

MINDS, MACHINES, AND THE LAW: THE CASE OF VOLITION IN COPYRIGHT LAW

Mala Chatterjee & Jeanne C. Fromer***

The increasing prevalence of ever-sophisticated technology permits machines to stand in for or augment humans in a growing number of contexts. The questions of whether, when, and how the so-called actions of machines can and should result in legal liability thus will also become more practically pressing. One important set of questions that the law will inevitably need to confront is whether machines can have mental states, or—at least—something sufficiently like mental states for the purposes of the law. This is because a number of areas of law have explicit or implicit mental state requirements for the incurrance of legal liability. Thus, in these contexts, whether machines can incur legal liability turns on whether a machine can operate with the requisite mental state. Consider the example of copyright law. Given the long history of mechanical copying, courts have already faced the question of whether a machine making a copy can have the mental states required for liability. They have often answered with a resounding, unconditional “no.” But this Essay seeks to challenge any generalization that machines cannot operate with a mental state in the eyes of the law. Taking lessons from philosophical thinking about minds and machines—in particular, the conceptual distinction between “conscious” and “functional” properties of the mind—this Essay uses copyright’s volitional act requirement as a case study to demonstrate that certain legal mental state requirements might seek to track only the functional properties of the states in question, even ones which can be possessed by machines. This Essay concludes by considering how to move toward a more general framework for evaluating the question of machine mental states for legal purposes.

* Law Clerk for the Honorable Judge Robert D. Sack, United States Court of Appeals for the Second Circuit; Ph.D. Candidate in Philosophy, New York University; Fellow, Engelberg Center for Innovation Law and Policy; Visiting Fellow, Information Society Project, Yale Law School; J.D., New York University School of Law 2018.

** Professor of Law, New York University School of Law. We thank Arnaud Ajdler, Jack Balkin, Barton Beebe, Colin Bradley, Christopher Buccafusco, Dan Burk, David Chalmers, Colleen Chien, Rebecca Crootof, Brian Frye, Eric Goldman, Michael Grynberg, Nathan Gusdorf, Scott Hemphill, Ben Holguín, Bert Huang, Arden Koehler, Andrew Lee, S. Matthew Liao, Jeffrey Lipshaw, Michael Meurer, Liam Murphy, Betsy Rosenblatt, Olga Russakovsky, Matthew Sag, Erick Sam, Jason Schultz, Jule Sigall, Christopher Sprigman, Jeff Stein, Katherine Strandburg, Olivier Sylvain, Jacob Victor, Patrick Winston, Felix Wu, Gideon Yaffe, and participants at workshops at New York University School of Law, Yale Law School, the 2019 Works in Progress in Intellectual Property Colloquium, and the *Columbia Law Review* Symposium on “Common Law for the Age of AI,” for choosing to share knowledgeable comments. Thanks to Moses Dyckman for helpful research assistance. Jeanne Fromer gratefully acknowledges support from the Filomen D’Agostino and Max E. Greenberg Research Fund.

INTRODUCTION

With the increasing prevalence of ever more sophisticated technology—which permits machines to stand in for or augment humans in a growing number of contexts—the questions of whether, when, and how the so-called actions of machines can and should result in legal liability will become more practically pressing.¹ Although the law has yet to fully grapple with questions such as whether machines are (or can be) sufficiently humanlike to be the subjects of law, philosophers have long contemplated the nature of machines.² Philosophers have considered, for instance, whether human cognition is fundamentally computation—such that it is in principle possible for future artificial intelligences (AI) to possess the properties of human minds, including consciousness, semantic understanding, intention, and even moral responsibility—or if humans and machines are instead fundamentally different, no matter how sophisticated AI becomes.³ It is thus unsurprising that, in thinking through how the law should accommodate and govern an increasingly AI-filled world, the lessons and frameworks to be gleaned from these philosophical discussions will have undeniable relevance.

One important set of questions that the law will inevitably need to confront is whether machines can have mental states, or—at least—something sufficiently like mental states for the purposes of the law. This is because a wide range of areas of law have explicit or implicit mental state requirements for the incurrence of legal liability.⁴ Consider, for example, questions of intent and recklessness versus negligence in tort law; *mens rea* and *actus reus* in criminal law; offer and acceptance in contract law; and, as we will see, infringement and authorship in copyright law. In each of these contexts, the law either implicitly or explicitly asks for the presence of some particular mental state on the part of the actors in question. Whether the operations of machines can incur legal liability—and what kind of liability they can incur—would thus often seem to turn on whether a machine is regarded as operating with the mental state required.

In some contexts, the decision already seems to have been made that machines can never possess the mental states required for liability. Consider copyright law's volitional act requirement for infringement. Copyright law has generally claimed that machines making copies of protected material lack the requisite volition for this conduct to give rise to legal liability on the part of those responsible for the machine, even when the machine has been designed to make copies, often of copyrighted works.⁵ In other contexts, such as criminal and tort law, the

-
1. See *infra* section I.B.
 2. See *infra* Part III.
 3. See *infra* Part III.
 4. See *infra* section I.A.
 5. See *infra* Part II.

question of machines' capacity for mental states remains open and underexplored.⁶

This Essay aims to challenge any hasty and blanket generalization that machines cannot have mental states as a legal matter, drawing on philosophical thinking surrounding mental states and using copyright's volitional act requirement as a case study. In so doing, this Essay concludes that—as a matter of copyright doctrine—a copying technology might be sufficiently “volitional” for the technology provider to be held directly liable for the technology's so-called actions in producing copies; and—as a matter of general legal theory—machines in some contexts might be capable of being sufficiently “mental” to count as agents of the humans behind them, depending on the aims of the area of law in question.⁷ This conclusion is thus not merely of philosophical interest but one with practical implications for determinations of legal liability. In the context of copyright law, this Essay's chosen case study, this conclusion has implications for who is and is not directly accountable for the copying of protected material and for the law's ability to effectuate its goals of encouraging the creation and dissemination of expressive works.

To mount this Essay's challenge, after giving an overview of mental states in the law and the puzzle raised by technological advancement in Part I, as well as the specific challenges posed by copyright law in Part II, Part III of the Essay recounts two of the most influential philosophical discussions on minds and machines, and the resulting theoretical distinction between the conscious and functional properties of mental states. Using this distinction as a framework, this Essay argues that it is an open question whether the law's mental state requirements seek to track the conscious or merely functional properties of the particular mental state in question,⁸ and the analysis depends on the ultimate aims of the relevant area of law. Part IV then defends the view that copyright law's volitional act requirement might be interested in merely functional

6. See generally Mark A. Geistfeld, *A Roadmap for Autonomous Vehicles: State Tort Liability, Automobile Insurance, and Federal Safety Regulation*, 105 *Calif. L. Rev.* 1611 (2017) (tort liability); Gabriel Hallevy, “I, Robot—I, Criminal”—When Science Fiction Becomes Reality: Legal Liability of AI Robots Committing Criminal Offenses, 22 *Syracuse Sci. & Tech. L. Rep.* 1 (2010) (criminal law); Ignatius Michael Ingles, Note, *Regulating Religious Robots: Free Exercise and RFRA in the Time of Superintelligent Artificial Intelligence*, 105 *Geo. L.J.* 507, 516 n.67 (2017) (criminal law).

7. See *infra* Parts IV–V.

8. Note that there are arguably nonconscious mental states aside from functional mental states, such as intentional and computational states. In this way, the distinction on which we focus—consciousness versus functionality—is not exhaustive, as one could similarly ask whether the law cares about intentionality, computation, and so forth. Nonetheless, the Essay focuses on consciousness versus functionality not only for the sake of simplicity, but also because this distinction is plausibly the most important one for legal purposes. The Essay otherwise leaves the question of whether intentionality (or other nonconscious, nonfunctional properties of mental states) should ever matter to the law for exploration in future work.

properties, which could—in principle—be replicated by machines. Next, Part V considers which functional properties copyright law might seek to track and what a machine might have to look like to be “functionally volitional” under copyright law, to count as the technology provider’s agent, and thereby to give rise to direct liability. These relevant functional properties include the ability to pause and analyze the nature of the work in question before “choosing” to undertake an act of copying, one which might cause exposure to liability. On the basis of this framework, this Essay concludes that machines with the appropriate functionality might satisfy copyright law’s volitional act requirement, thus forming the basis for holding technology providers directly liable for infringement. Finally, generalizing this Essay’s framework, Part VI offers preliminary thoughts on machines and mental state requirements in the contrasting contexts of criminal law and copyright authorship doctrine, as well as a general hypothesis regarding when the law is interested in conscious versus merely functional properties of the mental states in question.

I. MENTAL STATES, TECHNOLOGY, AND THE LAW

This Part explores the intersection of mental states and technology under the law. It first provides an overview of the law’s mental state requirements, and then surveys how businesses might use machines in lieu of humans to perform various operations that could—or would—incur liability if performed by a human, such that technological advancement inevitably raises the legal question of machine mental states.

A. *Mental State Requirements in the Law*

Mental state requirements for legal liability are pervasive. The most familiar include requirements of purpose (or intent), knowledge, and recklessness, in contrast to negligence (which is not itself a mental state but might be understood as distinguishable from, say, recklessness by the absence of such a state).⁹ Each relates in differing ways to beliefs or desires.¹⁰ Volition—which might be defined as the cause of willful actions, and which thus distinguishes actions from involuntary bodily

9. See, e.g., Model Penal Code § 2.02 (Am. Law Inst. 1985); see also Kyron Huigens, *On Commonplace Punishment Theory*, 2005 U. Chi. Legal F. 437, 453 (“In negligence and the other non-intentional fault doctrines, fault is found not in a discrete mental state, but in a broader set of facts surrounding the offense.”).

