditions and

new audiences for art? What will art ‘after’ AI look like? Who will be its recipient? (Zylinska 2020).

Eight years prior to his above statement on photography, Jean-Léon Gérôme painted *Truth Coming Out of Her Well to Shame Mankind* (*La Vérité sortant du puits armée de son martinet pour châtier l'humanité*). First attributed to the ancient philosopher Democritus by Diogenes Laertius et al. (1931), the figure of truth as the goddess Veritas being cast or lost down an abyssal well re-emerged and has endured since the early modern period (Oxford 2009). In Gérôme’s painting, Veritas is unleashed from her well brandishing a vicious-looking whip, chastising and embarrassing the viewer. In this painting, the goddess does not hold her customary hand mirror; the subtlety and comparative gentility of self-reflection is discarded in favour of a powerful, resurgent Truth who accuses and rebukes. Her nudity betrays her identity—the naked form is honest, if anything—but also follows in the tradition of academic painting. Veritas is an example of the ‘ideal’ female nude in its proportion and modelling. The evolution of the artistic nude is one of the reasons Gérôme eulogised photography towards the end of his life (Bowyer 2010), as he believed in the power of photography as a tool for achieving perfection. When readers of Gérôme and Émile Bayard’s nude photograph collection

✉ John McLoughlin
mcloughlinj@cardiff.ac.uk

¹ School of English, Communication and Philosophy,
Cardiff University, 90a Kimberly Road, Pen-Y-Lan,
Cardiff CF23 5DN, UK

Le Nu Ésthetique complained of bodily imperfections in the models, Gérôme replied that each, though imperfect, was an example of perfection in at least one respect (Cate 2000).

The relationship between photography and truth may seem a natural outcrop of the form; photographed images are intuitively ‘true’ in a way that paintings are not. The mechanical and chemical processes of early photography, whilst far from inviolable, did not invite the same decision-making processes which characterise painting or digital editing. Gérôme did not, however, proselytise about the virtues of photography just for its accuracy as a representative medium, but for its impact on those very decision-making and representational procedures. Representation of the figure of the horse in full gallop, for example, was a problem which had challenged artists for centuries. Though well-grounded assumptions could be reached from analysis of animal musculature, it was not until the advent of photography that ‘phases of locomotion were revealed which lay beyond the visual threshold’ (Scharf 1986). It was Gérôme who encouraged Lieutenant-Colonel Émile Duhoussset to publish a set of ‘preliminary discussions’ on the problem, wherein Duhoussset used Étienne-Jules Marey’s chronophotographic recordings of galloping horses to produce sketches of the various stages of motion (Scharf 1986).

Gérôme’s was, however, far from the only position on photography. The unveiling of the Daguerreotype led to uncertainty and outright snobbery over the place this new technology would assume. This reactionary element included an often-scathing critique of the very effect Gérôme would later celebrate: increased fidelity in the representative power of painting. In his review of the 1859, Paris Salon Charles Baudelaire characterises the obsessive fidelity to ‘real life’ engendered by photography as something obscene and deleterious:

A vengeful God has granted the wishes of this multitude. Daguerre was his messiah. And now the public says to itself: ‘Since photography gives us all the guarantees of exactitude that we could wish (they believe that, the idiots!), then photography and art are the same thing.’ From that moment squalid society, like a single Narcissus, hurled itself upon the metal, to contemplate its trivial image [...] In arranging and grouping together buffoons male and female, tricked up like butchers and washer-women in the carnival, in begging these heroes to be so good, during the time necessary for the operation, as to hold their smiles for the occasion, the photographer flatters himself that he is rendering scenes of ancient history tragic or noble. (Baudelaire 1880)

where Gérôme saw an opportunity to enhance painting, Baudelaire saw debasement and narcissism. A product of mere machinery and chemical interactions, photography

is rendered the domain of the incompetent: ‘As the photographic industry,’ writes Baudelaire, ‘was the refuge of all *peintres manqués*, of too slender talent or too lazy to complete their studies, this universal craze bears not only the mark of blindness and imbecility, but also has the flavour of vengeance’ (Baudelaire 1880). The painters of the Salon, spoiled by prizes, privileges, fame, and the popularity of realism, are producing a banal form of painting which no longer requires a true aesthetic imagination (Raser 1989). Timothy Raser phrases this eloquently:

What Baudelaire detects in the art of 1859 is interest: its emphasis on reality, its effort to acquire the visible world indicates desire, not aesthetic pleasure. The Salon artists produce realistic images of those things their buyers desire, and such pandering has nothing to do with beauty (Raser 1989)

Of particular note is the difference between interest and aesthetic pleasure as a response to an artwork. Raser identifies the rhetorical and aesthetic trends which would develop around photography as nascent in Baudelaire’s vitriolic response to it. For Roger Scruton, this is the difference between an ‘intentional’ and ‘causal’ relationship to the subject; paintings are intentional insofar as they represent the intentional act of the painter and their decision-making process.

Photographs, in contrast, are causal: ‘if a photograph is a photograph of a subject, it follows that the subject exists, and if x is a photograph of a man, there is a particular man of whom x is the photograph’ (Scruton 1981). This notion gets ahead of the natural counter-argument that photographers do in fact make artistic decisions: the posing of models, the theming and dress of the image, the size of the photograph. UK legal scholarship accounts for this conflict: Whilst all of this is true, none of it changes the fact that the resulting photography is a photograph of something, rather than the product of an artist’s intentionality towards something. Baudelaire’s denunciation of photography, writes Raser:

indicates how quickly Baudelaire understood the rhetoric of photography [...] before a photograph, naïveté returns, and one regains trust in a “natural” meaning: one forgets the difference between signifier and signified, and this allows one to seek in the photograph the satisfaction of the real world (Raser 1989).

As photography becomes more popular and its practices more intertwined with those of traditional painting, the aesthetic quality of those practices becomes irreversibly compromised. The availability of unremittably accurate representations of nature disabuses artists of their creative imagination in two ways: artists who use photographs are simply reproducing photographs, and their paintings can be judged retroactively against photographs.

The result is the diminishment of the narrative capacity of paintings; naiveté is not conducive to imagination. You cannot argue with a photograph, so a photograph cannot be allegorical or symbolic in the same way that a painting can: ‘Philosophy is idolatrous,’ writes Raser, ‘because it conflates signifier and signified, refusing to admit the autonomous existence of the latter. Further, photography takes the place not simply of objects, but also of those forms of art it most resembles, painting and drawing’ (Raser 1989). Raser then quotes Baudelaire himself, translated here: ‘If photography is allowed to replace art in some of its functions, it will soon have quite corrupted or supplanted it’ (Baudelaire et al. 1975).

This is the other great anxiety concerning the relationship between art and photography: replacement. By virtue of its novel technological processes and its remarkably, even magically accurate results, photography was threatening to render painting obsolete. The 1859 Salon was the first to allow photographs alongside paintings and sculptures, and 2 years before the Société Française de Photographie organised an exhibition in the salons of Gustave le Gray (de Font-Réaulx 2012). Portraiture quickly became ‘subordinated to the aesthetics of photography’ (de Font-Réaulx 2012), and artists were forced to pivot away from the scale and representational qualities of traditional portraiture to distinguish themselves. What followed was a period of rapid artistic innovation; portraiture was forced away from the traditional aesthetic markers which photography had adopted: pose, attitudes, attire and décor. Likeness instead ‘retreated before a symbolist or expressionist portraiture, the real elements receding before pure pictorial innovation’ (de Font-Réaulx 2012). Daguerre had pulled truth from her well, and representational art was scrambling to innovate and to justify its continued existence.

