



UNIVERSITY OF
OXFORD

Creativity, Artificial Intelligence, and the Requirement of Human Authors and Inventors in Copyright and Patent Law

JULY 2024



UNIVERSITY OF
OXFORD

Creativity, Artificial Intelligence, and the Requirement of Human Authors and Inventors in Copyright and Patent Law

CHRISTIAN E. MAMMEN,
MICHAEL COLLYER, RON DOLIN,
DEV GANGJEE, TOM MELHAM, MAGGIE
MUSTAKLEM, PIREENI SUNDARALINGAM,
VINCENT WANG-MAŚCIANICA

JULY 2024

Abstract

Copyright and patent law require the identification of an author or inventor, and further require the author or inventor to be human. We explore this requirement primarily with reference to U.S. law and provide additional illustrations from U.K. and E.U. law. A key rationale underlying the requirement of a human author or inventor is that there is something special and important about human creativity.

As AI, particularly generative AI, becomes more capable of producing outputs that look like they could have been human-created, arguments have increasingly been raised that the AI-generated outputs should be afforded copyright and patent protection, on the same basis as those made by human authors and inventors. And there have been arguments that these AI-generated outputs exhibit sufficient creativity, novelty, or innovativeness, to satisfy the laws' underlying creativity rationale.

We examine the concept of creativity from a multidisciplinary perspective, and identify three conceptually distinct components, all of which are necessary for a complete account of creativity. The *external* component refers to whether an artifact (or idea, or other thing) exhibits the qualities of being novel, valuable, and (on some accounts) surprising. The *subjective* component focuses on the psychological process of a creative act, which appears to involve a dance between task-focused and mental-wandering states, mediated by a salience functionality, where the person recognizes and selects novel, appropriate ideas. Third, embedded in the analysis of both the external and subjective components is a (largely-implicit) recognition that the *social context* is integral to creativity; it plays a role in determining whether an artifact has value (or is "appropriate"), and influences the subjective psychological process of plucking certain ideas or conceptions out of the flow of

mental activity.

With this enriched account of creativity, we examine how copyright and patent law value not only the creativity of the artifact, but also (to varying extents) social context's role as part of creativity. We then consider some ways in which arguments that AI-generated artifacts should be eligible for IP protection (e.g., because they are "just as good as" at least some human-generated and IP-eligible artifacts) are insufficient to satisfy the enriched understanding of the creativity requirement underlying the IP laws.

Our investigation additionally reveals some themes that may warrant further and deeper examination. First, copyright has not historically concerned itself with the quantum or quality of creativity in a human-created work; it just has to be non-zero. Some scholars have raised concerns that this standard has been too indulgent of copyright protection and have argued that the bar of creative quality should be raised. Placed in the context of generative AI's influence, such an argument raises intriguing possibilities, insofar as there is evidence that AI tends to narrow the band of creativity (both human creative thought and expression as well as the "creativity" of AI-generated outputs), hindering or suppressing what might otherwise be a wider range of creative possibility. Second, as AI becomes ever-more-capable, displacing not only traditionally human activities but also traditional human decisions and judgments, it may be necessary to consider fundamental revisions to some of our legal frameworks to accommodate this change.

There is a debate on the horizon over whether the laws should be changed to accommodate IP for AI-generated outputs.

Contents

Abstract.....	3
I. Introduction	5
II. The Requirement of Human Authors and Inventors Under Copyright and Patent Law, and the Underlying Value of Creativity.....	8
A. Core Aspects of Copyright Doctrine	8
1. United States	8
2. United Kingdom and European Union.....	11
B. Core Aspects of Patent Doctrine in the U.S.	12
III. A Conceptual Account of Creativity	14
A. External Creativity (Creativity of Artifacts).....	15
B. Subjective Creativity (Creativity of Mental Processes).....	17
C. Social Context	20
IV. Which Parts of Creativity Are Important to U.S. Intellectual Property Law?	21
A. Copyright Law.....	21
B. Patent Law	23
C. Summary.....	23
V. Risks to Creativity in an AI World.....	24
A. How AI Affects Human Creativity	24
B. Algorithmic Flattening: How the Design of AI Models Constrains the Quality and Diversity of “Creative” Outputs.....	27
VI. Revisiting the Question of AI “Creativity”	28
VII. Looking Ahead	31
Bibliography	35
Contributors.....	40

I. Introduction

What comes to mind when you are asked to think about “creativity”? Perhaps your first thought is of someone who makes something that has never been made before, or does something in a completely new way. But then, after a bit more consideration, you realize that cannot quite be right—if something is *completely* unlike anything that is already known to a society, it might be considered too foreign, too weird, too inappropriate, to be a recognizable exercise of creativity. So then you progress to thinking that creativity involves a modification of a known thing, or a combination of two or more known things in a new way, or a new (and elegant) solution to a previously-unsolved problem. This often involves building on the expertise and insights of predecessors—“standing on the shoulders of giants,” as is sometimes said. See Goldberg 2018 p. 43. But can you “brute force” your way to creativity by iteratively and non-selectively working through combinations and variations until you get something “good”? How do you select the problem or task? How do you identify the permutations to work through? How do you know when you get something “good”?

These are some of the core questions about creativity—human creativity. They have gained new relevance in the modern age of generative artificial intelligence (AI), which is an old field of AI research but surged into the public consciousness with the November 2022 public release of OpenAI’s ChatGPT (or, arguably, in the summer of 2019 with the publication of patent applications claiming to have been invented solely by an AI algorithm called DABUS).¹ The textual outputs of ChatGPT and the image outputs of Dall-E and Midjourney, to name a few examples, are amazing feats of natural language

processing (NLP) in converting the lexical semantics of human-supplied prompts into computer code, and then generating outputs – text or images – that *look like* human-generated outputs, (mostly) responsive to the prompt, and (mostly) constituting a finished product that did not exist before. Are those generative AI outputs “creative”? Are the generative AI platforms “creative”? What does that even mean?

In answering—or perhaps refereeing—these questions, intellectual property law plays an increasingly relevant role, as the legal field that protects creative outputs. That is, are AI outputs eligible for copyright and/or patent protection? Under U.S. law, current doctrine is clear that copyright law requires an author, which must be human, and patent law requires an inventor, which likewise must be human. These U.S. holdings have applied regardless of whether the AI is claimed to have acted independently or generated text or images in response to human-entered prompts. In part, these rulings have depended on a formal positive-law analysis—meaning that the statutes are worded in such a way as to make clear that “author” and “inventor” refer to human beings. Copyright laws in the U.K. and E.U. largely track this observation, as does the U.K. Supreme Court’s ruling requiring a human inventor.

Analysis to date begs the question whether the requirement of human authors and inventors is merely an artifact or is based on good and sufficient reasons to limit IP protection to humans.

¹ E.g., <https://www.surrey.ac.uk/news/world-first-patent-applications-filed-inventions-generated-solely-artificial-intelligence> (visited 13 June 2024); <https://news.sky.com/story/chatgpt-turns-one-the-first-year-of-the-chatbot-that-changed-the-world-13014185> (visited 13 June 2024).



To an extent, however, the analysis to date begs the question whether this statutory requirement of human authors and inventors is merely an artifact of historical assumptions built into the existing law or whether, instead, there are good and important reasons why intellectual property protections should be limited to humans. In part, the rationales stated in those rulings have implicitly endorsed the view that there is something special about the human creative contribution.

But some have argued that if the AI outputs are “just as good as” some IP-eligible human-generated counterparts, then it does not make sense to grant IP protection to the human-generated outputs while denying IP protection to the AI-generated outputs. *E.g.*, Abbott 2020; *cf* Yu 2017. Arguments such as this tend to focus on the quality and characteristics of the output—the text or image being indistinguishable from that of a human author or artist, for example. As discussed below, we believe that a complete account of the role of creativity in IP eligibility requires a broader examination of not just the outputs but also the process by which they were made, as well as the importance of the social or societal milieu in which the process takes place and the output, or artifact, is made.

Despite the relatively one-sided outcomes of recent cases (discussed below) concerning authorship and inventorship—a human creator remains necessary—we anticipate that there is a

debate on the horizon over whether the laws should be changed to accommodate IP protections for AI-generated outputs.²

There are many profound policy arguments underlying both sides of this debate, directed to the underlying purposes of the IP laws. These span economic policy, distributive justice, utilitarianism, and non-instrumentalist accounts of the value of human engagement in creative enterprises. *See generally, e.g.*, Merges 2011; O’Callaghan 2022 pp. 314-320. While these normative or justificatory arguments are important, they are beyond the scope of this paper. Instead, we focus on taking a closer look at “creativity,” which is frequently mentioned in the case law as a justification (or criterion) for limiting IP protection to certain human-generated outputs.

Authors and inventors must be humans, the reasoning goes, because only humans are creative, and creativity is the sine qua non for intellectual property protection.

In the meantime, another front has opened in the skirmish between copyright-owning human creatives and generative AI platforms. As of early 2024, approximately twenty lawsuits have been filed by copyright owners and other creatives against various generative AI platforms,³ asserting that the

² In talking about AI-generated outputs, we intentionally avoid question-begging characterizations by not referring to these outputs as “artifacts,” “inventions,” or similar terms normally and traditionally associated with human outputs. Nor do we limit ourselves to what some in the literature have called “emergent works,” which has been defined as “works of apparently creative expression that arise from the operation of a program but cannot be traced directly to a human source.” Boyden 2016 p. 379; *see also* Epstein 2023 pp. 1110-1111. Arguably, outputs of generative AI tools, output in response to a human-entered prompt, fall outside the definition of “emergent works” even though they are clearly within the scope of AI-generated outputs addressed in this paper.

³ *Thomson Reuters Enterprise Centre GmbH v. Ross Intelligence Inc.*, D.Del. Case No. 20-cv-00613-SB; *Getty Images (US), Inc. v. Stability AI, Inc.*, D.Del. Case No. 23-cv-00135-JLH; *DOE 1 v. GitHub, Inc.*, N.D. Cal. Case No. 22-cv-06823-JST; *DOE 3 v. GitHub, Inc.*, N.D. Cal. Case No. 22-cv-07074-JST; *Andersen v. Stability AI Ltd.*, N.D. Cal. Case No. 23-cv-00203-WHO; *P.M. v. OpenAI LP*, N.D. Cal. Case No. 23-cv-03199-TLT; *Kadrey v. Meta Platforms, Inc.*, N.D. Cal. Case No. 23-cv-03417-VC; *J.L. v. Alphabet Inc.*, N.D. Cal. Case No. 23-cv-03440-AMO; *Tremblay v. OpenAI, Inc.*, N.D. Cal. Case No. 23-cv-03223-AMO; *Silverman v. OpenAI, Inc.*, N.D. Cal. Case No. 23-cv-03416-AMO;



use of copyrighted material as AI training data infringes copyright and/or other rights held by the creators. In one case involving the use of human-drafted source code that was posted in an online repository as AI training data,⁴ the judge commented, “I have some difficulty in understanding how the training aspect of this product injures anybody’s rights, because if [I] can get a GitHub account and go on and read all of it until [my] heart’s content, why can’t a software program do the same thing? Why does that violate any right?”⁵ The judge’s question hints at a core issue underlying patent and copyright protection. These IP doctrines protect exercises of creativity, but creativity is not entirely out-of-the-blue and acontextual. Creative work is often (perhaps always) influenced by prior works, and our instinctual understandings of what is or isn’t creative are apparently influenced by an understanding of the concept. See, e.g., Buccafusco 2022.

What, then, is “creativity”? Within the broader, global, question about creativity is a narrower set of questions about what kind of creativity is necessary and sufficient to render an artifact IP-eligible. If nonhumans (including AI) engage in something that ticks all the relevant boxes for “creativity,” does that negate the reasons for distinguishing between humans and those nonhumans in terms of whose outputs can get IP protection? Or, conversely, if we can clearly see that, however much the nonhuman’s outputs “look just like” human outputs, they are deficient in some important criterion, does that bolster the reasons for continuing

to deny IP protection to those outputs?

The authors of this white paper, a multidisciplinary group of lawyers, philosophers, computer scientists, mathematicians, design theorists, and psychologists, aim to provide a preliminary account of creativity that could productively inform debates about whether and under what conditions IP rules should recognize AI-produced artifacts.

This paper broadly follows the following structure. First, Section II introduces the proposition that IP laws are founded on encouraging and rewarding creativity. The legal materials drawn on to support this view primarily relate to U.S. law, interspersed with some E.U. and U.K. materials as well. We wish to demonstrate that this view is widely accepted and conventional.⁶ Then, Section III provides an account of creativity that draws from both neuroscience and philosophy. Third, Section IV compares our account of creativity with the aspects of creativity that appear to be particularly valued by copyright and patent law. Section V considers how generative AI may hamper the quality and diversity of both human creativity and AI-generated outputs. Section VI turns to the question whether generative AI and AI outputs can be said to satisfy the components of creativity that are relevant to the IP laws. And finally, Section VII concludes with an overview of some of the questions to come about AI and law—questions that are explicitly not part of the AI-authorship and AI-inventorship debates, but are likely to arise soon, and perhaps inevitably will arise sooner than we expect.

T. v. OpenAI LP, N.D. Cal. Case No. 23-cv-04557-VC; *Chabon v. OpenAI, Inc.*, N.D. Cal. Case No. 23-cv-04625-AMO; *Authors Guild v. OpenAI Inc.*, S.D.N.Y. Case No. 23-cv-08292-SHS; *Concord Music Group, Inc. v. Anthropic PBC*, M.D. Tenn. Case No. 23-cv-01092; *The New York Times Company v. Microsoft Corp.*, S.D. N.Y. Case No. 23-cv-11195-SHS; *Huckabee v. Meta Platforms, Inc.*, N.D. Cal. Case No. 23-cv-06663-VC; *Basbanes v. Microsoft Corp.*, S.D. N.Y. Case No. 24-cv-00084-SHS.