10. See Kenneth W. Simons, *Rethinking Mental States*, 72 B.U. L. Rev. 463, 464–65 (1992) (“Properly understood, the principal mental state concepts do not reflect a single hierarchy of legal significance. Rather, they conceal two distinct mental state hierarchies, of desire and belief, as well as a third hierarchy, of conduct, which does not essentially involve mental states.”).

movements¹¹—can be understood as a mental state as well.¹² Thus, in addition to any further mens rea requirements, any area of law requiring a willful action for liability is implicitly asking for a mental state as well, because the presence of volition is what makes a movement count as a willful action (rather than, say, a muscle spasm) in the first place.¹³ Mental state requirements thus exist in nearly every area of law, including criminal law, torts, and contract.¹⁴ Indeed, these requirements are so prevalent that there is even a legal category arguably defined in terms of an absence of any mens rea beyond volition itself: namely, strict liability.¹⁵ These requirements are premised on the assumption that the mind—and not just the body—matters to the law.¹⁶ In other words, when such requirements exist, the body might move to do something prohibited, but only when this is conjoined with the corresponding illicit mental state is this a prohibited action.

As an evidentiary matter, discerning the presence of a mental state in a human requires “mind reading,” so to speak, because people cannot directly observe or measure a mental state.¹⁷ Nonetheless, the law typically feels comfortable—though perhaps it should not¹⁸—answering the question of whether a human had the required mental state. In light of these requirements, as machines become more pervasive in performing operations that humans traditionally performed, the law will find itself needing to assess not just the permissibility of machines’ operations but also whether they have operated with an illicit mental state.

11. See Michael S. Moore, *Act and Crime: The Philosophy of Action and Its Implications for Criminal Law* 113–65 (1993) [hereinafter Moore, *Act and Crime*] (defending a theory of volition as the mental state that causes actions).

12. See, e.g., *id.* at 115 (“Volition’ names a state or an event within the mind of the actor.”).

13. See *id.* at 113–65.

14. See, e.g., Kent Greenawalt, *A Pluralist Approach to Interpretation: Wills and Contracts*, 42 *San Diego L. Rev.* 533, 575–82 (2005) (contracts); Simons, *supra* note 10, at 468–73 (criminal law and torts).

15. See Simons, *supra* note 10, at 464.

16. See, e.g., Keren Shapira-Ettinger, *The Conundrum of Mental States: Substantive Rules and Evidence Combined*, 28 *Cardozo L. Rev.* 2577, 2579–81 (2007) (“[C]riminal law has adopted the vague metaphysical dualistic vision between a forbidden act and a state of mind that accompanied it.”). Some have criticized this assumption, suggesting it ought to be replaced with an integrated actus reus and mens rea. See, e.g., Douglas N. Husak, *Philosophy of Criminal Law* 126 (1987) (advocating for this integration “as an indivisible product of both what one thinks and what one does”).

17. See Teneille Brown & Emily Murphy, *Through a Scanner Darkly: Functional Neuroimaging as Evidence of a Criminal Defendant’s Past Mental States*, 62 *Stan. L. Rev.* 1119, 1129–30 (2010) (“Because we cannot presently read someone’s mind to determine her mens rea at the time of the crime, the jury is often told it can rely on the objective circumstances surrounding the criminal’s conduct to draw inferences about her state of mind.”).

18. See, e.g., James A. Macleod, *Belief States in Criminal Law*, 68 *Okla. L. Rev.* 497, 502–03, 514–34 (2016) (drawing on experimental epistemology to criticize how juries likely decide on the presence of a mental state).

B. *The Present and Future of Technology*

Increasingly, tasks once performed by only humans are being carried out or augmented by machines, which often perform better than humans ever could. In the copyright space alone—on which the Essay elaborates in the next Part—there are devices that can now recognize songs and other expressive content by listening to them,¹⁹ virtual assistants and bots that can locate and play user-requested content,²⁰ and software that can use machine learning techniques to create artwork based on a model derived from 15,000 portraits painted over the past six centuries.²¹ A piece of art created using this software recently sold at auction for over \$400,000.²²

Thus, questions of so-called machine liability are becoming more pressing. Legal scholars have already been puzzling over a tort liability regime for self-driving cars.²³ Plausibly, we might soon find ourselves asking whether a bot producing defamatory content about a public figure can itself have actual malice; whether an algorithm assessing risk can have discriminatory intent; or whether the price-setting systems of competing businesses can collude from the perspective of antitrust law. And in the copyright space, we might wonder whether technology creators or owners can be directly liable for copyright infringement when a bot fetches an infringing copy of a song in response to a user's request for that song or when software taught on portraits produces an artwork that is copied from and substantially similar to an existing portrait on which the software was trained.

19. E.g., Trent Gillies, *Shazam Names That Tune, Drawing in Money and Users*, CNBC (June 14, 2015), <https://www.cnbc.com/2015/06/14/shazam-names-that-tune-drawing-in-money-and-users.html> [<https://perma.cc/6DEK-L2KV>].

20. E.g., Taylor Martin, *9 Alexa Tips for Music Lovers*, CNET (Jan. 22, 2019), <https://www.cnet.com/how-to/alexa-tips-for-music-lovers> [<https://perma.cc/4ATW-C6HX>].

21. *Is Artificial Intelligence Set to Become Art's Next Medium?*, Christie's (Dec. 12, 2018), <https://www.christies.com/features/A-collaboration-between-two-artists-one-human-one-a-machine-9332-1.aspx> [<https://perma.cc/R6NB-RU5F>].

22. *Id.*

23. See, e.g., Geistfeld, *supra* note 6, at 1691–94 (arguing that a combination of state products liability law and federal regulations can provide an effective framework for self-driving cars); Gary E. Marchant & Rachel A. Lindor, *The Coming Collision Between Autonomous Vehicles and the Liability System*, 52 Santa Clara L. Rev. 1321, 1335–39 (2012) (suggesting “legal and policy tools that may help protect manufacturers [of autonomous vehicles] from liability,” including the assumption of risk defense, legislative limitations on liability, and federal preemption of state tort actions); Bryant Walker Smith, *Automated Driving and Product Liability*, 2017 Mich. St. L. Rev. 1, 2 (“[T]he current product liability regime, while imperfect, is probably compatible with the adoption of automated driving systems.”); Harry Surden & Mary-Anne Williams, *Technological Opacity, Predictability, and Self-Driving Cars*, 38 Cardozo L. Rev. 121, 178–80 (2016) (describing the potential for tort liability to encourage autonomous car manufacturers to program more predictable movements, as well as the ability for autonomous cars to transform the issue of fault in car accidents by providing a “‘black box’ record”).

II. THE COPYRIGHT EXAMPLE: A LONG HISTORY OF MECHANICAL COPYING

This Part uses copyright infringement as this Essay's case study for the challenge posed for the law by mental states and machines. In particular, this Part recounts copyright law's extensive history of mechanical copying, which has long provoked courts to explore whether and when machines and their owners can be directly liable for infringement. This history has led courts to develop a volitional act requirement for copyright infringement, while suggesting that this requirement—though *always* satisfied by human actions—can never be satisfied by machines. This Part also explains why the volitional act requirement ought to be understood as a mental state. For these reasons, the requirement provides a good test bed to explore whether machines should ever possess mental states as a legal matter.

A. *Background*

By way of background, American copyright law protects “original works of authorship fixed in any tangible medium of expression,” including literary works, sound recordings, and movies.²⁴ A copyright holder receives, among other things, the exclusive right to reproduce the work, distribute copies of it, and prepare derivative works,²⁵ typically until seventy years after the author's death.²⁶ Copyright protection extends to the expression of particular ideas rather than to the ideas themselves.²⁷ Yet protection actually reaches well beyond the literal work to works that are copied and substantially similar,²⁸ “else a plagiarist would escape by immaterial variations.”²⁹

The most widely embraced theory of copyright law in America is utilitarian and, in particular, economic.³⁰ According to this theory,

24. 17 U.S.C. § 102(a) (2012).

25. *Id.* § 106.

26. *Id.* § 302(a).

27. See *id.* § 102(b); *Nichols v. Universal Pictures Corp.*, 45 F.2d 119, 121 (2d Cir. 1930).

28. *Corwin v. Walt Disney Co.*, 475 F.3d 1239, 1253 (11th Cir. 2007) (citing *Herzog v. Castle Rock Entm't*, 193 F.3d 1241, 1249 (11th Cir. 1999)).

29. *Nichols*, 45 F.2d at 121.

30. See, e.g., *Harper & Row, Publishers, Inc. v. Nation Enters.*, 471 U.S. 539, 558 (1985) (embracing an economic theory of copyright, and stating that “[b]y establishing a marketable right to the use of one's expression, copyright supplies the economic incentive to create and disseminate ideas”); Shyamkrishna Balganesh, *Foreseeability and Copyright Incentives*, 122 *Harv. L. Rev.* 1569, 1576–77 (2009) (“[C]opyright law in the United States has undeniably come to be understood almost entirely in utilitarian, incentive-driven terms.”); Jeanne C. Fromer, *Expressive Incentives in Intellectual Property*, 98 *Va. L. Rev.* 1745, 1750–52 (2012) (“The Supreme Court, Congress, and many legal scholars consider utilitarianism the dominant purpose of American copyright and patent law.”); William M. Landes & Richard A. Posner, *An Economic Analysis of Copyright Law*, 18 *J. Legal Stud.* 325, 326 (1989) (proposing an “economic model of copyright protection”).

copyright law provides the incentive of exclusive rights for a limited duration to authors to motivate them to create and distribute culturally valuable works.³¹ Without this incentive, the theory goes, authors might not invest the time, energy, and money necessary to create and distribute these works because they might be copied cheaply and easily by free riders, eliminating authors' ability to profit from their works.³² By allowing a copyright holder to recover damages from and enjoin an infringer that breaches the copyright holder's exclusive rights—thereby undermining copyright's pecuniary incentive—the law preserves the copyright incentive.³³