The above is not designed to minimise or disregard the field of photography; photographers ‘deliberately and purposefully manipulate their apparatus in such a way as to generate objects with appearances which in some way lead spectators to recognise their subjects’ (Brook 1983): a description of artistry as useful as any other. The intention is to highlight an omission: analysis of where the actual artistic work occurs. Understanding this, and accurately dissecting the aesthetic, technical and creative implications of artistic labour will allow us to better understand AI art and to predict—or at least prepare for—its eventual place in the paradigm of human creative output. Such an understanding is complex and intimidating: AI lies at the bleeding edge of modern software engineering and asks multifaceted questions about aesthetics and the epistemology of creativity. Too much knowledge in one field and a shortfall in another can leave the researcher looking like a luddite or a naïve futurist. It is possible, however, to take lessons from history into the uncertain future.

2 Creative labour in the age of AI

One hundred and sixty-two years after the Salon of 1859, researchers at the OpenAI artificial intelligence laboratory in San Francisco unveiled DALL-E and CLIP, machine learning models for the generation and ranking of algorithmically produced images. An updated version was then made available to 1 million people as a ‘freemium’ beta; users given credit to produce a set of images on a monthly basis with the option to purchase additional credits (OpenAI 2020). The same model was adopted by others including David Holz’s Midjourney—used by The Economist to produce the cover image for its June 11 2022 issue. The same issue declared that ‘foundation models’ (of which DALL-E and its like are examples) promise a revolution in ‘high-status’ brainwork of the kind left unaffected by the industrial revolution:

For years it has been said that AI-powered automation poses a threat to people in repetitive, routine jobs, and that artists, writers and programmers were safer. Foundation models challenge that assumption. But they also show how ai can be used as a software sidekick to enhance productivity. This machine intelligence does not resemble the human kind, but offers something entirely different. Handled well, it is more likely to complement humanity than usurp it. (*The Economist* 2022)

Without considerable time, cooperation and funding, even approaching these neural networks from a position in the arts and humanities is an imposing task. This is also a phenomenal opportunity for interdisciplinary cooperation, despite significant obstacles. In the case of DALL-E and OpenAI, such analysis is made difficult by lack of access. Whilst they have released the code for DALL-E’s variational autoencoder (VAE), OpenAI have so far not seen fit to reveal the actual image transformer (Ramesh et al. 2021) which turns the underlying GPT (Generative Pre-trained Transformer) module into one which can generate images. From its announcement paper, we know that DALL-E is an autoregressive, decoder-only sparse transformer (Ramesh et al. 2021). Other groups are more forthcoming, including Craiyon (formerly DALL-E Mini) and Stable Diffusion, the latter of which undergirds multiple ‘freemium’ image-generation APIs. Technical, collaborative analysis of these models and their version of artistic work thus remains a possibility for those without the resources to engage big business.

Though tinged with optimism, the Economist article nevertheless touches on some of the deep-rooted existential anxieties engendered in the artistic community by the release and rapid—seemingly exponential—development

of artistic neural networks. Early technical analysis of DALL-E 2's output described its results as 'stunning', and researchers were particularly impressed with the model's ability to produce various perspectives and a wide variety of artistic styles (Marcus et al. 2022). AI-generated images already allow artists to rapidly test almost every aspect of a composition; poses, colours, values, perspectives. Just as Baudelaire believed photography was shackling painters to realist principles, so too may AI art engender homogeneity in artistic composition if, as discussed later in the essay, the underlying algorithmic principles remain unchanged.

The limitations of current systems mean that artists can struggle to achieve desired compositions from natural language prompting alone. They are bottlenecked by the functionality of the models and the capacity for natural language to succinctly express dense visual requirements: it is a translation issue as well as a technical one. As a result, the artist can either compromise, by accepting the received compositions, or perform the work of digitally altering an image to better suit their original vision. Both constitute a melding of artistic work with that of the algorithm and all its attendant labourers (original artists, programmers, mathematical functions), with the only difference being the extent to which the end-of-line artist can be said to have contributed to the resultant object. What occurs is an advanced form of the creative compromise induced by the advent of photography: the intentionality of the artist is subsumed by the collective intentionality of the neural network and its antecedents. Practically speaking, improvements to this process will likely follow along both axes: better parsing by the LLMs and more informed prompting. Such improvements, however, may never overcome the fundamental challenges of 'converting' natural language into images. If they can be said to truly exist, sites of mutual untranslatability between language and art represent a kind of 'hard' problem of AI art. McCormack et al. outline that problem neatly: text-to-image (TTI) systems with a 'literal' understanding of image must, when prompted by language, produce images which are themselves comprehensible in that language. Extraneous 'meaning, intention or encoded information assigned to a specific textual construct, be it metaphorical or culturally charged, will be lost in translation' (McCormack et al. 2023). McCormack concludes convincingly that current AI art enjoys a parasitic relationship with human creativity, one which not only fails to benefit, but actively harms human artists. They also, like this author, consider the artistic work of prompt generation to be essentially minimal—of which more shortly. Not within the purview of McCormack's study, however, is the artistic work performed in the construction of TTI generation. Tracing the nature and experience of work throughout the development cycle of a natural language model will allow us to better understand and delineate the extent of that parasitism.

If one assumes a future where such models proliferate throughout the creative industries—not itself a foregone conclusion, particularly considering the neo-Ludditism spreading in the visual arts—then the delineation of artistic contributions will enter the world as deeply impactful remunerative models. These are, pessimistically, the likely 'new conditions' of art; Zylinska observes how AI can exacerbate unfair labour conditions, particularly the precarity in the digital economy (Zylinska 2020). The likely result of AI art is an intense deepening of this precarity. As with the advent of photography, AI threatens to completely upend the art industry, particularly in fields like concept art and digital design where the product of an artist's work is rapidly becoming less distinguishable from that of the AI models. Cloud and dispersed computing have made the computing power required to run large neural networks available to end-users at scale. Indeed, users whose lap- or desktop computer has a compatible graphics card or—at a push—CPU can already run Stable Diffusion on their own systems. There is a growing sense of inevitability—and not a little despair—in the public perception of AI and creative work; after winning the 2022 Colorado State Fair's annual art competition with his impishly named Midjourney creation *Théâtre D'opéra Spatial*, artist Jason Allen was unapologetic: "This isn't going to stop," Mr. Allen said. "Art is dead, dude. It's over. A.I. won. Humans lost" (Roose 2022). In the same month, Rappler.com performed a series of interviews with creative professionals, all of whom expressed real concerns about the effects of easy-access AI on their industry. Emil Mercado, listed as a 'Creative Director', paints a compelling—if distressing—picture:

"From the client's perspective, why would they go to the trouble of hiring an artist?" he continued. "They can be difficult to talk to, temperamental, can't understand instructions, aren't able to meet the deadline, etc. But with AI, they just punch in the prompt words and they can have results in less than a minute. No more back-and-forth for a period of several weeks, and best of all, it's cheap. This could also mean letting go of in-house artists and design studios. No overhead costs. No benefits. Nothing." (de Leon 2022)

These responses each make an implicit assumption about the relationship between technology and the artwork: namely that the artistic work is occurring in the deep, iterative processes of the neural network itself. This assumption is a deeply problematic one. As rapid and seemingly magical as the generative process is for the end-user, the fact is that a huge amount of asynchronous, distributed work has occurred in the lifetime of each image. Photography has advanced far, far beyond the capabilities of the daguerreotype, but one would be hard-pressed to find someone who believes that the labour performed in producing a hi-resolution image of

a painting was performed exclusively by the photographer. Insofar as a photograph is always a photograph of something—a causal rather than an intentional relationship—it stands to reason that the original painter represents the predominant site of artistic effort performed in the production of that image.