⁴ *DOE 1 v. GitHub et al.*, N.D. Cal. Case No CV-22-06823-JST (May 4, 2023 hearing, Dkt. No. 91).

⁵ *Id.* Tr. pp. 27-28.

⁶ Reflecting the training and experience of the working group’s members, we focus heavily on U.S. law, with some discussion of U.K. law. Other nations’ legal systems, to the extent they differ from these two, are outside the scope of this paper.

II. The Requirement of Human Authors and Inventors Under Copyright and Patent Law, and the Underlying Value of Creativity

In this section, we provide a brief outline of U.S. copyright and patent law, showing that the requirement for human authors and inventors is related to the proposition that the IP laws have as one aim to support and reward human creativity. We also address selected analogous provisions in the IP laws of the U.K and E.U.

Cases from both patent and copyright law are clear that the U.S. Constitution’s Patents and Copyright Clause, Article I, Section 8, has as a main purpose to encourage and reward *creative effort* by authors and inventors. The IP Clause “is intended to motivate the creative activity of authors and inventors by the provision of a special reward ...” *E.g., Bouchat v. Baltimore Ravens Ltd. Partnership*, 737 F.3d 932, 936 (4th Cir. 2013) (citing *Sony Corp. Of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 429 (1984)). Below we address in more detail case law specific to the enacting statutes for copyrights and patents.

A. Core Aspects of Copyright Doctrine

1. United States

It is well-established in U.S. law that in order for an artifact to be copyrightable, it must be expressed in tangible form⁷ and must have an author (or artist, composer, etc.). See 17 U.S.C. § 102 (extending copyright protection to “original works of authorship fixed in any tangible medium”). The term “author” refers exclusively to humans. “An original work of authorship is a work that is independently created by a human author and possesses at least some minimal degree of creativity.” U.S. Copyright Office, *Compendium of U.S. Copyright Office Practices* (3d ed. 2021) § 306. See also *Naruto v. Slater*, 888 F.3d 418, 426 (9th Cir. 2018) (crested macaque in “monkey selfie” case lacks standing because only humans have statutory standing under Copyright Act; “author” must be a human). This has led at least one commentator to flatly assert that protection of “emergent works” is “impossible” because “the whole system of copyright law ... operate[s] in the notation of human creativity.” *Blaszczyk* 2023 pp.2-3.

Three recent determinations illustrate the U.S. Copyright Office’s policy concerning AI-generated images and creativity. In *Thaler v. Perlmutter*, 687 F.Supp.3d 140 (D.D.C. 2023), the D.C. District Court upheld the Copyright Office’s denial of a copyright registration⁸ for “A Recent Entrance to Paradise,” which was alleged to be created solely by DABUS, an AI. This image is below:


⁷ Copyright “protection is given only to the expression of the idea—not the idea itself.” *Mazer v. Stein*, 347 U.S. 201, 217 (1954).

⁸ U.S. Copyright Office Review Board, *Decision Affirming Refusal of Registration of A Recent Entrance to Paradise* (Feb. 14, 2022), <https://www.copyright.gov/rulings-filings/review-board/docs/a-recent-entrance-to-paradise.pdf>.



In February 2023, the Copyright Office revoked a previously-issued copyright certificate for a graphic novel called *Zarya of the Dawn*, based on newly-discovered information that the human author, Kristina Kashtanova, had used the Midjourney generative AI tool to create the images in the graphic novel. Ultimately, the Copyright

Office allowed a registration for the text (which Kashtanova wrote) and the arrangement of text and images in the final product, but not for the images alone.⁹ Images from *Zarya of the Dawn* are reproduced below:



⁹ U.S. Copyright Office, *Cancellation Decision re: Zarya of the Dawn* (VAu001480196) (Feb. 21, 2023), <https://www.copyright.gov/docs/zarya-of-the-dawn.pdf>.

Third, in September 2023, the Copyright Office denied copyright registration to Jason Allen’s work, “Théâtre D’opéra Spatial.” The final work involved an AI-generated work that Allen then spent substantial time refining. Allen refused to disclaim the AI-generated material, so the Copyright Office denied registration.¹⁰ Allen’s work gained notoriety in 2022 after it won first place in its category at the Colorado State Fair.¹¹ This image is below:



On March 16, 2023, the U.S. Copyright Office issued a statement of policy concerning works containing material generated by AI. 88 FR 16190 (March 16, 2023). This policy statement reiterated that “it is well-established that copyright can protect only material that is the product of human creativity.” *Id.* at 16191. The statement explained, in the case of generative AI, “when an AI received solely a

prompt from a human and produces complex written, visual, or musical works in response, the ‘traditional elements of authorship’ are determined and executed by the technology—not the human user.” *Id.* at 16192. Accordingly, human “users do not exercise ultimate creative control” over the AI-generated works and they are not copyright-eligible. *Id.* Varying conclusions about copyrightability may be possible on a case-by-case basis, depending on “how the AI tool operates and how it was used to create the final work.” *Id.*

As the Copyright Office’s policy statement emphasizes, the human authorship requirement is rooted in protecting and rewarding creativity. Justice Sotomayor echoed this in *Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508, 526 (2023), “The Copyright Act encourages creativity by granting to the author of an original work ‘a bundle of exclusive rights.’” This has long been the case.¹² To

qualify as a “work of authorship” requires “on statutory and constitutional grounds [that there be] a modicum of creativity.”¹³

However, the threshold of creativity required is low, but non-zero. The Supreme Court stated in 1991, “A work satisfies the originality requirement as long as it possesses some ‘creative spark,’ no matter how crude, humble or obvious it might be.” *Feist Publications, Inc. v. Rural Tel. Serv. Co.*,

¹⁰ U.S. Copyright Office, *Decision re Refusal to Register Théâtre D’opéra Spatial* (SR # 1-11743923581) (September 5, 2023), <https://www.copyright.gov/rulings-filings/review-board/docs/Theatre-Dopera-Spatial.pdf>.

¹¹ See <https://www.smithsonianmag.com/smart-news/artificial-intelligence-art-wins-colorado-state-fair-180980703/> (visited June 25, 2024).

¹² *Mitchell Bros. Film Group v. Cinema Adult Theater*, 604 F.2d 852, 856 (5th Cir. 1979) (“The purpose underlying the constitutional grant of power to Congress to protect writings is the promotion of original writings, an invitation to creativity.”).

¹³ *Atari Games Corp. v. Oman*, 979 F.2d 242, 244 (D.C. Cir. 1992) (citing *Feist Publications, Inc. v. Rural Tel. Serv. Co.*, 499 U.S. 340 (1991)). *In re Trade-Mark Cases*, 100 U.S. 82, 94 (1879) (writings are “founded in the creative powers of the mind” and “the fruits of intellectual labor ...”); see also *Burrow-Giles Lithographic Company v. Sarony*, 111 U.S. 53, 58-59 (1884) (“[A]uthor, in its constitutional sense, has been construed to mean an ‘originator,’ ‘he to whom anything owes its origin.’”).



499 U.S. 340, 358 (1991). Under this standard, living gardens and the teachings of non-human spirits are not copyrightable.¹⁴ In other words, “The standard of originality is low, but it does exist.” *Feist*, 499 U.S. at 362. The originality requirement can be satisfied by making “non-obvious choices” from “among more than a few options,” evincing “some inventiveness and imagination,” which may include “choices about style and setting” or “decisions about what materials to include and how to organize them.”¹⁵

And the Copyright Office, which registers copyrights in the U.S., does not engage in substantive evaluation about “how good” or “how creative” a work is, provided it meets the minimum threshold of creativity. U.S. Copyright Office, *Compendium of U.S. Copyright Office Practices* (3d ed. 2021) § 310.2 (Copyright Office “does not consider the aesthetic value, artistic merit, or intrinsic quality of a work.”).¹⁶ “The pursuit of creativity requires freedom to explore into the gray areas, to the cutting edge, and even beyond.”¹⁷

2. United Kingdom and European Union

Both E.U. law and consequently U.K. law, which still follows the E.U. copyright *acquis*, also require human authors for creative works. Citing to the Berne Convention on

Copyright, the Court of Justice of the European Union observes declared that “copyright within the meaning of [relevant E.U. law] is liable to apply only in relation to a subject-matter which is original in the sense that it is its author’s own intellectual creation.” [Case C-5/08 *Infopaq International* [2009] ECR I-6569, at [37]]. Under E.U./U.K. law, once it has been decided that the subject matter is of the kind which allows free and creative choice to be made, the analysis then turns to whether the presumptive author sufficiently exercised these choices in a manner which reflects the author’s personality. In a dispute over whether a portrait photograph qualifies for protection, the Court emphasized that the test was whether that work “is an intellectual creation of the author reflecting his personality and expressing his free and creative choices in the production of that photograph.” [Case C-145/10, *Painer v. Standard Verlags GmbH*, 2011 E.C.R. I-12533, at [99]]. “By making those various choices [relating to background, lighting, framing, the viewing angle etc], the author of a portrait photograph can stamp the work created with his ‘personal touch’.” (Ibid, [92]). For broadly similar reasons, the [non-binding] Advocate General’s Opinion concluded: “only human creations are . . . protected” under

¹⁴ See *Urantia Foundation v. Maaherra*, 114 F.3d 955 (9th Cir. 1997) (book containing the alleged teachings of non-human spiritual beings was copyrightable; requirement that there be “some element of human creativity” satisfied by humans’ formulation of questions to the spiritual beings, selection and arrangement of their revelations, and organization and order of the resulting book); *Oliver v. Saint German Foundation*, 41 F.Supp. 296 (S.D. Cal 1941) (no copyright protection for text allegedly dictated to copyright claimant by spirit of a deceased person); *Kelley v. Chicago Park Dist.*, 635 F.3d 290, 303 (7th Cir. 2011) (living garden designed by landscape architect not copyrightable).

¹⁵ *Premier Dealer Services, Inc. v. Allegiance Administrators, LLC*, 93 F.4th 985, 989 (6th Cir. 2024) (cleaned up); compare *Ragan v. Berkshire Hathaway Automotive, Inc.*, 91 F.4th 1267, 1270 (8th Cir. 2024) (basic customer intake sheet lacks sufficient originality to be copyrightable; it fails to “exhibit some degree of creativity”).

¹⁶ See also *Mitchell Bros. Film Group v. Cinema Adult Theater*, 604 F.2d 852, 856 (5th Cir. 1979) (Congress determined that “constitutional goal of encouraging creativity would not be best served if an author had to concern himself not only with the marketability of his work but also with the judgment of government officials regarding the worth of the work”).

¹⁷ *Id.*

copyright law.¹⁸ This underscores the requirement for human authorship. Copyright legislation in the E.U. is also premised on legal protection fostering this human creativity.¹⁹

However, there is a curious exception in the U.K. statute, authorizing copyrights for shortened terms for computer-generated works. Copyright, Designs and Patents Act (1988) §§ 12(7) (“If the work is computer-generated ... copyright expires at the end of the period of 50 years from the end of the calendar year in which the work was made”), 178 (“‘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”), 9(3) (“[T]he author [of a computer-generated work] shall be taken to be the person by whom the arrangements necessary for the creation of the work are undertaken.”). This statutory provision, enacted in 1988, has not been tested in the current era of generative AI. It is, however, beginning to attract scholarly attention in this regard. *E.g.*, Lee 2021.

B. Core Aspects of Patent Doctrine in the U.S.

Patent law in the U.S. requires that there be an identified inventor. 35 U.S.C. § 100

(“The term ‘inventor’ means the individual ... who invented or discovered the subject matter of the invention.”), § 101 (“Whoever invents or discovers ...”), § 102 (“A person shall be entitled ...”). The inventor must be human. *Thaler v. Vidal*, 43 U.S. 4th 1207 (Fed. Cir. 2022) (DABUS, an AI, cannot be an inventor).²⁰ Several earlier U.S. rulings, considering whether corporations could be inventors, also concluded that an inventor must be a human.²¹

In requiring a human inventor, U.S. patent law similarly serves the aim of human creativity. See *Bonito Boats, Inc. v. Thunder Craft Boats, Inc.*, 489 U.S. 141, 162, 167 (1989) (referencing the “*quid pro quo* of substantial creative effort required by the federal [patent] statute” as part of the patent bargain, or “careful balance between public right and private monopoly to promote certain creative activity”).²² The association between invention and creativity has long been part of the animating reasoning for patenting. *E.g.*, *Calkins v. Oshkosh Carriage Co.*, 27 F. 296, 298-299 (E.D. Wis. 1886) (“a device which displays only the expected skill of the maker’s calling, and involves only the exercise of ordinary faculties of reasoning upon materials supplied by special knowledge and facility of manipulation resulting from habitual intelligent practice, is

¹⁸ [Opinion of Advocate General Trstenjak, Case C-145/10, *Painer v. Standard Verlags GmbH*, 2011 E.C.R. I-12533, at [121]].

¹⁹ *E.g.*, Directive 2001/29, Recital 4 (“A harmonised legal framework on copyright and related rights, through increased legal certainty and while providing for a high level of protection of intellectual property, will foster substantial investment in creativity and innovation, including network infrastructure. ...”, Recital 9 (“Any harmonisation of copyright and related rights must take as a basis a high level of protection, since such rights are crucial to intellectual creation...”). Stech 2021 p. 235 (“While ‘author’s own intellectual creation’ is European verbiage describing the floor for authorial originality [for the E.U.], it is equally serviceable in the United States and elsewhere.”).

