Spring 2018

A Legal Perspective on the Trials and Tribulations of AI: How Artificial Intelligence, the Internet of Things, Smart Contracts, and Other Technologies Will Affect the Law

Iria Giuffrida
William & Mary Law School, igiuffrida@wm.edu

Fredric Lederer
William & Mary Law School, filede@wm.edu

Nicolas Vermeys

A LEGAL PERSPECTIVE ON THE

TRIALS AND TRIBULATIONS OF AI:

HOW ARTIFICIAL INTELLIGENCE,

THE INTERNET OF THINGS, SMART

CONTRACTS, AND OTHER

TECHNOLOGIES WILL

AFFECT THE LAW

Iria Giuffrida, Fredric Lederer, and Nicolas Vermeys[†††]

DEDICATION AND APPRECIATION

Paul Giannelli and I first met when we, along with Ed Imwinkelried and Fran Gilligan, were colleagues on the faculty of what today is The Judge Advocate General's School and Legal Center. We then became co-authors of *Courtroom Criminal Evidence.*[1] As a teacher of Evidence and Criminal Procedure at William & Mary, I followed Paul's career with admiration. One of our leading evidence scholars, Paul has combined creative, scholarly thinking with pragmatic realism to become our nation's leading scientific evidence authority. He has always been ready and able to engage with important scientific advances.

In this short article, we address some of the legal issues that may flow from the combination of Artificial Intelligence, the Internet of Things, Smart Contracts, and related technologies. In doing so, we

---

[†††] Iria Giuffrida is Visiting Assistant Professor of Law and Associate Director for Research, Center for Legal and Court Technology, William & Mary Law School; Fredric Lederer is Chancellor Professor of Law and Director, Center for Legal and Court Technology, William & Mary Law School; Nicolas Vermeys is Associate Director of the Cyberjustice Laboratory, Professor, Université de Montréal's Faculty of Law and Visiting Associate Professor of Law, William & Mary Law School. The authors' work is supported by a grant from the Silicon Valley Community Foundation, funded in turn by Cisco. Inc.

1. EDWARD J. IMWINKELRIED, PAUL C. GIANNELLI, FRANCIS A. GILLIGAN & FREDRIC I. LEDERER, COURTROOM CRIMINAL EVIDENCE (4th ed. 2005).

acknowledge Paul's outstanding career and seek to follow his lead as we explore the legal implications of our world-changing technology.

-Fred Lederer, Chancellor Professor of Law and Director, Center for Legal and Court Technology

CONTENTS

INTRODUCTION

Imagine the amazement that a time traveler from the 1950s would experience from a visit to the present. Our guest might well marvel at:

∞   Instant access to what appears to be all the information in the world accompanied by the virtual elimination of personal privacy;

∞   Personal worldwide communication via voice, text, and images;

∞   Decisions and recommendations made by computers whether in the form of instantly implemented stock trades, recommended medical diagnosis, or criminal case bail release;

∞   Crypto-currencies such as Bitcoin implemented by blockchain, a distributed and decentralized electronic ledger held by all users that updates instantly;

∞   Electronic commerce based in significant part on what computers anticipate and persuade consumers to purchase;

∞   Manufacture by robots;

∞   Semi-autonomous and, soon, fully autonomous self-driving vehicles of all types.

And so much more . . .

As history has shown us, every technological advance is accompanied by legal questions.[2] We believe that our modern high-technology era will be faced by an unusual number of such questions growing out of what we will undoubtedly term, "artificial intelligence" ("AI"), but which in fact is the combination of advanced algorithms, important pools of data, usually referred to as "big data," and the many technologies that exploit these. Some questions are versions of traditional issues, such as tort liability for semi-autonomous or autonomous automobile collisions. Others may be termed novel: when, if at all, might a "computer" statement be hearsay or a "computer" be liable for tortious injury—or even murder[3]—or might it be sued for breach of copyright because the "computer" is considered a "person"? How will we define a "smart contract;" what knowledge and skills will a responsible lawyer need to know to avoid a successful malpractice suit?

With the assistance of our student colleagues at William & Mary Law School's Center for Legal and Court Technology, and faculty and supporting staff of the University of Montreal's Cyberjustice Laboratory, the three of us are engaged in trying to predict the nature of the legal issues that exist, that will clearly grow out of, and those that might stem from AI and related technologies. This Article is only an introduction to that task. It aims to add to the already numerous publications and journal articles written on the topic of law[4] and AI by

---

2.   *See, e.g.,* ETHAN KATSH, THE ELECTRONIC MEDIA AND THE TRANSFORMATION OF LAW (1989).

3.   *See, e.g.,* GABRIEL HALLEVY, WHEN ROBOTS KILL: ARTIFICIAL INTELLIGENCE UNDER CRIMINAL LAW 38 (2013) (relating that in 1981 a Japanese motorcycle employee was killed by a robot working next to him after the robot's algorithm determined that the employee was a threat to its mission, and that pushing the employee into an adjacent machine would remove the problem, which it did, as it killed the employee).