This notion is enshrined in UK copyright law: a photograph of an extant artwork is considered a copy of that artwork, be it physical or digital (Copyright, Designs and Patents Act 1988). Copyright law, particularly in the digital age, represents more of a compromise between conflicting interests than an endpoint in the discussion about artistry and ownership (Stokes 2022), but nevertheless represents a valuable launching-off point for critiquing the notion of artistic work in the digital age. The Copyright, Designs and Patents Act 1988 makes clear that the artist has the ‘moral right’ to be cited and compensated appropriately if their original artwork is represented almost anywhere. It is also the ‘moral right’ of an artist to protect the integrity and reputation of their work by preventing unauthorised copying in any format (Greenberg and Reznicki 2015). The Act thus restricts secondary representations: if a photograph of an artwork appears in a film, then copyright for the image of the image remains with the original artist (Copyright, Designs and Patents Act 1988). The 1988 act does tackle the notion of ‘computer-generated’ artworks—though only briefly:

“computer-generated”, in relation to a work, means that the work is generated by computer in circumstances such that there is no human author of the work. (Copyright, Designs and Patents Act 1988)

This addendum leans rather heavily on the phrase ‘no human author’. In the place of a traditional ‘author’, computer-generated works are considered by the Act to be authored by ‘the person by whom the arrangements necessary for the creation of the work are undertaken’ (Copyright, Designs and Patents Act 1988). Despite the internal contradictions apparent in the absence or presence of human authorship, the Act is moving in the right direction in its search for authorship: it just does not go far enough.

This is perhaps a result of age: 1988 was—in technological terms—many lifetimes ago, and the quoted definitions have seen no amendments since. At this time, algorithmic art still meant fractal mathematics and physically executed algorithmic exercises: in 1982, Atari invested significant sums of money into fractal computer graphics, and procedurally generated environments remain a feature of modern games, none of which has eliminated artists as a category (du Sautoy 2019). Nevertheless, the motion is a useful one: to peer back into the ‘arrangements necessary for the creation of the work’ and identify who is responsible for each. Simon Stokes notes the ambiguity of this wording (Stokes 2022)

and identifies a useful precedent in *Express Newspapers plc v. Liverpool Daily Post & Echo plc*, wherein the defendant (the Post) was accused of stealing computer-generated grids of letters: not the program itself, just the resulting grids. The plaintiff’s case hinged on the notion that the grids constituted, by virtue of the labour and skill which went into making them, a ‘literary work’, based on which the Post was being sued for copyright infringement. The defence countered that, whilst the computer program itself might be protected under copyright, the word grids it produced were not a work of which it could truly be said that Mr. Ertel, the programmer, was the author (Whitford 1985). The judge rejected this notion, offering the following explanation:

The computer was no more than a tool. It is as unrealistic (to suggest that the programmer was not the author) as it would be to suggest that, if you write your work with a pen, it is the pen which is the author of the work rather than the person who drives the pen. (Whitford 1985)

Consider our photograph of a painting. It would seem facetious to claim that neither the painter nor the photographer but rather that the camera itself, the tool, is responsible for an artwork. However, a significant amount of technical choices have undeniably been made by the camera manufacturer in order to facilitate the production of images. This is the position taken by those who attribute these artworks, implicitly or explicitly, to ‘the AI’. If we are to come to terms with AI art, we have to discern whether the algorithm performs the same kind of artistic work as the photographer and painter, or whether it simply represents images like a tool.

3 Plagiarism and programming: evolving ideas

There are substantive differences between the ways artists and programmers view contribution: usage of another person’s work is baked into both practices, albeit in different ways and with different moral frameworks. Ryan Donovan at Stack Overflow describes it as ‘an open secret among coders’ that publicly available code—often given as answers to user questions—ends up in commercial production (Donovan 2020), and researchers have described the cross-pollination of code in positive terms:

Code reuse has well-known benefits on code quality, coding efficiency, and maintenance. Open Source Software (OSS) programmers gladly share their own code and they happily reuse others’. Social programming platforms like GitHub have normalized code foraging via their common platforms, enabling code search and

reuse across different projects. Removing project borders may facilitate more efficient code foraging, and consequently faster programming. (Gharehyazie et al. 2019)

Codebases can be extremely large, and sometimes the best solution to a problem in a particular programming language has already been discovered. Nevertheless, the detection of cloned code has been an active area of research (Juergens et al. 2009) within the academy and the corporate world. Software engineering necessarily brings students into contact with extant code: this is both necessary and productive for teaching the fundamentals. Student programmers use the same open-source platforms as their professional counterparts, and viable solutions to introductory problems are limited. As a result, plagiarism becomes difficult to define (Modiba et al. 2016). The ready and necessary use of the Internet in educational contexts muddies our understanding of plagiarism (Dominguez et al. 2019), with the result being a potentially widening gap between the way software engineers and those of us in the humanities view authorship.

This divergence can already be seen in current research:

The decisions of any AI system generating art or design are independent of humans and must be judged only on the ground of the final outcome. A work of art will be judged ‘artistic’ to the extent that humans will recognize an artistic intent in the work itself. In particular, the introduction of random processes makes an AI-generated artwork fully independent of human creativity. (Terzidis et al. 2022)

The principle seems sound enough; the introduction of stochasticity places the designer (artist or programmer) at a sufficient intentional remove from the resultant image that it can no longer reasonably be described as their labour. In an article for *The New Yorker* speculative fiction star Ted Chiang described Chat-GPT as ‘a blurry JPEG of the web’ (Chiang 2023), highlighting the extent to which such LLMs, though they contain stochastic processes, are nevertheless models of compression and recall, subject to the relatively rigid laws of compression algorithms. Chiang explains the ‘hallucinations’ of such models and their inability to perform complex mathematical calculations, in terms of ‘lossy’ compression and mathematical averaging, seemingly far removed from traditional creativity. His intention is not to deflate enthusiasm about the potential of the models, but rather ‘offers a useful corrective to the tendency to anthropomorphize large language models’ (2023); this is absolutely necessary if the artistic contributions—and rights—of the artists in the training data are to remain intact. In fact, the idea that certain types of data such as images and audio have a larger acceptable threshold of alteration is baked into approaches to compression like run-length encoding and

transformers (Salomon 2008). To some extent, the intentionality of such systems has hard-coded limits.