A utilitarian theory of copyright law rests on the premise that the benefit to society of creators crafting valuable works offsets the costs to society of the incentives the law offers to creators.³⁴ To prevent excessive rights that would undercut the goals of dissemination of works and of creation that builds on preexisting works, copyright law therefore limits copyright's duration and scope in certain ways.³⁵ For example, copyright law excuses some third-party uses that would otherwise be infringing by deeming them to be "fair use."³⁶ The fair use doctrine enables third parties to create culturally valuable works that must borrow from the original work in some capacity in order to succeed, often transforming it.³⁷

Moreover, copyright infringement is understood to be a strict liability offense. At the extreme, a person can infringe another's copyright even if they copy from the third party's work without any awareness of the

31. Stewart E. Sterk, *Rhetoric and Reality in Copyright Law*, 94 Mich. L. Rev. 1197, 1197 (1996).

32. See *id.*

33. See Roger D. Blair & Thomas F. Cotter, *An Economic Analysis of Damages Rules in Intellectual Property Law*, 39 Wm. & Mary L. Rev. 1585, 1617–46 (1998) (“[A] simple model of intellectual property rights suggests that the prevailing plaintiff in a . . . copyright . . . infringement action should be able to recover the greater of her lost profit attributable to the infringement, or the defendant’s profit so attributable . . .”); Jeanne C. Fromer & Mark A. Lemley, *The Audience in Intellectual Property Infringement*, 112 Mich. L. Rev. 1251, 1299–1300 (2014) (discussing the “multiple vantage points” used when assessing a copyright infringement as a way to structure when there is infringement liability and thus preserve copyright’s incentive).

34. See Mark A. Lemley, *The Economics of Improvement in Intellectual Property Law*, 75 Tex. L. Rev. 989, 996–97 (1997).

35. See *id.* at 996–98.

36. 17 U.S.C. § 107 (2012).

37. See *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 577 (1994) (“The fair use doctrine thus ‘permits [and requires] courts to avoid rigid application of the copyright statute when, on occasion, it would stifle the very creativity which that law is designed to foster.’” (alteration in original) (quoting *Stewart v. Abend*, 495 U.S. 207, 236 (1990))); Pierre N. Leval, *Toward a Fair Use Standard*, 103 Harv. L. Rev. 1105, 1111–16 (1990) (“Quotation can be vital to the fulfillment of the public-enriching goals of copyright law. The first fair use factor calls for a careful evaluation whether the particular quotation is of the transformative type that advances knowledge and the progress of the arts . . .”).

fact that they have copied.³⁸ For example, singer Michael Bolton was found liable for infringement for subconsciously copying the Isley Brothers' song "Love Is a Wonderful Song" decades later in his song of the same name.³⁹ As Judge Learned Hand explained,

Everything registers somewhere in our memories, and no one can tell what may evoke it. . . .

. . . Once it appears that another has in fact used the copyright as the source of his production, he has invaded the author's rights. It is no excuse that in so doing his memory has played him a trick.⁴⁰

B. *The Player Piano Roll*

In light of a consistent stream of advancements in copying technologies, copyright law has already had to grapple with whether and when copies made by machines constitute copyright infringement.⁴¹ One of the most striking illustrations of this dates back to the early twentieth century, when copyright law faced player piano rolls: rolls of paper with perforations in accordance with musical works.⁴² When installed on a player piano, these rolls cause the piano to play notes in sequence as determined by the position and length of the perforations, thereby performing the song encoded therein. In 1908, the Supreme Court considered in *White-Smith Music Publishing Co. v. Apollo Co.* whether the piano rolls—which would be "read" by a machine to play the encoded musical composition rather than by a human—were "copies" of the musical composition, thereby constituting copyright infringement.⁴³ The plaintiff in the case owned copyrights in certain musical compositions, and the

38. See, e.g., *Three Boys Music Corp. v. Bolton*, 212 F.3d 477, 482–85 (9th Cir. 2000) (stating that "[s]ubconscious copying has been accepted" alongside proof of widespread dissemination to satisfy proof of the reasonable access element of copyright infringement); *ABKCO Music, Inc. v. Harrisongs Music, Ltd.*, 722 F.2d 988, 998–99 (2d Cir. 1983) ("It is not new law in this circuit that when a defendant's work is copied from the plaintiff's, but the defendant in good faith has forgotten that the plaintiff's work was the source of his own, such 'innocent copying' can nevertheless constitute an infringement.").

39. *Three Boys Music*, 212 F.3d at 484–85.

40. *Fred Fisher, Inc. v. Dillingham*, 298 F. 145, 147–48 (S.D.N.Y. 1924).

41. Copyright law would likely not exist in the first place without the printing press, which made the large-scale copying of written material plausible. See *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 430 (1984) ("Indeed, it was the invention of a new form of copying equipment—the printing press—that gave rise to the original need for copyright protection.").

42. Zhengshan Shi, Kumaran Arul & Julius O. Smith, Modeling and Digitizing Reproducing Piano Rolls, in *Proceedings of the 18th International Society for Music Information Retrieval Conference 197*, 197 (Xiao Hu, Sally Jo Cunningham, Doug Turnbull & Zhiyao Duan eds., 2017), https://ismir2017.smcnus.org/wp-content/uploads/2017/10/25_Paper.pdf [<https://perma.cc/W9EH-JLJL>].

43. 209 U.S. 1, 17–18 (1908). Under the copyright statute in place at the time—and continuing through its current version—copyright law deemed copying of copyrighted works to be infringement. *Id.* at 9.

defendant was in the business of making and selling player pianos and piano rolls.⁴⁴

The Supreme Court ultimately held that the piano roll was not a copy of the musical composition it represented (and therefore the plaintiff could not prohibit this type of reproduction by the defendant).⁴⁵ In particular, the Court reasoned that something could not count as an infringing use unless it was “put in a form which [humans] can see and read.”⁴⁶ Because people did not read piano rolls as they read sheet music, piano rolls did not satisfy this requirement. The Court thought it irrelevant that “[t]hese perforated rolls are parts of a machine which, when duly applied and properly operated in connection with the mechanism to which they are adapted, produce musical tones in harmonious combination.”⁴⁷

In its ruling, the Court thus adopted the view that machines were unlike humans for purposes of copyright infringement: Machine-read materials did not constitute copyright infringement unless humans can read the same material as well.⁴⁸ However, Congress evidently did not share the Supreme Court’s broad view on this distinction between humans and machines.⁴⁹ Although there are arguably justifications for a focus on human readability, *White-Smith’s* formalism provoked severe criticism.⁵⁰ Even if a person could not read or hear the musical composition encoded in a piano roll, that same person could still consume the work with the help of a player piano.⁵¹ As a practical matter, *White-Smith* meant that copiers could circumvent copyright protections by creating copies of a work that were unreadable by humans, but could be made comprehensible with the aid of a machine.⁵²

The following year, Congress overturned the specific holding of *White-Smith* by granting copyright holders in musical works the right to control the mechanical reproduction of their works and instituting a compulsory license scheme for manufacturers of piano rolls and other

44. *Id.* at 8–9.

45. *Id.* at 18.

46. *Id.* at 17.

47. *Id.* at 18.

48. *Id.* at 17–18.

49. See Yvette Joy Liebesman, Redefining the Intended Copyright Infringer, 50 Akron L. Rev. 765, 790 (2016) (stating that “Congress amended the Copyright Act to include these works under its purview” (citing An Act to Amend and Consolidate the Acts Respecting Copyright, ch. 320, 35 Stat. 1075, 1081–82 (1909))).

50. See, e.g., H.R. Rep. No. 94-1476, at 52 (1976) (criticizing *White-Smith* for its “artificial and largely unjustifiable distinction[] . . . under which statutory copyrightability . . . has been made to depend upon the form or medium in which the work is fixed”).

51. *White-Smith*, 209 U.S. at 8–10.

52. See Liebesman, *supra* note 49, at 787–90 (finding that the Supreme Court’s decisions “confine[ing] copies of musical works . . . to those specific mediums of expression defined by Congress . . . resulted in a larger reach of legal copying and subsequently a smaller cohort of who was an intended infringer”).

mechanical reproductions.⁵³ And almost seventy years later, Congress changed its definition for copyright law of “copies” to include not only “material objects” that can be read or perceived “directly” by humans but also those “from which the work can be perceived, reproduced, or otherwise communicated . . . with the aid of a machine or device.”⁵⁴ With that definition, Congress took an expansive view of machine-readable forms of works as “copies,” so long as humans could perceive or read them via the machine.

C. *The Internet*

Nonetheless, further questions as to machines’ ability to engage in copyright infringement subsequently arose, especially as the internet era dawned in the 1990s. For the first time, machines—computers—interconnected on a vast network around the world were copying and transmitting material to one another (and ultimately often to people using these machines). Any human posting or emailing material that infringed another’s copyright would therein provoke countless interconnected machines to make copies of this material as well. Some frustrated copyright holders sued certain of these users and machine owners—typically, internet service providers—for copyright infringement.

The foundational case of *Religious Technology Center v. Netcom On-Line Communication Services, Inc.* addressed the liability of internet server owners.⁵⁵ *Netcom* was a suit by the Church of Scientology against both former minister Dennis Erlich, for uploading messages to Usenet containing copyrighted church texts and criticism of the church, and internet service providers, including BBS and Netcom, whose servers created copies of those messages.⁵⁶ The Northern District of California viewed the liability of the entities deploying these servers as turning on “whether possessors of computers are liable for incidental copies automatically made on their computers using their software as part of a process initiated by a third party.”⁵⁷ But the court refused to assign liability to the server owners: “Although copyright is a strict liability statute, there should still be some element of volition . . . which is lacking where a defendant’s system is merely used to create a copy by a third party.”⁵⁸ The *Netcom* court thought that because the defendants’ “systems can operate without any human intervention, . . . the mere fact that Netcom’s system incidentally makes temporary copies of [the church’s] works does not

53. An Act to Amend and Consolidate the Acts Respecting Copyright § 1(e).

54. Copyright Act of 1976, Pub. L. No. 94-553, § 101, 90 Stat. 2541, 2542 (codified as amended at 17 U.S.C. § 101 (2012)).

