2023

Artificial Intelligence and Contract Formation: Back to Contract as Bargain?

John Linarelli
University of Pittsburgh School of Law, jlinarelli@pitt.edu

ARTIFICIAL INTELLIGENCE AND CONTRACT FORMATION: BACK TO CONTRACT AS BARGAIN?

John Linarelli*

INTRODUCTION

I just booked a reservation at my favorite sushi restaurant using Google Assistant. Assuming a contract exists when one books a table at a restaurant,[1] Google Assistant can now enter contracts for humans using a natural language interface. The restaurant maître de (if sushi restaurants have these) heard a bot with a natural language interface. Listen: https://www.youtube.com/watch?v=-RHG5DFAjp8. Could you tell if this bot[2] is a human? For the limited function of booking a table and other clearly defined tasks, Google Assistant probably passes the Turing Test. This bot is the result of Google's Duplex software.[3] This technology is not artificial general intelligence. It is artificial narrow intelligence – AI[4] designed for the task of restaurant booking. But advances of AI of this kind now commonplace.

This chapter makes two claims. One is methodological—about how to think about the participation of AI in contracting processes. The other is a substantive claim about how AI

* Professor of Law, University of Pittsburgh School of Law.

[1] The elements of offer and acceptance do appear to exist in such cases, though it is unclear whether an intention to create a binding contract exists, given that it is customary that restaurants do not demand compensation for no shows. The reason for not seeking compensation probably has to do with transaction costs and the lack of loss because restaurant tables tend to get filled anyway at the sorts of places that accept reservations.

[2] Listen to this clip. Magikmaker, *Google Assistant Calling a Restaurant for a Reservation*, YOUTUBE (May 9, 2018), https://www.youtube.com/watch?v=-RHG5DFAjp8 [https://perma.cc/3Q9P-SRWG].

[3] *See* Yaniv Leviathan & Yossi Matias, *Google Duplex: An AI System for Accomplishing Real-World Tasks Over the Phone*, GOOGLE RESEARCH, https://ai.googleblog.com/2018/05/duplex-ai-system-for-natural-conversation.html (last visited Nov. 20, 2022) [https://perma.cc/CGF9-E6ML].

[4] From here on, "AI" will refer to "artificial intelligence."

has the potential to change the nature and institution of contract itself, though we have a way to go before major changes occur.

Part I sets forth the methodological claim. It explains how to approach questions about when AI agency is relevant. "Agency" here refers to human agency, not legal agency, or perhaps more accurately, the agency of a being or entity. This sort of philosophical or psychological agency is key to understanding contract law. Without agency of this sort, no contracting can occur. An agent in a philosophical sense is a being whose mental states (beliefs, attitudes, intentions) play a causal role in making that being act in a particular way according to those mental states.[5] Although this conception of agency appears in analytical philosophy to describe human agency, it is broad enough to describe agency for non-human animals and some AI. The methodological claim is this: the best way to determine how to produce legal rules to govern how AI may or may not interact with humans, or have consequences for human flourishing, is to start from the bottom up—that is, starting with a conception of the agent known as an AI and the kind of agency that produces action by an AI. All moral philosophy does this too and so it is puzzling why so many who write on AI take a top-down approach, starting the discussion with the sorts of ethical properties an AI should have or the AI "rights" or the rights of humans affected by AI decision making. Some may argue that top-down approaches to AI ethics makes AI ethics useless.[6] The same may be true for examining the role of AI from a legal perspective without exploring the nature of AI itself.

[5] This is at least accepted by many philosophers and psychologists. Of course, there is always disagreement and lots of discussion about such questions in philosophy and related fields and here is not the place to hash out these disagreements and discussions. *See* Juan S. Piñeros Glasscock & Sergio Tenenbaum, "Action," *Stanford Encyclopedia of Philosophy* (2023), https://plato.stanford.edu/entries/action/#Rel.

[6] *See* Luke Munn, *The Uselessness of AI Ethics*, AI ETHICS, Aug. 2022 (arguing that AI ethical principles are useless but making a different argument about what is useful).

Part II provides a brief history of the evolution of AI in contract practices. It sets the stage for Part III that examines the possible directions AI may go in terms of contract formation practices.

Part III explores an important, recent advance in AI that has the potential to push AI in contracting to a new level: the large language model. Such models are a significant advance. Part III explores the limits on the use of large language models in contract practices.

Part IV explores the next steps for AI to develop a more sophisticated agency by allowing for more interactions with human agents in an effort to come closer to what humans understand as acting with human-like intentionality. Further advances are needed for AI to come to contract practices in a way that will release it from its current purgatory of what Margaret Jane Radin calls "contract as product" to the older vintage "contract as bargain."[7] While Radin further distinguishes between contact as product and contract as consent, this chapter does not deal with consent because it presents a host of complex questions about power and cognition unnecessary for the purposes of this chapter. The focus in part IV is only on whether and how the move to more advanced forms of AI might open a window for aligning contract law to its pre-information age archetype of "contract as bargain," as AI becomes more "human" in its interactions with us. Welcome back contract law. Get lost click wrap. This is clearly an exaggeration, for now. Part IV deals with the potential for a transformative change as AI develops its capacity to interact with humans at the transactional level.

[7] Margaret Jane Radin, *The Deformation of Contract in the Information Society*, 37 OXFORD J. OF LEGAL STUD. 505, 530-31 (2017); Margaret Jane Radin, *Humans, Computers, and Binding Commitment*," 75 IND. L. J. 1125, 1126 (2000) [hereinafter Radin on Humans and Computers].

I. FROM THE BOTTOM UP

All contract theory presupposes a conception of how humans think. In philosophy, this presupposition is sometimes understood as a "conception of the person," but it is partly about how close human thought reflects the cognitive capacity to act rationally, with rationality conceptualized in different ways. Rational choice theory defines its field with an explicit reference to assumptions about human thought. It offers a theory of human action based on instrumental rationality.[8] Behavioral economics incorporates insights from cognitive psychology on predictable properties of human thought differing from the rational choice model and that have consequences for human action.[9] Moral theories about contract law depend on a conception of human thought too. Kantians accept a much different kind of human rationality than the simple instrumental rationality of economics to argue that humans as free beings must accept the moral law as a rational necessity.[10] Although oftentimes theories and philosophies about contract law may not clearly elucidate presuppositions about how humans think, such presuppositions, as a matter of necessity, have to exist. We cannot really say anything about either how the law affects human behavior or why humans justify the law without holding some view, even if unexamined, about how humans think, or to use a fancier word, cognition.