4.   *See e.g.,* RYAN CALO ET AL., ROBOT LAW (2016); MIREILLE HILDEBRANDT, SMART TECHNOLOGIES AND THE END(S) OF LAW (2016); JOHN FRANK

honing into what we believe to be the crux of the issue: AI-enabled devices exist in a technological ecosystem. Therefore, we cannot simply address the impact of a given technology without establishing how it will interact with others, more importantly how data will be generated, shared, used, and monitored by AI-enabled devices. The aim of this Article is to contribute further to a basic and useful understanding of the legal problems to be generated by that ecosystem, leaving to later articles more detailed discussions of those problems and related ones such as the critical and numerous privacy issues raised by these and related technologies.

Of course, anticipating the future does not easily lend itself to exhaustive prediction. What is absolutely sure is that the combination of the technologies addressed in this Article will change the world beyond anything most of us can anticipate and that the legal professions are unprepared for the legal consequences.[5]

Initially, this Article will define the relevant terms, such as "Artificial Intelligence," which can mean very different things. Emphasizing the impact of the combination of the related technologies, the Article will then survey the legal risks that can stem from algorithms, arguably the heart of AI. Next, this Article will briefly address Smart Contracts and some of their implications. Finally, this Article will discuss the need to create an environment where AI can flourish while co-existing with a society of rights.

WEAVER, ROBOTS ARE PEOPLE TOO: HOW SIRI, GOOGLE CAR, AND ARTIFICIAL INTELLIGENCE WILL FORCE US TO CHANGE OUR LAWS (2013); SAMIR CHOPRA & LAURENCE F. WHITE, A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS (2011). In fact, a quick Westlaw search for the expression "artificial intelligence" brings up over 2,500 journal and law review articles.

5.   This is not to say that governing bodies are ignoring the subject. Both federal and state organizations, for example, are attempting to encourage and regulate self-driving cars. *See generally Autonomous Vehicles/Self-Driving Vehicles Enacted Legislation*, NAT'L CONFERENCE STATE LEGISLATURES, http://www.ncsl.org/research/transportation/autonomous-vehicles-self-driving-vehicles-enacted-legislation.aspx [https://perma.cc/RJG5-436D] (last visited Jan. 25, 2018) ("Since 2012, at least 41 states and D.C. have considered legislation related to autonomous vehicles," and twenty-one states have enacted legislation). Advisory panels are being created to define problems and solutions. *See, e.g.*, Andrew Burt, *Leave A.I. Alone*, N.Y. TIMES (Jan. 4, 2018), https://www.nytimes.com/2018/01/04/opinion/leave-artificial-intelligence.html [https://perma.cc/44F8-FXQ6] ("[A] bipartisan group of senators and representatives introduced the Future of A.I. Act, the first federal bill solely focused on A.I. It would create an advisory committee to make recommendations about A.I.").

I. Artificial Intelligence, the Internet of Things, and Smart Contracts: What Does It All Mean?

One of the main issues that must be faced when addressing the legal underpinnings of technological innovations is rooted in the vocabulary used by those developing and marketing these tools. Information Technology ("IT") professionals, like lawyers, have developed a somewhat dense and opaque lexicon that is undeniably complex to master for the uninitiated. That being said, unlike legal terms which are expected to have a single definition unless otherwise stated in the statute, most technological constructs benefit from shifting meanings depending on the author.[6] This adds to the confusion of those who try to predict how the law should treat AI, for example, as authors cannot agree on what AI represents conceptually. Therefore, to borrow the language from Canadian author Hugh MacLennan, lawyers and IT specialists very much represent "two solitudes" who speak different languages, yet often using the same words.[7]

Given, however, that an Article like this one relies on a common understanding of somewhat novel concepts in order to carve out a legal framework, it is important to at least try and offer a general outline of the main terms popping up in the media which will undoubtedly find their way into the courtroom. Although the list of terms to choose from is long and ever-growing as new concepts seem to emerge daily, this Article will focus on the three interlinked, yet distinct, notions that have titillated the legal community in the last few years: AI,[8] the Internet of Things,[9] and Smart Contracts.[10]

A. What Is Artificial Intelligence?

According to common knowledge, the term "Artificial intelligence" may first have been coined by John McCarthy, Marvin L. Minsky, Nathaniel Rochester, and Claude E. Shannon,[11] in a 1955 paper, *A*

---

6.   *See* Statistics Can., *A Reality Check to Defining eCommerce*, Gov't Can. (1999), http://publications.gc.ca/collections/Collection/CS88-0006-99-06E. pdf [https://perma.cc/UCE8-6MNH] ("[A]s with any new concept, the understandings of what the terminology means are as diverse as the individuals involved. Hence it is often confused or misused.").

7.   Hugh MacLennan, Two Solitudes (1945).

8.   *See infra* Part I.A.

9.   *See infra* Part I.B.

10. *See infra* Part I.C.