55. 907 F. Supp. 1361 (N.D. Cal. 1995).

56. See *id.* at 1365–66.

57. *Id.* at 1368.

58. *Id.* at 1370.

mean Netcom has caused the copying.”⁵⁹ The court emphasized the risk of establishing a contrary rule:

[A contrary rule] would also result in liability for every single [internet] server in the worldwide link of computers transmitting [the ex-church minister’s] message to every other computer. These parties, who are liable under [the church’s] theory, do no more than operate or implement a system that is essential if [internet] messages are to be widely distributed. There is no need to construe [copyright law] to make all of these parties infringers.⁶⁰

Thus, the *Netcom* court strongly suggested that—although a human using a machine to make a copy is thereby volitionally infringing a copyright—a machine itself cannot possess the requisite volition to be regarded as an infringer, or as thereby “acting” on behalf of the technology provider.⁶¹

Building on *Netcom* and its progeny,⁶² the Second Circuit further stressed the differential treatment of humans and machines with regard to volition and copyright infringement in *Cartoon Network LP v. CSC Holdings, Inc.*⁶³ In that case, the court held that a cable company’s remote-storage digital video recording system did not directly infringe the copyrights of a cable television company when cable company customers requested or played back recordings on this system.⁶⁴ For one thing, the court dismissed the possibility that the cable company satisfied the volitional act requirement for infringement liability by virtue of its “conduct in designing, housing, and maintaining a system that exists only to produce a copy . . . made automatically upon [a] customer’s command.”⁶⁵ For even though the copying was instrumental to the function of the recording system, the court held that it was the customer requesting the recording—rather than the system or its owner—who made the copy.⁶⁶ The court thought that it would have been a different situation, however, had the customer requested a human employee of the cable system—rather than the machine itself—to make the copy: “In determining who actually ‘makes’ a copy, a significant difference exists between making a request to a human employee, who then volitionally operates the copying system to make the copy, and issuing a command

59. *Id.* at 1368–69.

60. *Id.* at 1369–70. The court left open the possibility that the internet service providers would instead be liable for contributory infringement. *Id.* at 1369, 1373–75.

61. *See id.* at 1370.

62. Cases in the intervening years on this issue include *CoStar Grp., Inc. v. LoopNet, Inc.*, 373 F.3d 544 (4th Cir. 2004); *Field v. Google Inc.*, 412 F. Supp. 2d 1106 (D. Nev. 2006); *Playboy Enters., Inc. v. Russ Hardenburgh, Inc.*, 982 F. Supp. 503 (N.D. Ohio 1997); *Marobie-FL, Inc. v. Nat’l Ass’n of Fire Equip. Distribs.*, 983 F. Supp. 1167 (N.D. Ill. 1997).

63. 536 F.3d 121 (2d Cir. 2008).

64. *See id.* at 123.

65. *Id.* at 131.

66. *Id.*

directly to a system, which automatically obeys commands and engages in no volitional conduct.”⁶⁷ The Second Circuit seemed to state categorically that machines—circa 2008—always lack the requisite volition to be infringers acting on behalf of technology providers, whereas humans, including human employees, always possess it.⁶⁸

While some courts were denying the possibility that machines could volitionally infringe on behalf of technology providers, others seemed to ignore the volitional act requirement entirely, instead readily assuming—without analysis—that computers’ owners had infringed when their machines automatically copied protected content. For example, in a series of cases, courts generally found businesses operating search engines not liable for copying infringing works found online to index and make them available for user searching.⁶⁹ But these courts never paused to question whether the machines had *volitionally* copied, proceeding instead to decide that there was in fact a *prima facie* case of copyright infringement by the search engine operators but that their copying was nonetheless fair use.⁷⁰ Similarly, the Supreme Court, in *American Broadcasting Companies v. Aereo, Inc.*, made no mention of volition before finding the owner of many small internet-connected antennae liable for streaming (that is, publicly performing) broadcast television programming to subscribers.⁷¹

67. *Id.*

68. In some ways, a prior decision by the Fourth Circuit had already muddied the volition waters further. The Fourth Circuit found that an internet service provider lacked volition when the company had its human employees take a quick look at whether commercial real estate photographs posted by users seemed to infringe on third parties’ copyrighted material and its computers copied the infringing material to check it against any new material uploaded by that user. See *CoStar Grp. v. LoopNet, Inc.*, 373 F.3d 544, 556 (4th Cir. 2004). The court elaborated:

The employee’s look is so cursory as to be insignificant, and if it has any significance, it tends only to lessen the possibility that [the provider]’s automatic electronic responses will inadvertently enable others to trespass on a copyright owner’s rights. In performing this gatekeeping function, [the provider] does not attempt to search out or select photographs for duplication; it merely *prevents* users from duplicating certain photographs. . . . [The provider] can be compared to an owner of a copy machine who has stationed a guard by the door to turn away customers who are attempting to duplicate clearly copyrighted works. [The provider] has not by this screening process become engaged as a “copier” of copyrighted works who can be held liable under . . . the Copyright Act.

Id.

69. See *Perfect 10, Inc. v. Amazon.com, Inc.*, 508 F.3d 1146, 1176–77 (9th Cir. 2007); *Kelly v. Arriba Soft Corp.*, 336 F.3d 811, 822 (9th Cir. 2003).

70. *Perfect 10*, 508 F.3d at 1168; *Kelly*, 336 F.3d at 822. Perhaps the courts never considered volition because the machines’ owners in these cases provoked the copying in the first instance. Cf. Robert C. Denicola, *Volition and Copyright Infringement*, 37 *Cardozo L. Rev.* 1259, 1279–80 (2016) (“[I]f no third party has participated in the alleged infringement, defendants rarely invoke the volition requirement; when they do, the issue is quickly resolved in favor of the plaintiffs.”).

71. See 134 S. Ct. 2498, 2498–511 (2014).

In dissent, Justice Scalia lambasted the majority for failing to consider whether volition was present as a prerequisite to finding infringement:

Although we have not opined on the issue, our cases are fully consistent with a volitional-conduct requirement. . . .

The volitional-conduct requirement is not at issue in most direct-infringement cases; the usual point of dispute is whether the defendant's conduct is infringing (*e.g.*, Does the defendant's design copy the plaintiff's?), rather than whether the defendant has acted at all (*e.g.*, Did this defendant create the infringing design?). But it comes right to the fore when a direct-infringement claim is lodged against a defendant who does nothing more than operate an automated, user-controlled system. Internet-service providers are a prime example. When one user sends data to another, the provider's equipment facilitates the transfer automatically. Does that mean that the provider is directly liable when the transmission happens to result in the "reproduc[tion]" of a copyrighted work? It does not. The provider's system is "totally indifferent to the material's content," whereas courts require "some aspect of volition" directed at the copyrighted material before direct liability may be imposed. The defendant may be held directly liable only if the defendant *itself* "trespassed on the exclusive domain of the copyright owner." Most of the time that issue will come down to who selects the copyrighted content: the defendant or its customers.

.....
The distinction between direct and secondary liability would collapse if there were not a clear rule for determining whether the *defendant* committed the infringing act. The volitional-conduct requirement supplies that rule; its purpose is not to excuse defendants from accountability, but to channel the claims against them into the correct analytical track.⁷²

Thus, *Aereo* has caused some to wonder whether the majority had implicitly rejected a volitional act requirement for copyright infringement,⁷³

72. *Id.* at 2513–14 (Scalia, J., dissenting) (citations omitted) (first quoting 17 U.S.C. § 106(1) (2012); then quoting *CoStar Grp.*, 373 F.3d at 550–51; then quoting *id.* at 550). There is a conceptual connection between a volitional act requirement and certain forms of secondary liability in copyright law. In particular, the Supreme Court has held—with respect to secondary liability for a provider of peer-to-peer file-sharing software—that “one who distributes a device with the object of promoting its use to infringe copyright, as shown by clear expression or other affirmative steps taken to foster infringement, is liable for the resulting acts of infringement by third parties.” *Metro-Goldwyn-Mayer Studios Inc. v. Grokster, Ltd.*, 545 U.S. 913, 936–37 (2005). Just as the presence of volition indicates that a technology provider has gone beyond merely deploying its automated system to copy, inducement of third-party infringement indicates that a technology provider has gone beyond merely providing a system or device that can be used by others to infringe copyright.

73. *E.g.*, Bruce E. Boyden, *Aereo* and the Problem of Machine Volition, 2015 Mich. St. L. Rev. 485; Kyle A. Brown, Comment, Up in the *Aereo*: Did the Supreme Court Just Eliminate the Volitional Conduct Requirement for Direct Copyright Infringement?, 46 Seton Hall L. Rev. 243 (2015).

although the Second and Ninth Circuits have affirmed the requirement's continuing relevance.⁷⁴ Owing to the ongoing relevance of volition in copyright law, it is worth making sense of this requirement, to which the next section now turns.