How one approaches this first question about understanding the nature of cognition will shape how one's consideration of AI and its "permission" in a legal system to interact with humans will proceed. One way to approach the AI-human interface is to argue that AI must

---

[8] For a classic exposition, *see* GARY S. BECKER, THE ECONOMIC APPROACH TO HUMAN BEHAVIOR (University of Chicago Press 1976).

[9] *See* Daniel Kahneman & Amos Tversky, *Choices, Values, and Frames,* 39 AMERICAN PSYCHOLOGIST 341 (1984).

[10] ANDREWS REATH, AGENCY AND AUTONOMY IN KANT'S MORAL THEORY (Cambridge University Press 2011).

meet a particular external ethical standard to qualify for interaction with humans beyond the very basic level of interaction.[11] The argument goes like so: humans have the capacity to be moral creatures. We, therefore, have to design AI that interacts with humans to share these same features. Pick your theory. We want an AI "programmed" to comply with Aristotelian ethics. Or we want it programmed to comply with Kant's categorical imperative. We also need it programmed to respect human rights. Or it would be nice, if in the appropriate cases, AI could weigh costs and benefits of a particular action and act in the best interests of humans. Or we need AI to be designed to take a contractual approach, or a consequentialist one. Or we do not see ethics as one big item and are fine with a mix of principles to be "programmed into" AI.

These kinds of "top-down" investigations matter in articulating values that AI designers should appreciate; however, they tend to be more in the nature of general directives with no real focus on how AI might make decisions or even interact with humans in some form of natural language interface. In the non-ideal world, humans, moreover, do not "comply" with the dictates of the above moral theories either. These theories speak more to the capacity of humans to act morally, but, if we want to understand actual human behavior, we have to turn to psychology.

This chapter takes a "bottom up" approach. Its focus is cognitive and behavioral, not normative or moral. That is, it examines the nature of the person or entity asked to respond to moral principles or legal commands. The focus here is on the way that AI does and can operate in the world. Moral disagreement is pervasive among humans. How we would "program" AI to do something we have not figured out for ourselves would be quite a task,

[11] E.g., Nesibe Kantar & Terrell Ward Bynum, *Flourishing Ethics and Identifying Ethical Values to Instill into Artificially Intelligent Agents*, 53 METAPHILOSOPHY 599, 602 (2022).

which I leave for others. We need to understand how forms of communication between human and artificial forms of life will interact with one another in a contractual context. This is a finer grained approach, influenced by analytical philosophy and less by competing philosophical approaches.[12] Once the analytics are understood, the moral questions will become more tractable and apt for concrete solutions in the actual interactions between AI and humans in a contract (or any) context.

II. GETTING CLEAR ON AI IN CONTRACT PRACTICES

What role has AI had in contract practices? An area of confusion seems to stem from conceptualizing an essential feature of AI as "computational" or "automated." Such an overly broad understanding of AI means that it extends to any computer programming or "coding" embedded in contract practices. The concepts of "computation" and "automation" do not provide sufficiently precise distinctions between traditional if-then software coding at the symbolic level and algorithms associated with machine learning. Coding is logic. Machine learning is statistical. AI, in a machine learning sense, is a Bayesian learning machine that writes its own "code" based on consuming vast quantities of data.[13] AI is not human-directed at the single-decision level, which means that the blockchain or distributed ledger is not AI nor are many smart contracts.

Another way to approach this question is to attempt to define or at least describe AI. For this chapter, we can use a conception of AI as "getting machines"[14] to act in ways that

[12] Of course, these fields are not entirely distinct. There is a significant body of work in ethics in the analytical tradition. For a survey, see STEPHEN P. SCHWARTZ, A BRIEF HISTORY OF ANALYTICAL PHILOSOPHY: FROM RUSSELL TO RAWLS (Wiley-Blackwell 2012).

[13] For a basic review of these points, *see* PEDROS DOMINGOS, HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD (BASIC BOOKS 2015).

[14] NILS J. NILSSON, THE QUEST FOR ARTIFICIAL INTELLIGENCE: A HISTORY OF IDEAS AND ACHIEVEMENTS xiv (Cambridge University Press 2010) (on how the computer has greatly expanded our notion of what a

depend on cognitive functions, such as learning and problem-solving, in a way humans could recognize.[15] Note the focus on action. Trying to conceptualize whether an entity is "thinking" may lead to a variety of problems and corresponding solutions that do not assist us very much in understanding AI. In his influential work on the history of AI, Nils Nilsson describes AI as "that activity devoted to making machines intelligent," with "intelligence" being "that quality that enables an entity to function appropriately and with foresight in its environment."[16] A being acting with foresight in its environment would seem capable of offering the sort of manifestation of mutual assent and intent necessary for an entity to constitute a contracting agent if they can act with foresight in transacting with humans in a contractual context. Nilsson's definition works well for understanding how AI might work in contracting – and differ from other sorts of computer technologies deployed in the contracting process – because of its focus on functionality and foresight, or what one might consider agency in contract formation. But we should note that it takes a particular kind of AI to be a transactional agent, a subject I have addressed elsewhere,[17] and what "acting with foresight" means is itself a contested concept.

Symbolic AI, more colloquially known as "good old-fashioned artificial intelligence" or GOFAI for short, is an early form of AI.[18] It depends on symbolic (human readable)

machine is, how software alone is often referred to as a "machine," and how the distinction between hardware and software has become blurred).

[15] S. Matthew Liao, *A Short Introduction to the Ethics of Artificial Intelligence*, *in* ETHICS OF ARTIFICIAL INTELLIGENCE 3 (S. Matthew Liao ed., Oxford University Press 2020) (suggesting that "we can broadly understand AI as getting machines to do things that require cognitive functions such as thinking, learning, and problem-solving when done in intelligent beings such as humans").

[16] *See* NILSSON, *supra* note 14, at xiii.

[17] John Linarelli, *A Philosophy of Contract Law for Artificial Intelligence: Shared Intentionality*, *in* CONTRACTING AND CONTRACT LAW IN THE AGE OF ARTIFICIAL INTELLIGENCE 59 (Martin Ebers, Christina Poncibò & Mimi Zou eds., Hart 2022) [hereinafter Linarelli on Shared Intentionality]; *see also* John Linarelli, *Artificial General Intelligence and Contract*, 24 UNIF. L. REV. 330 (2019) [hereinafter Linarelli on AGI and Contract].