²⁰ Ruling on the U.K. version of the same patent application, the U.K. Supreme Court recently reached the same conclusion. *Thaler v. Comptroller Thaler v. Comptroller-General of Patents, Designs and Trademarks* [2023] UKSC 49. (<https://www.supremecourt.uk/cases/docs/uksc-2021-0201-judgment.pdf>)

²¹ *E.g.*, *Beech Aircraft Corp. v. EDO Corp.*, 990 F.2d 1237, 1248 n.23 (Fed. Cir. 1993) (“only natural persons can be ‘inventors’”); *Karrer v. United States*, 152 F.Supp. 66, 69 (Ct. Cl. 1957) (“In the United States a patent application can be filed only by a natural person, the inventor ...”).

²² Under the so-called patent bargain, “inventors are awarded a limited monopoly through a patent grant to incentivize their creative effort, but after that limited time expires, the invention becomes available to the public.” *Suppes v. Katti*, 710 Fed. Appx. 883, 888 (Fed. Cir. 2017).



in no sense a creative work of inventive faculty, such as the constitution and patent laws aim to encourage and reward.”). In *In re Beineke*, 690 F.3d 1344, 1351 (Fed. Cir. 2012), the Federal Circuit emphasized that provisions in the Patent Act covering plant patents should be interpreted to reflect “Congress’s understanding that patent protection was available only for plants resulting from human creative efforts by the patent applicant, and not for found plants.”²³

Under the Patent Act, an invention must be “new and useful.” 35 U.S.C. § 101. Novelty is measured in two ways. First, if the invention was disclosed in a single “prior art” reference (e.g., a printed publication that was publicly available before the patent application’s effective filing date), it is said to be “anticipated” by the prior art. 35 U.S.C. § 102. Second, a claimed invention is not novel if it “would have been obvious before the effective filing date of the claimed invention to a person having ordinary skill in the art [(“POSITA”)] to which the claimed invention pertains.” 35 U.S.C. § 103. That is, if the claimed invention would have been obvious to a POSITA at the time of the patent application, then it is not novel.

Until the mid-20th century, this non-obviousness requirement in patent law was often articulated using a turn of phrase referencing a “flash of creative genius” as the act of invention. See *Cuno Corp. v. Automatic Devices Corp.*, 314 U.S. 84, 91 (1941) (to be patentable, an invention “must

reveal the flash of creative genius”). But this was supplanted by a more structured analysis of obviousness by the Supreme Court’s rulings in *Graham v. John Deere*, 383 U.S. 1 (1966) and *KSR Int’l Co. v. Teleflex, Inc.*, 550 U.S. 398 (2007).²⁴

As subsequent rulings noted, the *Graham/KSR* cases reframe the standard for non-obviousness so that it may be achieved either by a “spark of creative genius” or by more methodical investigation. “Generally, a patent may be obvious if it lacks skill and ingenuity that characterizes a patentable invention. However, a nonobvious invention can arise from systemic experimentation as well as from a flash of creative genius.” *Braintree Laby’s, Inc. v. Novel Laby’s*, No. 11-cv-1341-PGS, 2013 WL 2970739, *21 (D.N.J. June 4, 2013).²⁵

But the shift from the poetic “flash of creative genius” rhetoric to the standards articulated in *Graham* and *KSR* did not eliminate the creativity requirement. Invention as an act of *creativity* remains part of the discourse. For example: “Inventors are impelled to invest in creative effort by the expectation that, through procurement of a patent, they will obtain [a right to exclude others].” *Biotechnology Industry Organization v. District of Columbia*, 496 F.3d 1362, 1372 (Fed. Cir. 2007).²⁶ This expansion also accords with evolving understandings of creativity. Fromer 2010 pp. 1462-1463 (“Contrary to popular images of serendipitous discoveries in the sciences

²³ Similarly, Dolly the sheep clone was not patentable because she was just a time-delayed version of her exact genetic replica donor mammal. In *re Roslin Institute (Edinburgh)*, 750 F.3d 1333, 1337 (Fed. Cir. 2014).

²⁴ See also *Picard v. United Aircraft Corp.*, 128 F.2d 632, 640 (2d. Cir. 1942) (Frank, J., concurring) (discussing progression of inventive research, from “creative genius” of people like Thomas Edison, to “modern large-scale research laboratories and modern scientific techniques”).

²⁵ *PharmaStem Therapeutics, Inc. v. ViaCell, Inc.*, 491 F.3d 1342, 1377-1378 (Fed. Cir. 2007) (Newman, J., dissenting) (“The patent law recognizes that advances of great power may be based as much on persistent and skilled investigation as on the flash of creative genius, for both serve to transcend that which was previously achieved.”).

²⁶ See also *Intellectual Ventures I LLC v. Capital One Financial Corp.*, 280 F.Supp.3d 691, 698 (D. Md. 2017) (“...we value innovation that leads to new inventions that advance science and technology, protecting that creative effort by issuing patents.”); *Telebrands Direct Response Corp. v. Ovation*, 802 F.Supp. 1169, 1179 (D.N.J. 1992) (“This Court ... recognizes the importance of rewarding inventors for their creative genius and protecting their intellectual property rights from infringers.”).



and instantaneous aesthetic breakthroughs, creativity is hard work.”).

Further, the creativity characterization of invention can be reconciled with the notion that methodical investigation can yield non-obvious inventions. Not all methodical investigations will result in patentable inventions. For example, if the investigative techniques or experiments would have been “obvious to try” to a POSITA at the time of the invention, that can be tantamount to obviousness. In other words, even if the results are not known, if there was a recognized problem or need in the art, and there were a finite number of identified, predictable alternatives, combined with a reasonable expectation that one of those alternatives will provide a successful solution, then the claimed invention could be deemed obvious and therefore unpatentable. See Manual of Patent Examining Procedure § 2143(E). “The rationale to support a conclusion that the claim would have been obvious is that ‘a person of ordinary skill has good reason to pursue the known options within his or her technical grasp. If this leads to the anticipated success, it is likely that product [was] not of innovation but of ordinary skill and common sense. In that instance the fact that a combination was obvious to try might show that it was obvious under § 103.’” *Id.* (quoting *KSR*, 550 U.S. at 421). Other cases have distinguished patentably inventive contributions from activities that are unpatentably obvious combinations of known teachings on the basis that a “motivation to combine” known teachings (which would render the combination obvious) may be found in the “background knowledge, creativity, and common sense of the [POSITA].” *Intercontinental Great Brands, LLC v. Kellogg N.A. Co.*, 869 F.3d 1336, 1345 (Fed. Cir. 2017) (citing *KSR*, 550 U.S. at 418-421); *Plantronics, Inc. v. Aliph, Inc.*, 724 F.3d 1343, 1354 (Fed. Cir. 2013) (same). Thus,

just as creativity is an animating feature of copyright, it also underlies patent law as well.

III. A Conceptual Account of Creativity

Having established the centrality of creativity to the requirement that authors and inventors be human, particularly under U.S. copyright and patent law, we turn now to an in-depth analysis of creativity.

Core aspects of creativity are well-established in philosophical, neuroscientific, and psychological literature.²⁷ Most accounts of creativity include both an external component (e.g., manifestations of ideas, including made artifacts) and a mental component (e.g., a person’s thought process). As we discuss below, both components are not only necessary elements for a complete account of creativity, but both are also incorporated into IP law’s requirements. See *generally* Boden 2004; Paul and Kaufman 2014. And there is a third requirement, sometimes discussed separately, and sometimes embedded and only implicit in discussions of the external and internal components: the social milieu in which creativity happens.

Creativity includes three core elements: external (artifacts), subjective (psychological), and social context.

Additionally, there are numerous different continuums along which the magnitude of creativity, or some quality of creativity is assessed. For example, Marcus du Sautoy (discussing the work of Margaret Boden) identifies three types of creativity, ranked from common to rare. The most common,

²⁷ Here, we are focused on developing the concept in sufficient detail to address questions of creativity in IP and AI, without undertaking to provide an exhaustive literature survey.



he asserts, is “exploratory creativity,” where the creative act still adheres to the established rules and paradigms, and only nibbles at the edges. Less common is “combinatorial creativity,” in which two different, often unrelated, constructs are combined to provide some interesting new insight. And finally, rarely, comes the category of “transformational creativity,” which Du Sautoy refers to as “phase changes,” where the creative act introduces some new and completely different paradigm. Du Sautoy 2019 pp. 9-11; *see also* Boden 2004.

Another continuum is the community for which a particular act or artifact is deemed creative. Most narrowly, something may be creative only for the single person involved. Consider, for example, a drawing made by a young child; the drawing may reflect the child’s creativity even as it is highly similar to other drawings made by other children. Most profoundly, an act or insight may be creative vis-a-vis society at large. Boden folds this into her distinction between “historical creativity” (“H-creativity”) and “psychological creativity” (“P-creativity”). Boden 2004 pp. 2-3, 43-44. In Boden’s account, H-creativity is that which is new to “the whole of human history,” while P-creativity is that which is new to the person having the creative idea. Boden 2004 p. 43. At an intermediate waypoint, something may be creative as to a particular community, for example, one defined in terms of geography, interest, or specialized domain expertise. See Abraham 2018 p. 9. These acts and insights may often be considered transformational, though transformational acts of creativity may also be situated within particular communities of interest. Abraham 2018 pp. 9-10.

A. External Creativity (Creativity of Artifacts)

Discussions of creativity often begin with


reference to a thing—a text, image, or performance, for example—as “creative.” Some things that we might want to call creative in a broad sense are not human-generated. One often-repeated example is the products of evolution in nature. These could be said to be “creative” in some broad sense. But of course it is not possible to obtain patents on laws of nature or to get copyrights on naturally-occurring phenomena, such as snowflake patterns.

As such, we know there must be something more besides just being new and useful (and surprising). To get at this, we focus on “ideas” and “artifacts” as subsets of “things.” Margaret Boden emphasizes that “ideas” and “artifacts” are, by definition, the products of psychological processes. Boden 2014 p. 227, 233-234. Following Boden’s nomenclature, we use “thing” to refer to a physical manifestation, “idea” as the product of a mental process that is not physically manifested, and “artifact” as a physical manifestation that was generated at least in part through psychological processes. In short, a tree and a cloud are things but not artifacts. Figuring out how to repair a motorcycle is an idea. A painting is an artifact. Although it might be argued that some intangibles, such as an oration or a dance performance, are also artifacts, for the purpose of this paper we will limit our use of the term to tangible items.²⁸

Common definitions of creativity in reference to artifacts focus on two, sometimes three elements: that it is new, valuable, and (sometimes) surprising. Margaret Boden is often credited with articulating this definition. See Boden 2004 pp 1, 40-41. Newness (or novelty) and value are both highly variable terms, suggesting both variable frames of reference (new to the creator or society, valuable in a sentimental, scientific or monetary sense, for example) and society-dependent assessments (there must be some social

²⁸ Both patent and copyright law include a requirement of tangible items—copyright requires expression in a tangible form and patent requires reduction to practice.

milieu at a particular point in time making the assessment of value). Anna Abraham has provided a more expansive presentation of the first two elements in the following Venn diagram:



Abraham 2018 p. 8.

Abraham’s account highlights the range of ways that something can be considered original. She also subsumes “valuable” within a broader category that she labels “appropriate.”²⁹ This highlights the extent to which this element of creativity depends on social context. In addition to “value,” which suggests monetary or market-based considerations, Abraham has also included terms like “relevant,” “suitable,” and has labeled the overall category “appropriate.” Thus, to be considered “creative,” things cannot simply be randomly generated or randomly recombined versions of existing things even if the combination happens to be novel. At the same time, some artifacts are considered “creative” because of some particular relevance or salience (to the relevant audience, whether just the creator or some broader group or community), even if there is little or no market value for the

artifact.

Like value, originality is also sometimes assessed in relation to a relevant community of interest. Sometimes, the originality is just whether it is new to the person creating it; sometimes it is a group; sometimes it is society at large. As we have seen above, Boden folds this continuum into her accounts of creative things versus creative processes, calling something that is novel to society “historically creative” or “H-creative,” and calling something that is novel to the individual creator “psychologically creative” or “P-creative.” She

acknowledges that for something to be H-creative, it must also be P-creative. Boden 2004 pp. 2-3, 43-44. Others, such as Bence Nanay, adopt a different framework, one that distinguishes between the things or artifacts, on the one hand, and the people or mental processes, on the other hand. Nanay 2014 pp. 18-23. Drawing on earlier scholarship, Nanay defines subjective creativity as “a property of persons or their minds,” and contrasts this with what he calls “objective creativity”³⁰ as “a property ... of created works.” Nanay 2014 p.18.

At least initially, Marcus Du Sautoy appears to regard this external aspect of creativity as independently sufficient to provide an account of AI-creativity. In particular, he discusses several experiences of AI algorithms devising new strategies for winning various games, including breakout (a paddle-and-bricks video game), chess, and Go. In each case, acting within the rules of the game, the algorithm employed new strategies or sequences that proved more

²⁹ See also Fromer 2010 p. 1459.

³⁰ Because we emphasize the extent to which external creativity implicates the subjective assessment of the community (e.g., in determining value), we believe it is more appropriate to label this category as “external” rather than “objective.”

successful than the accepted conventional strategies. Du Sautoy 2019 pp. 16-39.³¹ Identifying the ways in which the AI strategies were new, valuable and surprising, he concludes that they were creative. However, limiting the definition of creativity in this manner would allow for a conclusion that some natural phenomena, such as adaptive genetic mutations, are also creative. See Boden 2019 pp. 173 et seq. But if Du Sautoy's argument is that the AI algorithm AlphaGo is creative in the same way that nature is creative, that is unrelated to the inquiry in this paper, concerning what kind of creativity is necessary to support IP eligibility. Natural phenomena are ineligible for IP protection. In later chapters, where he discusses *art*, Du Sautoy seems to come around to a more expansive viewpoint, at least in some circumstances: "Art is ultimately an expression of human free will and until computers have their own version of this, art created by a computer will always be traceable back to a human desire to create." Du Sautoy 2019 p. 98.

