



TOURO UNIVERSITY
JACOB D. FUCHSBERG LAW CENTER
Where Knowledge and Values Meet

**Digital Commons @ Touro Law
Center**

Scholarly Works

Faculty Scholarship

2022

A Philosophy of Contract Law for Artificial Intelligence: Shared Intentionality

John Linarelli

Touro Law Center, jlinarelli@tourolaw.edu

Follow this and additional works at: <https://digitalcommons.tourolaw.edu/scholarlyworks>



Part of the [Contracts Commons](#), and the [Law and Philosophy Commons](#)

Recommended Citation

John Linarelli, A Philosophy of Contract Law for Artificial Intelligence: Shared Intentionality, in *CONTRACTING AND CONTRACT LAW IN THE AGE OF ARTIFICIAL INTELLIGENCE* (Martin Ebers, Cristina Poncibò, & Mimi Zou eds., 2022).

This Book Chapter is brought to you for free and open access by the Faculty Scholarship at Digital Commons @ Touro Law Center. It has been accepted for inclusion in Scholarly Works by an authorized administrator of Digital Commons @ Touro Law Center. For more information, please contact lross@tourolaw.edu.

A PHILOSOPHY OF CONTRACT LAW FOR ARTIFICIAL INTELLIGENCE: SHARED INTENTIONALITY

John Linarelli*

INTRODUCTION

The aim of this chapter is to offer a theory of contract law to account for the inclusion of artificial intelligence in contract practices. It offers a new direction for a philosophy of contract law to be able to give significance to mental states or psychological attributes increasingly relevant to contract practices and which come in the form of machine learning, deep learning, and neural networks associated with artificial intelligence. The objective is to produce a general theory about contract to accommodate both human and artificial agency. The chapter identifies what makes contractual obligation distinctive, and the core concepts of contract law that are most relevant, when contracting involves the interaction of human and artificial agency.

Any practical philosophy – moral, political, or legal – either starts with or presupposes a conception of the person. A philosophy of contract law must account in some way for the capacities of contract parties as more or less responsive to reasons relevant to contractual obligation. With few exceptions, this focus on the conception of the person has been missing in legal philosophy in general and in a philosophy of contract law in particular. The relevance of artificial intelligence to contracting brings this gap to our attention because, as we shall see, so much of what makes a contract a matter of legal obligation depends on what the law identifies as a particular form of intention of persons to form a contract. An intent to enter contractual relations is what makes contracting possible in most legal systems.¹ Persons or other entities that cannot form or be understood to form an intent to contract lack the legal recognition to enter legally enforceable contracts. Intention, then, leads us directly to questions about the mental states of contract parties or their agents and how other contract parties come to recognize and accept that intent.

* Professor of Law, Touro College, Jacob D. Fuchsberg Law Center. I am grateful for comments on a presentation of this paper in the Contracting and Contract Law in the Age of Artificial Intelligence conference sponsored by the University of Turin Observatory on Economic Law and Innovation, the Robotics and AI Law Society, and the European Law Institute. I am also appreciative of comments at two Touro Law works-in-progress workshops. All errors are mine.

¹ See, e.g., Hugh Beale, Bénédicte Fauvarque-Cosson, Jacobien Rutgers, & Stefan Vogenauer, *Contract Law Cases and Materials (Jus Commune Series)* (3rd ed. Hart/Bloomsbury 2019): 301-303. My focus is on the objective theory of contract in Anglo-American contract law.

Artificial intelligence in contract practices brings out that what makes contract law a distinctive form of legal obligation is shared intentionality. I refer to this insight as the shared intentionality thesis. Shared intentionality is the psychological capacity of one agent to share and pursue a joint goal with another agent. It is an attribute of human thought empowering human planning and the ability to share agency with others. To have this form of shared intention, one must have the mental capacity for a second personal point of view, to be able to know and give due recognition to the views and interests of others, and to attribute the same point of view to others. Shared intentionality leads to a focus on objective intent to enter contractual relations in Anglo-American contract law as the primary concept in understanding what is distinctive about contractual obligation.

In demonstrating the shared intentionality thesis, theories such as “contract as promise” and “contract as consent” reduce to a theory about intent and mental states. When we add the role of artificial intelligence in contracting, conceptualizing contract around promise or consent do little or no work on their own. These concepts cannot direct us to what is distinctive about contractual obligation. Because we are dealing with artificial intelligence and in some future cases artificial beings with their own abilities to enter contracts independent of human agency, we must get underneath concepts like promise and consent. When artificial intelligence enters the mix of transacting, these intermediate concepts underdetermine what is going on.

Part I explains how artificial intelligence is involved in the process of contract formation today and how it may be involved in the future. The focus in Part I will be on contract formation because we need to understand the role of intention to create legal relations in contract.

Part II explains the relevance of objective intent in contract law to developing a philosophy of contract law to account for artificial intelligence. An intention to create a contract is key to understanding what is distinctive about contractual obligation. It is a necessary condition to persons cooperating in the form of a contract. The doctrine of objective intent in contract law operates as a Turing test for determining whether a contract has been formed. Part II further explains that it does not matter which “mind” is thought to produce this intent – human or artificial. The issue is attribution of the right sort of mental states by one contract party to another. Philosophy and cognitive science can assist in developing this argument, most prominently Daniel Dennett’s notion of the intentional stance,² as well as work on “theory of mind,”³ and “mindreading.”⁴ The bottom line here is understanding how folk psychology can lead humans to the sort of recognition that is

² Daniel C. Dennett, *The Intentional Stance* (MIT Press 1987).

³ Alvin I. Goldman, “Theory of Mind” in Eric Margolis, Richard Samuels, & Stephen Stich, *Oxford Handbook of Philosophy of Cognitive Science* 402 (Oxford University Press 2012).

⁴ Shaun Nichols & Stephen P. Stich, *Mindreading: An Integrated Account of Pretence, Self-Awareness, and Understanding Other Minds* (Oxford University Press 2003).

needed to engage with increasing levels of artificial intelligence in contracting. What gives artificial intelligence intent is not some internal workings of its programming but *us* – the ascription of intentionality to artificial intelligence by humans.