[18] John Haugeland, *Farewell to GOFAI?*, *in* SPEAKING MINDS: INTERVIEWS WITH TWENTY EMINENT COGNITIVE SCIENTISTS 101, 105 (Peter Baumgartner & Sabine Payr eds., Princeton University Press 1995).

representations of problem-solving and the rules of formal logic.[19] Symbolic AI requires good old-fashioned programming in the form of if-then statements for every step in a chain of reasoning to solve a problem. It is the realm of the logician,[20] in which code was designed to produce deductive forms of reasoning for every conceivable task or problem asked of the AI. Symbolic AI was the main form of AI until the mid-1980s. It does not involve machine learning. The machine or system does not learn from data, but instead must have all its decisions explicitly programmed in advance. Logic comes first, data comes second: the opposite direction of most AI today, for which data comes first and logic second. The most advanced forms of AI today rely on the notion of inductive inferences drawn from massive amounts of data.[21]

Symbolic AI, now understood to be an "old" form of AI vastly outperformed by machine learning, is clearly not symbolic, deductive, or logical in approach. However, symbolic AI still has many uses today in contracting. For instance, one application of symbolic AI in use today is the smart contract.[22] Of course, it might be more accurate to say that some smart contracts rely on the most rudimentary form of symbolic AI, while others may incorporate some limited aspects of machine learning. Smart contracts clearly do not, however, rise to the level where we might see AI as a transactional agent capable of serving as a tool for human agents or contract parties.

---

[19] NILSSON, *supra* note 9, at 53-54.
[20] Nilsson *supra* note 14, at 331-346; DOMINGOS, *supra* note 13, at 30, 49, 80-83 (Basic Books 2015).
[21] *Id*.
[22] "At present, the input parameters and the execution steps for a smart contract need to be specific and objective. In other words, if 'x' occurs, then execute step 'y.' Therefore, the actual tasks that smart contracts are performing are fairly rudimentary, such as automatically moving an amount of cryptocurrency from one party's wallet to another when certain criteria are satisfied." Stuart D. Levi & Alex B. Lipton, *An Introduction to Smart Contracts and Their Potential and Inherent Limitations*, HARV. L. SCH. F. on CORP. GOVERNANCE (May 26, 2018), https://corpgov.law.harvard.edu/2018/05/26/an-introduction-to-smart-contracts-and-their-potential-and-inherent-limitations/ [https://perma.cc/7ZPL-5FGG].

The term "smart contract" is ambiguous. It is not a legal concept. Very simply, a smart contract is a contract for which some or all contract performance will have both digital execution and enforcement without the need for human intervention except at the level of writing code to automate contract performance.[23] Distributed ledger technology has advanced substantially the ability of contract parties to write and use smart contracts. The combination of the distributed ledger, the network, and the consensus mechanisms built into distributed ledger technology facilitate trust between contract parties and replace humans in institutions operating as intermediaries.[24] Note that this trust in the blockchain is a very limited form of trust: trust only that the transaction will execute in accordance with the instructions of the contracting parties. This system differs from a broader form of trust in terms of the probity or honesty of the contracting agent or value of the contract itself.[25] In short, smart contracts substitute algorithms for human contract performance and enforcement.

Symbolic AI cannot rise to the level of "agent" if we understand an agent to be a being with the capacity to act with foresight or intentionality. The reasons why pertain to the

[23] Various authors have offered definitions of a smart contract, but Nick Szabo is credited with inventing the phrase. Kevin Werbach & Nicolas Cornell, *Contracts* Ex Machina, 67 DUKE L. J. 102 (2017). Szabo defines a smart contract as a "set of promises, specified in digital form, including protocols within which the parties perform on these promises." NICK SZABO, SMART CONTRACTS: BUILDING BLOCKS FOR DIGITAL MARKETS (1996), http://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart_contracts_2.html [https://perma.cc/7D7D-MSSQ]. Max Raskin describes smart contracts as "agreements wherein execution is automated, usually by computers." Max Raskin, *The Law and Legality of Smart Contracts*, 1 GEO. L. TECH. REV. 305, 306 (2017); Werbach and Cornell define a smart contract as an "agreement in digital form that is self-executing and self-enforcing." Werbach & Cornell at 320. Jeffery Lipshaw describes a smart contract as "simply computer code that automatically execute agreed-upon transactions." Jeffrey M. Lipshaw, *The Persistence of 'Dumb' Contracts*, 1 STAN. J. BLOCKCHAIN L. & POL'Y 1, 4 (2019).

[24] Werbach & Cornell, *supra* note 23, at 330.

[25] As recent debacles in the crypto world confirm. *See* Erin Griffith, *Why the Crypto Collapse Matters: The Failure of the Cryptocurrency Exchange FTX Put the Entire Industry Under Scrutiny*," N.Y. TIMES, (Nov. 17, 2022), https://www.nytimes.com/2022/11/17/briefing/crypto-collapse-ftx.html [https://perma.cc/VST5-QHSX].

limits of symbolic logic. Contracts in the classical or traditional sense come from the perspective of mutual assent based on the shared meanings of contractual language.[26] This conception of contract is now under considerable threat in the form of automated contracts of adhesion between consumers and firms. After all, such contracts move the point of normative significance for contract formation from mutual assent to constructive notice, or in Radin's terminology, have resulted in a shift from contract as consent to contract as product.[27] But GOFAI cannot reach the level of acting with foresight or intentionality necessary to provide the sort of shared intentionality required for a contract to form.[28]

Machine learning is AI that can learn and adapt without following explicit coding or instructions by using learning algorithms and statistical methods to draw inferences from patterns in data. The result is that a learning algorithm produces another algorithm to perform a particular task or solve a particular problem. In machine learning, a computer program writes its own computer program in an iterative process through the study of large amounts of data.