In sum, there is a sense, external to the maker's state of mind, in which we can refer to things and ideas and artifacts as creative. In so doing, under commonly-adopted definitions, we are saying that the thing (or idea or artifact) is creative because it is new and valuable—and possibly also that it is surprising. As we will see below, at least value and surprisingness cannot be assessed in a vacuum; they require assessment within a social context.

B. Subjective Creativity (Creativity of Mental Processes)

In addition to the externally-focused conception of creativity in relation to artifacts, much has been written about the mental process of creativity. Sometimes it is expressed as a "flash of genius," or "divine inspiration" or other involuntary and sudden

occurrences. But sometimes creative insight is the product of considerable effort, work, and iterative refinement. Nanay adds the observation that a person who experiences an idea as creative "experiences it as something she has not taken to be possible before." Nanay 2014 p.23 (emphasis omitted). As a consequence of Nanay's additional gloss, the process of engaging in research via the performance of "brute force" iteration through a large number of permutations that were imaginable but perhaps not tried before is *not* creative. Nanay 2014 p.25.

The literature is replete with discussions of various aspects of human psychology as they relate to creativity. For example, is it a conscious or unconscious process? See *generally* Baumeister 2014. Is an act of creativity an intentional act? See Kieran 2014 pp.126-127; Boden 2014 pp. 233-234. What role does domain knowledge or expertise play? See, e.g., Kieran 2014 pp. 130-131, 141-142; Abraham 2018 p. 57 ("increased fluency is linked to an increased likelihood of more remote or unusual associations").

The accounts we find most persuasive involve interplay between what is sometimes referred to as divergent and convergent thinking. As Elizabeth Picciuto and Peter Carruthers explain: "Divergent thinking is associated with defocused attention, which involves more defocused thought. Out-of-left-field generation of ideas comes about via divergent thinking. Convergent thinking is more rigorous and analytical and is associated with focused attention." Picciuto 2014 p. 207. See *also* Abraham 2014 pp.34-38, 81. Neurologically speaking, divergent thinking is associated with the "default mode network" (DMN), which is the part of a brain that is active when we're not engaged in any particular task, such as when we are daydreaming or

³¹ But see <https://www.vice.com/en/article/v7v5xb/a-human-amateur-beat-a-top-go-playing-ai-using-a-simple-trick> (accessed June 27, 2024) (reporting that humans have more recently identified a tactic enabling the human to beat an AI at Go).

sleeping. Goldberg 2018 pp. 53-54. Convergent thinking is associated with the “central executive network” (CEN), which is the part of our brain we use when focused on a particular task, leading the CEN to be characterized as “task-positive.” *Id.* p. 53. See also Bartholomew p. 387 n.168 & accompanying text.

At the risk of oversimplification, the default mode network is free to wander, sometimes calling up ideas or knowledge from disparate subject areas into overlapping or adjacent brain regions. Much of the time, the activity of the DMN is entirely routine, or the overlaps it presents are nonsensical, irrelevant, fleeting. But sometimes something that is novel and relevant will emerge from the DMN that gets noticed. Elkhonon Goldberg describes this noticing function as being performed by the “salience network” (SN) of the brain. Goldberg 2018 pp. 54-56. Getting noticed may then trigger the CEN to focus and pay attention. *Id.* Additionally, the DMN and SN may be primed to notice particular things, based on current or recent CEN-directed efforts to solve a particular problem, and in the context of particular domain expertise. Hence there are many stories of someone working very hard on a problem, going to bed, and waking up in the morning with the solution suddenly in mind. Goldberg refers to this as “directed wandering” of the DMN. Goldberg 2018 p. 132 (“There are good reasons to believe that mental wandering saves the day when a conscious, systematic effort to solve a problem comes up short. *But in order for this to happen, it must be preceded by a period of a ‘hyperfrontal’ deliberate go at the problem.*” (emphasis in original)). As Goldberg puts it all together:

Deliberate effort plays an important role in the creative process in synergy with mental wandering. It is precisely the combination of the two processes, one deliberate and guided by the ‘hyperfrontal’ frontal lobes, the other one spontaneous and ‘liberated’ from frontal-lobe control in hyperfrontal states, and

going back and forth between these processes, that makes the creative process productive and ultimately successful.

...

A creative process usually begins with a conscious idea of what needs to be accomplished, however vague and imprecise. ... [A]n innovative idea usually does not occur to someone who has never pondered the subject matter before; even when the subjective experience is one of the idea ‘appearing out of nowhere,’ it occurs to a prepared mind. The birth of a creative idea begins through the frontal lobe-driven process by activating certain regions [of the brain]. The brain is in a state of task-specific (‘task-positive’) hyperfrontality. The activated regions are likely to be quite disparate within the cortex ... But these disparate regions activated during the hyperfrontal state will constrain the ‘mental wandering’ that will come later.

Goldberg 2018 pp. 132-133.

Thus, creative mental processes involve cycling between task-positive activity in the CEN (convergent thinking) and mental wandering in the DMN (divergent thinking), mediated by novelty- and relevance-filtering, carried out, e.g., by the SN.

This account appears to be consistent with the experiments conducted by Baumeister et al. to test the role of conscious attention in creativity. Baumeister, 2014. They conducted three experiments. In the first, called the guitar experiment, they asked jazz musicians to improvise three guitar solos along with chord progressions they had not been exposed to in advance, while also performing a mental task. In one, they were subjected to a higher cognitive load (counting backwards by 6 from 913), in another they had a lower cognitive load (counting up by 1 from 15), and in the third, they were given no task. Independent experts judged the relative creativity of the three solos. While the low-load and no-load

solos were roughly equivalent, the high-load solos were deemed much less creative (though still competently in-rhythm and on-key) in that they tended toward repetition and simple, stilted phrasing. Baumeister 2014 p. 191. The second, a picture drawing experiment, the subjects were asked to draw pictures while listening to two different songs. For one of the songs, they were asked to count the number of times the word “time” appeared in the lyrics; in the other, they were given no special instructions. Again, independent judges rated the low-load pictures as significantly more creative than pictures drawn under high cognitive loads. There was no difference in the coherence of the pictures, their mood, or even the number of colors used. Baumeister 2014 pp. 192-193. In the third experiment, the authors sought to increase creativity via conscious and unconscious cues. The task involved coming up with a title for a short story about popcorn in a frying pan. To trigger conscious engagement, the instructions about writing a title simply included the word “creative.” To trigger unconscious engagement, the participants were first asked to participate in a word-association exercise that included pro-creativity phrases. The authors concluded that the conscious instruction “significantly increased creativity,” but the exercise aimed at the unconscious “had no effect.” Baumeister 2014 p. 195. This assessment of magnitude of creativity has support in the literature. As outlined by Anna Abraham, it can be useful to categorize the magnitude of creativity into four tiers: mini-c, little-c, Pro-c, and Big-C. Abraham 2018 pp. 11 et seq., p.28. In this framework, mini-c corresponds to “personally meaningful interpretations of experiences, actions and events,” while Big-C corresponds to “monumental and lasting” demonstrations of genius. *Id.* The other two, little-c and Pro-c, reflect interpersonal levels of creativity corresponding to varying social groups and varying levels of domain expertise. *Id.* Abraham makes a further observation that

aligns with the results of Baumeister’s experiments: “[H]igh cognitive load negatively impacts the ability to retrieve remote associations by narrowing attentional control whereas, under conditions of low cognitive load, the activation of wider associations is an exploratory process by default.” Abraham 2018 p. 57.

Both the guitar experiment and the picture experiment demonstrated that if a person’s conscious attention (the CEN) is impaired by another unrelated task, then the relative creativity of the primary task (improvising a jazz guitar solo or drawing a picture) is reduced (but not completely eliminated). The experiments also align with another of Goldberg’s observations that, as particular tasks become familiar, routine, mastered, the part of the brain responsible for them shifts from the right hemisphere to the left hemisphere of the brain. Goldberg 2018 p. 97 (“When a task is novel, it is supported mostly by the right hemisphere, but the left hemisphere takes over as the task is being mastered and becomes increasingly familiar.”). Accordingly, the guitarists could stay in tune, in time, and play familiar (even repetitive) riffs—because, as accomplished musicians, these are familiar enough to be performed using (left hemisphere) parts of the brain, which are not associated with creativity-seeking for conscious tasks.

Within the framework outlined above, some of the important roles of domain expertise begin to come into view. For example, in the task-positive state, a person’s level of domain expertise can help focus the mind on the particular challenges at hand. And those who are expert in a particular field – for which certain parts of the analysis are routine – will experience those routine parts being performed by the left hemisphere, freeing the right hemisphere to perform the novelty-seeking aspects of the project without congestion from the routine or known details.

C. Social Context


The notion that creativity and the creative process is situated within a social context permeates the analysis of both the external and subjective aspects of creativity. “Any individual creative act can only be understood in the cultural context in which it occurs.” Goldberg 2018 p. 44.

Commentators appear fond of invoking a 1957 quotation from artist Marcel Duchamp: “The creative act is not performed by the artist alone; the spectator brings the work in contact with the external world by deciphering and interpreting its inner qualification and thus adds his contribution to the creative act.” Nanay 2014 p. 27; Carroll 2014 p. 62. But Duchamp’s observation is ultimately under-determinative. Not only can an audience find meaning in a urinal hung on the wall and called “Fountain,” (1917) but an audience can also find meaning in particular formations of clouds (e.g., that they look like hippopotamuses) or burned toast (e.g., that it contains the face of Elvis). The role of the community and social context is more complex than simply reacting and responding to a phenomenon encountered in the world.

Looking first at the components of the working definition for creativity in the “external” sense, that a creative artifact is new, valuable and (perhaps) surprising, the question of novelty or originality depends heavily on what is already known in the relevant community. As Abraham explains, “Frames of reference can be at the level of (a) an individual, which is akin to P-creativity, (b) a group, or (c) mankind, which is akin to H-creativity.” Abraham 2018 p.9. So in one sense, the measure of novelty (and, for that matter, surprise), is against that frame of reference. Something may be not-novel to mankind, but nonetheless novel to the individual creator or to a particular group. To be sure, things that are mini-c creative in the sense that they are new to the individual creator may reflect asocial insights. But if the novelty is assessed with respect to a

relevant community (e.g., Miss Shannon’s algebra class, plasma physicists, or classical composers) (i.e., little-c or Pro-c, respectively) or with respect to society at large (Big-C), the community’s knowledge is central. See Abraham 2018 p. 11. Legal scholars have likewise recognized the importance of social context to the external component of creativity. Buccafusco 2014 pp. 1932-1933 & nn.63-66 (“[A]ppropriateness indicates that some community recognizes the contribution as socially valuable.”) (collecting sources); Fromer 2010 pp. 1460-1461 (same).

Similarly, the question of whether an artifact is valuable depends on the community’s assessment. This assessment occurs not just on a positive scale – a new painting is worth millions, or a stock photo is worth a few dollars per use – but also in relation to the ways in which things may lack value. Something routine or well-known may lack value, but similarly something that is too unconventional, too weird, too inappropriate may also lack value—whether it is simply “before its time” or is just considered to be nonsense. Du Sautoy illustrated this, at least in part, by the following diagram:



Du Sautoy 2019 p. 131.

The “internal,” psychological components of novelty also interpose a certain degree of social context, in several senses. The notion of “directed wandering”—oscillating between



task-oriented convergent thinking and free-associative divergent thinking means that the divergent thinking is oriented or preconditioned to work on the problems coming from the task-positive mode, and the brain's filters are likely to flag as relevant those things that are relevant and appropriate. Relevance and appropriateness could come from any number of sources, not just including the particular problem under consideration, but also assumptions or unchallenged propositions, notions of rules and norms governing the exercise, norms and rules governing the acceptability of the outcome, social acceptance of what has already been decided or what is out-of-bounds, and the like.

In sum, the community plays important roles in providing both constraints or boundaries, as well as encouragement or facilitation, thus both enabling and funneling human endeavors into what we consider to be creativity. The community participates in setting rules, in teaching rules as well as habits, in fostering a culture that encourages exploration and synthesis, and in stimulating creativity. Professor Acemoğlu has also emphasized the extent to which our collective interaction as agents and social learning tends to increase the quantity and complexity of our accomplishments. Acemoğlu 2024 at 5:50-6:30, 8:00-10:30. All told, collective learning process is different in kind from merely encouraging learning through imitation. See Du Sautoy 2019 pp. 11-15.

All three core elements of creativity are important to both patent and copyright law.

IV. Which Parts of Creativity Are Important to U.S. Intellectual Property Law?

Having established that copyright and patent law require a human author or inventor, and that a substantial reason for that requirement has something to do with humans' exercise of creativity, we explored the concept of creativity in greater detail. Now we turn back to IP law to take a closer look at which aspects of this expanded concept of creativity are particularly germane.

Both patent and copyright law clearly require a tangible artifact, matching the external aspect of creativity. For copyright, the work must be expressed in tangible form, via the fixation or recordation requirement. For patents, the invention must be actually or constructively reduced to practice. Curiously, however, copyright does not require novelty (other than requiring that the work not have been copied from someone else) or value.³² Patent law, on the other hand, does require that an invention be "new" and "useful," practically mirroring the standard definition of external creativity, which includes novelty and value.

In addition, both patent and copyright include requirements rooted in the internal, or psychological, aspect of our description of creativity. This is discussed in more detail below.

A. Copyright Law

While the threshold for creativity in copyright is very low, it does require some degree of intellectual effort in producing the work. The measure of originality is author-

³² In copyright, "original" "means only that the work was independently created by the author (as opposed to copied from other works), and that it possesses at least some minimal degree of creativity." *Feist*, 499 U.S. at 345.



centric, primarily focusing on the work not being copied from other existing works to which the author had access. This may be analogized to the “mini-c” type of creativity discussed above. One court even went so far as to comment that two poets, working independently of each other could theoretically each write the same poem and it would be copyrightable as to each.³³

These doctrinal requirements lead to some interesting conclusions.