Doctrines about capacity to contract are ruled out as a relevant locus for this discussion. The law on capacity to contract involves a set of policy decisions, developed along social preferences or conventions, for determining when the state will or will not enforce contracts. Capacity doctrines deal mainly with exceptional cases through rules of positive law external to whether intent to enter contractual relations can be formed. We want to cover the central cases of contractual intent. It may come to pass that forms of artificial intelligence may be regulated as not entitled to recognition as a question of capacity to contract or in a simple prohibition on their operation as an agent for contracting purposes, but such determinations usually depend on some policy justification beyond our bounds here.

The intent to enter a contract means more than the recognition of a capacity in an agent (human or artificial) to take actions directed by some end the agent has. That sort of intentionality is certainly a necessary condition for agency. But agents also must be planning agents. They need to be able to reciprocate intentions. The intentional stance on its own does not offer an adequate explanation of a special kind of intent that is needed for contract formation. It does not inform us about the sort of “we” intentionality or the ability to engage in future directed intentions as elements of stable plans of action in the form of contracts. Part III fills in this argument. It argues that shared intentionality is the core of contractual obligation. It does not matter whether this shared intentionality comes from human or artificial minds. The role of intention in contract law comes in two steps. First, we must understand the role of attribution of intentional mental states from one agent to another. This attribution is the subject of Part II. Next, we need to determine if there is a role for a psychological capacity that only humans have possessed, to be able to share intent – to represent and act as “we.”⁵ The insight to explore here is whether agents can combine their future directed plans or intentions to produce mutual obligations in the form of a contract.

This chapter will go beyond the state of the possible for artificial intelligence right now. The theory in this chapter will rest on some general assumptions about the state of machine learning today, but the goal is a theory that is sufficiently general, to remain relevant as the state-of-the-art for artificial intelligences changes. Artificial intelligence will change.⁶ Many have

⁵ See n. 73-79.

⁶ See Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford University Press 2019).

tried to predict its future and the predictions will continue.⁷ If one is constrained to write only on what is possible right now for artificial intelligence, what is written will be out of date as soon as the ink is dry on the page. Artificial intelligence makes our common sense about what is “practical” for discussion wrong: it is more “practical” to focus on what will be rather than on what is.⁸ This chapter offers no predictions, but it is an attempt at an account that is resilient to change as artificial intelligence develops into the future.

I. ARTIFICIAL INTELLIGENCE in CONTRACTING: PAST, PRESENT, AND FUTURE

How is artificial intelligence involved in contracting? The strategy deployed to answer this question is to proceed through the development of artificial intelligence, from its early days as symbolic artificial intelligence, on to machine learning and big data, and finishing with some possibilities for the future of artificial intelligence. The various ways in which artificial intelligence interjects into contract practices can be understood by mapping contract practices to the various stages of the development of artificial intelligence. As this analysis proceeds, I will try to identify plausible uses of artificial intelligence in contracting that either have not yet occurred or which have not reached widespread use. The idea here is to lead from the technological discussion, to explore how artificial intelligence can be involved in contract practices and not necessarily how it has been involved in those practices.

A. *The Search for Agency*

A common presupposition in discussions about artificial intelligence and contractual obligation is the idea that at some point a boundary is crossed in which artificial intelligence becomes an “agent.” The reason why this boundary exists is that a common understanding of contractual obligation it is that it is a form of obligation in private law this is chosen, unlike obligations in tort.⁹ One cannot choose unless one has agency. An agent is a being with the capacity to act with intentionality.¹⁰

⁷ Stuart Armstrong & Kaj Sotala, “How We’re Predicting AI – or Failing to” in Jan Romportl, Eva Zackova, & Jozef Kelemen eds., *Beyond Artificial Intelligence. Topics in Intelligent Engineering and Informatics*, Vol 9 11 (Springer 2015).

⁸ This approach to artificial intelligence research is common. See, e.g., Anthony J. Casey & Anthony Niblett, “Self-Driving Contracts” (2017) 43, *the Journal of Corporation Law* 101; F. Patrick Hubbard, “Do Androids Dream?: Personhood and Intelligent Artifacts” (2011) 83 *Temple Law Review* 405; Lawrence B. Solum, “Legal Personhood for Artificial Intelligences” (1992) 70 *North Carolina Law Review* 1231.

⁹ Daniel Markovits, “Theories of the Common Law of Contracts”, *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/win2019/entries/contracts-theories> (last accessed June 24, 2021).

¹⁰ Markus Schlosser, “Agency”, in *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/win2019/entries/agency/> (last accessed May 13, 2021).

Can artificial intelligence offer the sort of agency needed to be normatively significant actor in contract formation? In other words, how is artificial intelligence any different from any other technology we might use to facilitate the formation of a contract? Pen and paper is a technology. Why might we consider artificial intelligence to be different from these other technologies?

To understand how artificial intelligence is or might be involved in contacting, we first have a working conception or “definition” of artificial intelligence. No agreed-upon definition of artificial intelligence exists. This is not a problem. There is no way to “capture the essence” of artificial intelligence in a definition.¹¹ Think of artificial intelligence, rather, as a set of practices or methods associated with relevant computer and engineering domains. Stuart Russell and Russell Norvig, authors of the definitive textbook on artificial intelligence, say that artificial intelligence can be understood along four dimensions: (a) acting humanly, (b) thinking humanly, (c) thinking rationally, and (d) acting rationally.¹² They focus on category (d), “on general principles of rational agents and on the components for constructing them.”¹³ The focus on action in category (d) is the right focus as trying to conceptualize whether an entity is “thinking” leads to a variety of problems whose solutions do not aid us very much in understanding artificial intelligence. But category (d) is problematic for artificial agency in contract law because humans are not rational or at least not always rational and so how could perfectly rational and “boundedly”¹⁴ rational agents cooperate to produce a contract that is sensible or not putting the human side of the bargaining process in a disadvantageous position. The fix to the rationality problem is to think of rationality as an archetype or idealized notion of human thought. For our purposes, we can rely on a minimalist conception of artificial intelligence as getting “machines”¹⁵ to act in ways that depend on cognitive functions like learning and problem-solving, in a way that would be

¹¹ Dictionaries offer something that differs substantially from, say, the work of semantics, which is about the uncovering of the meaning of sentences. Dictionary definitions are only one kind of definition. They aim to define words with enough information to provide a language user with sufficient information to be able to use a term in sentences. An example of this distinction is at work in the difference between the definition of a word in *Black’s Law Dictionary* and the development of the meaning of a word by frequent interpretation over time by common law courts.