Machine learning is in ubiquitous use in commercial and financial contexts to detect fraud, to conduct automated trading, to fill in price terms with dynamic pricing, to provide financial advisory services, to identify risks and predictive analytics for construction

[26] *See* RESTATEMENT (SECOND) OF CONTS. §§ 18-20, 205 (AM. L. INST. 1981).
[27] *See* Robin Bradley Kar & Margaret Jane Radin, *Pseudo-Contract and Shared Meaning Analysis*, 132 HARV. L. REV. 1135 (2019); MARGARET JANE RADIN, BOILERPLATE: THE FINE PRINT, VANISHING RIGHTS, AND THE RULE OF Law (Princeton University Press 2014). The recent debates about the new Restatement of the Law of Consumer Contracts focused on the move from mutual assent to notice as an animating principle for contract formation and enforceability; for a summary of the legal (as opposed to empirical) issues associated with this move, see Melvin Eisenberg, *The Proposed Restatement of Consumer Contracts, if Adopted, Would Drive a Dagger Through Consumers' Rights*, YALE J. LAW ON REGUL. NOTICE AND COMMENT BLOG (Mar. 20, 2019) https://www.yalejreg.com/nc/the-proposed-restatement-of-consumer-contracts-if-adopted-would-drive-a-dagger-through-consumers-rights-by-melvin-eisenberg/ [https://perma.cc/C8KJ-TL2P].
[28] Linarelli on Shared Intentionality, *supra* note 17.

among other uses.[29] Perhaps most significantly, machine learning has made firms, or what we know as "merchants" in the Uniform Commercial Code,[30] ever more powerful contract parties because they can exploit the use of big data to take advantage of significant information asymmetries when contracting with consumers.[31] In these contexts, humans use machine learning algorithms to determine who to contract with and on what terms. For example, Amazon and Uber use dynamic pricing.[32] Perhaps the most sophisticated machine learning ongoing right now in the contracting context is Google's ad exchange, which fills in ad space on websites in milliseconds as the website loads.[33] Still, the framework for these contracts is human-produced.

Based on information available on the most recent AI innovations, machine learning has not progressed to the point where it can possess the sort of agency required for contract formation on its own without instruction from human agents. The simple example in the Introduction of making a restaurant reservation with Google Assistant suggests that we are close, but only for very simple contracts. In that example, a human restaurant customer delegates the task of natural language interface to a narrow form of AI designed only to make restaurant reservations. The contract terms are basic: date, time, restaurant, and number of parties. One could argue that this is too simple a task to suggest any significant

[29] Lauren Henry Scholz, Algorithmic Contracts 20 STAN. TECH. L. REV. 128 (2017); OECD, ARTIFICIAL INTELLIGENCE IN SOCIETY, (OECD Publishing 2019); Martin Ebers, *Regulating AI and Robotics: Ethical and Legal Challenges, in* ALGORITHMS AND LAW (Martin Ebers & Susana Navas Navarro eds., Cambridge University Press 2020).
[30] U.C.C. § 2-104.
[31] *See* Stacy-Ann Elvy, *Contracting in the Age of Internet of Things: Article 2 of the UCC and Beyond* 44 HOFSTRA L. REV. 839 (2016).
[32] Alexander Shartsis, *Dynamic Pricing: The Secret Weapon Used by the World's Most Successful Companies*, FORBES (Jan. 8, 2019), https://www.forbes.com/sites/forbestechcouncil/2019/01/08/dynamic-pricing-the-secret-weapon-used-by-the-worlds-most-successful-companies/?sh=4024be35168b [https://perma.cc/Z9K9-XDL3].
[33] Dina Srinivasan, *Why Google Dominates Advertising Markets Competition Policy Should Lean on the Principles of Financial Market Regulation*, 24 STAN. TECH. L. REV. 55, 78 n. 45 (2020).

progress towards an AI with the sort of contract intentionality required. But there is progress on the horizon in AI technology that may lead to movement towards richer forms of collaboration between humans and AI in contract formation. Part III will address these advances, which have come in the form of large language models (LLMs).

III. WHERE ARE WE HEADED WITH AI IN CONTRACT FORMATION?

A common cognitive error that humans make is to assume that more advanced forms of self-awareness and acting with intention presupposes the use of natural language. When we hear natural language, we attribute "smart," "intelligence," "sentience," or "consciousness" to the speaker. In contract law, we can avoid these more complex concepts in favor of the simpler concept of intention in an objective sense.[34] But still, humans are prone to

---

[34] The objective theory of contract tells us that intention to be bound to or form a contract is determined by evidence external to the actual intentions of the parties. Judge Learned Hand has said:

> A contract has, strictly speaking, nothing to do with the personal, or individual, intent of the parties. A contract is an obligation attached by the mere force of law to certain acts of the parties, usually words, which ordinarily accompany and represent a known intent. If, however, it were proved by twenty bishops that either party, when he used the words, intended something else than the usual meaning which the law imposes upon them, he would still be held, unless there were some mutual mistake, or something else of the sort.

Hotchkiss v. National City Bank, 200 F. 287, 293 (S.D.N.Y. 1911), *aff'd,* 201 F. 664 (2d Cir. 1912), *aff'd,* 231 U.S. 50 (1913). Judge Frank Easterbrook has explained that intention to be bound "does not invite a tour through [a contract party's] cranium." Skycom Corp. v. Telstar Corp., 813 F.2d 810, 814 (7th Cir. 1987). Instead, it "must necessarily be derived from a consideration of the words, written and oral," and "actions" of the parties. *Id*. (quoting Household Utils., Inc. v. Andrews Co. 236 N.W.2d 663, 669 (Wis. 1976)). Often quoted on the objective theory of contract formation and interpretation is the New Hampshire Supreme Court, itself quoting Oliver Wendell Holmes Jr.:

> A contract involves what is called a meeting of the minds of the parties. But this does not mean that they must have arrived at a common mental state touching the matter at hand. The standard by which their conduct is judged and their rights are limited are not internal but external. In the absence of fraud or incapacity, the question is: What did the party say and do? "The making of a contract does not depend upon the state of the parties' minds; it depends upon their overt acts."