11. *See, e.g.*, Gill Press, *Artificial Intelligence (AI) Defined*, Forbes (Aug. 27, 2017, 12:00 PM), https://www.forbes.com/sites/gilpress/2017/08/27/ artificial-intelligence-ai-defined/#45cc151f7661 [https://perma.cc/Z6NY-4U C4; *see also*, Chris Smith et al., *The History of Artificial Intelligence* 5

Proposal for the Dartmouth Summer Research Project on Artificial Intelligence.[12] The authors explained that:

> An attempt will be made to find how to make machines use language, form abstractions and concepts, solve kinds of problems now reserved for humans, and improve themselves. . . . For the present purpose the artificial intelligence problem is taken to be that of making a machine behave in ways that would be called intelligent if a human were so behaving.[13]

Fast forward to 2018, and although AI is talked about in the media almost every day, there is still no generally accepted definition of the term. Individual definitions run the gamut from a super-intelligent, humanoid, sapient, world-conquering robot to an app that suggests that the weather justifies wearing a coat. According to the Merriam-Webster dictionary, "Artificial Intelligence" can be defined as "[a] branch of computer science dealing with the simulation of intelligent behavior in computers," or [t]he capability of a machine to imitate intelligent human behavior."[14] This definition is at best misleading and functionally useless.[15]

Rather than taking this approach, some have defined AI by its components.[16] For example, while giving a lecture to the Council of Bars and Law Societies of Europe, Andrew Arruda, co-founder of Ross Intelligence,[17] presented AI as a blanket term encompassing four types

---

(2006), https://courses.cs.washington.edu/courses/csep590/06au/projects/history-ai.pdf [https://perma.cc/RKY3-CR53].

12.   The article was re-published in 2006. *See* John McCarthy et al., *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence: August 31, 1955*, AI Magazine, Winter 2006, at 12, 12.

13.   *Id.* at 2, 11.

14.   *Artificial Intelligence*, MERRIAM-WEBSTER ONLINE DICTIONARY, https://www.merriam-webster.com/dictionary/artificial%20intelligence [https://perma.cc/3BXT-QKEL] (last visited Mar. 4, 2018).

15.   This Article later asserts that AI based on machine learning basically exists when a computer, via its algorithms, can modify its implementing algorithms in order to better carry out the goals set by its major algorithms. *See infra* notes 19–22 and accompanying text.

16.   CCBE, *Presentation ROSS Intelligence by Andrew Arruda*, YOUTUBE (Nov. 18, 2016), https://www.youtube.com/watch?v=hJk-dQnn4M8.

17.   *See generally* ROSS, ROSSINTELLIGENCE.COM [https://perma.cc/YQ87-T6FS] (last visited Mar. 29, 2018); John Manes, *ROSS Intelligence Lands $8.7M Series A to Speed Up Legal Research with AI*, TECHCRUNCH, https://techcrunch.com/2017/10/11/ross-intelligence-lands-8-7m-series-a-to-speed-up-legal-research-with-ai/ [https://perma.cc/B2F4-XZSR] (last visited Feb. 11, 2018).

of technologies: machine learning, speech recognition, natural language processing, and image recognition.[18] Although the Authors of this Article would agree that these four concepts fall within the boundaries of AI, it could be argued that they do not actually represent distinct technologies as speech recognition and natural language processing could be seen as two sides of the same coin. Furthermore, both these technologies—as well as image recognition—can, and often do, rely on machine learning algorithms.

Of course, this begs the question: what are machine learning algorithms?

Let us first address the simpler of the two terms. In its most basic form, an algorithm is the set of software rules that a computer follows and implements. Put slightly differently, an "algorithm" is a program that evaluates data and executes given instructions. For example, in today's world, much of the day's stock market trading is conducted by highly complex algorithms rather than by people. The algorithm is the key to AI.[19] A computer's ability to function sufficiently well to carry out its programmed texts requires sufficiently adequate hardware. However, *what* the computer does is the result of the algorithms running in the computer's hardware.

Machine learning can be summarized as the ability of a computer to modify its programming to account for new data and to modify its operations accordingly.[20] It "uses computers to run predictive models that learn from existing data to forecast future behaviors, outcomes, and trends."[21] Machine learning therefore is dependent on data. The more data it can access, the better it can "learn." However, the quality of said data, the way the data is inputted into the system, and how the system is "trained" to analyze the data can have dire effects on the validity, accuracy, and usefulness of the information generated by the algorithm.[22]

---

18. CCBE, supra note 16, at 2:00-2:19.

19. A sufficiently well executed extraordinarily complex algorithm might well pass the "Turing test": can a remote human being distinguish a machine from a person? *Cf. ELIZA*, WIKIPEDIA, https://en.wikipedia.org/wiki/ELIZA [https://perma.cc/7YLP-MM7V] (last visited Mar. 4, 2018).

20. *E.g.*, a computer monitoring a factory assembly line determines that employees are more efficient in the afternoon in cooler temperatures than usual and drops the line temperature to 67 degrees from 69 degrees.

21. Jonathan Sanito et al., *Deep Learning Explained*, EDX, https://www.edx.org/course/deep-learning-explained-microsoft-dat236x-1 [https://perma.cc/G94C-9EG5] (last visited Mar. 4, 2018)].

22. *See* Pedro Domingos, *A Few Useful Things to Know About Machine Learning*, COMM. OF THE ACM, Oct. 2012, at 78, 78.