¹² Stuart Russell & Peter Norvig, *Artificial Intelligence: A Modern Approach* (Person 3rd ed. 2009): 1-5.

¹³ *Ibid.*, 5.

¹⁴ I use “boundedly” here in a broad sense to refer to any departures in the social and behavioral sciences from conceptions of persons as perfectly rational. This is the clear direction now of the sciences that study human thought and action. See Lex Fridman Podcast, “Daniel Kahneman: Thinking Fast and Slow, Deep Learning, and AI, <https://lexfridman.com/daniel-kahneman/> (last accessed May 18, 2021); Tim Adams, “Daniel Kahneman: ‘Clearly AI is going to win. How people are going to adjust is a fascinating problem’,” *The Guardian*, May 16, 2021.

¹⁵ Nils J. Nilsson, *The Quest for Artificial Intelligence* (Cambridge University Press 2010): xiv (on how the computer has greatly expanded our notion of what a machine is, how software

recognizable to humans, particularly if they are shorn of the psychology that makes them prone to depart from rational thought.¹⁶

The best description of artificial intelligence might be Nils Nilsson's, in his definitive work on the history of artificial intelligence: artificial intelligence is "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."¹⁷ Nilsson's definition works well for understanding how artificial intelligence might work in contracting – and differ from other sorts of technologies deployed in the contracting process – because of its focus on functionality and foresight, or what might be understood as intentionality in contract formation, explored in parts II and III below. Nilsson's focus on agency is what is needed to take the discussion of artificial intelligence into the future of contracting, where an agent is understood not in its conventional legal sense as the representative of a principal but psychologically or philosophically as a system or organism with the capacity to act with intentionality.¹⁸

The ability to possess or be seen by others to be able to act as an agent with intentionality is key for understanding if artificial intelligence could have a distinctive role on its own independently of human agency. An agent is a being with the capacity to act, with action requiring intention. Very simply, an agent is an entity that can take actions on its own and that are directed to some end the agent has.¹⁹ Agents take actions for reasons. These reasons can be based on beliefs, desires, and attitudes. That an agent has reasons for actions means that we ascribe intentions to the agent. Some philosophers, most notably Donald Davidson, have argued that the intentions of agents are the causes of the actions of the agent.²⁰

But we must be cautious here: an agent may be acting because it was programmed or directed to act by some external force, such as through the direction of a principal or a computer programmer. Does the agent have a reason for action in such a case? In a sense, yes. Its reasons for actions are the direction of its programming, though it might be difficult to describe programming as a "reason." The agent may intend to comply with the direction of another. But what if it has no choice but to comply with its programming? We do not know whether human agents face this same difficulty, or whether

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

¹⁶ S. Matthew Liao, "A Short Introduction to the Ethics of Artificial Intelligence" in S. Matthew Liao ed., *Ethics of Artificial Intelligence* (Oxford University Press 2020): 3 ("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").

¹⁷ Nilsson, n. 15, at xiii.

¹⁸ Schlosser, n. 10.

¹⁹ Samir Chopra & Laurence F. White, *A Legal Theory for Autonomous Artificial Agents* (University of Michigan Press 2011): 5-28.

²⁰ Donald Davidson, *Essays on Actions and Events* (Oxford University Press 1980).

their “programming” comes from the evolution of their psychologies. This brief discussion is meant to expose the difficulties in trying to know whether an agent acts with the requisite “aboutness,” self-originating direction, or reasons independent of the influence of external actors or technologies.²¹

B. *GOF AI Contracts*

An early form of artificial intelligence is symbolic artificial intelligence, more colloquially known as “good old-fashioned artificial intelligence” or GOF AI for short.²² GOF AI is based on symbolic (human readable) representations of problem solving and the rules of formal logic. The animating idea behind GOF AI was that a physical system would provide the necessary and sufficient conditions for intelligent action.²³ Symbolic artificial intelligence required 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 logicist²⁴ or logician, in which code is to be designed to produce deductive forms of reasoning for every conceivable task or problem the artificial intelligence in question was meant to do or solve. Symbolic artificial intelligence was the main form of artificial intelligence until the mid-1980s. It does not involve any machine learning. The machine or system does not learn from data. Rather, all its decisions must be explicitly programmed in advance. Logic comes first data comes second, the very opposite direction of most artificial intelligence today in the post-symbol area, for which data comes first and logic second. The most advanced forms of artificial intelligence are based on the notion of inductive inferences drawn from massive amounts of data.²⁵

Symbolic artificial intelligence offers a comfortable zone for traditional coders, but it is now understood to be an “old” form of artificial intelligence vastly outpaced by machine learning, clearly not symbolic or logicist in ap-

²¹ See the discussions in Chopra & White, n. 19; Samir Chopra & Laurence White, “Artificial Agents and the Contracting Problem: A Solution Via Agency Analysis (2009) *University of Illinois Journal of Law Technology & Policy* 363.

I have avoided any discussion of free will here as that leads us to a larger set of metaphysical questions that will add more questions than answers. The same goes for consciousness.

²² John Haugeland, “Farewell to GOF AI?,” in Peter Baumgartner & Sabine Payr, *Speaking minds: Interviews with Twenty Eminent Cognitive Scientists* (Princeton University Press 1995): 101.

²³ Nilsson, n. 15, at 331.

²⁴ *Ibid.* 331-346; Pedro Domingos, *The Master Algorithm* (Basic 2015): 30, 49, 80-83.

²⁵ Domingos, n. 24.

proach. But it has many uses today in contracting. One application of symbolic artificial intelligence in use today is the smart contract.²⁶ It is well-accepted that smart contracts do not rely on artificial intelligence.²⁷ It might be more accurate to say that some smart contracts rely on the most rudimentary form of symbolic artificial intelligence, but many do not. 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 is executed and enforced digitally and without the need for human intervention except at the level of writing code to automate contract performance.²⁸ 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.²⁹ In short, smart contracts substitute algorithmic for human contract performance and enforcement.

Symbolic artificial intelligence will probably never be able to develop to the level of “agent,” if we understand an agent to be a being with the capacity to act with intentionality. It can only be a tool for agents.³⁰ It therefore cannot take artificial intelligence very far into contract practices. The reasons why have to do with the limits of symbolic logic. Contracts in the classical or traditional sense are understood from the perspective of mutual assent

²⁶ “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,” *Harvard Law School Forum on Corporate Governance*, <https://corpgov.law.harvard.edu/2018/05/26/an-introduction-to-smart-contracts-and-their-potential-and-inherent-limitations/>, May 26, 2018 (last accessed May 13, 2021).