First, because of the requirement that a copyrightable work be fixed in tangible form, there is a superficial temptation to think that the measure of creativity required should draw guidance solely from the “external” component of the definition of creativity. However, this would not be correct.

Instead, the case law and Copyright Office statements of policy make an explicit linkage between the required “creativity” and “intellectual labor,” meaning that the subjective, or psychological, sense of creativity is required—even if only by crossing a low threshold.

Some commentators have suggested that there ought to be a heightened threshold of creativity as a condition of copyrightability. See Buccafusco 2014 pp.1923, 1930-1931 (discussing “growing debates about whether copyright law’s creativity threshold is set too low”); Subotnik 2011 p. 1491 n.19 (collecting authorities); Parchomovsky 2009 pp. 1517, 1550. One proposal is to adopt something like patent law’s non-obviousness standard. See Miller 2009 pp. 11-13; Bartholomew 2021 pp. 373-376. Bartholomew has also argued that authorial intent-to-create is relevant (*id.* pp. 368-373), as is the author’s domain-specific expertise (*id.* pp. 376-382). Others, however, have argued that no change in the law is needed, and that any problems generated by

a low bar can be handled by varying the scope of protection – from so called “thin” to “thick” copyright – depending on the degree of creativity in the underlying work. Stech 2021 p. 271.

Similarly, copyright’s “originality” requirement, namely that the work not have been copied from another, reflects a low but non-zero situatedness in a social context. Professor Craig makes the case for a strong social, relational aspect of authorship., arguing that copyright exists not just to incentivize outputs, but also to incentivize the very activity of authorship. This includes, on her account, the dialogic processes and exchange of meaning that constitute authorship and reflect the creative agency vital to relational autonomy and human flourishing Craig 2022 pp. 16, 23. She further argues, “It must be urged that the protection of AI-generated works would not advance the kind of ‘creative progress’ with which copyright is concerned—but worse, it could cause copyright to defeat its own ends, stultifying creative practices in a thicket of privately owned algorithmic productions.”

Finally, the Supreme Court has rejected a “sweat-of-the-brow” version of copyrightability, not considering how much effort or labor went into a tangible work, but continuing to emphasize the creativity or intellectual labor. *Feist*, 499 U.S. at 359-360. This tends to undermine any argument that a “brute force” version of generating works—for example, the proverbial thousand monkeys banging away on typewriters to produce the works of Shakespeare (see Goldberg 2018 p.120)—can result in copyrightability.

³³ *Sheldon v. Metro-Goldwyn Pictures Corp.*, 81 F.2d 49, 54 (2d. Cir. 1936) (“[I]f by some magic a man who had never known it were to compose anew Keats’s Ode on a Grecian Urn, he would be an ‘author,’ and, if he copyrighted it, others might not copy that poem, though they might of course copy Keats’s [because Keats’s copyright would have expired].”)

B. Patent Law

For patents, there appears to be a more robust alignment between the patentability analysis and the elements of creativity.

Not only is an external artifact required, but it must be “new” and “useful,” satisfying the novelty and value components of the external aspect of creativity.

Subjective/psychological creativity is also required. The literature on creativity acknowledges that some creative efforts may come all in a flash, while others may come slowly over years of painstaking effort. See, e.g. Goldberg 2018 pp. 132-140. This aligns with the case law’s observation that a patentable, nonobvious invention may arise from either a flash of genius or from

painstaking incremental research, so long as the research plan involved some creative insight and was not merely working through known permutations (*i.e.*, was not obvious to try). In that regard, “brute force” efforts to work through known permutations may not satisfy the creativity requirement – even if they can be done with a computer at significantly reduced cost and effort.

Furthermore, the social context is important. In addition to the social aspect of determining usefulness or value, obviousness is measured in terms of what would have been known to a person of ordinary skill in the art. This aligns roughly with the “Pro-C” type of creativity discussed above.

C. Summary

We summarize this comparison in the following table:

	Copyright	Patent
External	Required – expression in tangible form	Required – reduction to practice (actual or constructive)
New	Irrelevant (unless “copied” from existing work)	Must be “new”
Valuable	Irrelevant	Must be “useful”
Surprising	Irrelevant	Must be “non-obvious”
Internal	Required – must be “fruit of intellectual labor” that “are founded in the creative powers of the mind”	Required – either as “flash of genius” or methodical research (but not obvious-to-try and not a combination that a POSITA using ordinary creativity could identify)
Convergent	(implicitly required?)	(implicitly required?)
Divergent	(implicitly required?)	(implicitly required?)
Salience	(implicitly required?)	(implicitly required?)
Social Factors	Minimal – not copied from another; additionally, the process of authoring may be inherently social	Role of level of skill in the relevant art at the time of the invention

V. Risks to Creativity in an AI World

Picking up several themes noted in the previous section, that copyright law does not currently require a quantum or quality of creativity (a low threshold that has been subject to criticism), and that patent law requires something more than the ordinary creativity of a POSITA, it appears that there is an undercurrent of creative *quality* that forms part of the IP-creativity discussion. As we prepare to turn our attention back to AI “creativity” as viewed through the lens of IP law’s requirements, it is worth pausing to note how AI appears to be impacting the quality of both human creativity and AI “creativity.”

A. How AI Affects Human Creativity

During our examination of the role of creativity in IP law, we have seen that copyright, in particular, does not require a qualitative threshold that must be exceeded for works to be copyright-eligible. And we have seen that some legal scholars have expressed concern about this, arguing for an increased qualitative threshold. This debate has important ramifications for the broader discussion about AI and creativity as well, because there is emerging evidence that the increasing prevalence of machine learning platforms may be downgrading the nature of human creative capacity.

Advances in neuroscience have clearly demonstrated that the “wetware” of our neurological systems changes in response to our physical surroundings, including the spaces in which we dwell and the tools with which we engage. The resulting impact on our cognitive capacities can take place either at a short-term behavioral level (e.g., “state-dependent learning” influences the way we access memories depending on our physical environment, as first demonstrated by Godden & Baddeley in 1975) or at more fundamental levels leading to sustained changes in our neural pathways after even

just a few minutes of interaction (Pascual-Leone 2001). For example, experimental data indicates that interacting with AI-generated spaces, such as the social media or search engines found on a typical smartphone, triggers a downgrading in human cognitive processing. The conditioned response is so strong that the mere presence of our smartphones is enough to catalyze an effective short-term drop in fluid intelligence (*i.e.*, a reduction in the type of divergent thinking associated with insight and complex problem-solving). Ward et al. 2017.

Given this sensitivity of human neural systems, there are real concerns that AI deployments can negatively impact humans’ own long-term creative capacity, ostensibly in exchange for the short-term rewards (such as speed) offered by AI. While human creativity can be deployed across a variety of functional levels as noted by Kaufman & Beghetto (2009), ranging from learning new things (“mini c” creativity) and problem-solving in the workplace (“Pro-C” creativity) to pioneering eminent new achievements (“Big C” creativity), AI platforms currently only support a few of these creative processes, (namely Pro C rather than mini c and Big C). In other words, ongoing expansion of AI deployments may compromise the powerful, multi-level nature of human creativity – including processes such as the “mini c” creativity needed to learn new things, and even processes such as Big C creativity (LLMs are necessarily limited to the set of domains found in the training data, thus tending to disadvantage revolutionary insights typically associated with Big C creativity).

There are several ways in which interactions with AI downgrade human cognitive capacity for creativity. First, experimental research indicates that an individual’s creative ability expands significantly when exposed to conditions which prompt meta-awareness of their own thinking patterns and mental models: when participants are challenged to rethink the

“convergent thinking” frame they use in a particular domain, and to practice cognitive flexibility (such as toggling between convergent and divergent routes of thinking), this led to a significant and sustained increase in creative capacity. The long-term impact on participants’ creativity could even be seen in other domains, as measured by tests such as the Alternative Uses Task, Gibson’s Functional Fixedness Test, and the Remote Associates Task (Maddux & Galinsky, 2010). In contrast, working with AI to produce creative products involves working in a way that emphasizes speed and instant answers (as well as becoming the passive consumer of such answers), rather than self-reflection or toggling between convergent and divergent thinking frames. Landmark experiments more than 20 years ago clearly demonstrated that exposure to an environment of fast-paced, reactive “shallow focus” digital prompts is enough to temporarily downgrade human creative problem-solving capacity dramatically, the equivalent of dropping 10-15 IQ points after just one hour of exposure (Jackson, Dawson, & Wilson 2003). Moreover, the human neural system is particularly vulnerable to informational overload: exposure to excess data, with its increased processing demands, reduces the cognitive resources we have available for deeper processing and memory consolidation. It leads to swift changes in our nervous system, narrowing our attentional focus, and reducing our ability to see the type of larger patterns or widely-distributed associations necessary for creative thought. Using EEG mapping, neuroscientists have shown that even a few minutes spent trying to navigate through the wealth of data offered up by AI-generated recommendations is enough to shrink the brain’s effective field of focus, downgrading the alpha brain waves needed for deeper thought and imagination (Peng et al, 2018).

Second, as Vinchon and others have pointed out, in order for humans to have a stable capacity for creativity, they need not

just to be able to generate novel ideas, but also to have a driving underlying curiosity, a desire to seek out new and unexpected options (Vinchon et al, 2023). Yet as researchers have shown, humans interacting with AIs tend to lose confidence in their own creative skills (Habib et al, 2024). Furthermore, their curiosity plummets and they start to restrict the range of their own creative repertoire in favor of creating “mash-ups”, made from assembling pieces from pre-existing image/ music, rather than experimenting with independently creating something novel and surprising. Harvard psychologist Howard Gardner has shown that over the 20 year period in which teenagers started using the web and becoming able to access a wealth of idea prompts online, the quality of the teenagers’ exercise of imagination plummeted: while in the 1990s, 64% high schoolers were able to write short stories with unconventional plots, non-linear narratives and were given to blurring the boundaries of different genres, imagining different time and space than where they were, by 2011 the number of high schoolers using unconventional plots had fallen to 14% and only 5% used a different locale for their story, compared to 32% in the 1990’s. Teachers reported that by 2011, even teens who had won places at Art & Design Colleges were no longer willing to imagine a story idea or visual image in their “mind’s eye”, and tended to rely on searching online first (Gardner & Davies, 2013).

Recent research with adults suggests that while using an LLM such as Chat GPT4 to aid in creating a short story helped less creative individuals produce stories that were judged as more interesting, and more creative, such help had no effect for those in this group who tended to score highly on creative tests already (note that this study sample did not even include highly accomplished “creatives”). Furthermore, independent evaluation of the stories produced after using ChatGPT4 for just a few minutes indicated that participants significantly reduced the diversity of ideas

incorporated into the stories, leading to a greater homogeneity between the stories (Doshi & Hauser, 2024). This has important implications for the deployment of AI: if interacting with AI simply serves to increase the creativity of humans who find creativity difficult at the cost of decreasing creativity for humans who are functioning adequately, and also at the cost of narrowing all human creative output down to a lackluster, limited set of ideas, the question arises as to whether this is a worthwhile net sacrifice to human creativity. Further, Professor Acemoğlu argues that the ways in which AI takes a narrow view of human talent will lead to less experimentation and fewer new discoveries (*i.e.*, diminished creativity). Acemoğlu 2024 at 10:45-12:30. He further argues that AI leads to a bias toward following AI-supplied recommendations without intervening exercises of human judgment – what he calls “conformity bias” and “informational herding.” *Id.* at 40:00-44:50.

The nature of the types of algorithms underlying LLMs makes it inevitable that the more that one works with platforms such as Midjourney or DALL-E, the more that one becomes acclimatized to mash-up, novel-yet-not-too-surprising solutions. LLM's are built to find solutions that are "plausible", doing so by using predictive AI algorithms, steering themselves towards the most likely outcome, and away from novel or unusual possibilities; any outlier data points have little impact on training and will "most likely be ignored" (Esling & Devis, 2020, p.7). Consequently, although humans engaging in such AI-based environments may well be exposed to a greater *absolute* number of "creative solutions", such solutions will tend to cluster within the same narrow range, effectively acclimatizing participants to a narrower range of possibilities, and flattening their own future creative repertoire. Additionally, the "pixelated" nature of the AI creative process (where creative products are made out of a sequence of short-term decisions), means that our exposure to AI-driven creativity

leaves us starved of examples of the type of longer, larger scale patterns found in the best (“Big-C”) human creative achievements (Brandt, 2023), which further catalyzes a recalibration of our imagination toward a greater propensity to accept shallower types of creative products.

Human imagination and creative capacity is further being downgraded by the fact that LLM's have limited access to the realm of embodied knowledge. To be sure, all the text on the Internet, or all images available in digital form, are incredibly broad selections. But computer engineers have yet to design code that captures the large swathes of human knowledge not digitized and online, such as information derived from the somatic experience of an organic body, from the emotional impact of a piece of music to our somatic experience of the physical environment, or of our flesh-based interaction with other life-forms. As a result, AI-generated creative products are tilted away from such considerations. For example, in April 2024, artist Sarah Bird projected an image of a redwood tree onto an iconic urban landmark (the San Francisco Ferry Building clock tower), an art piece that both provokes emotions (by confronting the viewer with the physical impact of encountering such a large life-form in a city environment) as well as pushing us to rethink our assumptions about scale, assumptions that arise from being small life-forms with very specific visual apparatuses. The piece (“Being/ Tree”) challenges the viewer's experience of seeing and thinking from within an organic body. This represents the type of creative work and philosophical questions not easily accessed by LLMs, and that are thus vulnerable to being sidelined A.I.-driven decisions about creativity.