In short, an otherwise perfect algorithm can not only fail to accomplish its set goals but may prove affirmatively harmful. For example, the algorithm used by Google to answer user questions erroneously declared that former president Barack Obama, a Christian, was a Muslim.[23] In that case, the algorithm was not at fault. It simply gathered data from the Internet, "feeding" on websites that propagated false information. Its data pool was polluted, and the algorithm could not discern between "good" and "bad" data. Another example is that of the Microsoft chatbot, "Tay," which learned to interact with humans via Twitter.[24] Within twenty-four hours, the chatbot "became racist," for lack of a better word, because "Internet trolls"[25] had bombarded it with mostly offensive and erroneous data, i.e. inflammatory tweets, from which the Chatbot had "learned."[26]

Even when the data is accurate, the individual "training" the AI could infuse his or her own biases into the system. This may have been a factor in crime-predicting software that has led to the arrest of an unjustifiably high number of African Americans and other minorities,[27] as well as sentencing tools that predict higher rates of recidivism for these same individuals.[28]

---

23. Jack Nicas, *Google Has Picked an Answer for You—Too Bad It's Often Wrong*, WALL ST. J. (Nov. 16, 2017, 10:58 AM), https://www.wsj.com/articles/googles-featured-answers-aim-to-distill-truthbut-often-get-it-wrong-1510847867 [https://perma.cc/2PV2-HDA2].

24. Daniel Victor, *Microsoft Created a Twitter Bot to Learn from Users. It Quickly Became a Racist Jerk.*, N.Y. TIMES (Mar. 24, 2016), https://www.nytimes.com/2016/03/25/technology/microsoft-created-a-twitter-bot-to-learn-from-users-it-quickly-became-a-racist-jerk.html [https://perma.cc/JF5A-VXCH].

25. This is a slang term used to identify an Internet user who:

    sows discord on the Internet by starting quarrels or upsetting people, by posting inflammatory, extraneous, or off-topic messages in an online community (such as a newsgroup, forum, chat room, or blog) with the intent of provoking readers into an emotional response or of otherwise disrupting normal, on-topic discussion, often for the troll's amusement.

    See, for lack of a better source, *Internet Troll*, WIKIPEDIA, https://en.wikipedia.org/wiki/Internet_troll (last visited Jan. 25, 2018) [https://perma.cc/68V4-6A6P].

26. Victor, supra note 24.

27. CATHY O'NEIL, WEAPONS OF MATH DESTRUCTION 85–87 (2017).

28. Julia Angwin et al., *Machine Bias*, PROPUBLICA (May 23, 2016), https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing [https://perma.cc/45JM-QM5P]. If a software decision-making assistance tool for judges erroneously predicts racial minority members are more likely than others to violate pretrial release terms or to re-offend, then

Accordingly, the effective accuracy of an algorithm is dependent on both the programming and the data. This dictates a further, legally-troubling conclusion. If there are doubts about the results of an algorithm, one can at least theoretically inspect and analyze the programming that makes up the algorithm. Given the volume of data available on the Internet, however, it may be impossible to adequately determine and inspect the data used by the algorithm.

Machine learning should be envisioned as a spectrum that ranges from relatively simple algorithms to complex self-teaching systems that could eventually mirror the human brain in their complexity—if not their structure. This later subset of machine learning is usually presented as "deep learning," i.e. "a sub-field of machine learning, where models inspired by how our brain works are expressed mathematically, and the parameters defining the mathematical models, which can be in the order of few thousands to 100+ million, are learned automatically from the data."[29] Deep learning relies on what is referred to as neural networks. an interconnected group of nodes said to be modeled after the human brain.

This overview leads to two somewhat obvious, yet essential, observations. First, use of AI does not require or imply self-aware technology. It does, however, encompass technology that can substantially change sub-goals when necessary to maximize accomplishment of a larger goal, a possibility that permits unanticipated and potentially deadly consequences.[30] Second, although AI requires hardware, it should not be understood as such—hence, why the term "robot," although technically accurate, is somewhat misleading because of how robots have been depicted in science-fiction books and movies. AI should be understood as software incorporated in or installed on hardware to implement the designer's goals. This is essential to understand because, as is common knowledge, software can be hacked, pirated, or otherwise corrupted. Third, and most importantly, the vocabulary associated with AI can be somewhat misleading. Because of the use of terms like "intelligence," "learning," "teaching," etc., AI can sometimes be seen as a form of sentient being, a belief reinforced by Hollywood blockbusters. However. as stated bv Justice Mahonev in the Canadian case, *Apple Computer, Inc.* v. *Mackintosh Computers Ltd.*:[31]

---

that could be because of defects in the algorithm, defects in the underlying data, or an accurate reflection of a racially biased criminal justice system.

29.   Sanito et al., *supra* note 21.

30.   And as Hamlet mused, *"ay there's the rub."* WILLIAM SHAKESPEARE, HAMLET act 3, sc. 1; *see also infra* notes 48–49 (briefly discussing AI behavior when reacting to an unexpected situation).

31.   [1988] 1 F.C. 673 (Can.).

> The principal difficulty which this case has given me arises
> from the anthropomorphic character of virtually everything
> that is thought or said or written about computers.
> Words like "language", "memory", "understand",
> "instruction", "read", "write", "command", and many others
> are in constant use. They are words which, in their primary
> meaning, have reference to cognitive beings. Computers are
> not cognitive. The metaphors and analogies which we use to
> describe their functions remain just that.[32]

Although, as this Article addresses further on, there are those who
would grant AI legal personhood[33]—which might well be proven nec-
essary in the future—this Article posits that Justice Mahoney's wise
words remain valid when considering current AI algorithms. AI is not
simply R2-D2 or C-3P0, it also encompasses simple algorithms that
help you pick which movie to watch on Netflix or that protect your
mailbox from spam emails.