²⁷ See “What are Smart Contracts on the Blockchain?,” <https://www.ibm.com/topics/smart-contracts> (last accessed May 13, 2021); Michael Mylrea, “AI Enabled Blockchain Smart Contracts: Cyber Resilient Energy Infrastructure and IoT,” *The 2018 AAAI Spring Symposium Series*.

²⁸ Various authors have offered definitions of a smart contract. Nick Szabo is credited with inventing the phrase. Kevin Werbach & Nicolas Cornell, “Contracts Ex Machina” (2017) 67 *Duke Law Journal* 102. 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), available at http://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart_contracts_2.html (accessed Feb. 7, 2019). 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 & Cornell define a smart contract as an “agreement in digital form that is self-executing and self-enforcing.” Werbach & Cornell, n. 28, at 108. Jeffrey Lipshaw describes a smart contract as “simply computer code that automatically execute agreed-upon transactions.” Jeffrey M. Lipshaw, “The Persistence of ‘Dumb’ Contracts” (2019) 1 *Stanford Journal of Blockchain Law and Policy* 1, 4.

²⁹ Werbach & Cornell, n. 28, at 118.

³⁰ I am grateful to Mimi Zhou for this distinction. See Lauren Henry Scholz, “Algorithmic Contracts” (2017) 20 *Stanford Technology Law Review* 128.

based on the shared meanings of contractual language.³¹ This conception of contract is now under considerable threat in the form of automated contracts of adhesion between consumers and firms that move the point of normative significance for contract formation from mutual assent to constructive notice.³² Regardless of this shift, there will have to be some use of natural language, either for purposes of assent or notice, for a contract to be binding on a human agent. Symbolic artificial intelligence, however, focuses on the meaning of sentences, presupposing that symbols can have self-contained or concrete meanings without the need for any further information. It suffers from the classic philosophical problem of reference. Human users of language rely on the meanings of speakers, on what philosopher Paul Grice calls conversational implicature, to be distinguished from logical implicature.³³ The Gricean insight is that humans rely on social contexts and cooperative norms about language to convey meanings, and not on the literal meanings of terms and sentences. For example, using the contract in the famous case of *Raffles v. Wichelhaus*,³⁴ if contract parties say that a shipment of cotton is to arrive “to arrive ex Peerless from Bombay” what is necessary to understanding the meaning of this sentence is a set of norms about context. No amount of focus on the meaning or logic of the sentence will solve the problem of shared meaning that is necessary for a contract to exist, at least a contract in the traditional sense. Artificial intelligence that can reliably deal with conversational implicature will not be symbolic. It will more likely be probabilistic in approach, as it will have to learn through many repeat instances the cooperative norms that humans use when they use natural language. It will be machine learning.

C. Machine Learning and Contract

Machine learning is artificial intelligence 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

³¹ See Restatement (Second) Contracts §§18-20, 205.

³² See Robin Bradley Kar & Margaret Jane Radin, “Pseudo-Contract and Shared Meaning Analysis” (2019) 132 *Harvard Law Review* 1135; 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 Journal of Law and Regulation 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/> (last accessed May 13, 2021).

³³ H. Paul Grice, “Logic and Conversation,” in Peter Cole & Jerry L. Morgan eds., 3 *Syntax and Semantics* (Academic Press 1975) 41.

³⁴ [1864] EWHC Exch. J19.

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, in automated trading, in filling in price terms with dynamic pricing, to provide financial advisory services, to identify risks and predictive analytics for construction, and in many other uses.³⁵ Perhaps most significantly, machine learning has been used to make firms or what we know as “merchants” in the Uniform Commercial Code³⁶ 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.³⁷ 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. 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 is loading.³⁸ Still, the framework for these contracts is human produced.

Artificial intelligence can be used to assist in contract drafting.³⁹ Here is an example that to my knowledge has yet to be developed. Choice of law clauses are in ubiquitous in contracts. They are often coupled with an alternative dispute resolution clause mandating arbitration. They are in frequent use in contracts with consumers, usually because the merchant (or its lawyers) known that the law of, say, South Dakota, is more favorable to credit card companies than say, the law of California.⁴⁰ This is a relatively easy call for contract planners who represent the credit card company. But what about negotiated contracts, for, say, the purchase of the assets of a business, or in

³⁵ Scholz, n. 30; OECD, *Artificial Intelligence in Society*, Jun 11, 2019, <https://www.oecd.org/publications/artificial-intelligence-in-society-ceedfee77-en.htm> (last accessed May 13, 2021); Martin Ebers, “Regulating AI and Robotics: Ethical and Legal Challenges” in Martin Ebers & Susana Navas Navarro eds., *Algorithms and Law* (Cambridge University Press 2019).

³⁶ U.C.C. §2-104.

³⁷ See, e.g., Stacy-Ann Elvy, “Contracting in the Age of Internet of Things: Article 2 of the UCC and Beyond” (2016) 44 *Hofstra Law Review* 839

³⁸ Dina Srinivasan, “Why Google Dominates Advertising Markets Competition Policy Should Lean on the Principles of Financial Market Regulation” (2020) 24 *Stanford Technology Law Review* 55.

³⁹ See, e.g., Kathryn D. Betts & Kyle R. Jaep, “The Dawn of Fully Automated Contract Drafting: Machine Learning Breathes New Life into a Decades Old Promise” (2017) 15 *Duke Law and Technology Review* 216; Irene Ng, “The Art of Contract Drafting in the Age of Artificial Intelligence: A Comparative Study Based on US, UK, and Austrian Law, TTLF Working Papers (2017), <https://law.stanford.edu/publications/the-art-of-contract-drafting-in-the-age-of-artificial-intelligence-a-comparative-study-based-on-us-uk-and-austrian-law/> (last accessed May 13, 2021). For an investment treaty drafting example, Wolfgang Alschner & Dmitry Skougarevskiy, “Can Robots Write Treaties? Using Recurrent Neural Networks to Draft International Investment Agreements” in F. Bex & S. Villata eds., *JU-RIX: Legal Knowledge and Information Systems* (IOS Press 2016): 119.