Randomness has, however, also long been built into artistic intentionality: see Jackson Pollock. Terzidis and the others get ahead of this by arguing that irreversible emergent properties within a neural network override this ‘decide-to-delegate’ (Floridi and Cowsls 2019) model, as the designer has no choice but to relinquish artistic intentionality to the ‘unintentional intentionality’ of the AI (Terzidis et al. 2022). There is an attempt to dismantle the intentional connection between the generated image and its authors by introducing stages of randomness between them. However, any noise simply constitutes further iteration on an extant image whose intentionality remains with the original artist: you can make it as fuzzy as you like, that labour still occurred. As long as this emergence is incidental and not ‘intended’ by the AI as a result of task-agnostic cognitive processes, then the causal properties (Searl 1980) of the artwork remain with the AI’s designers and with the original artist. DALL-E 2 even introduced a post-generation image editor (Strickland 2022), meaning that intentionality can be instantly restored once the generated image appears; if AI researchers believe that randomness breaks the chain of intentionality, they must also concede that this in-turn negates ‘unintentionality’.

Markus du Sautoy sees in the development of such algorithms a reflection of the way human beings have always created art. ‘The idea of learning from what artists have done in the past and using that knowledge to push into the new,’ he writes, ‘is of course the process which most human artists go through’ (du Sautoy 2019); art has always been an ‘evolutionary’ model, and it is this evolution which the algorithms have picked up and accelerated (du Sautoy 2019). Sautoy is not the only one; in 2006, Patrick Janssen proposed a method of design which takes advantage of evolutionary principles, using computers to discover ‘inspiring or challenging design alternatives for ill-defined design tasks’ (Janssen 2006). In this ‘schema method’, design teams were required to participate actively in the construction of a generative algorithm regardless of programming ability, helping the programmer with initial design parameters and responding to emergent ideas (Janssen 2006). Computer programming, particularly at a high level of complexity, falls somewhere between creative practice and mathematical exercise. The demands of a compiler are paradoxically more liberal and more stringent than a human reader of fiction; a computer program simply will not run if it is written *wrong*, but it will run if it is ungainly or counter-intuitively arranged. The need for creativity is baked into engineering workflows and pedagogy, stemming from the versatility inherent to the mathematical foundations of programming itself (Howard et al. 2008). Computational thinking denotes a complex of interrelated skills whose overlap with what we consider

artistic thinking is significant: decomposition, generalisation, abstraction, and algorithmic thinking (Romero et al. 2017). Like artists, engineers are mediated by their materials: subject to and inspired by the limitations of their tools. In AI art, this includes choices around training data and image processing. Adobe Firefly's introduction of image 'intensity' sliders signifies that such ideas are at the forefront of commercial AI solutions.

Galit Wellner discusses the result of this creative programming practice and its mediations:

This detailed description of the creation of AI-based works of art leads us to conclude that both parties, the human and the technological, are engaged in the creation process. The participation of human and technological actors changes the situation and both actors. This is known as the co-shaping process between humans and technologies (2021).

If we accept that 'co-shaping' is a constituent property of current AI art then it is important that we analyse the creative practice of programming, at the very least in this context. This distributive creativity introduces 'fragmentation of the imagination' of artistic labour (Wellner 2021): software engineers become artists. Alternately, artists must 'take an expanded role within the feedback loop developing between them and the generative system' (Elgammal and Mazzone 2020); this sort of collaboration counts for the critics, too. The article the previous quote is taken from is one such example of collaboration between the arts and sciences, and intriguing aesthetic perspectives are emerging from the intersection.

Economies of scale have already introduced fractures into the notion of artistic authorship. Pop artists such as Andy Warhol and David Hockney, whose wild popularity turned their creative practice into one of industrial scale, often used assistants and technicians to manufacture or—in the case of performance and conceptual art—otherwise manifest their ideas (McClellan 2018). Where legal structures struggle to define authorship, such as with performance art and temporary sculptural installations, the law simply refuses to classify the work appropriately (Stokes 2001). Generally, however, it is the central artist upon whose brand the work is built that is credited with its creation. The idea behind the artwork is considered the fundamental labour and rights are allocated accordingly. Painters have been using assistants to perform grunt work for centuries, but concept art and printmaking take this further: artists need not touch a piece to be its creator. 'The specificity of contemporary creation', writes Nadia Walravens, 'changes the whole basis of the law. The work may be executed by the artist, or by a person who is not predetermined. Or it may not be executed at all: for only the concept is really important' (Walravens 2022). By contrast, testers, designers, technical specialists, and creators

are all rejected as authors of computer programs in lieu of the programmer (Aplin 2005).

In the United Kingdom, copyright can only be granted for an algorithm once the creator converts it into code, and that copyright does not protect against the development of alternate code which effects the same process. This stems from implementation of EU Directive 2009/24/EC (2009) on the legal protection of computer programs, however, and as such may be subject to change in light of the U.K.'s exit from the European Union; the directive asserts 'to the extent that logic, algorithms and programming languages comprise ideas and principles, those ideas and principles are not protected'. In the United States, where copyright protections have also proven unfavourable to algorithms (Bonadio and McDonagh 2020), Federal and State legislatures have nevertheless been wrestling with the definition of 'algorithm', with cases in the late nineties serving to decouple copyright from *physical* processes or natural scientific laws, and instead to enhance the patentability of algorithms with 'nonobvious' utility (Saladi 1999). The patentability of algorithms parallels our discussion of artistic labour: both require the delineation of boundaries between the strictly 'mathematical' and the decision-making of the designer.

Both art and programming privilege the original idea-author with the supreme moral right: though programming takes a more rigorous approach to intellectual contribution in this regard. It is often designers, executives or entrepreneurs who have the initial idea for a piece of software (Bandey 1996). The modern elevation of the artistic 'idea' over the artefact has led to a decoupling of execution from authorship; the only difference between an AI reproduction of a painting from a model trained on the original and a photograph is the complexity of the intermediary tool and the fact that the 'author' of an AI reproduction—in this case the LLM design team—actually *made* the tool. Nevertheless, 'multimedia works' which result from a program currently fall under the authorship of the programmer who writes the software underlying them (Aplin 2005). In the United Kingdom the output of a program is thus considered the legal property of the programmer just as much as the code itself. Programmers will, in turn, sign their output over to a parent company as part of an employment contract: DALL-E images, for example, are the copyright of OpenAI. This works for self-contained code like fractals which, by their constituent processes—mathematical or otherwise deterministic—output content not subject to copyright.

Problems arise when code draws upon copyrighted material as a part of the AI 'training', and outputs iterations on that material. At that point the network is essentially reprinting a painting, and the resultant muddying of copyright represents a further blurring of the boundaries between artists and programmers. Large language models like GPT have a demonstrated tendency to 'leak' the examples from which

they are trained verbatim (Carlini et al. 2021), and researchers were able to ‘attack’ GPT-2 by generating a diverse set of high-probability samples and sorting them to detect the sources from which they were drawn. These results were achieved with ‘black-box’ access to the GPT-2 API (Carlini et al. 2021), meaning researchers were able to identify exact duplication of GPT-2 training materials without access to the neural network itself. If GPT-3 and DALL-E are subject to the same weaknesses, then this dents the notion that the AI is ‘creating’ new artworks: would a human being, somehow able to know, precisely compute, and perfectly execute another artist’s process, be said to ‘make art’ if they used this capacity to produce replicas? Francis Bacon’s pope paintings, for example, are not in the same class of image as a photograph of the Velasquez painting which inspired them—legally or artistically. Bacon’s work involves compression, loss, and addition.