B. Algorithmic Flattening: How the Design of AI Models Constrains the Quality and Diversity of “Creative” Outputs

Similarly, AI outputs are not all champion-beating stratagems for Go or chess. Indeed (even as generative AI continues to be remarkable in its production of outputs bearing human-seeming verisimilitude), there is a proliferation of what in some circles is derided as “AI slop,” or vast swathes of junk content output by generative AI. This raises the question how the AI models themselves influence the quality and diversity of creative outputs.

In *Filterworld* (2024), Kyle Chayka notes that the technology embedded in the platforms used for visual inspiration is narrowing our choices, creating an “algorithmic flattening” of visual culture. The notion of algorithmic flattening predates generative AI, broadly applying to machine learning in social media systems. Algorithmic platforms are endogenous, ranking potential outputs from within self-contained systems (Ursu, 2018) where they moderate and curate the flow of content (Gillespie et al., 2014). This contributes to a homogenization of cultural influences within filtered communities (Schulman, 2013). Furthermore, algorithmic systems are often evaluated or trained with data from users already exposed to algorithmic recommendations; creating an insidious feedback loop that amplifies the homogenization of user behavior (Chaney et al., 2018).

The algorithmic structures of machine learning systems that create this flattening are complex, inclusive of the data, the engineers who design algorithmic systems, and the structural functionality of the algorithmic models themselves. LLMs and diffusion models are built upon earlier machine learning systems, with many of the same structural issues. It is quickly becoming clear that generative AI outputs demonstrate many of the shortcomings present in earlier machine learning models.

This algorithmic flattening can tend to amplify the limitations of the AI’s training data, resulting in, for example, racial bias in image generators, leading Midjourney and Dall-E to distort portraits of Black women (Small, 2023). Gender bias is also an issue, as it would be challenging to train models on inclusive data. Only 1% of works in the National Gallery in London by women and only 11% of works at MoMA in New York. The biases embedded in training data and “digital negative space,” (Scolere et al., 2018) or what is absent from the training data, may distort or influence the quality and diversity of generative AI outputs.

Furthermore, when AI-generated content feeds future models, there is the potential to create a “self-referential aesthetic flywheel” (Epstein et. al. 2023), where successive models are defined and constrained by the outputs of the prior models, without the kind of diverse new inputs that comes from human-generated training data. Generative AI outputs have the potential to map onto Chayka’s assertions about algorithmic flattening. Chayka highlighted that every coffee shop looks the same worldwide because of Instagram inspiration. Some have argued that AI images already reflect this flattening, with video game aesthetics being highly prevalent in AI images. As AI-generated images reflect the sociocultural context in which they were created and on which they were trained, they have the potential to accelerate biases and flattening, perpetuating AI-driven cultural norms. In fact, research has shown when LLMs are trained on recursive data, model decay sets in, with model performance decreasing over time (Shumailov et al, 2024). In extreme cases, where AI is trained and retrained only on AI-generated outputs, without the input of human creativity, there quickly follows an “irreversible model collapse.” Acemoğlu 2024 at 40:00-44:50 (discussing Shumailov). In other words, Shumailov’s findings on model decay suggest that generative AI models will, for the time being, require original creative content in order to perform.

These observations suggest that there may be limitations to AI “creativity”—it may tend to regress to the median of its training data (which, in many circumstances, may include skewed or biased sample sets), and without the “refresh” of *human* creative input, may not be self-sustaining. If AI creativity becomes *less* “creative” over time by virtue of the design of generative AI systems, what does this mean for AI, creativity and IP-eligibility?

VI. Revisiting the Question of AI “Creativity”

We now turn back to the question of AI “creativity” as it relates to IP-eligibility. As we have seen, the concept of creativity—particularly in relation to IP law—involves ideas clustered around at least three distinct elements: the artifact itself that is supposed to be the product of creativity, the mental processes by which creativity is performed, and the social milieu within which creativity is assessed and perhaps even made possible. As we suggested at the beginning of this paper, most popular discourse—and a significant proportion of the scholarly literature—focuses solely on the external aspect of creativity, whether the text or image or other generative AI output is new, valuable and surprising. *E.g.*, Fromer 2010 p. 1459. A 2023 article that attempted to quantify the creativity of ChatGPT-4 explicitly adopted such a construction, framing the question as “Is AI creative according to the external evaluation of its actual output?” Guzik 2023. At the same time, Margaret Boden, who has for decades produced widely cited scholarship at the intersection of creativity and computers, concluded as recently as 2014, “[T]he question of whether a computer could ever ‘really’ be creative is currently unanswerable, because it involves several highly contentious philosophical questions.” Boden 2014 p. 242.

That said, both our anecdotal observations and emerging scholarship support the ready conclusion that generative AI outputs can meet or exceed the purely artifactual or “external” *originality* of human-created outputs. Erik Guzik et al. administered the standardized Torrance Tests of Creative Thinking (TTCT) to ChatGPT version GPT-4 (eight separate submissions) and a control group of 24 humans. As measured by the TTCT, the average of the GPT-4 submissions consistently outscored the average of the human control group, with the top-scoring humans on par with the GPT-4 submissions. Guzik 2023. In an indication of how rapidly generative AI’s capabilities are advancing, Guzik cites a 2022 study of GPT-3, just one year earlier, which concluded that “GPT-3’s ability to generate unexpected and novel ideas ... did not match that of humans.” Guzik 2023.

An even more recent study concluded that GPT-4 scored higher than humans on a standardized battery of tests to measure the identification of “divergent thinking” tasks (*i.e.*, the identification of multiple solutions to a problem rather than a single optimal solution). Hubert 2024. The Hubert paper also describes a number of other recent studies, broadly supporting the conclusion that GPT-4 can outperform humans on standardized tests designed to measure certain types of human creativity. However, Hubert notes, these experiments focus on the external aspect of creativity, emphasizing the character of the outputs produced. Hubert employs the external-creativity definition: “To comprehensively examine creativity requires not only an assessment of *originality*, but also of the *usefulness* and *appropriateness* [such as accounting for sociocultural and historical contexts] of an idea or product.” Hubert 2024 (emphasis added). (Recall that above, we suggested that usefulness and appropriateness both fall within the broader rubric of “value.”) Hubert noted that in their study, even though GPT-4’s originality was higher than that of humans, “the feasibility



or appropriateness of an idea could be vastly inferior to that of humans.” *Id.* Further, Hubert noted, “humans came up with a wider range of responses” than GPT-4, highlighting that “flexible thinking may be the strong point in human-centered divergent thinking.” *Id.*³⁴ Finally, Hubert notes the dependence of AI on human prompting: “AI creative potentials are dependent on the assistance of a human user to elicit responses. Therefore, the creative potential of AI is in a constant state of stagnation unless prompted.” *Id.*

These observations and studies have all focused on the external aspect of creativity. But recall that the intellectual property laws, while requiring an artifact or output, also value the creative process by which that output was generated, as well as the social context in which the activity occurs. It remains difficult to map AI’s internal operations onto cognitive processes that we recognize as analogous to those involved in human creativity. Likewise, modern generative AI such as GPT and Diffusion models lack an explicitly engineered, internal model of the world. It is therefore hard to see how they possess the social situatedness that we have seen is involved in human creativity.

A number of GPT-based text-generation platforms now offer user-selectable settings to be “more focused” or “more creative.” But this setting is not actually able to increase psychological creativity as we have explained it. Rather, the setting merely alters the probability and randomness of next-token selection during the text generation process. Language models based on the transformer architecture output a *probability* for each possible next token – in essence a numerical indication of how likely each next token is to appear immediately after the block of text that was entered as the input, based on an internal statistical model derived from the training data. The token that is presented at the

output is chosen according to these probabilities, with more likely ones, according to the model, more likely to be selected. This process is iterated, with the output token appended each time to the input text and fed back into the model, to generate sentences and paragraphs.

With respect to the question of AI “creativity,” it is important to know that these probabilities are modified by the so-called “temperature” parameter of the model, selected by the user. At low temperatures, the likelihood of selecting unlikely tokens *decreases*, while the likelihood of selecting likely tokens *increases*. Conversely, at high temperatures, the likelihood of selecting unlikely tokens *increases* and that of sampling likely tokens *decreases*. At very low temperatures there is less randomness in outputs, as probability congeals around the most-likely or “best” token. And at very high temperatures there is more randomness, as probability becomes more evenly distributed over tokens. This difference in randomness underlies the appearance of “focused” model responses at low temperatures and “creative” or “unexpected” responses at high temperatures.

Similarly, diffusion models, such as those used in AI image-generation, work by iteratively adding detail, starting from random so-called “noise,” until a desired output is reached. Ho 2020. To achieve this for images, first a series of “noising” maps are obtained, each of which perturbs images in a random (but predetermined, and input-dependent) manner. The effect of the series of noising maps on the space of all possible images is intended to be mathematically analogous to the physical process of how a gas diffuses over time to fill a space in an unstructured way. For each stage of noising, a corresponding “denoiser” model is assigned, tasked with undoing the small perturbation at that stage. When the series of denoisers are well-trained to clean up

³⁴ See *supra* § V.B.

noise in this way, they may be applied serially, starting with truly random noise, to obtain a well-formed image (such as one that fulfills the description provided in the generative AI prompt), and this finished serial composite of denoisers is called a diffusion model. “Stable Diffusion” models work similarly, except on a more compact encoding of images in the so-called latent space of an autoencoder, rather than directly on the images themselves. Rombach 2021.

Nonetheless, implausible as it may seem from what we know of the architecture of these systems, in order to test the assertion that AI models can be sufficiently “creative” to qualify as authors or inventors under IP laws, we undertake here to tentatively examine some of the claims made in the published literature about these models through the lens of the Convergence-Divergence-Saliency framework of creativity outlined above. To be clear, we conclude that it is not presently possible to “map” onto AI models the psychological and social-context aspects of our robust conception of creativity. Because those aspects matter, in varying ways, to copyright and patent law, AI models cannot be deemed analogously creative to human authors and inventors.

Regarding convergence, generative AI can generate what the human audience experiences as coherent and contextually appropriate outputs in response to human-entered prompts. However, it is unclear whether these models can reliably identify the best or most optimal solution, which is a characteristic of the reasoning associated with convergent, or task-focused, behavior. Reliably selecting optimal solutions arguably requires world models and the ability to evaluate the consequences of outcomes with respect to such models, and generative AIs do not yet possess either capacity. As the Stanford Human-Centered AI’s 2024 Artificial Intelligence Index Report states, AI “cannot reliably deal with facts, perform complex reasoning, or explain its

conclusions.” Stanford HAI Report 2024 at p.3; see *also id.* pp. 112 et seq. In other reports, researchers have recently made claims that LLMs might possess human-coherent conceptual representations (Søgaard 2023) and simple world models (Nanda et. al 2023), that they might induce complex representations and computations that go beyond mere memorization and retrieval of training data patterns (Millière, Raphaël, and Buckner 2024a, 2024b), and that they might exhibit some degree of systematicity and compositional generalization in their linguistic behavior (Mahowald et. al 2024). However, other research has concluded that they might not consistently apply the abilities they may have to plan or reason about outcomes in a human-like manner (Wang et. al 2023).

In terms of divergence, generative AI models can display a rudimentary form of “divergent” exploration through the injection of randomness. Language models based on the transformer architecture predict a probability distribution over the next token in a sequence, and by adjusting the so-called “temperature” parameter, they can be made to select either high-probability tokens (low temperature) or low-probability tokens (high temperature) as they generate a sequence of output words. The effect is that the output as a whole can appear more unexpected or “creative” when a high temperature is stipulated. Similarly, diffusion models used for image generation can incorporate randomness in the initial seed or the denoising process, leading to a diversity of generated images. As we note above, the Hubert study demonstrated that AI models could generate lots of different outputs in response to prompts from standardized tests, thereby generating a high score on the “divergent” task batteries, but there was not necessarily much feasibility or practical common sense applied in the responses generated.

And there are, of course, aspects to human divergent-mode thought that randomness will fail to reliably emulate,

being arguably at a higher level of informational representation. For the time being, anyway, these might seem to be essentially human. For example, randomness alone does not naturally integrate the cultural nuances or the deeper contextual relevance that human creativity can incorporate, and which—at least on the social axis—we have seen are fundamental to some forms of human creativity. Without situatedness in the social world, or the embodied experience we all share as we physically mature in the world, it seems unplausible that AI could achieve the highest levels of emotional depth and aesthetic appeal integral to some forms of human creativity.

The salience aspect, which involves filtering for novelty and mediating between convergent and divergent modes of thought, is perhaps the most challenging to postulate in current generative AI models. On a basic analysis for text-models, any notion of autonomous salience detection is mutually exclusive with how closely model behavior adheres to prompts from the human user, and current commercially available LLMs optimize heavily for such compliance (Ziegler et. al 2020, OpenAI 2024). In this respect, one of the key measures of salience remains human-supplied. While there have been efforts to make LLMs more automatic by introducing feedback loops so that human prompting is not required at every step (Nakajima 2023), generative AI cannot effectively filter for novelty or assess the relevance of its own outputs in a goal-directed fashion. Echoes of this theme also appear in Hubert’s observations about the limitations of usefulness and appropriateness of the AI-generated responses in its standardized testing. Hubert 2024.

Looking forward, there are ongoing efforts to bridge the gap between AI and cognitive science, both theoretically and practically. For example, more sophisticated forms of memory and information retrieval are being developed and refined for

generative AI, such as the use of vector databases for "retrieve-and-generate" architectures (Gao et. al 2024). As our understanding of the computational principles underlying cognition advances, it is conceivable that future AI systems may achieve a closer structural and operational resemblance to human creative processes, at least within the constraints of well-defined domains and evaluation criteria. Benchmark reports such as that published by the Stanford University Institute for Human-Centered AI, show that AI’s capabilities have been steadily progressing on a number of measures. For now, however, there are key aspects of human creativity—ones which matter to IP-eligibility—which are both absent and lack analogues in AI models.