⁴⁰ Amy Sullivan, “How Citibank Made South Dakota the Top State in the U.S. for Business,” *The Atlantic*, July 10, 2013.

construction contracts, or in contracts that cross national borders? Artificial intelligence has much more computing power to determine which law (and courts) are best for a contract party. Lawyers often insert these clauses into contracts without much deliberation or based on intuition or “judgment” that a particular jurisdiction has the more favorable law, but machine learning can probably inform us with a much higher degree of accuracy which is the better law and courts, if it has access to the right sort of data on the law and courts of the relevant jurisdiction.

Still, humans control this assistive role in contract drafting to the point where artificial intelligence does not rise to the level of an agent in the philosophical sense in which that concept is understood here. Except in some limited and rudimentary ways, as in speech recognition technologies, artificial intelligence is not yet at a point where it can interact directly with humans using a natural language interface in any way close to what might be needed for contract formation for contracts that are negotiated or what Karl Llewellyn called “dickered.”⁴¹ The natural language interface may be less important for “notice” type contracts that do not rely on negotiation, though the boilerplate terms still must be in a natural language for a contract to be formed.

C. *The Future*

Is there some possibility that artificial intelligence can rise to the level of an actual agent in the contract formation process, or at least share some of that agency, properly delegated, with a human contract party? What has to happen is for algorithms to produce algorithms to draft and negotiate contracts. An artificial intelligence would learn without human intervention to select contract parties and contract terms. It requires another step towards automation towards agency in some form of action by artificial intelligence unsupervised by humans. Promising areas might be coded or partially coded contracts in which artificial intelligence is responsible for the “drafting” of the code and the code translates into a natural language equivalent for human understanding, a natural language – code interface of some sort. Humans express intention through linguistic communication and any interaction with humans in a contracting context would seem to require some use of language. Intelligent assistants, disembodied agents that assist humans when mobility is not required, are already in use.⁴² Artificial intelligence may plausibly develop natural language abilities. This would be task-level artificial intelligence with natural language abilities.⁴³ Finally, in some far future,

⁴¹ Karl N. Llewellyn, *The Common Law Tradition: Deciding Appeals* (Little Brown 1960): 363-372.

⁴² Nilsson, n. 15 522,

⁴³ Nilsson, n. 15, at 525-534.

there may be an artificial general intelligence with natural language abilities and full-blown agency.⁴⁴

At some point, humans will attribute mental operations and intentionality to artificial intelligence in the contracting process. The intent needed for contract formation will either be shared with a human-level principal or co-party or humans will attribute the ability to produce such intent entirely to artificial intelligence. How this will occur will depend on the folk psychology at work when humans typically attribute intentionality to others. We are not there yet but will be. It is to these issues that I will now turn.

II. OBJECTIVE INTENT AND THE ATTRIBUTION OF INTENTIONALITY

Returning to Nilsson's definition of artificial intelligence, how do we go about determining for contracting purposes whether an entity can function appropriately and with foresight in its environment? In other words, when can we say when an artificial intelligence can have the needed "aboutness" that will give humans reasons to believe that an artificial intelligence involved in contacting has the intentionality needed to contribute to the formation of a contract? This is the question for this Part II. Part III deals with the question of whether this intentionality can be shared between a human and an artificial intelligence.

A. *Turing Test Intentionality in the Common Law of Contracts*

In his 1950 article, *Computing Machinery and Intelligence*, Alan Turing starts by stating his question as "can machines think?"⁴⁵ He eventually finds this question "too meaningless to deserve discussion"⁴⁶ and replaces it with "are there any imaginable digital computers which would do well in the imitation game?"⁴⁷ The imitation game proceeds as follows. The players in the game are a human being and a machine. The game also includes a human interrogator. The interrogator is in a separate room from the human and the machine. The interrogator knows the others are labelled *X* and *Y*, knows one is human and the other machine, but does not know which. The object of the game is to test whether the interrogator can tell the difference between the human and the machine through a series of questions. The interrogator is to ask questions to the machine and the individual through a text channel to avoid revealing which is the human and which the machine. If the interrogator cannot reliably distinguish the human from the machine, then the machine passes the test. The test is meant to assess whether a machine can give answers that consistently resemble those a human would give such that

⁴⁴ John Linarelli, "Artificial General Intelligence and Contract" (2019) 24 *Uniform Law Review* 330; Bostrom, n. 6.

⁴⁵ Alan Turing, "Computing Machinery and Intelligence" (1950) LIX *Mind* 433.

⁴⁶ *Ibid.*, 442.

⁴⁷ *Ibid.*

humans cannot tell the machine apart from a human, in terms of the cognitive performance of the machine.⁴⁸

The Turing test has been effectively embedded into Anglo-American contract law, in the objective theory of contract.⁴⁹ The focus of this tradition in contract law is on outward appearances – on what can be proven as a matter of evidence independent of the mental states the parties may or may not have. 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.⁵⁰

Judge Frank Easterbrook has explained that intention to be bound “does not invite a tour through [a contract party’s] cranium” but must necessarily be derived from a consideration of the words, written and oral, and actions of the part parties.⁵¹ 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.”⁵²

The U.S. Restatement (Second) of Contracts, intended to reflect a consensus about contract law in the United States, does not contain any section

⁴⁸ See Stevan Harnad, “The Turing Test Is Not A Trick: Turing Indistinguishability Is A Scientific Criterion” (1992) 3 *SIGART Bulletin* 9.

⁴⁹ See Joseph M. Perillo, “The Origins of the Objective Theory of Contract Formation and Interpretation” (2000) 69 *Fordham L. Rev.* 427; Timothy A. O. Endicott, “Objectivity, Subjectivity, and Incomplete Agreements,” in Jeremy Horder ed., *Oxford Essays in Jurisprudence Fourth Series* (Oxford University Press 2000): 151.

⁵⁰ *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).

⁵¹ *Skycom Corp. v. Telstar Corp.*, 813 F.2d 810, 814 (7th Cir. 1987).

⁵² *Woburn National Bank v Woods*, 77 N.H. 172, 89 A 491, 492 (1914)(citation omitted), quoting Oliver Wendell Holmes, Jr., *The Common Law* (Little Brown 1881): 307.

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. Restatement (Second) section 21 provides that “neither real nor apparent intention that a promise be legally binding is essential to the formation of a contract. . . .”⁵³ but the objective theory of contract still prevails. 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.⁵⁴ Mutual assent is objectively determined. While English law does not reflect this Restatement (Second) language of manifestation of mutual assent, it is substantially similar in adhering to an objective theory of contract formation and interpretation.⁵⁵ 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.