If legal authorship of AI images falls to the responsible programmers or the hiring entity, then a hierarchy will emerge between the ‘original’ artwork used for training and its generative sibling, one which offers a renewal of Walter Benjamin’s famous ‘aura’: the power of a unique artwork embodied in its scarcity, physicality and narrow market circulation (Zylinska 2020); this is especially relevant as AI art scales to industrial usage and becomes truly commodified. Benjamin makes the following observation:

the technology of reproduction detaches the reproduced object from the sphere of tradition. By replicating the work many times over, it substitutes a mass existence for a unique existence. And in permitting the reproduction to reach the recipient in his or her own situation, it actualizes that which is reproduced. These two processes lead to a massive upheaval in the domain of objects handed down from the past—a shattering of tradition which is the reverse side of the present crisis and renewal of humanity. (Benjamin 2006)

There are a number of democratising principles tied into Benjamin’s essay, and we should take a moment to consider the actualisation Benjamin is talking about. It is unclear at present whether easy AI reproducibility of artistic styles will diminish the novelty and impressiveness of those styles, or whether the ‘aura’ of the original artwork will be emphasised by the phenomenal complexity of the networks needed to mimic them accurately.

4 Algorithmic homogeneity

The notion of ‘original’ art, so important to understanding the value and position of AI-generated artworks, comes to us from the legal frameworks which emerged out of the nineteenth century (Walravens 2022). Outside of these legal and

sociological studies of art, where the principle of aesthetic neutrality offers insulation from the muddying waters of influence, the idea has become increasingly compromised by the emergence of new technologies and historiographies. Art is as much a pedagogical practice as a creative one; artist’s incomes have been supplemented by training for as long as there have been students willing to learn. Artists influence one another as peers, and not just in the visual arts; Harold Bloom famously dissects the thorny relationships which make up, and for him are indistinguishable from, the history of creativity. Bloom talks about poetry, but his observations are equally applicable to painting:

Poetic history, in this book’s argument, is held to be indistinguishable from poetic influence, since strong poets make that history by misreading one another, so as to clear imaginative space for themselves [...] Weaker talents idealize; figures of capable imagination appropriate for themselves. But nothing is got for nothing, and self-appropriation involves the immense anxieties of indebtedness. (Bloom 1997)

The result of these nuanced interactions is style; aesthetic choices which converge to make an image look a certain way, sometimes regardless of subject matter.

Whether or not unauthorised appropriation of style constitutes an infringement of legal rights is an area of debate. Arjun Gupta argues:

Art that appropriates content (“appropriation art”) from other works and sources of visual culture renders inadequate current interpretations of copyright law in its exclusion of alternate meanings in the act of copying. Gupta 2005)

Artists influence each other. They copy, defy, mock, revere and reply; where such relationships are concrete enough one can identify the formation of an artistic school or circle. AI models do not just reproduce paintings verbatim, they are creating novel images in a particular artistic style by finding the ‘average’ between the chosen image and the chosen style, and this is an area where artists have fewer legal rights. Settled U.S. law dictates that, ‘where the only similarity between two works relates to uncopyrightable elements, there can be no infringement’ (Stokes 2001). This spells bad news for artists whose work is used to train AI in a particular style; in the UK, however, legal thinking has come to understand that originality ‘presupposes the exercise of substantial independent skill, labour, and judgement, which offers more room for the original artist, but retreats even further from programming where functional support work is legally irrelevant.

There are moments in history where sets of artistic norms become enshrined within institutions; the Paris of Jean-Léon Gérôme was one in open revolt against Academicism. With

its careful curation, criticism, and celebration of artworks the Académie des Beaux-Arts (prior to the revolution the Académie Royale de Peinture et de Sculpture) had developed a very specific sense of taste, and routinely rejected art which failed to match its standards of artistic or moral virtue. As the Salons were the foremost vehicle for art in the country, and in wider continental Europe, this in turn defined the canon of art for the period. Similar academies sprung up all over Europe during and after the Renaissance and had largely retained their predilection for classical themes and idealised subjects (Williams 2009). Paris in particular produced a significant chunk of theoretical material, as the goal of the academy system was systematisation: art training and execution was to be ‘rationalised’, so that successive generations of artists would continue to iterate on the principles of perfection handed down from the ancient world and revitalised during the Renaissance. As early as the mid-to-late 1600s artists like Charles le Brun and Henri Testelin were offering comprehensive models of drawing and painting from the Academy, with the explicit intention of formalising artistic education (Harrison et al. 2000).

In the context of AI art, it is useful to think of this not just as a social–historical or art–historical process, but as an algorithmic one. The foundation of the European academies established principles of selection, upon which subsequent sub-systems (prizes, scholarships and exhibitions) evolved, with generational iteration built into the system at a pedagogical level. Whilst this algorithm achieved its goal—some truly magnificent neo-classical and academic art—there was nevertheless a serious problem, one which manifested itself in the nineteenth century and which may manifest itself again if neural networks become the primary source of artistic output: academic art got stale. Whilst the Salons continued to grow throughout the nineteenth century, there was a ‘sense that the work it produced and encouraged had become irremediably formulaic and outmoded, and that truly creative work had to be done outside it and in pointed opposition to the system of art-theoretical values it embodied’ (Williams 2009). By the mid-nineteenth centuries, artists and writers in Europe were witnessing and partaking the ‘breakdown of a common language of classicism, the dissipation of revolutionary idealism, and the growing division between artists and public’ (Eisenman 1994). Photography both exacerbated this stagnancy and spearheaded the avant garde emerging from its dissolution. Artists and writers like Gustave Courbet and Charles Baudelaire were rejecting the idealised forms of the Classical period in favour of ‘gross wrestlers, drunken priests, peasants, prostitutes, and hunters [...] common scribes, pharmacists, journalists, students, and adulterers’ (Eisenman 1994). The aesthetic retrenchment of the Academy and its *Secrétaire Perpétuelle de l’Ecole des Beaux-Arts Quatrième de Quincy* led in the 1960s to the *Salon des Refusées*: a show for rejected works which

would feature luminaries like Édouard Manet (Raser 1989). Romantics, impressionists and otherwise non-academic artists flourished in the new Salons, and by the fin de siècle the cultural significance of the Academy was on the wane.

Despite the phenomenal range of subject matter, the academic algorithm had exhausted the attention of its viewing public and the patience of its critics. ‘Thus’, writes Baudelaire in his review of the 1848 Salon:

...that ideal is not that vague thing—that boring and impalpable dream—which we see floating on the ceilings of academies; an ideal is an individual put right by an individual, reconstructed and restored by brush or chisel to the dazzling truth of its native harmony. (Baudelaire et al. 2021)

The underlying algorithmic principles had become too narrow; the set of moral, aesthetic and historical rules which underpinned the production of art led to mundanity. If one imagines that AI art will come to replace human art, and that a small set of powerful models will dominate the field, then one must imagine the same problem. Regardless of subject matter, approach, or artistic intention, if all art is the result of a single process, sameness is built into the system.