D. *What Is Volition in Copyright Law?*

What exactly is this “volition” mental state required for copyright infringement liability? As this Essay noted earlier, volition might be understood as the mental state that causes willful actions.⁷⁵ In other words, the question of whether some event counts as volitional is the question of whether it is something genuinely willed or chosen by the so-called actor. The presence of a volitional mental state as a cause thus distinguishes involuntary bodily movements—such as those during a seizure—from voluntary ones.⁷⁶ With that distinction, a volition requirement coheres with the intuition that individuals should be held responsible for, and only for, that which was under their control.⁷⁷ As the Restatement (Second) of Torts explains, “Some outward manifestation of the defendant's will is necessary to the existence of an act which can subject him to liability.”⁷⁸

Note that copyright's volitional act requirement is asking for volition or control in something very specific: the production of the infringing copy itself. After all, technology providers have chosen—that is, willfully acted—in providing copy-making technologies, such that holding them responsible for resulting infringements would not constitute responsibility for something entirely out of their control.⁷⁹ Nonetheless, volitionally providing the technology is not sufficient for satisfying copyright law's volitional act requirement. Instead, copyright requires that the instance of infringing copying itself be volitional—or itself count as a willful action on the part of the technology provider—and that the infringing

74. See *BWP Media USA Inc. v. Polyvore, Inc.*, 922 F.3d 42, 49 (2d Cir. 2019) (per curiam) (“[W]e have reaffirmed post-*Aereo* . . . that “[v]olitional conduct is an important element of direct liability.” (quoting *EMI Christian Music Grp., Inc. v. MP3tunes, LLC*, 844 F.3d 79, 96 (2d Cir. 2016))); *Perfect 10, Inc. v. Giganews, Inc.*, 847 F.3d 657, 667 (9th Cir. 2017) (explaining that one element of a direct infringement claim is volitional conduct).

75. See, e.g., Moore, *Act and Crime*, supra note 11 (canvassing and assessing different philosophical conceptions of volition); Robert Audi, *Volition, Intention, and Responsibility*, 142 U. Pa. L. Rev. 1675, 1680 (1994) (“Moore sees conflict as a pervasive element in our desire and belief systems. Action cannot occur without resolution of such conflicts; volition here plays the role of reconciler, or, at least, of referee.”).

76. See, e.g., Restatement (Second) of Torts § 2 cmt. a (Am. Law Inst. 1965) (“There cannot be an act without volition. Therefore, a contraction of a person's muscles which is purely a reaction to some outside force, such as a knee jerk . . . , are not acts of that person. . . . So too, movements of the body during sleep . . . are not acts.”).

77. Moore, *Act and Crime*, supra note 11, at 48.

78. Restatement (Second) of Torts § 2 cmt. a.

79. See Denicola, supra note 70, at 1265 (explaining that courts have found volition when defendants made a choice to deploy systems that made infringement possible).

conduct can be attributed to the provider rather than the technology user alone.⁸⁰ This is to say that copyright law asks for volition at a specific point on the causal chain: not simply the instance of providing copying technology but the particular instance of copying.⁸¹ This requirement is plausibly motivated by the policy that it would be bad to hold technology providers responsible for all infringements resulting from their technologies—including ones proximately caused by someone else's actions—when these technologies are capable of value-adding, non-infringing uses and are therefore not ones that the law seeks to disincentivize entirely. For copyright, such technology providers thus must have volitionally “committed” the infringing action themselves, perhaps with opportunity to pause, evaluate, and then choose whether to proceed with the particular infringing action, in order to be held responsible for it.⁸²

All in all, given that the legal attention to machine operations has been relatively extensive in the context of copyright's volitional act requirement, it provides a good test bed for exploring machine mental states more broadly across the law. For in copyright law, many courts have treated liability for human and mechanical, or automated, acts of copying dichotomously: Humans always have volition, even when they are copying subconsciously, whereas machines can—and, to some courts, always—lack volition, even when carrying out acts of copying for which they are centrally designed.⁸³ Indeed, the particularly strong language of

80. See *id.* at 1272 (describing a hypothetical in which a customer uses a provider's machine to reproduce a copyrighted work to demonstrate that “[t]he volition requirement . . . defines the connection between the owner of a copying system and the copied work that is sufficient to justify attributing the copying of that work to the owner”).

81. An alternative way of describing copyright's volitional act requirement is that it requires that the actions of the technology provider be the proximate cause of the production of the copy for the technology provider to be liable for infringement. See, e.g., *BWP Media USA Inc. v. Polyvore, Inc.*, 922 F.3d 42, 61–67 (2d Cir. 2019) (Newman, J., concurring in the result) (“Infringement is a tort . . . ‘Volition’ . . . is best understood to mean a concept essentially reflecting tort law causation. . . . ‘[C]ausation,’ in the context of copyright infringement, is tort law ‘proximate cause,’ rather than ‘but for’ causation.”). Note that this interpretation of the volitional act requirement is ultimately equivalent to the interpretation we favor according to which volitions are the mental states causing actions, for it is asking whether the proximate cause of infringement is the action of the technology provider. Furthermore, what determines whether something counts as the technology provider's actions (rather than someone else's) is whether it is the result of the technology provider's (or its machine's) volitional mental state (which causes actions rather than mere movements).

82. See, e.g., Moore, *Act and Crime*, *supra* note 11, at 111–65.

83. Cf. James Grimmelmann, *Copyright for Literate Robots*, 101 *Iowa L. Rev.* 657, 657 (2016) (“Almost by accident, copyright law has concluded that it is for humans only: reading performed by computers doesn't count as infringement. Conceptually, this makes sense: Copyright's ideal of romantic readership involves humans writing for other humans.”). Professor Matthew Sag has observed that whether machines or their owners are liable for copyright infringement ought to turn on whether the machines are copying

Cartoon Network seems to entail that if we imagine an (inefficient) internet whose computers—servers and all—are each replaced with a human given the task to copy received material and pass it on toward the specified destination, then this imagined internet would count as having volition under copyright law at each node, whereas the currently automated internet lacks it entirely.⁸⁴ This implication is notwithstanding the fact that both variations of the internet—by stipulation—would be functionally identical systems. But this thought experiment is reminiscent of those deployed by philosophers in their efforts to understand the nature of human minds and machines, to which we now turn.

III. THE PHILOSOPHY OF MIND AND MACHINES

This Part surveys two of the most influential philosophical discussions on the mind—namely, John Searle’s “Chinese Room” argument and David Chalmers’s two concepts of mind—in order to explicate the important conceptual distinction between “conscious” and “functional” understandings of mental states. It then explains the implications of this philosophical distinction for the question of whether any of the law’s mental state requirements, such as copyright law’s volitional act requirement, can or should be satisfied by machines.

A. *John Searle and the “Chinese Room” Argument*

Philosophers of mind have long contemplated whether there is any fundamental difference between human and artificial minds. Perhaps the most well-known challenge to the possibility of computers with truly human-like mental states is John Searle’s “Chinese Room” argument. This argument has shaped much of the course of philosophical thinking on these questions since its publication in 1980, spurring continuing debate about the possibility of so-called “strong” AI—purely computational systems that possess *conscious* mental states like those of humans—versus “weak” AI, which merely *functionally* simulates the human mind.⁸⁵ In particular, Searle asks us to consider the following thought experiment:

Suppose that I’m locked in a room and given a large batch of Chinese writing. Suppose furthermore (as is indeed the case) that I know no Chinese Now suppose further that after this first batch of Chinese writing I am given a second batch of

works for expressive or nonexpressive uses. See Matthew Sag, *Copyright and Copy-Reliant Technology*, 103 Nw. U. L. Rev. 1607, 1624–44 (2009).

84. See *Cartoon Network LP v. CSC Holdings, Inc.*, 536 F.3d 121, 131 (2d Cir. 2008) (“In determining who actually ‘makes’ a copy, a significant difference exists between making a request to a human employee, who then volitionally operates the copying system to make the copy, and issuing a command directly to a system, which automatically obeys commands and engages in no volitional conduct.”).

85. See Paul M. Churchland & Patricia Smith Churchland, *Could a Machine Think?*, *Sci. Am.*, Jan. 1990, at 32, 32–34 (noting that “Searle’s paper provoked a lively reaction from AI researchers, psychologists and philosophers alike”).

Chinese script together with a set of rules for correlating the second batch with the first batch. The rules are in English, and I understand these rules as well as any other native speaker of English. They enable me to correlate one set of formal symbols with another set of formal symbols, and all that “formal” means here is that I can identify the symbols entirely by their shapes. . . . Suppose also that after a while I get so good at following the instructions for manipulating the Chinese symbols . . . that from the external point of view—that is, from the point of view of somebody outside the room in which I am locked—my answers to the questions are absolutely indistinguishable from those of native Chinese speakers.⁸⁶

In other words, Searle asks us to imagine that he is performing computational operations on the Chinese characters in accordance with formal rules, thereby instantiating a computer program.⁸⁷ Although the program that he is operating has the same input–output structure as a human fluent in Chinese, such that it is computationally *equivalent* to a Chinese speaker, Searle argues that he—and the program—nonetheless lack the *conscious experience* of a Chinese speaker who genuinely understands the language.⁸⁸ In other words, he explains, there is a fundamental difference between what goes on in the Chinese Room and an alternative scenario in which Searle responds to English inputs with outputs on the basis of formal rules.⁸⁹ In the case of English, Searle is not solely *functionally* instantiating the English program but also consciously understands.⁹⁰ In the Chinese Room, however, he merely simulates a conscious Chinese speaker.⁹¹

Searle’s thought experiment challenged both the view that it is possible for there to be an artificial system with conscious mental states resulting from purely computational processes⁹² and the view that human

86. John R. Searle, *Minds, Brains, and Programs*, 3 *Behav. & Brain Sci.* 417–18 (1980). Note that Searle himself originally put forth the “Chinese Room” argument as a challenge to the possibility of computation-based *understanding* rather than *consciousness*. But some philosophers have subsequently interpreted the argument as actually challenging the possibility of an artificial computer *experiencing* understanding, which is ultimately the question of artificial consciousness. See, e.g., David Chalmers, *The Conscious Mind: In Search of Fundamental Theory* 322–23 (1996). For this Essay’s purposes, we follow these philosophers’ interpretation of Searle’s argument. Nonetheless, we flag the alternative interpretation and note that the choice of interpretation ultimately has no bearing on this Essay’s thesis.