Therefore, to summarize, AI covers a gamut of technologies from
simple software to sentient robots, and everything in between, and
unavoidably includes both algorithms and data. This is why the current
buzz surrounding AI is somewhat inaccurately portrayed by many, as
it is not AI as a whole that has people talking, but rather advances in
machine learning and related technologies. At least initially, this may
simplify matters such as tort liability for injuries caused by AI.

B. What Is the Internet of Things?

The "Internet of Things" ("IoT"), or as Cisco's Maciej Kranz calls
it, "The Internet-of-Everything,"[34] describes the way in which so many
electronic devices communicate with each other, sharing data and some-
times even operations. More specifically, it refers to "the networking
capability that allows information to be sent to and received by objects
and devices, such as fixtures and kitchen appliances, using the
Internet."[35]

Conceptually, the IoT is much easier to grasp for those of us who
are less technologically inclined than the more complex notion of AI. A
refrigerator automatically reordering milk when its sensors notice that

---

32. *Id.* at 27.

33. *See* Lawrence B. Solum, *Legal Personhood for Artificial Intelligences*, 70
     N.C. L. REV. 1231 (1992); *see also* F. Patrick Hubbard, *"Do Androids
     Dream?": Personhood and Intelligent Artifacts*, 83 TEMP. L. REV. 405
     (2011).

34. *See* MACIEJ KRANZ, BUILDING THE INTERNET OF THINGS 12, 15 (2017).

35. *Internet of Things*, MERRIAM WEBSTER ONLINE DICTIONARY, https://
     www.merriam-webster.com/dictionary/Internet%20of%20Things   [https://
     perma.cc/9L3R-K27H] (last visited Feb. 9, 2018).

you are about to run out is simpler to explain than how the algorithms
that allows it to do so were programmed.

Yet, the IoT poses its own terminological challenges, and, like AI,
is lacking in a universally accepted definition. At the core of the con-
fusion is the link, or lack thereof, between the IoT and other similar
concepts such as machine-to-machine communications, or cyber-physi-
cal systems. For example, according to the National Science Founda-
tion, "Cyber-physical systems (CPS) are engineered systems that are
built from, and depend upon, the seamless integration of computational
algorithms and physical components."[36] Using this definition, one could
suggest that all IoT devices are actually cyber-physical systems, while
cyber-physical systems are not necessarily IoT devices as they are not
all connected to the Internet. However, the International Organization
for Standardization ("ISO") offers a different analysis:

> It became clear right from the outset that the notions of IoT,
> Cyber Physical Systems (CPS), and Machine to Machine
> Communications (M2M) were quite similar. This conclusion was
> reached based on observing M2M standardization activities . . . as
> well as academic research in the CPS area. Therefore, [the
> working group] expanded the scope of its search and identified
> about two dozen definitions for IoT, M2M, and CPS that were
> regarded as better than many others that were found. More
> definitions were found and added to the list of reasonable
> definitions later. Over the past two years, other notions such as
> the Industrial Internet, Internet of Everything, and Industrial IoT
> have been proposed [but were regarded] as too similar to the IoT.
> Hence, [the working group] decided to define the IoT in such a
> way that it would include the characteristics of all these similar
> notions. It is unlikely that in the long run more than one of these
> terms would survive.[37]

So how does ISO define the IoT? Although the organization admits
that "there is no way of capturing all the complexities of the IoT in a
2 to 3 line definition."[38] it still proposed one for the sake of discussion.
According to the ISO, the IoT is "[a]n infrastructure of interconnected
objects, people, systems and information resources together with intel-

---

36. *Cyber-physical Systems*, NAT'L SCI. FOUND., https://www.nsf.gov/funding/
    pgm_summ.jsp?pims_id=503286 [https://perma.cc/4V9B-UBPW] (last
    visited Feb. 9, 2018).

37. ISO/IECJTC1, INTERNET OF THINGS (IOT) PRELIMINARY REPORT 2014,
    at 2, https://www.iso.org/files/live/sites/isoorg/files/developing_standards/
    docs/en/internet_of_things_report-jtc1.pdf [https://perma.cc/AH9Q-SP7B]
    (last visited Mar. 29, 2018).

38. *Id.* at 3.

ligent services to allow them to process information of the physical and the virtual world and react."[39]

Arguments could obviously be made for or against this or other definitions of the IoT. However, for the purpose of this Article, the interesting aspect of the ISO definition is that it emphasizes the undeniable link between AI and the IoT, or rather the fact that the IoT relies on AI algorithms. And, to once again paraphrase Shakespeare: "there's the rub," for AI data likely comes to the algorithm via the Internet potentially originating from countless different sources. Consider, for example, that a computer responsible for stock market trades almost certainly is monitoring and responding to Internet-derived data describing financial transactions from all over the world. Given the immense number of devices and the amount of data available on the Internet, a computer that uses Internet-derived data can yield unpredictable results. As we will see, one of the most difficult issues inherent in AI is how to assure that the data used by a computer is in fact accurate. Not only is information originating on the Internet, such as on social media, often inaccurate, but the Internet also contains intentionally false data often spread extensively by "bots" and similar technologies that run automated tasks—such as spreading inflammatory content—at a higher rate than humanly possible.[40]

From a legal standpoint, this issue is rarely addressed when discussing the IoT. Most authors usually center their analysis on the legal implications of connected devices collecting and sharing personal data about their users.[41] Although the Authors fully recognize the need to study the effects of the IoT and AI more generally on privacy, the topic

---

39. *Id.* at 4.