Woburn Nat'l Bank v. Woods (quoting 1 OLIVER WENDELL HOLMES, JR., THE COMMON LAW 307 (Little, Brown, & Co. 1881)). The U.S. Restatement (Second) of Contracts, intended to reflect a consensus about contract law in the United States, does not contain any section explicitly titled on intention to form a contract. It advises us that American contract law has likely abolished the idea of intention to be legally bound. Its section 21 provides that "neither real nor apparent intention that a promise be legally binding is essential to the formation of a contract," but the objective theory of contract still prevails. RESTATEMENT (SECOND) OF CONTS. § 21 (AM. L. INST. 1981). American contract law relies on what is known in American law as manifestation of mutual assent, which requires each party either to commit, objectively understood, or perform. *Id*. Again, mutual assent is objectively determined. *Id*. While English law does not reflect the language in the

accept the unreliable view that something speaking to us in a natural language must know the topic about which it speaks and have some capacity to take responsibility for its actions based on the intentions expressed in natural language terms. We do the same with robots and AI more generally that converse or look like us or move in particular ways like us.[35] In other words, we are fooled. This deception problem partly explains why the Turing Test has increasingly fallen into disuse as a test for a machine's ability to exhibit intelligent behavior indistinguishable from that of a human.[36] Turing himself called the test an imitation game.[37]

One of the more recent AI technology innovations that has the potential to fool us if we are not more reflective about our interactions with AI is the large language model. A large language model is AI that has access to very large datasets of texts that it uses to express natural language in remarkably human-like sequences.[38] A language model is, in short, a

---

restatement on manifestation of mutual assent, it is substantially similar in adhering to an objective theory of contract formation and interpretation. *See, e.g.*, EDWIN PEEL, TREITEL ON THE LAW OF CONTRACT 1 (Sweet & Maxwell 14th ed. 2015) (in particular, section 1-002, which discusses "[t]he objective principle"). In English law, intention to create legal relations is traditionally used to distinguish promises the parties want the law to enforce and promises they do not want the law to enforce. *Id*. The intent or manifestation of mutual assent to enter a contract is thus not a matter of investigating some inner mental operations of a contract party, but about whether one contract party can reasonably conclude that the other contact party has the requisite objective intent. *Id*.

[35] There is significant research on this point. *E.g.*, Cindy M. Grimm, T*he Danger of Anthropomorphic Language in Robotic AI Systems*, TECH STREAM (June 18, 2021), https://www.brookings.edu/techstream/the-danger-of-anthropomorphic-language-in-robotic-ai-systems/ [https://perma.cc/X2H9-W4AU]; Arleen Salles, Kathinka Evers, & Michele Farisco, *Anthropomorphism in AI*, 2020 AJOB NEUROSCIENCE 88 (2020).

[36] Matthew Sparkes, *Google Wants to Challenge AI with 200 Tasks to Replace the Turing Test*, THE NEW SCIENTIST (June 14, 2022), https://www.newscientist.com/article/2323685-google-wants-to-challenge-ai-with-200-tasks-to-replace-the-turing-test/ [https://perma.cc/JHY9-GR7B]; Will Oremus, *Google's AI Passed a Famous Test – and Showed How the Test Is Broken*, THE WASH. POST (June 17, 2022), https://www.washingtonpost.com/technology/2022/06/17/google-ai-lamda-turing-test/ [https://perma.cc/BR7H-9EYD]; Gary Marcus, *The Search for a New Test of Artificial Intelligence*, SCI. AM. (Mar. 1, 2017), https://www.scientificamerican.com/article/the-search-for-a-new-test-of-artificial-intelligence/ [https://perma.cc/N482-QHLB]; Gary Marcus, Francesca Rossi, & Manuela Veloso, *Beyond the Turing Test*, AI MAG., Apr. 2016, at 3; Gary Marcus, *What Comes After the Turing Test?*, THE NEW YORKER (June 9, 2014), https://tinyurl.com/yrnb68ac [perma.cc/DYK7-GBJ8].

[37] Alan M. Turing, *Computing Machinery and Intelligence*, 49 MIND 433, 433 (1950) ("The new form of the problem can be described in terms of a game which we call the 'imitation game.'").

[38] *See* Blaise Agüera y Arcas, *Do Large Language Models Understand Us?*, 151 DAEDULUS 183 (2022).

probability distribution over a sequence of words. Some AI researchers have expressed sanguine views on the potential of large language models to reach the level of artificial general intelligence. Ray Kurzweil, for example, has expressed the view that large language models have made a significance advance towards artificial general intelligence.[39]

Artificial general intelligence refers to the ability of an AI agent to understand or learn any task that a human can, across a range of domains.[40] It differs from artificial narrow intelligence, which expresses intelligence across a specific domain.[41] For example, Google's AlphaGo can defeat humans at the game of Go, and according to Google "is arguably the strongest Go player in history,"[42] but AlphaGo cannot play chess or even checkers. As I have argued elsewhere, artificial general intelligence is probably needed for an AI to become a full transactional agent on its own in a contract setting.[43]

The main large language models in existence today are Google's Language Model for Dialogue Applications, known as LaMDA—a sophisticated natural language "conversational" AI;[44] OpenAI's Generative Pre-Training Transformer—known as GPT-3;[45] and now very recently, OpenAI's ChatGPT and advances on it.[46] Some would say that Google

---

[39] Lex Fridman, *Ray Kurzweil: Singularity, Superintelligence, and Immortality*, YOUTUBE (Sept. 17, 2022), https://www.youtube.com/watch?v=ykY69lSpDdo [https://perma.cc/24UQ-26XY]; CHIP at Boston Children's Hospital, *CHIP Landmark Ideas: Ray Kurzweil*, YOUTUBE (Dec. 20, 202), https://www.youtube.com/watch?v=KklEmSBlUcM [https://perma.cc/BB2Z-T6MQ].

[40] *See* Ben Goertzel, *Artificial General Intelligence: Concept, State of the Art, and Future Prospects*," 5 J. ARTIFICIAL GEN. INTEL., 1 (2014). However, the definitive text defining these concepts comes from a Russell and Norvig book. STUART RUSSELL & PETER NORVIG, ARTIFICIAL INTELLIGENCE: A MODERN APPROACH (Pearson 4th ed. 2020).

[41] *Id*. at 98.

[42] AlphaGo, *What is Go*?, DEEPMIND, https://www.deepmind.com/research/highlighted-research/alphago (last visited Mar. 19, 2023) [https://perma.cc/5BXZ-YD4V].

[43] Linarelli on AGI and Contract, *supra* note 17.

[44] Eli Collins & Zoubin Ghahraman, *LaMDA: Our Breakthrough Conversation Technology*, GOOGLE: THE KEYWORD (May 18, 2021), https://blog.google/technology/ai/lamda/ [perma.cc/9GLJ-FVZM].

[45] Kevin Roose, *How ChatGPT Kicked Off an A.I. Arms Race*, N.Y. TIMES (Feb. 3, 2023), https://www.nytimes.com/2023/02/03/technology/chatgpt-openai-artificial-intelligence.html [https://perma.cc/DS6S-R6ZD].