87. Searle, *supra* note 86, at 418.

88. *Id.*

89. *Id.*

90. *Id.*

91. *Id.*

92. As Searle explains,

Whatever else intentionality is, it is a biological phenomenon, and it is as likely to be as causally dependent on the specific biochemistry of its origins as lactation, photosynthesis, or any other biological phenomena. No one would suppose that we could produce milk and sugar by running a

consciousness is itself simply the product of computation.⁹³ In other words, Searle argued, because the functional processes of computation cannot give rise to conscious mental states and because our human minds clearly possess such mental states, it cannot be the case that our human minds are solely instantiating a program.⁹⁴

This argument triggered decades of discussion, including a slew of critical responses from philosophers, psychologists, and computer scientists. Some of these challenges reject Searle's conclusion about the Chinese Room, saying it in fact does experience understanding of Chinese, even if the person inside the room—who is only a part of the computational system—does not.⁹⁵ Others have said that even if the Chinese Room lacks such experience, this is only because it is running the wrong kind of program; if it were instead running, say, a program simulating all the intricacies of the human brain, then it would have the experience of a Chinese speaker.⁹⁶ But Searle himself has responded to these objections, even addressing many in his original paper;⁹⁷ and there thus remains a rift between those who find the Chinese Room to be compelling in showing that the human mind could not be a computer and those who regard the argument as fundamentally mistaken.

B. *David Chalmers and the Hard Problem of Consciousness*

Regardless of whether Searle's argument is successful, the conceptual distinction between conscious and functional properties of mental states—which is made particularly vivid by the Chinese Room argument—remains enormously important and is taken seriously by all such philosophers. Pointedly, even human mental states can be understood in

computer simulation of the formal sequences in lactation and photosynthesis, but where the mind is concerned many people are willing to believe in such a miracle because of a deep and abiding dualism

Id. at 424.

93. Id. (“Whatever it is that the brain does to produce intentionality, it cannot consist in instantiating a program since no program, by itself, is sufficient for intentionality.”).

94. Id.

95. For example, Daniel Dennett posits that

Searle, laboring in the Chinese Room, does not understand Chinese, but he is not alone in the room. There is also the System, . . . and it is to *that* self that we should attribute any understanding

This reply to Searle's example is what he calls the systems reply. It has been the standard reply of people in AI from the earliest outings of his thought experiment.

Daniel C. Dennett, *Consciousness Explained* 439 (1991).

96. See, e.g., Chalmers, *supra* note 86, at 323–25 (arguing that at least a system with the same functional organization or structure as a brain would mirror the “causal relations between neurons” and therefore have the same conscious properties); Churchland & Churchland, *supra* note 85, at 37 (arguing that a system mimicking a human brain might be conscious).

97. Searle, *supra* note 86, at 419–22.

terms of either conscious or functional properties. David Chalmers famously made this point in *The Conscious Mind*, which articulated what he called the “hard problem” of consciousness.⁹⁸ As Chalmers explains, the term “conscious” might be understood as synonymous with “phenomenal,” the idea being that if an entity is conscious, then there is something that it is like to be that being.⁹⁹ To illustrate this concept, consider the contrasting examples of a human and a thermometer. Although a human and a thermometer both possess functional attributes that enable them to detect heat, the human feels or experiences heat, whereas the thermometer does not.¹⁰⁰ This is the difference between beings—such as humans—that have the capacity for such subjective experiences and beings—such as thermometers—that do not: Only the former are conscious beings.

In his book, Chalmers demonstrates that individual human mental states can be analyzed either in terms of what he calls their *psychological* properties—their functional role in producing behavior, or what they do—or their *phenomenal* properties—their conscious quality, or how they feel.¹⁰¹ That is, according to Chalmers, the functional and the conscious concepts of the mind are distinguishable, even with respect to the human mind.¹⁰² Consider, for instance, Chalmers’s example of the “pain” mental state.¹⁰³ Pains have conscious aspects: There is something it is like to be in pain (indeed, it is unpleasant).¹⁰⁴ But pains also have entirely functional properties, which specify their structural roles in causal systems. For example, a pain has the functional properties of typically being the product of some damage to one’s body, leading to adverse reactions to the stimulus such as saying “ow,” recoiling, and so forth.¹⁰⁵ Upon separating the two concepts of mind, Chalmers ultimately argues for the conceivability of an entity that possesses human mental states understood entirely in terms of their functional properties, but which nonetheless lacks any conscious experience of those states.¹⁰⁶ As he explains, the Easy Problem (despite being difficult in its own right) is the question of the precise functional nature of mental states;¹⁰⁷ the Hard Problem is the

98. Chalmers, *supra* note 86, at xi–xii.

99. *Id.* at 285–86. See generally Thomas Nagel, *What Is It Like To Be a Bat?*, 83 *Phil. Rev.* 435 (1974) (explicating the philosophical difficulties surrounding the concept of consciousness).

100. At least, we plausibly suspect that it does not. An alternative view is offered by panpsychism, the idea that all objects possess conscious minds. See, e.g., Chalmers, *supra* note 86, at 297–301.

101. *Id.* at 11.

102. *Id.* at 17.

103. *Id.*

104. *Id.*

105. *Id.*

106. *Id.* at 17–18.

107. *Id.* at xi–xii.

question of why or how certain beings—such as humans—also have conscious experience.¹⁰⁸

This philosophical distinction between the conscious and functional properties of the mind has important implications for the law and its governance of machines. This is because, regardless of one's views on whether conscious AI is possible, most philosophers—including Searle—agree that machines (like the Chinese Room) can in principle replicate the functional properties of human minds.¹⁰⁹ Moreover, for each of the law's mental state requirements, it remains an open question whether the law ultimately seeks to track the conscious or functional properties of the states in question. Because the law has primarily been designed for human actors, for whom the conscious and the functional typically coincide, this is a question we have principally been able to avoid until now. But the increasing prevalence of ever-sophisticated machines requires us to take it seriously. If the law is concerned only with functional properties, then these properties could very well be possessed by the states of a nonhuman machine.¹¹⁰ In other words, then, it is far from settled that all the law's mental state requirements should be satisfied only by conscious minds. The remainder of this Essay challenges this assumption, analyzing the case of the aforementioned volitional act requirement in copyright law.

IV. VOLITION AND AI: IS CONSCIOUSNESS RELEVANT?

This Part argues that the volition requirement in copyright law ultimately does not seem interested in tracking conscious properties of the human infringer but instead functional ones, which could in principle be possessed by a machine.

The earlier analysis of the purpose of copyright's volitional act requirement¹¹¹ still leaves open the question of whether such "volition" at the instance of infringement must be conscious rather than some functional analogue, or whether such a purely functional state of a machine can result in something that, at least for the law's purposes, should be regarded as a "willful action" on the part of the technology provider. In

108. *Id.* at 4–5.

109. See Searle, *supra* note 86, at 418 (granting that the Chinese Room is functionally "indistinguishable . . . from native Chinese speakers").

110. According to one school of artificial intelligence, human-like intelligence in machines can emerge only from machines that are embodied with features that are human-like, such as the brain and eyes. See generally Rodney A. Brooks, *Cambrian Intelligence: The Early History of the New AI* (1999) (exploring how behavior-based robots can act in ways that appear intelligent); Andy Clark, *Being There: Putting the Brain, Body, and World Together Again* (A Bradford Book reprint ed. 1998) (1997) (theorizing how the brain is a controller for embodied activity, and deriving an action-oriented theory of the mind). To the extent that this school is correct, artificial intelligence will appear relatively human.

111. See *supra* section II.D.

other words, given the law's concerns and that a business's human employees almost always count as "acting" on the part of the business for the law's purposes, is there a reason for thinking, as a categorical matter, that the business's nonconscious machines—no matter their functions—never could? We think the answer is "no."

Consider first the general question: When should *any* area of law require a conscious rather than so-called functional volition? One might argue that a being should be held legally responsible for itself—or as a conscious, autonomous agent—only if that being is genuinely conscious. But this thesis would certainly need to be defended, for it would depend on the purpose of liability in the particular legal domain. If the purpose is entirely to produce the proper incentives—the dominant American view of copyright¹¹²—then it is not clear why the actor being held responsible must have consciousness, rather than simply the right functional responses to such incentives. On the other hand, at least for some areas of law, one might have the view that legal responsibility is meant to track moral responsibility.¹¹³ Such a theorist thus might argue that it is nonsensical to hold a nonconscious being morally responsible for its behavior, as such a being is not a moral agent. Underlying this claim is the premise that, for something to be a moral agent, it must have conscious experience. But even this supposition requires substantiation and is undoubtedly up for debate.¹¹⁴ For instance, imagine a machine with all the functional properties of a human. Such a machine would thereby have the capacity for something functionally equivalent to moral deliberation and judgment, and for choosing an action on the basis of such judgment, all despite lacking any conscious experience of this process. We might thus wonder why these functional capacities are not themselves sufficient for moral agency, or why their conscious quality (or lack thereof) would be relevant to the question at all.

In any event, even if one embraces the view that a being must be conscious for it to be held legally responsible for itself, this ultimately does not pose a challenge for the suggestion—say, in the context of copyright law—that the mere functionality of a technology provider's machine could suffice for holding that provider responsible. This is because holding a human or business entity responsible for its machine (or,

112. See *supra* text accompanying notes 30–37. See generally William M. Landes & Richard A. Posner, *The Economic Structure of Intellectual Property Law* (2003) (articulating and defending an economic understanding of the aims of intellectual property law).

113. See, e.g., Michael S. Moore, *Causation and Responsibility: An Essay in Law, Morals, and Metaphysics* 4 (2009) (“[C]riminal and tort liability must track moral responsibility, because justice is achieved only if the morally responsible are held liable to punishment or tort damages.”).