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 contract party has the requisite objective intent. If artificial intelligence is responsible for all or some of that intent, the issue therefore becomes when a human contracting party would come to recognize the artificial intelligence as providing the intent on the other side of the transaction. When will artificial intelligence reach the point of development that it can be considering as having or contributing to having the sort of intent necessary for contract to come into existence?

B. *The Intentional Stance*

The presence of artificial intelligence in contracting requires us to compare human and artificial mental operations. We have never had to do this before when we only had to think about human intention in contract relations. We can do this in functionalist terms – what makes something a mental state does not depend on its internal constitution but on the way it functions, on the role it plays, in the system in which it is a part.⁵⁶

A fruitful way to engage questions about intentionality is to look to what is known in the philosophy of mind as a theory about the content of mental

⁵³ Restatement (Second) of Contracts §21 (1981).

⁵⁴ *Ibid.*, §18. The “manifestation” language is pervasive in American contract law and reflects the notion of contracting making sense only in the form or external representations to other persons. *See, e.g.*, Restatement (Second) of Contracts §2 (1981), which defines core concepts such as a promise as “a manifestation of intention to act or refrain from acting in a specified way, so made as to justify a promisee in understanding that a commitment has been made.” Comment b explains that a manifestation of intention is an “external expression” as opposed to “undisclosed intention.” Restatement (Second) of Contracts §2 cmt. b (1981).

⁵⁵ *See, e.g.*, Edwin Peel, *Treitel on the Law of Contract* 1 (Oxford University Press 14th ed. 2015)(§1-002, ‘The objective principle’).

⁵⁶ For an accessible explanation of functionalism, see Daniel C. Dennett, *Consciousness Explained* (Penguin 1991): 30-32.

representations. A theory of content of mental representations is an explanation of how humans can form thoughts about concepts or things. For example, how do we explain how we think about “computer?” Think of the various propositional attitudes that can come to one’s mind. I may entertain a belief that I need to be at the law school every Monday for office hours. This belief may lead to a desire that I get in my car and travel to my law school every Monday. These belief-desire states reflect my intentions. My belief about going to the law school is about me intending to go to the law school. This is a very basic sketch that suppresses a literature and the many debates in philosophy of mind,⁵⁷ but it is enough for us here.

So far, we have evaluated my own beliefs from my own perspective, a first-person perspective. My own beliefs, desires, and attitudes may explain my own behavior. But a link is missing here, the attribution of my intentions to me by another person. 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 states to other humans, and explain the behavior of humans in terms of mental states.⁵⁸ Having this capacity is necessary to understand how intent to enter a contract operates, because that intent must be externally manifested and attributed, in the form of something like an imitation game, for the objective intent necessary for a contract to be formed to come into existence. Some capacity for belief attribution is necessary for contract formation.

When it comes to evaluating intentionality in contract law, Daniel Dennett’s theory of content, his so-called intentional stance, holds substantial promise.⁵⁹ Dennett argues that humans can take three explanatory stances about a complex system. We can take a physical stance, to predict behavior by understanding how a system or entity is built.⁶⁰ We might, for example, profitably use a physical stance to understand how coding for a smart contract works. We can take a design stance, to predict behavior by understanding how a system or entity is designed. 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” that a contract is formed in accordance with any automated terms included in the contract. Finally, the intentional stance is the position we take to predict the behavior of a system or entity that we believe that system or entity has beliefs and goals

⁵⁷ A place to start would be: Fred Adams & Ken Aizawa, “Causal Theories of Mental Content”, *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/spr2021/entries/content-causal> (last accessed May 17, 2021).

⁵⁸ See n. ____.

⁵⁹ Several have so found. Chopra & White, *A Legal Theory for Autonomous Artificial Agents*, n. 19, at 11-13; Chopra & White, “Artificial Agents and the Contracting Problem”, n. 21; Giovanni Sartor, “Cognitive Automata and the Law: Electronic Contracting and Intentionality of Software Agents” (2009) 17 *Artificial Intelligence and Law* 253.

⁶⁰ Dennett, *The Intentional Stance*, n. 2: 16-17. From the perspective of cognitive science, see Andrew Brook & Don Ross, “Dennett’s Position in the Intellectual World,” in Andrew Brook & Don Ross, *Daniel Dennett* (Cambridge University Press 2002): 3.

and the capacity on its own to achieve its goals based on its beliefs. In Dennett's words, we are looking for "true believers:" a system that acts or at least appears to act based on its beliefs.⁶¹

Think about the intentional stance from an adaptation of Dennett's own explanation of his account, in the context of humans. Humans may be unique in their use of natural language communication. We use sound linguistically in the form of syntax and semantics, which lead to many of our communications as being evaluable as true and false. This is a big move, according to Dennett, a "radical reconstrual of the data, and abstraction from its acoustic and other physical properties."⁶² Humans make sense of the sounds as words. It is a record of speech acts, not just sounds but assertions, declarations, questions, answers, promises, comments, requests for clarification, critiques and so on.⁶³ Dennett characterizes the evaluation of these utterances as a process of interpretation that involves the taking of an intentional stance. According to Dennett, "we must treat the noise emitter as an agent, indeed, as a rational agent, who harbors beliefs and desires and other mental states that exhibit intentionality or 'aboutness,' and whose actions can be explained (or predicted) on the basis of the content of these states."⁶⁴ The utterances are interpreted as propositions the entity wanted to make, based on reasons the entity holds as content in its thought.

There is no theoretical obstacle to extending the possibility of someone taking an intentional stance about an entity that does not "speak" words verbally or more generally does not use natural language to communicate. It is about belief-desire attribution to a "system" that may or may not be a human. The account seems deliberately left open to be able to accommodate artificial intelligence. The mode of communicating the belief-desire attribution may indeed matter in determining the reasonableness of the attribution but it will not make it impossible. But a system may be sufficiently rational or a true believer for the task at hand.⁶⁵

The intentional stance and similar theories about belief attribution⁶⁶ rely on what philosophers and cognitive scientists call "folk psychology," an array of mental concepts that humans have known since childhood and that they effortlessly deploy as part of being human, such as beliefs, desires,

⁶¹ Dennett, *The Intentional Stance*, n. 2, at 13.