VII. Looking Ahead

The foregoing analysis has shone light on the centrality of creativity, and assumptions about the innate human-ness of creativity, in determining who can be authors or inventors under the intellectual property laws of the U.S. and the U.K. But this deeper dive has also exposed potential tenuousness in the connection between creativity and the aspects of creativity that provide the hook for IP protection. This is all the more true in the age of generative AI, where the AI tools are already able to produce text, images, and audio that *seem* creative – at least as creative as the more modestly creative human artifacts that are copyright-eligible. In the near future, the generative AI capability will only increase and become even more impressively capable in its creative-*seeming* outputs.

If it turns out that the relevant constituent components of creativity—external (novelty, value and surprise), psychological (interplay of divergence, convergence and salience), and social context—have equivalents (or sufficient analogs) in generative AI, this may create pressure on the positivist legal position (the conclusion from reading



statutes and court rulings that all authors and inventors must be humans) to simply change the law. That would place us at a crossroads. On one path, firmly committed humanists may need to come up with better explanations or rationales as to why authorship and inventorship should be limited to humans. For example, perhaps (consistent with historical suggestions that the “bargain” embedded in Article I, Section 8 of the U.S. Constitution is intended to provide incentives to humans to be creative) the law should recognize and embody an intrinsic value of the human creative endeavor and the satisfaction that comes from making a creative connection—of having a “Eureka!” moment when the proverbial light bulb comes on.

On the other path, it may be necessary to expand IP regimes to acknowledge when generative AI outputs are equivalently “creative” (e.g., novel, valuable, and (perhaps) surprising) to human endeavors to warrant expansion of IP eligibility.³⁵ This would be but one aspect of a broader examination of how ever-more-capable AI should be handled in our legal framework. This would be a multifaceted, complex undertaking, with both immediate, pragmatic issues and deep philosophical implications. While this white paper focuses on an example of an immediate, pragmatic issue, it is perhaps worth outlining some of the issues we see coming out of the broader framework. As AI gains in capabilities, it moves along a continuum from simple tool to a projected endpoint of artificial general intelligence (AGI), approaching the point at which it might be considered autonomous and capable of acting as an omnipurpose agent. In a wide array of legal domains, the progression of growing capabilities raises issues around how to integrate such

systems into a legal framework oriented around human activities, human agency, and human responsibility.

In the millennia during which human society has had formal legal systems, we have not encountered anything quite as broadly paradigm-challenging as the prospect (and not-entirely-implausible possibility) of AGI, or even AI tools that simply begin to emulate AGI capabilities. Viewed through the lens of AI challenges, questions and assumptions regarding how the law arrived at its current state may yield contradictions, or highlight inconsistencies and biases that require a reexamination or reevaluation of the underlying principles.

As a simple example, some aspects of legal civil and criminal culpability are based on the notions of free will and moral agency. However those concepts are constructed, one may argue even the most autonomous-seeming current forms of AI are too deterministic and algorithmic to qualify for these forms of culpability. Indeed, this discontinuity—and the potential need for a legal framework to hold “robots” responsible—was the impetus for the E.U.’s 2017 proposal concerning “personhood” for robots and other computer-controlled technology.³⁶ Among other things, the E.U. proposal suggested ways to address liability for certain kinds of harms committed by robots. However, the proposal was not well-received at the time and did not move forward.

Yet as our understanding of neuroscience and psychology advances to provide ever more materialistic explanations for human mental processes, and at the same time AI, machine learning and neural networks become ever more complex, approaching AGI, one could argue that the convergence

³⁵ Though we should be mindful of Du Sautoy’s observation, “Often we respond to code that we don’t understand by assigning it some sort of agency. Back when people didn’t understand earthquakes or volcanoes, they created gods that were responsible for these elusive forces.” Du Sautoy 2018 p. 137.

³⁶ European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)) https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html#title1.



of the two trendlines is increasingly foreseeable. As we approach that time, what is the point at which we might be excluding AI from potential legal personhood—or at least legal and/or moral agency—based on incorrect premises and unwarranted bias?

The issue of potential legal status for future-capable AI raises questions about the legal definitions of “personhood,” “agency,” and related concepts. Some non-human organizations and institutions have legal personhood, as is concisely stated in the Dictionary Act of the U.S. Code (1 U.S.C. § 1), which states, “the words ‘person’ and ‘whoever’ include corporations, companies, associations, firms, partnerships, societies, and joint stock companies, as well as individuals.” Similarly, some humans, such as infants and those under conservatorship, may not have the full complement of rights and responsibilities that accompany full legal agency, though they do have some protected rights and interests. Further, some non-humans, such as animals, are referred to as having “rights” and have even been recognized as suffering legal harms sufficient to confer standing under Article III of the U.S. Constitution. *Cetacean Cmty. v. Bush*, 386 F.3d 1169, 1175 (9th Cir. 2004). But quaere whether, within the legal system, “animal rights” instead merely refers to a set of interests that the human legal system obliges humans and human institutions to respect. These questions, beyond the scope of this paper, provide a launching point for discussion of AI interests. To be sure, the ruling by the U.S. Ninth Circuit Court of Appeals in *Naruto v. Slater*, 888 F.3d 418 (9th Cir. 2018) is frequently cited in discussions of AI rights. See *Thaler v. Perlmutter*, 687 F.Supp.3d 140, 148 (D.D.C. 2023) (citing *Naruto*). In *Naruto*, the Ninth Circuit determined that Naruto, a crested macaque, lacked statutory standing to assert a

copyright in the now-famous “monkey selfie.” *Naruto*, 888 F.3d at 425-426.

Historically, computer software has been used merely as a tool by some human user. Resultant rights, liabilities and obligations have generally devolved to the humans and human institutions associated with the tool and its use. In the age of AI, with more autonomous or proto-autonomous behavior by the algorithms, some discussion has turned to whether the traditional allocation of rights, liabilities and obligations to those humans and institutions remains appropriate and sufficient. If not, should there be some attribution to the AI itself? And if that were to happen, what else would have to happen? Should talk of AI responsibility come first, before consideration of AI rights – as was hinted at in the E.U.’s 2017 proposal? Are current legal constructs sufficient to appropriately allocate responsibility when something goes wrong?

The literature regarding modern theories of “personhood” assigned to humans, animals, and organizations, focuses on the assignment of rights and duties within a specific context.³⁷ This paradigm extrapolates from classic personhood appropriate for an adult human with full capacity—someone who has, for example, the right to sign a contract, or the responsibility to pay taxes. An extension of legal personhood beyond adult humans also accommodates, within a consistent framework, variable (generally more limited) collections of rights and duties for, say, children, or non-humans. The legal fiction that a corporation is a person is not an assignment of humanness to an organization, but serves to address the ways in which such an organization may hold rights or be held responsible for carrying out

³⁷ One such analysis places actions along two axes – separateness and independence. Kurki 2019 pp. 182-187; Chopra 2011 pp. 153-191. Separateness describes the degree to which two parties’ resources are commingled (e.g. funds). Independence relates to the “exercise of competences”, such as with guardianship or power of attorney. The two axes are necessary but not sufficient measures of legal personhood—a full legal person would be viewed as both completely separate and independent.



duties.

Similarly, as the capabilities of AI systems continue to evolve, the law could in theory assign various rights and duties to them, depending on legal expediency, without having to go down the (discomfiting to many) path of addressing the degree to which such systems may have achieved various levels of adult-human cognition. As AI grows in capabilities, so too might the law allow such systems to be assigned a collection of rights and duties appropriate for their ability to act according to those rules. For example, AI might be capable of guiding an automobile, or executing certain commercial transactions. As such, it might be legally expedient to allow such systems to act with various types of autonomy so long as there exists a mechanism of handling associated liabilities. An AI-driven car, for example, might be made to carry insurance such that resolving an accident does not require an analysis of causality back to the designer or owner of the AI.³⁸ This assignment of responsibilities is not declaring AI to be a human, but simply facilitating certain types of actions with reduced legal friction, while also setting up systems of consequences and protections that allow humans who have been wronged to seek (human-world) meaningful recourse.

At the current and anticipated rate of development of AI, personhood issues will become increasingly imperative in the not-so-distant future.

³⁸ See, e.g., U.K. Automated and Electric Vehicles Act 2018 s.2 (allocating liability for accidents caused by autonomous vehicles to insurers or vehicle owners under particular circumstances).

Bibliography

- Abbott, R., *The Reasonable Robot* (Cambridge U. Press, 2020).
- Abraham A., *The Neuroscience of Creativity* (Cambridge 2018).
- Acemoğlu, D., Ethics in AI Colloquium “Redesigning AI – For Better Work, Human Agency and Democracy,” Lecture given at The Martin School, Oxford (17 May 2024) <https://www.oxfordmartin.ox.ac.uk/videos/redesigning-ai>.
- Bartholomew, M., “Copyright and the Creative Process,” 97 *Notre Dame L. Rev.* 357 (2021).
- Baumeister, R.F., B.J. Schmeichel, and C.N. Dewall, “Creativity and Consciousness: Evidence from Psychology Experiments” pp. 185 et seq. in Paul, E.S. and S.B. Kaufman, eds., *The Philosophy of Creativity: New Essays* (Oxford University Press 2014).
- Blaszczyk, M., “Impossibility of Emergent Works’ Protection in U.S. and EU Copyright Law,” 25 *North Carolina J. of Law & Tech.* 1 (2023).
- Boden M.A., “Creativity and Biology,” pp. 173 et seq. in Gaut, B. and M. Kieran, eds., *Creativity and Philosophy* (Routledge 2018).
- Boden, M.A., “Creativity and Artificial Intelligence: A Contradiction in Terms?” pp. 224 et seq. in Paul, E.S. and S.B. Kaufman, eds., *The Philosophy of Creativity: New Essays* (Oxford University Press 2014).
- Boden, M.A., *The Creative Mind: Myths and Mechanisms* (Routledge 2d ed 2004).
- Boyden, B., “Emergent Works,” 39 *Colum. J.L. & Arts* 377 (2016).
- Brandt, A.K. “Beethoven’s Ninth and AI’s Tenth: A Comparison of Human and Computational Creativity.” *Journal of Creativity* 33 (2023).
- Buccafusco, C., “There’s No Such Thing as Independent Creation, and It’s a Good Thing, Too” 64 *Wm. & Mary L. Rev.* 1617 (2022).
- Buccafusco, C., Z.C. Burns, J.C. Fromer, & C.J. Sprigman, “Experimental Tests of Intellectual Property Laws’ Creativity Thresholds,” 92 *Texas L. Rev.* 1921 (2014).
- Carroll, N., “The Creative Audience” pp. 62 et seq. in Paul, E.S. and S.B. Kaufman, eds., *The Philosophy of Creativity: New Essays* (Oxford University Press 2014).
- Chaney, A. J. B., B.M. Stewart, & B.E. Engelhardt, “How algorithmic confounding in recommendation systems increases homogeneity and decreases utility.” *Proceedings of the 12th ACM Conference on Recommender Systems*, 224–232 (2018). <https://doi.org/10.1145/3240323.3240370>.
- Chopra, S. & L.F. White, *A Legal Theory for Autonomous Artificial Agents* (Michigan 2011).
- Craig, C. J., “The AI-Copyright Challenge: Tech-Neutrality, Authorship, and the Public Interest” in Abbott, R. (ed.) *Research Handbook on Intellectual Property and Artificial Intelligence* (Edward Elgar Press 2022), <https://ssrn.com/abstract=4014811> or <http://dx.doi.org/10.2139/ssrn.4014811>.

- Doshi, A.R., and O.P. Hauser. "Generative Artificial Intelligence Enhances Creativity but Reduces the Diversity of Novel Content." <https://arxiv.org/pdf/2312.00506v3> (2023).
- Du Sautoy, M., *The Creativity Code: Art and Innovation in the Age of AI* (Harvard 2019).
- Epstein, Z., A. Hertzmann, M. Akten, H. Farid, J. Fjeld, M.R. Frank, M. Groh, L. Herman, N. Leach, R. Mahari, A. "Sandy" Pentland, O. Russakovsky, H. Schroeder, & A. Smith, "Art and the science of generative AI." *Science*, 380(6650), 1110–1111 (2023). <https://doi.org/10.1126/science.adh4451>.
- Esling, P., and N. Devis. "Creativity in the Era of Artificial Intelligence." (2020).
- Fromer, J.C., "A Psychology of Intellectual Property," 104 *Northwestern U. L. Rev.* 1441 (2010).
- Gao, Y., Y. Xiong, X. Gao, K. Jia, J. Pan, Y. Bi, Y. Dai, J. Sun, M. Wang, and H. Wang, "Retrieval-Augmented Generation for Large Language Models: A Survey" (arXiv, 27 March 2024) <https://doi.org/10.48550/arXiv.2312.10997>.
- Gardner, H., and K. Davies. *The App Generation: How Today's Youth Navigate Identity, Intimacy, and Imagination in a Digital World.* (Yale University Press, 2013).
- Gillespie, T., P.J. Boczkowski, & K.A. Foot, *Media technologies: Essays on communication, materiality, and society* (MIT Press 2014).
- Godden, D.R., and A.D. Baddeley. "Context-Dependent Memory in Two Natural Environments: On Land and Underwater." *British Journal of Psychology* 66, no. 3 (1975): 325-33.
- Goldberg, E., *Creativity: The Human Brain in the Age of Innovation* (Oxford University Press 2018).
- Guzik, E.E., C. Byrge, and C. Gilde, "The originality of machines: AI takes the Torrance Test," *Journal of Creativity* 33 (2023) 100065 (<https://www.sciencedirect.com/science/article/pii/S2713374523000249?via%3Dihub>).
- Habib, S., T. Vogel, X. Anli, and E. Thorne. "How Does Generative Artificial Intelligence Impact Student Creativity?" *Journal of Creativity* 34 (2024).
- Ho, J., A. Jain, & P. Abbeel, "Denoising Diffusion Probabilistic Models," 34th *Conference on Neural Information Processing Systems (NeurIPS 2020)* (2020), <https://arxiv.org/pdf/2006.11239>.
- Hubert, K.F., K.N. Awa, & D.L. Zabelina, "The current state of artificial intelligence generative language models is more creative than humans on divergent thinking tasks." *Sci Rep* 14, 3440 (2024). <https://doi.org/10.1038/s41598-024-53303-w>.
- Jackson, T., R. Dawson, and D. Wilson. "Reducing the Effects of Email Interruptions on Employees." *International Journal of Information Management* 23, no. 1 (2003): 55-65.
- Kaufman, J.C., and R.A. Beghetto. "Beyond Big and Little: The Four C Model of Creativity." *Review of General Psychology* 13, no. 1 (2009): 1-12.