[46] *Id*.

undersells LaMDA in its characterization of it as a "breakthrough conversation technology."[47] These are remarkable chat bots. ChatGPT has been all the rage as this chapter is being written. Deepmind, Google's AI subsidiary, claims that its GATO is a "generalist agent," which means that it goes beyond chatting and can play video games, write, and control a robotic arm.[48] While GATO reflects an attempt to demonstrate that the move towards artificial general intelligence may indeed be possible, at least one commentator has described it as "mediocre" at its tasks.[49]

In mid-2022, former Google engineer Blake Lemoine informed the *Washington Post* that he believed LaMDA was sentient and had achieved consciousness.[50] According to the *Post*, Lemoine said, "I know a person when I talk to it" and "[i]t doesn't matter whether they have a brain made of meat in their head. Or if they have a billion lines of code. I talk to them. And I hear what they have to say, and that is how I decide what is and isn't a person." According to the *Post*, Lemoine concluded LaMDA was a person in his capacity as a mystic Christian priest and not as a scientist.[51] He referred LaMDA to an attorney.[52]

---

[47] *See* Collins & Ghahraman, *supra* note 44; *see also* CNET Highlights, *Watch Google's AI LaMDA Program Talk to Itself at Length (Full Conversation)*, YOUTUBE (May 18, 2021) https://tinyurl.com/mznrd7b2 [https:///perma.cc/K9NZ-34WK]

[48] Scott Reed, et al., *A Generalist Agent*, TRANSACTIONS MACH. LEARNING RSCH. Nov. 2022, at 1.

[49] Tiernan Ray, *DeepMind's 'Gato' Is Mediocre, So Why Did They Build it?*, ZDNET (May 14, 2022), https://www.zdnet.com/article/deepminds-gato-is-mediocre-so-why-did-they-build-it/ [https://perma.cc/LM42-AC2K] (asserting that "DeepMind's program is a generalist, to test the notion that over time, greater computing power will win in AI.").

[50] Nitasha Tiku, *The Google Engineer Who Thinks the Company's AI Has Come to Life*, THE WASH. POST (June 11, 2022), https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/ [https://perma.cc/454B-23QM]. Google subsequently suspended and then eventually terminated Mr. Lemoine's employment. Nico Grant, *Google Fires Engineer Who Claims Its A.I. Is Conscious*, N.Y. TIMES (July 23, 2022), https://tinyurl.com/yhyv9d4k [https://perma.cc/3S4S-JATB].

[51] Tiku, *supra* note 50.

[52] Jon Christian, *Google Engineer Says Lawyer Hired by "Sentient" AI Has Been "Scared Off" the Case*, FUTURISM (June 22, 2022), https://futurism.com/google-engineer-lawyer-representing-ai-scared-off [https://perma.cc/28TC-PBYE].

A large language model is not your ordinary chat bot. Try running the query "what is the meaning of life" on Google Assistant on your cellphone. What you get in response is a canned answer of how Bill and Ted in the movie, Bill and Ted's Excellent Adventure, answered the question. We can readily recognize this answer as a joke by Google engineers. But if you ask ChatGPT or LaMDA this question, you will get a very credible response. Think of how poorly Google Assistant responds to simple commands. It fails rather dramatically with simple conversational implicature, a well understood aspect of human natural language.[53] Large language models have mastered conversational implicature.

Do large language models succeed at conferring capabilities onto AI that would make them acceptable as contracting agents? No, or at least not yet. They are, using the phrase from an article co-authored by two former Google AI ethics chiefs that Google terminated, "stochastic parrots."[54] A large language model does not understand the true meaning of words. It is "mindless,"[55] though one has to question what form of "mind" is at issue here. It analyzes the statistics of language in a disembodied context. Large language models lack the capacity to use their intelligence in a phenomenological sense of using multiple senses to engage in social forms of communication and interaction. As linguist Emily Bender has explained, "[w]e now have machines that can mindlessly generate words, but we haven't learned how to stop imagining a mind behind them."[56] It has no capacity for shared or collective intentionality. A large language model is not an agent with plans and intentions.[57]

---

[53] PAUL GRICE, STUDIES IN THE WAY OF WORDS 22-40 (Harvard University Press 1989)
[54] Emily M. Bender, Angelina McMillan-Major, Timnit Gebru, & Shmargaret Shmitchell, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, ACM CONF. ON FAIRNESS, ACCOUNTABILITY & TRANSPARENCY Mar. 2021, at 610.
[55] David Chalmers, *GPT-3 and Intelligence*, DAILYNOUS https://dailynous.com/2020/07/30/philosophers-gpt-3/#chalmers (last visited Nov. 22, 2022) [https://perma.cc/8SD2-WT6N].
[56] Tiku, *supra* note 50.
[57] Chalmers, *supra* note 55.

A large language model lacks a coherent identity. It does not hold beliefs across contexts.[58] Shobita Parthasarathy has provided a description that cuts through the fictions about the potential of large language models. She describes large language models as "artificial intelligence tools that can read, summarize and translate texts and predict future words in a sentence letting them generate sentences similar to how humans talk and write."[59]

In short, a large language model is a talker, not a doer. It is not a "true" agent in the philosophical sense of bearing responsibility for its actions. It does not engage in action that has practical or moral consequences. It uses language for conversation, but has no cognitive capacity to use language (or action) in the form of a binding commitment. Perhaps in the future it will have this capacity. These models might go in unexpected directions. But its mere ability to engage in more natural language processing cannot determine whether it does achieve it.

So, could AI have the potential to advance us to a contract as bargain context? What improvements to technology would need to happen for this paradigm change to occur? Right now, when computing power is involved in contracting, we tend to be in a standard form or adhesion context. This tendency is obviously not always true, as B2B contracting may deploy AI to engage in significant functions at this stage. In Radin's words, the "human/computer interface" has existed for some time in contract contexts.[60] But we have not yet reached the point where AI could somehow substitute for a human or collaborate with

---

[58] Amanda Askell, *GPT-3: Towards Renaissance Models*, DAILYNOUS https://dailynous.com/2020/07/30/philosophers-gpt-3/#chalmers (last visited Nov. 22, 2022) [https://perma.cc/8SD2-WT6N].

[59] Press Release, Shobita Parthasarathy, Professor of Pub. Pol'y, Univ. of Mich., Parthasarathy Discusses Implications of Large Language Models (Nov. 7, 2022), https://fordschool.umich.edu/news/2022/parthasarathy-discusses-implications-large-language-models [https://perma.cc/ZJ5D-3252].