⁶² Daniel C. Dennett, *Consciousness Explained* (Penguin 1991): 74-75.

⁶³ *Ibid.*, 76.

⁶⁴ *Ibid.*

⁶⁵ Dennett, *The Intentional Stance*, n. 2, at 21.

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 Nichols & Stich, n. 4.

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 intentional 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.

Another question that has not yet been asked is whether artificial intelligence can also possess intention attribution capacities. Shall we simply assume humans have intentions and that there is no need for an intentional stance to move from artificial to human agent. This seems wrong, because, as explained above, intent in contractual relations must work both ways. Of course, an algorithm in an artificial intelligence could inform it that humans possess the ready capacity to intend to enter a contract.

We do need another step to connect intentionality to contract – that is – can we share these intentions in the form of a cooperative effort? And if so, can we share them in a way that creates mutual obligation? I deal with these questions in Part III. Contractual intent is a special form of intent. They often involve planning about future conduct. It is routinely future directed about future conduct. It requires parties to be able to believe they have a legal obligation to do something in the future. How contract parties that include some kind of artificial agent can share such an intention is the subject of the Part III.

III. SHARED INTENTIONALITY AS THE CORE OF CONTRACTUAL OBLIGATION

Part II covered belief attribution capacities of humans to attribute intention to artificial intelligence. It also dealt tentatively with the need in contract formation for that intent to work in both directions. To the extent that an artificial intelligence would somehow have or share agency with a human in the process of contact formation, it may need to be able to take an intentional stance about the human on the other side of the transaction. This tells us something about the capacity of artificial intelligence to mind read or recognize the intentional stance taken by other entities, in this case humans.

This dual intentionality does not go far enough, however, to describe the sort of intentionality needed for contract formation. Contract formation requires a particular kind of intention. We need to understand its semantics. The special kind of intention needed for contract formation is the intention to share a goal or plan. The goal or plan is to complete contract performance, a cooperative venture between contract parties. At present, only humans have this facility.

⁶⁷ See Jerry Fodor, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind* (MIT Press 1987); Dennett, *The Intentional Stance*, n. 2, 7-11.

What is needed for agents to be able to engage in contracting is for them to be planning agents. A planning agent is an agent with the cognitive ability to have a future directed intention. There is a literature in philosophy and psychology that claims that future directed intentions play an important role in human practical reasoning.⁶⁸ An influential account of practical reasons is that of Donald Davidson, who claims that practical reasoning is the weighing of reasons, in the forms of beliefs and desires, for and against differing actions.⁶⁹ The idea here is that my intentions provide me with reasons for actions which lead to action. So, intentions are causes of actions. But what Michael Bratman and others have argued is that accounts like Davidson's do not go far enough because they do not account for intentions about the future. They only deal with intentions as reasons for actions for the moment of deliberation and the action to ensue. They do not account for the holding of a prior intention about achieving an end that would lead to further intentions about means or steps towards achieving those ends.

I do not see it as fatal to Davidson-type accounts that future directed intentions are not accounted for, but nevertheless the insight of Bratman and others on future directed intentions direct us to an important point of great relevance to contracting. And that is that a focus on future directed intention accommodates two different functions humans undertake in their lives, the need to be able to deliberate about future commitment and the need to plan our future actions together.

A. *Shared Intentionality*

That humans are planning agents means we engage in what Bratman calls shared cooperative activity or shared intentions.⁷⁰ Bratman's shared intentionality thesis is set forth below. Assume *J* to be a joint activity, such as contracting:

We intend to *J* if and only if

1. (a) I intend that we *J* and (b) you intend that we *J*.
2. I intend that we *J* in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b; you intend that we *J* in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b.

⁶⁸ See Michael E. Bratman, *Intentions, Plans, and Practical Reason* (CSLI 1999); Michael E. Bratman, *Faces of Intention: Selected Essays on the Intention and Agency* (Cambridge University Press 1999); Michael E. Bratman, "What is Intention?" in Philip R. Cohen, Jerry Morgan, & Martha E. Pollack, *Intentions in Communication* (MIT Press 1990): 15-31; Margaret Gilbert, *On Social Facts* (Princeton University Press 1989); Raimo Tuomela, *The Philosophy of Sociality: The Shared Point of View* (Oxford University Press 2010).

⁶⁹ Davidson, n. 2.

⁷⁰ Bratman, *Faces of Intention*, n. 68, Essay 5 on "Shared Cooperative Activity" and Essay 7, "Shared Intention and Mutual Obligation."

3. 1 and 2 are common knowledge between us.⁷¹

Stated in the context of contracting:

We intend to *contract* if and only if

1. (a) I intend that we *contract* and (b) you intend that we *contract*.
2. I intend that we *contract* in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b; you intend that we *contract* in accordance with and because of 1a, 1b, and meshing subplans of 1a and 1b.
3. 1 and 2 are common knowledge between us.

Bratman makes no appeal to a notion of obligation, such as a legal or moral obligation, but this is no obstacle. Contract law specifies a legal obligation when shared intent is present in the right context for contracting.

Shared intentionality is a uniquely human ability.⁷² It is the kind of intentionality that seems necessary for contracting is what anthropologist Michael Tomasello calls it joint intentionality, to distinguish it from the broader notion of collective intentionality.⁷³ The account from evolutionary anthropology is good for informing us about the sort of intentionality we are talking about for contracting and what might be needed for an artificial intelligence to be an entity with the agency to contract.

Individual intentionality is classic Davidsonian intentionality.⁷⁴ Interestingly, great apes have this capacity, though it will be expressed very differently because they do not use language. This first form of intentionality, known as “I” 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. “I” intentionality is adequate for attributing an intentional stance to an artificial intelligence, but it is not enough for contract formation.

Joint intentionality for Tomasello, or shared intentionality for Bratman, differs significantly from individual intentionality because it is about cooperation, albeit a cooperation for small groups. For humans, it is structured around linguistic communication. Evolutionary anthropologists trace its origins to small scale collaboration in human foraging. This kind of intentionality is commonly known as “we” intentionality. It is unique to humans. When humans engage in “we” 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

⁷¹ *Ibid.*, 131.

⁷² Michael Tomasello, *A Natural History of Human Thought* (Harvard University Press 2014): 35-36.

⁷³ *Ibid.*, 32-79.

⁷⁴ *See ibid.*, 7-31.