40. Consider the allegations that the United States and other national elections have been intentionally influenced by false data such as computer produced or bot social media communications. *See e.g.*, Scott Shane, *The Fake Americans Russia Created to Influence the Election*, N.Y. TIMES (Sept. 7, 2017), https://www.nytimes.com/2017/09/07/us/politics/russia-facebook-twitter-election.html [https://perma.cc/KE6X-LSYS]; Kai Kupferschmidt, *Social Media 'Bots' Tried to Influence the U.S. Election. Germany May Be Next*, SCIENCE (September 13, 2017, 3:45 PM), http://www.sciencemag.org/news/2017/09/social-media-bots-tried-influence-us-election-germany-may-be-next [https://perma.cc/98W7-8EFS].

    If one is familiar with the programming of a given algorithm, it is even possible to intentionally create data that the algorithm will interpret as something entirely different. *See, e.g.*, Mark Harris, *Researchers Find a Malicious Way to Meddle with Autonomous Cars*, CAR & DRIVER (Aug. 4, 2017, 11:06 AM), https://blog.caranddriver.com/researchers-find-a-malicious-way-to-meddle-with-autonomous-cars/ [https://perma.cc/527D-F3A4] (among other matters, directional street signs could be altered in such a way as to fool the algorithm into interpreting them as speed signs).

41. *See, e.g.*, Rolf H. Weber, *Internet of Things—New Security and Privacy Challenges*, 26 COMPUTER L. & SECURITY REV. 23, 24 (2010).

remains outside of the scope of this Article which focuses on the impacts of the quality and availability of data, rather than the legality of collecting, processing, and retaining data.

C. What Are Smart Contracts?

A "Smart Contract" can be defined as a legal agreement that contains or exists in the form of an algorithm. Unlike a traditional contract, which only lays out the terms of agreement for subsequent execution, a smart contract *autonomously executes* some or all of the terms of the agreement. A smart contract can be extraordinarily sophisticated and complicated, executing via the Internet, for example, transactions at different costs and dates depending upon data such as currency exchange rates, stock market prices, costs of given raw materials, and anticipated weather conditions.

Notwithstanding their names, smart contracts are actually fairly "dumb" as they ultimately rely on code that contains a set of instructions determining what happens when certain circumstances occur. In this sense, even though they self-execute—thus not requiring any human intervention or any other form of intelligence—they remain "computable contracts"[42] which rely on being provided with data relevant to compliance or performance.

From a programming point of view, smart contracts are generally based on blockchains, a technology "that permanently records transactions in a way that cannot be later erased but can only be sequentially updated, in essence keeping a never-ending historical trail."[43] Originally created to support crypto-currencies such as Bitcoin,[44] the distributed ledger technology behind blockchains is now being used in other fields,

---

42. Harry Surden, *Computable Contracts*, 46 UC DAVIS L. REV. 629, 636 (2012) ("When a contract term is "computable," the parties have arranged for a computer to make automated, prima-facie assessments about compliance or performance (i.e., as in the comparison of payment terms to payment data).").

43. WILLIAM MOUGAYAR, THE BUSINESS BLOCKCHAIN: PROMISE, PRACTICE, AND APPLICATION OF THE NEXT INTERNET TECHNOLOGY xxi (2016).

44. *See* Satoshi Nakamoto, *Bitcoin: A Peer-to-Peer Electronic Cash System*, BITCOIN, https://bitcoin.org/bitcoin.pdf [https://perma.cc/9K8Z-P557] (last visited Mar. 29, 2018).

such as the management of land registries,[45] securities transactions,[46] and even trademark registration.[47]

For lawyers, smart contracts are a cause for concern for several reasons. First, it remains unclear whether a traditional lawyer can competently "draft" or program a smart contract, even with the assistance of an IT expert. Second, and this brings us back to AI: what happens if a smart contract infused with learning capacities accesses contaminated pools of data? This could cause the smart contract to be executed in ways incompatible with the parties' intent, bringing into question whether a smart contract can even be considered a contract.

This and other questions and examples make clear that the key AI issue is not merely whether a given AI-enabled device is safe, but rather how the aforementioned technologies will impact one another. Discussions of AI tend to focus on the hardware, as if we are dealing with self-aware, reasoning artificial beings. However, as should now be evident, one cannot speak of AI without taking into account all of the other associated technologies and the ways in which they all interact. It is actually a technological ecosystem. As this Article will now address, for legal purposes, this complicates life extensively.

II. WHAT ARE THE LEGAL RISKS STEMMING FROM THESE "NEW" TECHNOLOGIES?

As discussed in Part I, AI is based on algorithms. These algorithms can be written by humans, *or* with sufficient AI ability, a computer system can create its own algorithms in order to accomplish goals set by the master algorithms.[48] Since a computer will always follow its algorithm-supplied goals, we must be careful to anticipate ways in which a computer might so comply. For example, *IF* a computer

---

45. This is notably the case in Sweden. *See* Gertrude Chavez-Dreyfuss, *Sweden Tests Blockchain Technology for Land Registry*, REUTERS (June 16, 2016), https://www.reuters.com/article/us-sweden-blockchain/sweden-tests-block chain-technology-for-land-registry-idUSKCN0Z22KV [https://perma.cc/M8S Y-ZXJJ].