[60] Radin on Humans and Computers, *supra* note 7, at 1125.

a human to produce the requisite intentional states that agents need to be contracting agents. Part IV will turn to what needs to change for AI to advance to this stage: it is a move from "tool" to "agent."

IV. GETTING TO CONTRACT AS BARGAIN

AI seemingly needs to make at least two advances to possess the cognitive capacity to become a true transactional agent.[61] AI will need to possess the capacity to form the requisite objective intent to create a contract. Intent is a necessary condition, but it is not sufficient on its own. AI will also need to possess the capacity to be subject to law; to have this capacity, it will have to internalize a notion of practical authority, the idea that legal rules provide reasons for an agent to act in a particular way because they are law. It will also need the capacity to recognize and accept its own accountability and the accountability of others for promises and expectations in the appropriate social and legal contexts. We explore these requirements here.

A. *Shared Intentionality as a First Step*

The presence of AI in contracting requires us to compare human and artificial mental operations, also known as functionally. We may need to accept that what makes something a mental state does not depend on its internal constitution, but instead, on the way it functions and the role it plays in the system in which it is a part.[62]

Philosophers and cognitive scientists have identified a set of mental capacities of humans, consisting of the capacity to explain and predict human behavior, attribute mental

[61] Avoided here is discussion of granting AI any form of "personhood." Such a question is unnecessary for our discussion and raises a host of issues beyond our scope here.
[62] For an accessible explanation of functionalism, see DANIEL C. DENNETT, CONSCIOUSNESS EXPLAINED 30-32 (Little, Brown, & Co. 1991).

states to other humans, and explain the behavior of humans in terms of mental states. Having this capacity is necessary to understand how intent to enter a contract operates, as that intent requires externally manifestation and attribution in the form of something like an imitation game, though as explained in Part III, the Turing Test is unreliable and new tests are being applied. . Some capacity for belief attribution is necessary to establish the objective intent necessary for contract formation.

Daniel Dennett's intentional stance holds substantial promise in explaining the intent needed for contract formation,[63] though we have to limit our reliance on the intentional stance to now avoid reliance on the Turing Test as the only or main method for assessment. Dennett argues that humans can take three explanatory stances about a complex system. Humans can take a physical stance to predict behavior by understanding the underlying structure of a system or entity.[64] We might, for example, profitably use a physical stance to understand how coding for a smart contract works. Second, humans can take a design stance to predict behavior by understanding the design of a system or entity. Again, using the smart contract example, a person may have no idea how the coding of a smart contract works, but she does know that when she presses "I agree," she has entered into a contract subject to any automated terms included therein. Finally, the intentional stance is the position that we predict the behavior of a system or entity based on our belief that that system or entity has its own beliefs and goals and the capacity on its own to achieve its goals based

[63] Several have so found. SAMIR CHOPRA & LAURENCE F. WHITE, A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS 11-13 (University of Michigan Press 2011); Samir Chopra & Laurence White, *Artificial Agents and the Contracting Problem: A Solution Via an Agency Analysis*, 2009 J. L., TECH, & POL'Y 363 (2009); Giovanni Sartor, *Cognitive Automata and the Law: Electronic Contracting and Intentionality of Software Agents*, 17 A.I. & L. 253 (2009).

[64] DANIEL C. DENNETT, THE INTENTIONAL STANCE 16-17 (MIT Press 1989). From the perspective of cognitive science, see Andrew Brook & Don Ross, *Dennett's Position in the Intellectual World*, *in* DANIEL DENNETT 3 (Andrew Brook & Don Ross eds., Cambridge University Press 2002).

thereon. In Dennett's words, humans look for "true believers:" a system that acts, or at least appears to act, based on its beliefs.[65]

The intentional stance and similar theories about belief attribution[66] rely on what philosophers and cognitive scientists call "folk psychology," an array of mental concepts that humans have known since childhood that they effortlessly deploy these, that is: beliefs, desires, knowledge, pain, fear, hope, expectation, intention, imagination, concern, and so on.[67] As society progresses toward more automation, our folk psychology will likely accommodate more leniency in the application of the international stance to artificial life. Humans routinely attribute intentions to non-human animals. Particularly with robotic forms of artificial intelligence, we often want to believe they express intention, particularly if they look or move like us.

The intentional stance must run both ways. AI must also possess these intention attribution capacities. Here is where the concept of shared intentionality comes into play. The special kind of intention needed for contract formation is the intention to share a goal or plan. With contract formation, that goal or plan is to complete performance, a cooperative venture between contract parties. At present, only humans have this facility.

For an agent to contract, it must be a planning agent, that is, one with the cognitive ability to have a future-directed intention. That humans are planning agents means we

[65] DENNETT, *supra* note 64, at 13.
[66] Another way to understand the capacities to predict and explain behavior is known as mindreading. Philosophers and psychologists have identified a set of cognitive capacities in humans, consisting of the capacity to explain and predict human behavior, attribute mental states to humans, and explain human behavior in terms of mental states. This research has focused on humans, but it is common for humans to attribute mental states to non-human animals (dogs and cats), machines, computers, toasters, etc. Some of these may be metaphorical, but there is substantial evidence that some primates other than humans might be able to predict behavior of others in their species. *See* SHAUN NICHOLS & STEPHEN P. STICH, MINDREADING (Oxford University Press 2003).
[67] *See* JERRY FODOR, PSYCHOSEMANTICS: THE PROBLEM OF MEANING IN THE PHILOSOPHY OF MIND (MIT Press 1987); DENNETT, *supra* note 64, at 7-11.

engage in what philosopher Michael Bratman[68] calls shared cooperative activity or shared intentions and what anthropologist Michael Tomasello calls joint intentionality.[69]

To understand what shared or joint intentionality is, distinguish it from intentionality in a strictly individual sense.[70] The strictly individualist form of intentionality is about an agent having the ability to self-regulate in situations in which an individual can recognize novel situations and deal reflectively with them with an understanding of the causal relations between intentions to actions. Joint or shared intentionality differs from individual intentionality because it is about cooperation for small groups. For humans, its structure centers around linguistic communication. It is unique to humans. When humans engage in joint or shared intentionality, they engage in cooperative activity. Think of it in the context of its ancient origins. Chimpanzees hunt in parallel. They will pursue prey on they own, considering the behavior and possibly the intentions of other chimpanzees. Each chimpanzee has an individual goal to separately capture the prey. But humans developed the ability to hunt cooperatively, that is, to capture it together with other humans as part of a joint goal.[71] Humans but not chimpanzees can enter into what only humans at present can conceptualize as the institution of contract because of the human capacity for shared intentionality. This capacity is not, however, the only capacity of the human mind required for the institution of contract to work in a society.