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, to capture it together with other humans as a joint goal.⁷⁵ Humans developed the second-personal point of view and to use that second-personal point of view to engage in a relationship of mutual recognition.⁷⁶

Finally, 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."⁷⁷ The cognitive capacities associated with collective intentionality have to do with the ability of entities to engage in self-governance that is responsive to a culture's norms of rationality.⁷⁸ Collective intentionality is a kind of intentionality that any entity must have to participate in a common culture. It is necessary for what Lon Fuller's characterizes as "subject to" law.⁷⁹ It concerns a broader set of issues I will leave for future discussion.

Are we at the point where an artificial intelligence could share intentionality with a humans to form a contract? Could two artificial intelligences share an intent to form a contract? These issues need substantial exploration beyond the scope of the theory set forth in this chapter. The question for the theorist is, is it plausible to theorize around these questions. I believe the answer to be yes, because it is plausible to expect the technology of artificial intelligence to reach the appropriate level of advancement.

B. *A General Theory of Contract*

Once we engage in the sort of ground clearing this paper does around the distinctive features of contractual obligation, it seems clear that other theories of contract reduce to questions about intent. I will make just a few brief remarks here, to lay some groundwork for future discussion.

Consider Charles Fried's theory of contract as promise.⁸⁰ It is probably the most influential theories about contract. But promise cannot describe what is salient or distinctive about contractual obligation. Many contracts are routinely formed without a basis in promise.⁸¹ Of particular relevance for us here, however, is that there are many promises that are not legally binding

⁷⁵ *Ibid.* 35-36.

⁷⁶ See Stephen Darwall, *The Second-Person Standpoint: Morality, Respect, and Accountability* (Harvard University Press 2000); TM Scanlon, *What We Owe to Each Other* (Harvard University Press 1998).

⁷⁷ Pascal Boyer, *Minds Make Societies: How Cognition Explains the World Humans Create* (Yale University Press 2018).

⁷⁸ Tomasello, n. __, at 80-123.

⁷⁹ Lon L. Fuller, *The Morality of Law* (Yale University Press 1966): 162-163.

⁸⁰ Charles Fried, *Contract as Promise: A Theory of Contractual Obligation* (Harvard University Press 1982).

⁸¹ Radin, n. 32; Kar & Radin, n. 32.

but which may be morally binding or binding as a matter of social consensus outside of the law. If we take a hard look at what makes a promise a feature of contracting, we will discover that it is the intention to create a legal relation with the promise, in the form of an objective manifestation of mutual assent to contracting, that makes the promise a contractual one, not the promise itself. We intend many promises to be not legal in nature, for which we reasonably can expect compliance. For example, if I promise my colleague to return a book I borrowed from her, we both may believe I have made a promise that is morally obligatory but not legally enforceable. We may share in an intent to create a moral obligation but not a legal one. It is the intent supporting the promise that provides the criteria for legal recognition of a particular kind of promise as a contractual obligation.

A similar argument can be made about Robert Barnett's theory of contract as consent.⁸² The consent theory of contract explains contract law as means to transfer property rights. It relies on the objective manifestation of an intention to be legally bound to transfer a property right as a core feature of contract. Its articulation appears to conflate consent with intent.⁸³ But certainly, consent and objective intent to contract are different. An agent can manifest an objective intent to contract and yet not consent to contract. Conversely, a person can consent to a contract and not manifest an objective intent. In fact, a theory of subjective intent about contract formation would seem to be required if actual consent to a contract is required, because there may be situations in which an agent provides sufficient objective indicia of intent but in truth does not want to enter a legally enforceable contractual obligation. More fundamentally for purposes of artificial agency, consent would take us too close to a requirement for what has been characterized as "strong AI": for an entity to qualify as intelligent in a strong AI sense, it actually has to think and have actual intentions associated with its actions and not just offer an intentional stance to humans.⁸⁴ What we have dealt with in this account is what may be understood as weak AI, which is reflected in Turing's imitation game and Dennett's intentional stance, and which is the standard conception of artificial intelligence by its developers.⁸⁵

⁸² Randy E. Barnett, "A Consent Theory of Contract" (1986) 86 *Columbia Law Review* 269; Randy E. Barnett, "Contract is Not Promise; Contract is Consent" in Gregory Klass, George Letsas, & Prince Saprai, *Philosophical Foundations of Contract Law* (Oxford University Press 2014): 43-57.

⁸³ Barnett, "Contract is Not Promise" n. 83, at 48.

⁸⁴ A longstanding disagreement in philosophy of mind is between those who argue that AI, or at least what is known as strong AI, is impossible. Strong AI is AI that *actually* thinks, is conscious, has a phenomenology of the particular experiences of life, and has the properties of intentionality that humans have. Weak AI is AI that acts *as if* it is thinking, conscious, and acting with intentionality. See Russell & Norvig, n. 12, at 1020-1033. While staking out a middle ground David Chalmers offers a good summary of the arguments. David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory* (Oxford University Press 1996): 313-332. This is an argument about whether artificial intelligence must have some form of inner life, some phenomenology of conscious experience, or true understanding or whether a simulation of these things will suffice, is beyond our scope here.

⁸⁵ Norvig & Russell, n 12, at 1020-1033.

For contracting purposes, it is enough that we be able to attribute intent and not actual consent to putative contract parties.

CONCLUSION

This chapter has set forth a philosophy of contract law formed around the most basic elements of the contracting process that of the manifestation of assent to contract based in the notion of objective intent to enter contractual relations. Objective intent has special qualities that make it the most direct route to understanding the core of contract formation and contract obligation. The focus on objective intent allows us to connect other minds – not just human minds – in melding both artificial and human agency in contract formation. As artificial intelligence progresses, so too, I believe, will my approach to understanding contractual obligation become more salient.

The philosophy in this chapter is analytical in the sense that I have not asked whether contract law and contract as an institution in a society *ought* to accept participation in its practices by artificial agency. This normative question is beyond our scope here. It is an important question that is in need of being addressed. It is a question of justification. Law makes a claim to authority or legitimacy to humans. It takes humans as having primacy of place as the subjects of law. How do we account for artificial intelligence if it reaches an appropriate level of agency and what would that level of agency be? That we can possibly share intention does not mean we should. These are questions that relate to the common rationality we share on a mass scale in the form of collective intentionality. This is perhaps our next discussion.