46. Michael Mainelli & Alistair Milne, *The Impact and Potential of Blockchain on the Securities Transaction Lifecycle* (May 9, 2016) (SWIFT Inst., Working Paper No. 2015-007), https://ssrn.com/abstract=2777404 [https://perma.cc/EM7E-PZ2X].

47. This type of service is one currently being offered by companies such as Cognate. *See generally Secure Your Company's Most Valuable Assets—Its Trademarks*, COGNATE, https://cognate.com/ [https://perma.cc/7R25-GZ2H] (last visited Feb. 9, 2018).

48. Cade Metz, *Building A.I. That Can Build A.I.*, N.Y. TIMES (Nov. 5, 2017), https://www.nytimes.com/2017/11/05/technology/machine-learning-artificia l-intelligence-ai.html [https://perma.cc/2C9Z-2TD8].

charged with keeping a sidewalk clean had the capacity to do so, absent
programming protections, it might well determine that human beings
cause trash and that to keep the sidewalk clean, it should remove all
people from the sidewalk.[49]

Arguably, the most important near-term legal question associated
with AI is who or what should be liable for tortious, criminal, and con-
tractual misconduct involving AI and under what conditions. This is
why it becomes essential to establish the risks stemming from an
overreliance on AI and identifying who can and should be held respon-
sible for adopting the counter-measures aimed at mitigating these risks.

A. A Survey of Legal Risks Stemming from an Overreliance on Algorithms

AI is already a part of many people's lives. But, to fully understand
what we usually mean when we refer to AI, we have to start with the
business world, which is rapidly adopting AI-enabled technologies to
enhance productivity and profit. Perhaps the most useful examples are
those that at first blush might appear to be highly limited, a far cry
from avenging computer intelligences. In his book, *Building the Internet
of Things*,[50] Maciej Kranz relates two examples from the mining indus-
try, which we now paraphrase:

∞   Rio Tinto, a global, open-pit mining concern "has the
largest fleet of giant autonomous trucks in the world"

---

49.   For a much more frightening real-world example, see HALLEVY, *supra* note
3, at xv, 38 (2013) (relating that in 1981 a Japanese motorcycle employee
was killed by a robot working next to him after the robot's algorithm
determined that the employee was "a threat to its mission," and that pushing
the employee into an adjacent machine would remove the problem, which it
did, killing the employee.). Science fiction fans will be familiar with Isaac
Asimov's Three Laws of Robotics which were intended to prevent such a
result:

> A robot may not injure a human being or, through inaction, allow a
> human being to come to harm.

> A robot must obey orders given it by human beings except where
> such orders would conflict with the First Law.

> A robot must protect its own existence as long as such protection
> does not conflict with the First or Second Law.

See, e.g., Isaac Asimov's "Three Laws of Robotics", AUBURN UNIVERSITY,
https://www.auburn.edu/~vestmon/robotics.html [https://perma.cc/222H-
6ZUQ] (last visited Feb. 9, 2018). As many have noted, the three laws will
not actually avoid the type of behavior that killed the Japanese employee.
See, e.g., HALLEVY, *supra* note 3, at 15–17 (2013) (asking what would a
police robot do when a perpetrator is threatening to kill a hostage and the
robot must kill the perpetrator to save the hostage?).

50.   KRANZ, *supra* note 34, at 13, 15.

transporting "more than 200 million tons of materials across approximately 3.9 million kilometers." Extreme conditions and extreme loads create major and expensive maintenance problems. Installation of sensors in the trucks connected via the Internet to computers able to evaluate the truck data permits preventive maintenance which forestalls breakdown, recovery, and repair.[51]

∞ Goldcorp operates an underground "connected" gold mine in Canada worked by more than 1,000 people. Implementation of multi-faceted technology allowed the company to "achieve real-time visibility, monitoring, and ventilation control" over a single wireless network that uses radio frequency identification (RFID) to provide live tracking of people and equipment. As a result, the company saves between $1.5 and 2.5 million dollars in energy costs for ventilation; in the event of emergency can locate people 45 to 50 minutes faster than before; and can locate and track its equipment.[52]

These types of operations combine sensors, connected equipment, and the data they supply to produce more efficient and profitable business. By incorporating machine learning algorithms, the system could, for example, redeploy miners to more productive areas of the mine. Of course, if, as we saw in Part I, the algorithm has been badly designed and poorly trained—i.e. the coder, for instance, did not test sufficiently rigorously the algorithm or the trainer incorporated a bias into the system—or has access to polluted pools of data, the safety of these miners could be put into jeopardy.

Similarly, a judge preparing to sentence an offender might consult an AI-enabled digital report and recommendation that will predict the probability of recidivism.[53]

Another example is that of now-anticipated autonomous vehicles. If we assume fully autonomous, self-driving cars, we might have the following: The user or passenger enters the car and speaks the destination. The car's internal computer communicates with multiple computers located elsewhere to determine the most efficient, safest and perhaps economical route. While *en route*, both the car's own sensors and those in other cars, on, above, below, and near the street monitor progress, automobile condition, and compliance with operational and traffic requirements. Mechanical and electronic functions—and if privately owned, perhaps the status of the owner's required payments—

---

51.   *Id.* at 47–49.