5 Generative originality

This sameness is visible today in weaker models which leave a distinct visual footprint: blending and image scaling artefacts which persist regardless of textual input. These are considered ‘flaws’ to be ironed out with further training and advances in machine learning, but the principle remains—the algorithm is producing images in the same way every time. The underlying operations never change. There are, however, a vast number of images used to train models (often in the millions or billions); this may mean that a commensurately gigantic number of generations are needed to identify deep visual repetitiveness beyond the immediate visual markers mentioned above, making creative exhaustion unlikely on a human timescale. Recognising such repetitiveness or its inverse—originality—itself presents a complex set of ontological questions. Daniel Dennett intuits the problems with humans defining such patterns within our own observational frameworks: ‘Other creatures with different sense organs or different interests, might readily perceive patterns that were imperceptible to us. The patterns would be *there* all along, but just invisible to *us*’ (Dennett 1991).

The question we ask is this: whether this algorithmic form of ‘creativity’ is substantially different from human ‘generalised’ problem-solving and decision-making, or if human originality is an emergent property of scale and complexity in fundamentally similar systems. In short: whether human artistic originality is similarly exposed to sameness over a

long enough timescale. Dennett, like Chiang, observes that understanding is ‘the greatest degree of compression’ (Chiang 2023); the former announces, interestingly, that with regards to the existence per se of ‘beliefs’ he likes Paul Churchland’s ‘alternative idea of propositional-attitude statements as indirect “measurements” of a reality diffused in the behavioural dispositions of the brain (and body)’ (1991). Belief statements, stated as such, bear a resemblance to Chiang’s description of Chat-GPT: ‘blurry JPEGs’ of a phenomenally large number of neuronal connections. If art is understood from such a position of eliminative materialism as Churchland’s, then the human-made artwork is the same: a ‘blurry JPEG’ of the human experience, mediated by the chosen materials. A human artist, producing an homage of someone else’s work, is bringing together two averages also: their own experience defined by the organisation of their synaptic connections and the preceding work in question. What we call ‘generalised’ problem-solving or ‘task-agnostic’ understanding is simply a matter of scale and embodiment, of the hundred billion neurons and hundred trillion synaptic connections where Churchland places the whole human conception of the world (Churchland 1995). Zylinska asks whether there is a difference between AI art and older computational artworks; this, in turn, invites us to ask what the difference is between computational and traditional artworks. All involve the use of a medium to produce a visual outcome, the only difference, if one were to accept the above, would be in the scale of computation.

OpenAI also seems to believe that the answer is scale. Foundation models like GPT-3 have begun to offer more generalised solutions to the complex problems associated with natural language processing (NLP). NLP systems are generally trained for the execution of a specialised task. The ideal system, however, is one which can use broad skills to approach novel problems (Tingiris 2021); DALL-E, trained with 250 million online images, was able to produce high-quality generative images with previously unseen material from the much smaller MS COCO (Microsoft Common Objects in Context) database without manual tagging of images. This form of ‘zero-shot’ image generation also allowed for rudimentary image-to-image translation, demonstrating the increasingly flexible and task-agnostic ways that pre-trained neural networks are achieving language-processing and production (Brown et al. 2020).

Like Janssen’s generative evolution schema, the goal is an algorithmic system able to identify and tackle ill-defined challenges, with complementary machine learning models replacing direct feedback and input from designers. The process is the same, the only difference is that the participating models (generative and classifying) are not part of a ‘general’ intelligence. Artists and engineers are already experimenting with ways to ‘merge’ the cognitive-compositional processes which go into making music and algorithms

alike, offering non-human movements as a form of inspiration-prompting (Pošćić and Kreković 2020). The harmonic mean or ‘f-score’ used in machine learning demonstrates the importance of nuanced feedback and human judgement:

$$F1 \text{ score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Whilst not strictly constituting a binary classification dataset (Shalev-Schwartz and Ben-David 2014), a vast number of training images will nevertheless not fit the requirements of the natural language prompt. As a result, success criteria are based simultaneously on the ability of a model to produce accurate images (recall), and fail to produce inaccurate ones (precision), in the eyes of its partner model and the human viewer. The f-score combines precision and recall into a more nuanced, holistic judgemental category.

Human intelligence is sufficiently task-agnostic that task-specific transformers and pre-training may never be able to replace the full gamut of human artistic labour. This is the position taken by McCormack and contributors:

human artists do not learn to create art exclusively from prior examples. They can be inspired by experience of nature, sounds, relationships, discussions and feelings that a GAN is never exposed to and cannot cognitively process as humans do. The training corpus is typically highly curated and minuscule in comparison to human experience. This distinction suggests GANism is more a process of mimicry than intelligence. (McCormack, Gifford and Hutchings 2019)

The researchers at OpenAI agree:

Humans do not require large supervised datasets to learn most language tasks—a brief directive in natural language (e.g. “please tell me if this sentence describes something happy or something sad”) or at most a tiny number of demonstrations (e.g. “here are two examples of people acting brave; please give a third example of bravery”) is often sufficient to enable a human to perform a new task to at least a reasonable degree of competence. (Brown et al. 2020)

However, one could easily argue that McCormack’s sights, sounds and smells constitute only a difference of input class or quantity rather than function. GPT researchers point briefly towards ‘meta-learning’ as a solution for reducing the pre-training requirements of current models. Meta-learning involves a model with ‘a broad set of skills and pattern recognition abilities’ utilising those skills at inference time to deduce the meaning of natural-language prompts and offer appropriate output in a few- or zero-shot scenario (Brown et al. 2020). Here, perspectives between the arts and sciences converge somewhat; ‘It is not unfair to say,’ writes José Hernández-Orallo, ‘that we evaluate the researchers

that have designed the system rather than the system itself' (Orallo 2017). Orallo outlines the various forms of evaluation discussed already—including human discrimination and peer confrontation—and emphasises a need for systemic evaluation of cognitive abilities as constructs rather than as properties: as skill-bases rather than task-specific outcomes.

The construction of natural language prompts is itself a rapidly developing commercial field; though such prompt-writing constitutes a form of artistic intentionality and can be considered a site of artistic work, the actual labour involved is so minute compared to the others mentioned that the legal rights of the prompt-writer can be commensurately minimised. Prompt-writers should not share significantly in attribution unless they make substantive subsequent visual changes to the resultant image. This is in line with current UK legal scholarship around photography which suggests that copyright for images be modified to consider the 'special technical nature' of the medium, removing copyright in cases which 'require no skill and involve no labour beyond pressing the trigger' (Tappin et al. 2018). One solution is Creative Commons-style licences for generated images with special attribution to those artists whose work was used to train the AI and which closely resembles the generated image. Just as the emergence of generative adversarial networks (GANs) led to current advances in image generation, the concurrent development of convolutional neural networks (CNNs) and item response theory (IRT) (Martínez-Plumed et al 2019) has enabled advanced automation of these classifying and categorisation tasks; including models which analyse visual elements like brushstrokes and textures which the generative models replicate. Computer vision and stylistics have advanced to the point where programs 'have the capacity to yield highly variegated meanings, including the ability to recover three-dimensional forms of representations from a two-dimensional image, the recognition of objects in an image, and the analysis of human activities, gestures, facial expressions, and interactions'. AI can be used to crawl through a training database and retroactively rank images according to similarity with the generated image.