114. See, e.g., S. Matthew Liao, *The Basis of Human Moral Status*, 7 *J. Moral Phil.* 159, 169 (2010) (arguing that the basis of human moral status is not the conscious properties of human beings but rather the fact that human beings possess the genetic basis for moral agency, and that nonhuman beings could also possess moral status).

indeed, its employee) does not seem to amount to treating said machine (or employee) as a conscious, autonomous agent; rather, it amounts to treating the human or business entity as responsible for the machine. In other words, whether or not machines themselves must have conscious mental states in order to be held responsible for their own so-called behavior, the question of whether a business entity can be held responsible for its machines—that is, whether these machines can be regarded as “acting” on said corporation’s behalf—does not seem like it should turn on whether the machine in question is conscious.

Moreover, the idea that copyright’s rules for infringement liability are ultimately unconcerned with consciousness is further suggested by the doctrine of subconscious copying, which has been widely criticized but nonetheless firmly remains a part of copyright law.¹¹⁵ Recall that, under existing law, a human who subconsciously copies the work of another—that is, without any awareness that he or she is doing so—is still liable for copyright infringement.¹¹⁶ Of course, the term “subconscious” as used in this doctrine is importantly different from the concept of phenomenal consciousness discussed earlier, for “subconscious” in the doctrine refers to the absence of awareness that an act of copying—rather than original creation—has occurred, whereas “unconscious” in the phenomenal sense refers to the absence of any phenomenal qualities whatsoever. But nonetheless, if copyright does not care about a potential infringer’s awareness of their infringement, the question arises: Why think that it cares about the presence of any conscious awareness or experience whatsoever, even awareness of action? It is hard to see a reason to think it would. Indeed, when we imagine the case of a human employee operating a technology provider’s copy machine—one whose mental states, we have seen, would always satisfy copyright’s volition requirement—it seems plausible that this requirement might ultimately be interested in tracking what action the employee does and the function of their mind in facilitating this action, rather than their phenomenology while doing it.

V. A FUNCTIONAL UNDERSTANDING OF MACHINE VOLITION

The preceding discussion suggests that copyright’s volition requirement may not demand consciousness and may instead be more concerned with functionality. The doctrinal upshot is that so-called “functional volition”—or functional properties that capture what the law

115. See, e.g., Olufunmilayo B. Arewa, *The Freedom to Copy: Copyright, Creation, and Context*, 41 U.C. Davis L. Rev. 477, 531–39 (2007); Jessica Litman, *Copyright as Myth*, 53 U. Pitt. L. Rev. 235, 240 (1991); Wendy J. Gordon, *Toward a Jurisprudence of Benefits: The Norms of Copyright and the Problem of Private Censorship*, 57 U. Chi. L. Rev. 1009, 1029–31 (1990) (book review); Carissa L. Alden, Note, *A Proposal to Replace the Subconscious Copying Doctrine*, 29 Cardozo L. Rev. 1729, 1743–52 (2008).

116. *Supra* text accompanying notes 38–40.

is ultimately interested in tracking here—may suffice for copyright, such that the operation of a machine could give rise to direct liability for the technology provider, rather than solely for the technology user. This upshot has undoubted practical significance: It makes a substantial difference to copyright owners who would otherwise be limited to attempting to hold only individual users directly liable, and it prevents technology providers from avoiding direct liability simply by replacing human employees with copy-making machines. But this framework raises the question of which functional properties the volitional act requirement might seek and what a machine would have to look like to possess them.

Consider the human reproducing copies of another's copyrighted work, whom copyright law says always possesses the requisite volition for infringement liability. Indeed, consider the human employee making copies of a protected work. Principal-agent liability would readily confer liability on the employer without a doubt as to the employee's volition.¹¹⁷ Although such humans have the full range of the functional properties of a human mind, in which of these properties or capacities is the law ultimately interested in findings of infringement? Plausibly, it is not all of them, because the hypothetical copy-making humans do not use this full range of capacities. Instead, perhaps it is simply the humans' capacity to evaluate whether what they have been asked to copy is likely to be material within the realm of copyright subject matter—putting aside for the moment more complicated determinations such as the fair use defense to infringement, which is addressed shortly¹¹⁸—and to decline to make the copy on the basis of this determination. This functional capacity would align coherently with copyright law's stated aim to disincentivize third parties from copying protected materials in order to preserve the corresponding incentive that copyright offers to authors to create.¹¹⁹ And more broadly, it would seem to comport with the volitional act requirement's general purpose of ensuring that the actor has had an opportunity to pause to evaluate whether to proceed in acting—and to decline to perform the action if she so chooses on the basis of this evaluation—before being held responsible for the action.¹²⁰

This functional capacity seems relatively basic. Although it is not possessed by, say, a rudimentary copy machine—which is “compelled” to make copies upon the pressing of a button and therefore has no “choice” regardless of what is being copied—a more sophisticated computer could plausibly be designed to “choose” whether to make a copy, despite lacking the full range of human functional properties. In other words, such a computer—despite being functionally subhuman—would

117. See *infra* text accompanying notes 128–132 (summarizing agency law).

118. *Infra* text accompanying notes 123–124.

119. See *supra* text accompanying notes 30–33.

120. See *supra* section II.D.

be equivalent in all the ways copyright law cares about to a human operating a copy machine, who we already know is always “volitional” for purposes of copyright infringement.¹²¹ For instance, bots fetching songs and software generating new art based on learning from existing artwork could readily possess volition in this sense of the word.¹²²

On the other hand, perhaps the volitional act requirement seeks to track a more sophisticated functionality, such as the capacity to determine whether an instance of copying is likely to be fair use and choose to act on this determination. As Professor Dan Burk suggests, it is difficult—if not impossible—to devise algorithms that appropriately decide questions of fair use: “[T]he cost structure of algorithmic content policing has created a largely impersonal process, in which the context-specific factors that should be taken into account in fair use analysis are absent and go unconsidered.”¹²³ In particular, Burk worries about the “human judgment” that must be baked into these systems *ex ante* or in evaluating machines’ outputs *ex post*, such as a model of the markets for copyrighted works to assess the effect of a use on the market for a copyrighted work and the significance of the part of the work used.¹²⁴ Thus, if copyright law is interested in tracking the functional capacity to make plausible fair use determinations, then it seems that a functionally volitional machine remains far off.

Ultimately, this Essay does not aim to settle the question of the right functionality in which copyright law ought to be interested. Instead, it hopes to show that this is the type of question scholars and policymakers need to be asking, rather than simply assuming that machines can never be volitional as a matter of law.

Moreover, it must be emphasized that this conclusion is not merely one of philosophical interest. Rather, whether and when machines can possess the requisite volition to infringe copyright has great practical import. The precise contours of the volitional act requirement have implications for who is and is not directly accountable for the copying of protected material. For copyright law to accomplish its goals of encouraging the creation and dissemination of expressive works, it must provide

121. *Supra* notes 64–68 and accompanying text.

122. *Supra* text accompanying notes 19–21.

123. Dan L. Burk, *Algorithmic Fair Use*, 86 *U. Chi. L. Rev.* 283, 290 (2019) (emphasis omitted).

124. *Id.* at 296. Similarly difficult, as Sonia Katyal and Jason Schultz point out, are questions of which parts of a work are protectable as original and whether the author has expressly or implicitly licensed uses of the work. See Sonia K. Katyal & Jason M. Schultz, *The Unending Search for the Optimal Infringement Filter*, 112 *Colum. L. Rev. Sidebar* 83, 96–101 (2012), <https://columbia-lawreview.org/wp-content/uploads/2016/05/Katyal-Schultz.pdf> [<https://perma.cc/V4EW-MNUT>]. Other scholars are more supportive of the possibility of algorithmic copyright enforcement so long as the machine providers are transparent about and accountable for their substantive determinations. See Maayan Perel & Niva Elkin-Koren, *Accountability in Algorithmic Copyright Enforcement*, 19 *Stan. Tech. L. Rev.* 473, 477–78 (2016).

sufficient incentive to creators with copyright's exclusive rights and concomitant disincentive to third parties from infringing those rights by holding them liable for infringement.¹²⁵ Holding the providers of machines that act with the requisite volition directly liable for infringement thus plays an important role in doing just that. Indeed, even if there is also a technology user—a so-called customer—to hold accountable for infringing uses of a technology, this should not rule out holding qualifying technology providers liable for infringement as well. And given that technology users in these cases might be judgment proof while the technology provider frequently is not, the ability to hold the technology provider liable can have significant practical import. Moreover, because of its intricate connection to copyright policy, an inquiry into machine volition as a matter of direct liability will frequently be more pertinent and straightforward than an investigation of secondary liability, in light of the law's relatively mystifying standards for the latter.¹²⁶

At this point, one might be concerned with the policy implications of a conclusion that machines can have functional mental states or that functionality is what matters for findings of copyright infringement. For instance, does this overly discourage innovation of more sophisticated technologies, ones which—unlike simple copy machines—possess functional volition, to the extent that technology providers will attempt to “design around” liability? Or should technology providers be required to employ functionally volitional machines? Perhaps it would be sufficient to require machines to flag certain (or all) material for review by a human—such as a lawyer—before copying it, and thereby introduce human volition at the instance of copying. Such a design would give the machine the ability to pause and evaluate before proceeding to copy protected material. But it might also incapacitate machines from automating many of the tasks we have come to expect from them, precisely as the *Netcom* court worried.¹²⁷

Thus, the reader might wonder whether the forgoing discussion on human and machine volition should move us to reconsider the volitional act requirement itself. For instance, we might ask whether (on the one hand) a technology provider's volition in providing copying technology should be sufficient for liability rather than requiring volition at the

125. See *supra* text accompanying notes 30–37.

126. See Mark Bartholomew & John Tehranian, *The Secret Life of Legal Doctrine: The Divergent Evolution of Secondary Liability in Trademark and Copyright Law*, 21 *Berkeley Tech. L.J.* 1363, 1409–10 (2006) (“The difficulty of pursuing direct infringers has never served as a doctrinal basis for . . . secondary liability. Such reasoning undermines the stability of legal guidelines, rendering them unreliable . . . and erod[ing] the principled bases for secondary liability.”); Lital Helman, *Pull Too Hard and the Rope May Break: On the Secondary Liability of Technology Providers for Copyright Infringement*, 19 *Tex. Intell. Prop. L.J.* 111, 123 (2010) (stating that Supreme Court case law on secondary liability for copyright infringement “may have actually sowed the seeds of confusion reflected in the area . . . to this day”).