- Kieran, M., "Creativity as a Virtue of Character" pp. 125 et seq. in Paul, E.S. and S.B. Kaufman, eds., *The Philosophy of Creativity: New Essays* (Oxford University Press 2014).
- Kurki, V., *A Theory of Legal Personhood* (Oxford University Press 2019).
- Lee, J.-A., "Computer-generated Works under the CDPA 1988," in Lee, J.-A, R. Hilty, & K-C Liu, eds., *Artificial Intelligence and Intellectual Property* (Oxford University Press 2021) <https://ssrn.com/abstract=3956911>.
- Maddux, W.W., H. Adam, and A.D. Galinsky. "When in Rome ... Learn Why the Romans Do What They Do: How Multicultural Learning Experiences Facilitate Creativity." *Personality & Social Psychology Bulletin* 36, no. 6 (2010): 731-41.
- Mahowald, K., A.A. Ivanova, I.A. Blank, N. Kanwisher, J.B. Tenenbaum, and E. Fedorenko, "Dissociating Language and Thought in Large Language Models" (arXiv, 23 March 2024) <https://doi.org/10.48550/arXiv.2301.06627>.
- Merges, R., *Justifying Intellectual Property* (Harvard U. Press 2011).
- Miller, J.S. , "Hoisting Originality," *Lewis & Clark Law School Legal Research Paper Series* No. 2009-4 (2009) <http://ssrn.com/abstract=1361040>.
- Millière, R., and C. Buckner, "A Philosophical Introduction to Language Models – Part I: Continuity With Classic Debates" (arXiv, 8 January 2024) <http://arxiv.org/abs/2401.03910>.
- Millière, R., and C. Buckner, "A Philosophical Introduction to Language Models - Part II: The Way Forward" (arXiv, 6 May 2024) <http://arxiv.org/abs/2405.03207>.
- Mumford, L. *Technics and Civilization*. New York: Harcourt, Brace, and Company, 1934.
- Nakajima, Y.. "Task-Driven Autonomous Agent Utilizing GPT-4, Pinecone, and LangChain for Diverse Applications" (28 March 2023) <https://yoheinakajima.com/task-driven-autonomous-agent-utilizing-gpt-4-pinecone-and-langchain-for-diverse-applications/>.
- Nanay, B., "An Experiential Account of Creativity," pp.17 et seq. in Paul, E.S. and S.B. Kaufman, eds., *The Philosophy of Creativity: New Essays* (Oxford University Press 2014).
- Nanda, N., A. Lee, and M. Wattenberg, "Emergent Linear Representations in World Models of Self-Supervised Sequence Models" (arXiv, 7 September 2023) <http://arxiv.org/abs/2309.00941>.
- Noble, S. U., *Algorithms of oppression: How search engines reinforce racism* (New York University Press 2018).
- O'Callaghan, C., "Can Output Produced Autonomously by AI Systems Enjoy Copyright Protection, and Should It? An Analysis of the Current Legal Position and the Search for the Way Forward," 55 *Cornell Int'l. L.J.* 305 (Fall 2022).
- OpenAI, J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F.L. Aleman, et al., "GPT-4 Technical Report" (arXiv, 4 March 2024) <https://doi.org/10.48550/arXiv.2303.08774>.

- Pascual-Leone, A. "The Brain That Plays Music and Is Changed by It." *Annals of the New York Academy of Sciences* 930 (2001): 315-29.
- Parchomovsky, G. & A. Stein, "Originality," 95 *Virginia L. Rev.* 1505 (2009).
- Paul, E.S. and S.B. Kaufman, "Introducing *The Philosophy of Creativity*" pp. 3 et seq. in Paul, E.S. and S.B. Kaufman, eds., *The Philosophy of Creativity: New Essays* (Oxford University Press 2014).
- Peng, M., X. Chen, and Q. Zhao. "Attentional Scope Is Reduced by Internet Use: A Behavioral and ERP Study." *PLoS ONE* 13, no. 6 (2018): e0198543.
- Picciuto, E. and P. Carruthers, "The Origins of Creativity" pp. 199 et seq. in Paul, E.S. and S.B. Kaufman, eds., *The Philosophy of Creativity: New Essays* (Oxford University Press 2014).
- Rombach, R., A. Blattmann, D. Lorenz, P. Esser, & B. Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models," (2021), <https://arxiv.org/pdf/2112.10752>.
- Schulman, S., *The Gentrification of the Mind: Witness to a Lost Imagination* (California University Press 2013).
- Scolere, L., U. Pruchniewska, & B.E. Duffy, "Constructing the Platform-Specific Self-Brand: The Labor of Social Media Promotion." *Social Media + Society*, 4(3), 205630511878476 (2018). <https://doi.org/10.1177/2056305118784768>.
- Shumailov, Ilia et al. (2024), "The Curse of Recursion: Training on Generated Data Makes Models Forget." arXiv. (<https://arxiv.org/pdf/2305.17493>)
- Small, Z., "Black Artists Say A.I. Shows Bias, With Algorithms Erasing Their History." *New York Times* (July 4, 2023). <https://www.nytimes.com/2023/07/04/arts/design/black-artists-bias-ai.html>.
- Søgaard, A., "Grounding the Vector Space of an Octopus: Word Meaning from Raw Text" *Minds and Machines* 33, no. 1 (1 March 2023): 33–54. <https://doi.org/10.1007/s11023-023-09622-4>.
- Stanford University Institute for Human-Centered AI, *Artificial Intelligence Index Report 2024* (2024) https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_2024_AI-Index-Report.pdf.
- Stech, M.T., "The Semantics of Authorial Originality: Four Pillars," 29 *Texas Int'l Prop. L.J.* 235 (2021).
- Subotnik E.E., "Originality Proxies: Toward a Theory of Copyright and Creativity," 76 *Brooklyn L. Rev.* 1487 (2011).
- Ursu, R. M., "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions." *Marketing Science*, 37(4), 530–552 (2018). <https://doi.org/10.1287/mksc.2017.1072>.
- Vinchon, F., T. I. Lubart, S. Bartlotta, V. Gironnay, M Botella, S. Bourgeois-Bougrine, and et al. "Artificial Intelligence & Creativity: A Manifesto for Collaboration." *The Journal of Creative Behavior* 57, no. 4 (2023): 474-84.



- Wang, L., C. Ma, X. Feng, Z. Zhang, H. Yang, J. Zhang, Z. Chen, et al., "A Survey on Large Language Model Based Autonomous Agents," *Frontiers of Computer Science* 18, no. 6 (March 2024): 186345. <https://doi.org/10.1007/s11704-024-40231-1>.
- Ward, A.F., K Duke, A Gneezy, and M.W. Bos. "Brain Drain: The Mere Presence of One's Own Smartphone Reduces Available Cognitive Capacity." *Journal of the Association for Consumer Research* 2, no. 2 (2017).
- Xu, D., S. Fan, & M. Kankanhalli, "Combating Misinformation in the Era of Generative AI Models." *Proceedings of the 31st ACM International Conference on Multimedia*, 9291–9298 (2023). <https://doi.org/10.1145/3581783.3612704>.
- Yildirim, I., and L.A. Paul. "From Task Structures to World Models: What Do LLMs Know?" *Trends in Cognitive Sciences* 28, no. 5 (May 2024): 404–15. <https://doi.org/10.1016/j.tics.2024.02.008>.
- Yu, R., "The Machine Author: What Level of Copyright Protection Is Appropriate for Fully Independent Computer-Generated Works?" 165 *University of Pennsylvania Law Review* 1245 (2017).
- Ziegler, D.M., N. Stiennon, J. Wu, T.B. Brown, A. Radford, D. Amodei, P. Christiano, and G. Irving, "Fine-Tuning Language Models from Human Preferences" (arXiv, 8 January 2020) <https://doi.org/10.48550/arXiv.1909.08593>.

Contributors

The authors of this white paper participated in a working group on AI and creativity convened by Christian Mammen (Academic Visitor, University of Oxford Faculty of Law) during Trinity Term (May-June) 2024. The working group met at University College, Oxford multiple times over the course of the term, and approved this white paper for release in July 2024. Dr. Mammen is the primary author; all other co-authors (listed in alphabetical order) substantially participated in the working group, contributed to the drafting of the white paper, and have approved the final draft. In addition to the listed co-authors, Bryce Goodman and Linda Eggert participated in one or more meetings of the working group.

Dr. Christian E. Mammen was an Academic Visitor with the University of Oxford Faculty of Law during Trinity Term 2024, while on sabbatical from the law firm Womble Bond Dickinson (US) LLP. At Womble Bond Dickinson, Mammen practices intellectual property litigation and is Managing Partner of the firm's San Francisco Office. He is an established thought leader in issues of Artificial Intelligence and Intellectual Property law.

Contact: cemammen@gmail.com

LinkedIn: <https://www.linkedin.com/in/christian-e-mammen/>

Michael Collyer is a Doctoral Candidate at the University of Oxford. He is a co-founder of the Oxford AI Network and Founders & Funders Foundation. He holds a Bachelors in Politics, Psychology, Law and Economics from the University of Amsterdam and a Masters in Social Science of the Internet from the University of Oxford.

Contact: michael.collyer@oii.ox.ac.uk

LinkedIn: <https://www.linkedin.com/in/mcollyer/>

Dr. Ron Dolin has degrees in math, physics, law, and computer science, and is a licensed attorney in California. He has taught about legal informatics at Stanford, Notre Dame, and Harvard law schools, and is co-editor and co-author of the textbook Legal Informatics (Cambridge University Press, 2021).

Contact: ron.a.dolin@radicalconcepts.com

LinkedIn: <https://www.linkedin.com/in/ron-dolin-3908246/>



Dr. Dev Gangjee is Professor of Intellectual Property Law at the Oxford Law Faculty. His research presently focuses on the adoption of machine learning technologies by Intellectual Property Offices around the world. He is also working on the legal regulation of data, which fuels machine learning.

Contact: dev.gangjee@law.ox.ac.uk

LinkedIn: <https://www.linkedin.com/in/dev-gangjee-641b298b/>

Dr. Tom Melham FRSE FBCS CEng is a Professor of Computer Science at the University of Oxford and a Fellow of Balliol College. Melham's research is focused on mathematical methods for assuring quality and correctness of digital hardware and software systems. Alongside this, his research interests include AI and technology in legal services and the justice system and testing and evaluation of AI-based systems. He also works in computer-aided dispute resolution and was the lead technical co-author of a landmark LawtechUK study on online dispute resolution for SMEs. Melham serves on the Board of Trustees of the Alan Turing Institute, the UK's national institute for AI and data science. He is a chartered engineer and a Fellow of the British Computer Society and of the Royal Society of Edinburgh.

Contact: thomas.melham@balliol.ox.ac.uk

LinkedIn: <https://www.linkedin.com/in/melham/>

Maggie Mustaklem is a design lead and doctoral researcher focusing on the implications of AI in design. Her research project, *Design Interrupted*, centres on the "everyday AI" in platforms like Pinterest, Instagram (and increasingly generative tools) that designers and architects use to search for inspiration. She is interested in how these tools may be flattening what designers see for inspiration, influencing what they ultimately produce. Mustaklem holds a Master of Arts in History of Design from the Royal College of Art and Victoria & Albert Museum and a Bachelor of Arts in Psychology from the University of Michigan.

Contact: maggie.mustaklem@oii.ox.ac.uk

LinkedIn: <https://www.linkedin.com/in/mustaklemmaggie/>



UNIVERSITY OF
OXFORD

Pireeni Sundaralingam is a cognitive scientist. Trained at Oxford and MIT. she has held posts as Principal Advisor for Human Potential for the United Nations initiative UNLive and led research at Silicon Valley's Center for Humane Technology. Her consulting firm Neuro-Resilience Consulting specializes in optimizing conditions for human innovation and decision-making. Sundaralingam is also an award-winning poet, whose work has been translated into 5 languages, and is currently College Poet Laureate and InterDisciplinary Catalyst at University College, University of Oxford.

Contact: pireeni.sundaralingam@univ.ox.ac.uk

LinkedIn: <https://www.linkedin.com/in/pireeni-sundaralingam-3734a8187/>

Dr. Vincent Wang-Maścianica is a research scientist affiliated with the Compositional Intelligence branch of Quantinuum and the Human Centred AI Lab at St. Catherine's College, both in Oxford, where he was trained in Mathematics and Computer Science. Wang-Maścianica develops and works with high-level yet intuitively accessible visual reasoning languages for complex systems, including models of human cognition, natural language, and computation both quantum and classical.

Contact: vincentwangsemailaddress@gmail.com

LinkedIn: <https://www.linkedin.com/in/vincent-wang-ma%C5%9Bcianica-052a0b188/>



UNIVERSITY OF
OXFORD