If the coming years witness the emergence of a general AI with cognitive properties analogous to the human, then the point is moot: much human resistance to AI art is based on the ego-bruising notion that a fundamentally unthinking machine or combination of processes can produce art as a result of natural language prompts at the first pass. Like the photograph, AI art promises produce art at a scale and pace beyond the human. Zylinska asks 'who' AI art is for: if such an intelligence comes to exist, it will no longer be making art for human aesthetic enjoyment—but for itself, within its own set of judgement criteria which may not include human experiences. From the arts and humanities, it appears that the problem of labour in the production of AI images remains decidedly unsolved; even the formulation

of questions which allow for its solution poses issues of its own. The chain of intentionality runs through a series of complicated and collaborative processes which adulterate extant artworks and utilise varying forms of cognition. This will inevitably provoke moral and legal disputes which themselves stem from epistemological and aesthetic challenges. The philosophy of art is interwoven with technological trajectories and it appears that the present paradigm primarily indicates a coming change in the way artists and the legal system think about the appropriation of style. The humanities also need to engage further with the creative elements of computer programming, especially as software engineers are now making decisions which directly affect the composition of artworks, and at least until AI becomes sufficiently intentional that we can shift our philosophical focus. The actual gaze of arts researchers should, in both instances, be levelled on whatever code and literature is available on the cutting edge: whether we like it or not, AI is making art and there is a distinct need for collaboration if the goals and perspectives of researchers and artists are not to diverge deleteriously.

Data availability The GPT-3 datasets generated during and/or analysed during the current study are—where publicly available—accessible in the following repositories: <https://commoncrawl.org>. <https://paperwithcode.com/dataset/webtext>. <https://wikipedia.org>. Books1 and Books2—not yet specified.

Declaration

Conflict of interest The author did not receive support from any organisation for the submitted work. The author is a practising artist; there are otherwise no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Aplin T (2005) Copyright law in the digital society: the challenges of multimedia. Hart, Oxford
- Artificial Intelligence's New Frontier (2022) The economist <https://www.economist.com/leaders/2022/06/09/artificial-intelligences-new-frontier>. Accessed 13 Sept 2022

- Bandey B (1996) International copyright in computer program technology. CLT Professional Publishing, Birmingham
- Baudelaire C, Pichois C (ed) (1975) *Ouvres Complètes* Editions Gallimard, Paris
- Baudelaire C, Mayne J (trans) (2021) *The Salon of 1846* David Zwirner Books, New York
- Baudelaire C (1880) Salon de 1859 in *Curiosités Esthétiques: Salon 1845–1859*, 3rd edn. M. Lévy, Paris, pp 245–358, translated in: Scharf J (1986) *Art and Photography*, 2nd edn. Penguin, Middlesex, pp 145–146
- Benjamin W (2006) The work of art in the age of its technological reproducibility. In: Eiland H, Jennings MW (eds) *Walter Benjamin: selected writings*, vol 3, 1935–1938. Belknap Press, Harvard, pp 101–133
- Bloom H (1997) *The anxiety of influence: a theory of poetry*, 2nd edn. Oxford University Press, Oxford
- Bonadio E, McDonagh L (2020) Artificial intelligence as producer and consumer of copyright works: evaluating the consequences of algorithmic creativity. *Intell Prop Q* 2:112–137. <https://ssrn.com/abstract=3617197>. Accessed 08 July 2023
- Bowyer E (2010) Monographic impressions. In: Allen S, Morton, MG (eds) *Reconsidering gérôme*. Getty, Los Angeles, pp 22–39
- Brook D (1983) Painting, photography and representation. *J Aesthetics Art Crit* 42:171–180. <https://www.jstor.org/stable/430661>. Accessed 15 Sept 2022
- Brown TB et al (2020) Language models are few-shot learners. [arXiv:2005.14165](https://arxiv.org/abs/2005.14165). <https://doi.org/10.48550/arXiv.2005.14165>. Accessed 17 Sept 2022
- Carlini N et al (2021) Extracting data from large language models. [arXiv:2012.07805v2](https://arxiv.org/abs/2012.07805v2). <https://doi.org/10.48550/arXiv.2012.07805>. Accessed 21 Sept 2022
- Cate PD (2000) Prints abound: Paris in the 1890s. In: Higman S (eds) *Prints abound: Paris in the 1890s* National Gallery of Art, Washington, pp 12–47
- Chiang T (2023) Chat-GPT is a blurry JPEG of the web. *The New Yorker*. <https://www.newyorker.com/tech/annals-of-technology/chatgpt-is-a-blurry-jpeg-of-the-web>. Accessed 08 July 2023
- Churchland P (1995) *The engine of reason. The MIT Press*, Cambridge
- Copyright, Designs and Patents Act 1988 (c.48). <https://www.legislation.gov.uk/ukpga/1988/48/section/17>. Accessed 17 Sept 2022
- Directive 2009/24/EC of the European Parliament and of the Council of 23 April 2009 on the legal protection of computer programs. <https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX%3A32009L0024>. Accessed 20 Sept 2022
- Dennet D (1991) Real patterns. *J Philos* 88(27–51):2023. <https://doi.org/10.2307/2027085>
- Dominguez C et al (2019) The effects of adding non-compulsory exercises to an online learning tool on student performance and code copying. *ACM Trans Comput Educ* 19:16. <https://doi.org/10.1145/3264507>
- Donovan R (2020) Good coders borrow, great coders steal. *Stack Overflow*. <https://stackoverflow.blog/2020/05/20/good-coders-borrow-great-coders-steal/>. Accessed 30 Sept 2022
- Juergens E et al (2009) Do code clones matter? In: Fickas S, Atlee J, Inverardi P (eds) *Proceedings of the IEEE 31st International Conference on Software Engineering IEEE*, New York, pp 185–495. <https://doi.org/10.1109/ICSE.2009.5070547>. Accessed 17 Sept 2022
- Eisenman SF (1994) The rhetoric of realism: courbet and the origins of the avant-garde. In: Eisenman F (ed) *Nineteenth century art: a critical history*. Thames and Hudson, London, pp 206–224
- Elgammal A, Mazzone M (2020) artists, artificial intelligence and machine-based creativity in playform. *Artnodes* 26:1–8. <https://raco.cat/index.php/Artnodes/issue/view/28774/i26>. Accessed 29 Sept 2022
- Floridi L, Cows J (2019) A unified framework of five principles for AI in society. *Harvard Data Sci Rev* 1(6–7):6–7. <https://doi.org/10.1162/99608f92.8cd550d1>[Accessed25September2022]
- Font-Réaulx D (2012) *Painting and photography: 1839–1914*. Flammarion, Paris
- Gérôme JL (1902) Preface. In: Bayard É (eds) *Le Nu Ésthetique* Bernard, Paris
- Gharehyazie M et al (2019) Cross-project code clones in GitHub. *Empir Softw Eng* 23:1538–1573. <https://doi.org/10.1007/s10664-018-9648-z>
- Greenberg EC, Reznicki J (2015) *The copyright zone: a legal guide for photographers and artists in the digital age*, 2nd edn. Routledge, London. <https://learning.oreilly.com/library/view/the-copyright-zone/9781138022577/>. Accessed 23 Sept 2022
- Gupta A (2005) I'll be your mirror – contemporary art and the role of style in copyright infringement. *University of Dayton Law Review* 31:45–82. <https://heinonline.org/HOL/P?h=hein.journals/udlr31&i=53> Accessed 23 Sept 2022
- Harrison C, Wood P, Gaiger J (eds) (2000) *Art in theory: 1648–1815*. Blackwell, Oxford
- Hernández-Orallo J (2017) Evaluation in artificial intelligence: from task-oriented to ability-oriented measurement. *Artif Intell Rev* 48:397–447. <https://doi.org/10.1007/s10462-016-9505-7> Accessed25September2022
- Howard T, Culley SJ, Dekoninck E (2008) Describing the creative process by the integration of engineering design and cognitive psychology literature. *Design Stud* 29(2):160–180. <https://doi.org/10.1016/j.destud.2008.01.001>
- Janssen P (2006) A generative evolutionary design method. *Dig Creat* 17:49–69. <https://doi.org/10.1080/14626260600665736> Accessed22September2022
- Laertius D, Hicks RD (trans) (1931) *Lives of the eminent philosophers*, books 6–10 Harvard University Press, Cambridge. https://www.loebclassics.com/view/LCL185/1925/pb_LCL185.i.xml. Accessed 09 Sept 2022
- de Leon M (2022) We asked artists how they felt about AI-generated art—and they had a lot of feelings. *Rappler.com*. <https://www.rappler.com/life-and-style/arts-culture/asked-artists-how-they-felt-ai-generated-art-lot-of-feelings/>. Accessed 15 Sept 2022
- Marcus G, Davies E, Aaronson S (2022) A very preliminary analysis of DALL-E 2. [arXiv:2204.13807](https://arxiv.org/abs/2204.13807) <https://doi.org/10.48550/arXiv.2204.13807>. Accessed 13 Sept 2022
- Martínez-Plumed F et al (2019) Item response theory in AI: analysing machine learning classifiers at the instance level. *Artif Intell* 271:18–42. <https://doi.org/10.1016/j.artint.2018.09.004>
- McClellan D (2018) Introduction. In: McClellan D (ed) *Artist, authorship & legacy: a reader*. Ridinghouse, London, pp 11–27
- McCormack J, Gifford T, Hutchings P (2019) Autonomy, authenticity, authorship and intention in computer generated art. In: Ekárt A, Liapis A, Luz Castro Pena M (eds) *Computational intelligence in music, sound, art and design: proceedings of the 8th international conference. EVOMUSART, Leipzig*, pp 35–50 https://doi.org/10.1007/978-3-030-16667-0_3. Accessed 23 Sept 2022
- McCormack J, Gambardella C, Rajcic N, Krol S, Llano M, Yang M (2023) Is writing prompts really making art? [arXiv:2301.13049](https://arxiv.org/pdf/2301.13049.pdf) <https://arxiv.org/pdf/2301.13049.pdf> Accessed 16 Oct 2023
- Modiba P, Pieterse V, Bertram H (2016) Evaluating plagiarism detection software for introductory programming assignments. In: Pieterse V, van Eekelen M (eds) *Proceedings of the computer science education research conference. CSERC, Pretoria*, pp 37–46 <https://doi.org/10.1145/2998551.2998558>. Accessed 17 Sept 2022
- OpenAI (2020) DALL-E now available in beta. <https://openai.com/blog/dall-e-now-available-in-beta>