[68] MICHAEL BRATMAN, FACES OF INTENTION: SELECTED ESSAYS ON INTENTION AND AGENCY 93-108, 130-141.
[69] MICHAEL TOMASELLO, A NATURAL HISTORY OF HUMAN THOUGHT 32-79 (Harvard University Press 2014).
[70] BRATMAN, *supra* note 69, at 7-31.
[71] *Id.* at 35-36.

B. *Responsiveness to Practical Authority*

A topic under-explored in the philosophy of law is the conception of the person subject to law. Law, from the standpoint of practical authority, addresses its commands presumably only to humans. Who else would be an appropriate subject for law's commands? Practical authority is the authority of an agent to provide reasons for another agent to act.[72] Law has this kind of authority. This section will briefly sketch out a challenge for AI to overcome to become a contracting agent, or, indeed, as an agent more generally subject to law: the psychological ability to respond to law as providing reasons for action.

There are several ways to understand practical authority. One is a relatively neglected account of Lon Fuller in his *Morality of Law*.[73] The focus here is on Fuller because he may be the only legal philosopher to have explicitly addressed the requirement that law requires a responsible agent as its subject.[74] A possible reason for the neglect of this agency discussion in philosophy of law is that law is addressed from a practical authority standpoint only to humans. Moreover, the implicit assumption of most philosophy of law is that humans are rational creatures responsive to reasons for action.

For Fuller, there is an internal morality of law that requires a responsible agent. There is no point to law without the subject of the law constituting a responsible agent. For Fuller, this point underscores the distinction of being "subject to" law rather than being "acted upon by" law.[75]

---

[72] *See* Leslie Green, *Authority*, *in* CONCISE ROUTLEDGE ENCYCLOPEDIA OF PHILOSOPHY 68 (Edward Craig ed., Routledge 2000).
[73] Lon L. Fuller, *The Morality of Law* (2d. ed. Yale University Press 1969).
[74] KRISTEN RUNDLE, FORMS LIBERATE: RECLAIMING THE JURISPRUDENCE OF LON L. FULLER (Hart Publishing 2012).
[75] *Id.* at 98.

Philosophical accounts of the practical authority of law are usually first personal: the focus is on the claims the law makes on its subject. But to understand the normative power of contact law, shared intentionality requires a second personal point of view for contract parties to internalize the obligation to become bound to a contract. A second-person standpoint is one in which an agent possesses the psychological capacity to make claims on other agents and acknowledge and respect the claims other agents make on the agent.[76] Humans have developed a second-personal point of view and use that view to engage in a relationship of mutual recognition.[77] The second-personal account of practical authority is likely more relevant for the practical authority of rules relating to personal interaction rather than compliance with law, but the law often requires persons to recognize and accept features of legal rules that rely on personal interaction to produce binding legal commitment. An agent must be able to understand that it might owe an obligation to another, to be able to fully accept a law that imposes just this sort of obligation.

While an unexplored topic, practical authority likely connects to collective intentionality, the notion that an agent can recognize the institution of contract. Collective intentionality in Tomasello's framework is a massive form of uniquely human cooperation we know of as states, societies, and communities. Common cultural social practices such as law, including contract law, derive from collective rationality. In the words of anthropologist Pascal Boyer, "minds make societies."[78] The cognitive capacities associated with collective intentionality center on the ability of entities to engage in self-governance responsive to a

[76] STEPHEN DARWALL, THE SECOND-PERSON STANDPOINT: MORALITY, RESPECT, AND ACCOUNTABILITY 3-11 (Harvard University Press 2006).
[77] *See id;* T.M. SCANLON, WHAT WE OWE TO EACH OTHER 162, 194 (Harvard University Press 1998). "Second personal" refers to the capacity of an person or agent to address others and to understand concepts like "you" and "yours."
[78] PASCAL BOYER, MINDS MAKE SOCIETIES: HOW COGNITION EXPLAINS THE WORLD HUMANS CREATE (Yale University Press 2018).

culture's norms of rationality.[79] Collective intentionality is a kind of intentionality that any entity must have to participate in a common culture.

Current AI models do not appear to meet these responsible agency requirements. To have the sort of psychology that the recognition of law as having practical authority on an agent demands seems to be currently impossible for an intelligence that canperform complex computations and use probability distributions only to engage in human-like conversation but not more. But AI has a way of advancing in unexpected ways. What may be more achievable, at least in the next several years, will be for humans and AI to be hybrid thinkers and actors[80] such that a human and an AI can share an intention to form a contract and engage in the action needed for contract formation.

CONCLUSION

This chapter tries to offer a way to deal with some problems in contract law that some might view as off in the "far future" for AI. We should learn from AI technologists, however, that describing something about AI as "far off" or speculative is not an exercise in pragmatism, but more in the nature of a gamble. We must be ready for the future. AI will change and it is difficult to predict the pace of that change.[81] Many have tried to predict its future and the predictions will continue.[82] If one writes only on what is possible right now for AI, what is written will be out-of-date as soon as the ink is dry on the page. AI makes our common sense about what is "practical" for discussion wrong: it is more "practical" to

[79] TOMASELLO, *supra* note 69, at 80-123.
[80] RAY KURZWEIL, THE SINGULARITY IS NEAR: WHEN HUMANS TRANSCEND BIOLOGY (Viking Penguin 2005).
[81] *See* NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES (Oxford University Press 2014).
[82] Stuart Armstrong & Kaj Sotala, *How We're Predicting AI – or Failing To*, *in* 9 TOPICS IN INTELL. ENG'G & INFORMATICS, 1, 11 (Jan Romportl, Eva Zackova, & Jozef Kelemen eds., Springer 2015).

focus on what will be rather than on what is.[83] This chapter offers no predictions, but it is an attempt at an account that is resilient to change as AI develops into the future.

[83] This approach to AI research is common. *E.g.*, Anthony J. Casey & Anthony Niblett, *Self-Driving Contracts*, 43 J. CORP. L. 1 (2017); F. Patrick Hubbard, *"Do Androids Dream?'": Personhood and Intelligent Artifacts*, 83 TEMP. L. REV. 405 (2011); Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70 N.C. L. REV. 1231 (1992).