127. *Supra* text accompanying note 60–61.

instance of copying itself, or whether (on the other hand) the doctrine of subconscious copying should be rejected. And, indeed, such skeptical musings are ones in which we ourselves are inclined to engage. But note that they bear on the question of whether the volitional act requirement is a good thing and not whether—given what it seems to be trying to do—machines of any kind would and should ever satisfy it. Questions of the latter sort, we have demonstrated, cannot be handled so indelicately as some courts seem to think, for the law—for better or for worse—very well might here be interested in tracking only functional properties. Thus, as we enter a world in which users ask bots to find particular songs online and software gathers existing artworks to learn to create new art, it is increasingly important that we address such questions with due care. Moreover, because copyright law's volitional act requirement has served only as a case study, note that—regardless of what should or does become of this particular requirement—the challenge posed by the rest of the law's countless mental state requirements remains. The presented framework offers a path forward in analyzing how to adapt these requirements to a technologically evolving world.

The preceding analysis has pushed back on the assumption that mental state requirements can be satisfied only by human minds, instead asking both whether the particular requirement in question is ultimately about conscious or functional properties, and what a machine would have to look like to possess the functional properties of interest. But this analysis has focused in detail on one example, considering the apparent aims of a specific mental state requirement in copyright law. The next Part thus moves to generalize a theory of machines' mental states.

VI. TOWARD A GENERAL THEORY

As noted at the outset, despite the analytical focus until now on copyright law, the point to be gleaned from the present Essay is ultimately general: In the case of each implicit or explicit mental state requirement in the law, legal scholars and policymakers will need to engage in a similar analysis while attending to the unique interests and values at stake with regard to that law, in order to determine whether consciousness or mere functionality is what matters.

Of course, even if machines can have functional mental states, they do not have money, rights, or status as legal persons (at least, for the time being). Thus, the consequence of our analysis is that—to the extent that machines might be understood as having mental states for the law's purposes—machines might cogently be understood as agents of the business principal that creates or deploys them, performing actions for which that principal can be held directly responsible.¹²⁸ As one scholar puts it, an

128. See Anat Lior, *The Artificial Intelligence Respondeat Superior Analogy* 54 (unpublished manuscript) (on file with the *Columbia Law Review*). By contrast, if machines cannot be understood to possess mental states in the view of the law, it is likely that they

agent “functions as the principal’s representative, as an extension of the principal, while retaining the agent’s own separate legal personality.”¹²⁹ Agency, as per the most recent Restatement of the Law on the topic, is “the fiduciary relationship that arises when one person (a ‘principal’) manifests assent to another person (an ‘agent’) that the agent shall act on the principal’s behalf and subject to the principal’s control, and the agent manifests assent or otherwise consents so to act.”¹³⁰ An agent can act with actual or apparent authority from the principal vis-à-vis third parties.¹³¹ When an agent does so, pursuant to principles of respondeat superior, the principal can be legally liable for the agent’s actions.¹³² Thus, by suggesting the possibility of machines with legally required mental states, we are ultimately suggesting that there are contexts in which such machines are (functionally) agents in all the ways that matter. For that reason, just as a business would be liable for the conduct of its human agents, a business that creates and deploys these machines should be liable as principals for the conduct of these machines.¹³³ The possibility of technology providers being directly liable for infringement by their functionally volitional copying technologies is only one example of how this might manifest.

would instead be perceived as instrumentalities of the businesses or individuals that create and deploy them. Cf. *id.* at 12 (discussing the possible analogy of artificially intelligent machines to property). The Restatement (Third) of Agency takes the position that computers circa 2006 cannot be agents on the ground that “[t]o be capable of acting as . . . an agent, it is necessary to be a person, which in this respect requires capacity to be the holder of legal rights and the object of legal duties.” Restatement (Third) of Agency § 1.04, cmt. e (Am. Law Inst. 2006). According to the Restatement, “a computer program is not capable of acting as . . . an agent At present, computer programs are instrumentalities of the persons who use them.” *Id.* In light of this Essay’s analysis and the trajectory of artificial intelligence technology, this position may warrant reconsideration.

129. Deborah A. DeMott, *The Contours and Composition of Agency Doctrine: Perspectives from History and Theory on Inherent Agency Power*, 2014 U. Ill. L. Rev. 1813, 1816.

130. Restatement (Third) of Agency § 1.01.

131. See *id.* §§ 2.01–2.02 (actual authority); *id.* § 2.03 (apparent authority).

132. *Id.* § 2.04 (respondeat superior); *id.* § 2.06 (liability of undisclosed principal); *id.* §§ 7.03–7.08 (principal’s liability for an agent’s actions). According to agency principles, “[a]n agent is [also] subject to liability to a third party harmed by the agent’s tortious conduct . . . although the actor acts as an agent or an employee, with actual or apparent authority, or within the scope of employment.” *Id.* § 7.01. What this might mean with regard to artificially intelligent machines is beyond the scope of this Essay.

133. Note that a complete defense of the idea that machines can and should sometimes be regarded as the agents of humans or corporations would also require an explication of what mental states (or other requirements) *humans* need to possess to count as a machine’s principal. We set aside consideration of this question for future work. An important point to note, however, is that requirements for technology providers to be regarded as the principals of their functionally volitional machines are plausibly different from, and perhaps weaker than, what existing courts require of technology providers under their present (and, in our view, mistaken) understanding of the volitional act requirement.

To move toward that more general enquiry, one might start by considering some preliminary thoughts on two very different mental state requirements: namely, volitional act requirements in criminal law rather than copyright¹³⁴ and copyright's requirements for authorship rather than infringement.¹³⁵ One could coherently embrace the view that, although functionality is all that matters for volition in the context of copyright infringement, consciousness matters in both of these alternative legal contexts. For instance, one might argue that the punitive aims of criminal law ultimately require that those engaging in criminal conduct have a conscious experience of the actions in which they have engaged.¹³⁶ One might also argue that, because status as an author under copyright involves possessing rights of ownership in one's creative work, it ultimately requires personhood,¹³⁷ something which—the argument would go—requires possessing a conscious mind.¹³⁸ We neither defend nor reject either such line of argument, as to do so would involve distinct projects in their own right. Rather, we invoke these two additional contexts to illustrate the way such analyses might go and how they might differ from our primary example of copyright infringement, owing to the distinct aims and considerations at play in each context.

At this point, one might wonder about the availability of a general theory regarding when the law cares about conscious versus purely functional properties of mental states such that this framework need not be applied on a painstakingly case-by-case basis. Perhaps the search for such a theory is precisely where this Essay should lead future work. Nonetheless, as a preliminary hypothesis—one reacting to, and consistent with, the examples we have here discussed—it might be that the law is interested in conscious properties of mental states when it seeks to treat the actor in question as a rightsholder (such as in copyright authorship) or an autonomous and responsible agent (such as in criminal punishment). But in contexts in which the law is seeking simply to protect the rights or

134. See Moore, Act and Crime, *supra* note 11, at 44–46.

135. See, e.g., Shyamkrishna Balganes, Causing Copyright, 117 *Colum. L. Rev.* 1, 11–47 (2017) [hereinafter Balganes, Causing Copyright] (defending and analyzing the idea of “authorial causation” as a requirement for copyrightability).

136. See Samuel W. Buell & Lisa Kern Griffin, On the Mental State of Consciousness of Wrongdoing, 75 *Law & Contemp. Probs.*, no. 2, 2012, at 133, 139–44 (exploring how blameworthiness can justify a requirement of conscious awareness of wrongdoing); cf. Shapira-Ettinger, *supra* note 16, at 2578 (“A normative theory [of guilt in criminal law] stands in contrast to the dominant psychological theory of guilt . . . prevailing . . . in legal systems today. The focus of the psychological approach to guilt is on . . . the internal state of mind that reflects the kind of consciousness with which one acts.”).

137. Balganes, Causing Copyright, *supra* note 135, at 27 (“Given that authorship was invariably tied to ownership and the assertion of legal rights, it made little sense to speak of nonhuman authorship.”).

138. See generally Christopher Buccafusco, A Theory of Copyright Authorship, 102 *Va. L. Rev.* 1229 (2016) (defending a theory of authorship that requires intent: namely, the intention to produce mental effects in an audience).

interests of others from the actor (such as copyright infringement), functionality might be all that matters.¹³⁹ A thorough exploration or defense of this preliminary hypothesis is reserved for future work. But we hope this Essay has impressed the need to engage in such explorations and to wrestle with the fundamental questions surrounding the law's aims, in order to adapt the law to an increasingly machine-filled world.

CONCLUSION

In sum, it is a mistake to assume that machines can or should never satisfy implicit or explicit mental state requirements, entirely by virtue of the fact that they are machines. The law is not always or necessarily concerned with the existence of conscious experience or even with the full range of human-level functionalities. Instead, it will always be a substantive question what the law's various mental state requirements are aiming to track, one which depends on the interests and values at stake in the particular legal domain. It follows from this that, in adapting the law to a world with increasingly sophisticated technologies replacing the actions of humans, the challenge for the law is not that mental state requirements exist. Rather, it is that scholars and policymakers must start asking the normative questions of what such requirements are designed to achieve and therefore what relevant mental states must be.

139. Thanks to Erick Sam for suggesting this hypothesis in conversation.

Reproduced with permission of copyright owner. Further reproduction prohibited without permission.