- Pošćić A, Kreković G (2020) On the human role in generative art: a case study of AI-driven live coding. *J Sci Technol Arts* 12:45–62 <https://doi.org/10.34632/jsta.2020.9488>
- Ramesh A et al (2021) Zero shot text-to-image generation. *arXiv:2102.12092v2* <https://doi.org/10.48550/arXiv.2102.12092>. Accessed 17 Sept 2022
- Raser T (1989) *A Poetics of Art Criticism: The Case of Baudelaire*. University of North Carolina Press, Chapel Hill
- Romero M, Lepage A, Lille B (2017) Computational thinking development through creative programming in higher education. *Int J Educ Technol Higher Educ* 14:42. <https://doi.org/10.1186/s41239-017-0080-z>
- Roose K (2022) An A.I.-generated picture won an art prize. artists aren't happy. *The New York Times* <https://www.nytimes.com/2022/09/02/technology/ai-artificial-intelligence-artists.html>. Accessed 15 Sept 2022
- Saladi I (1999) Computer software: patentable subject matter jurisprudence comes of age, 18 *J. Marshall J. Computer & Info. L.* 113. *J Comput Inform Law* 1:113–139. <https://repository.law.uic.edu/jitpl/vol18/iss1/2>. Accessed 09 July 2023.
- Salomon D (2008) *A concise introduction to data compression*. Springer, London
- du Sautoy M (2019) *The creativity code: how AI is learning to write, paint and think*. 4th Estate, London
- Scharf A (1986) *Art and photography*. Penguin, Middlesex
- Scruton R (1981) Photography and representation. *Crit Inquiry* 7:577–603 <https://www.jstor.org/stable/1343119>. Accessed 13 Sept 2022
- Searle JR (1980) Minds, brains and programs. *Behav Brain Sci* 3:417–457. <https://doi.org/10.1017/S0140525X00005756>
- Shalev-Schwartz S, Ben-David S (2014) *Understanding machine learning: from theory to algorithms*. Cambridge University Press, Cambridge
- Stokes A (2001) *Art and copyright*. Hart, Oxford
- Stokes A (2022) Copyright, art and digitisation: European and U.K. Perspectives. In: McClean D, Schubert K (eds) *Dear images: art, copyright and culture*. Institute of Contemporary Arts, London, pp 128–141
- Strickland E (2022) DALL-E 2's failures reveal the limits of AI. *Spectrum*, IEEE <https://spectrum.ieee.org/openai-dall-e-2>. Accessed 25 Sept 2022
- Tappin M et al (2018) *Laddie, Prescott and Vitoria: the modern law of copyright*, 5th edn. LexisNexis Butterworths, London
- Terzidis K, Fabrocini F, Lee H (2022) Unintentional intentionality: art and design in the age of artificial intelligence. *AI & Soc.* <https://doi.org/10.1007/s00146-021-01378-8>. Accessed 25 September 2022
- Tingiris S (2021) *Exploring GPT-3: an unofficial first look at the general-purpose language processing AP from OpenAI*. Packt, Birmingham
- Truth Lies at the Bottom of a Well. *The Oxford Dictionary of Proverbs*. <https://www.oxfordreference.com/view/https://doi.org/10.1093/acref/9780199539536.001.0001/acref-9780199539536-e-2323>. Accessed 09 Sept 2022
- Walravens N (2022) The concept of originality and contemporary art. In: McClean D, Schubert K (eds) *Dear images: art, copyright and culture*. Institute of Contemporary Arts, London, pp 170–195
- Wellner G (2021) Digital imagination fantasy AI art. *Found Sci.* <https://doi.org/10.1007/s10699-020-09747-0>
- Whitford J (1985) *Express newspapers v. Liverpool Daily Post (Ch.D.)*. *The Weekly Law Reports* 1:1089–1099, The Incorporated Council of Law Reporting for England and Wales, London
- Williams R (2009) *Art theory: an historical introduction*, 2nd edn. Wiley-Blackwell, Chichester
- Zylinska J (2020) *AI art: machine visions and warped dreams*. Open Humanities Press, London

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