52.   *Id.* at 49.

53.   *See supra* note 27.

are all monitored with instant corrections. Police no longer need or, perhaps even legally, can stop the vehicle. Instead, they have full data access from the vehicle which, of course, is not being "driven" by a human being.[54] In the event of a traffic violation, who is responsible? Logic would dictate that it is those responsible for the technology oversight, but who exactly? The car's manufacturer? The AI programmer? The trainer? This would obviously depend on the source of the oversight—hardware, software, data, data sources, instruction transmission, etc.

These and other examples serve to show the complexity of: 1) establishing how liability should apply to AI; and 2) who should ultimately be held responsible if an AI-enabled device fails to function in the manner it was supposed to.

B. How to Address the Liability Issues Linked to Algorithms—Initially A Status Question

There are essentially three ways to address legislatively the liability issues linked to AI. First, AI-enabled devices can be treated as property and therefore be the responsibility of their users, owners, or manufacturers.[55] Second, they could be treated as "semi-autonomous beings," and fall under a legal regime similar to that of children[56] or persons with mental disabilities, or even one similar to the notion of

---

54. If one assumes full use of the Internet-of-Things for police to monitor data transmission to and from the autonomous car, see, for example, JOHN S. HOLLYWOOD ET. AL., USING FUTURE INTERNET TECHNOLOGIES TO STRENGTHEN CRIMINAL JUSTICE 4 (2015), https://www.rand.org/pubs/research_reports/RR928.html [https://perma.cc/85KK-34TX] (providing that police justifications for ordinary traffic stops should vanish as there will be no need to stop the car to investigate when vehicle operation data is already and immediately available to the police and the responsible person, organization, or computer(s) will not be a "driver").

55. *See, e.g.*, F. Patrick Hubbard, *Allocating the Risk of Physical Injury from "Sophisticated Robots": Efficiency, Fairness, and Innovation, in* CALO ET AL., *supra* note 4, at 25, 39.

56. In the United States, tort law in this area will vary by state. The general rule is that courts will ordinarily determine the reasonableness of children's negligence by comparison against a reasonable child of the same age, assuming the child is not engaged in an adult activity such as flying a plane. Some states, however, exempt children below a given age from tort liability or establish rebuttable presumptions about liability exposure. *See, e.g.*, 1A STUART M. SPEISER ET. AL., THE AMERICAN LAW OF TORTS § 5:16 (2013).

agency.[57] Third, like corporations, they could be treated as fully autonomous beings.[58]

From a legislative standpoint, the first model is relatively simple to imagine and implement. It would also be the least strenuous to implement as it would require very little by way of legislative amendments. In fact, foreign laws are already drafted in a way that allows for this scenario. For example, in Quebec, as in most civil law jurisdictions, the Civil Code states that "[t]he custodian of an inanimate object is bound to make reparation for injury resulting from the autonomous act of said object, unless he proves that he is not at fault."[59] This would be akin to the common law doctrine of *res ipsa loquitur* under which negligence is presumed if one's property causes harm to a third party.[60] In cases where no negligence on the part of the custodian, owner, or user is established, liability could be transferred to the manufacturer of the AI-enabled device.[61] This does bring up an interesting question of how to apportion liability among the manufacturer, programmer and trainer of the AI.[62]

---

57. David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 WASH. L. REV. 117, 122 (2014) ("[T]he key conceptual question that autonomous thinking machines will pose is whether it is fair to think of them as agents of some other individual or entity, or whether the legal system will need to decide liability issues on a basis other than agency.").

58. Hubbard, *supra* note 33, at 407 (2011). Note that machines with "personhood" successfully sued in tort could pay damages from their earnings or via a universal insurance pool likely funded from a percentage of the initial cost of the machine. *See infra* note 64.

59. Civil Code of Québec, S.Q. 1991, c 64, art 1465 (Can.) (emphasis added).

60. Byrne v. Boadle, 159 Eng. Rep. 299 (Exch. 1863). Put differently, this would establish an evidentiary presumption shifting the burdens of production and proof to the custodian seeking to avoid liability.

61. Vladeck, *supra* note 57, at 141–42.

62. *See, e.g.*, FLA. STAT. § 316.86 (2016) (exempting automobile manufacturers from liability when third-party AI is installed: "The original manufacturer of a vehicle converted by a third party into an autonomous vehicle is not liable in, and shall have a defense to and be dismissed from, any legal action brought against the original manufacturer by any person injured due to an alleged vehicle defect caused by the conversion of the vehicle, or by equipment installed by the converter, unless the alleged defect was present in the vehicle as originally manufactured."); *see also* WEAVER, *supra* note 4, at 56.

The third model—granting legal personhood to AI—is also rel-atively simple to legislate. It would necessitate AI-insurance[63] or the creation of a regime of compulsory compensation[64], but these are schemes that legislators have dealt with before and do not pose unique challenges as such. This model does, however, raise the more phil-osophical question of whether we consider autonomous vehicles, bots and other AI-enabled technology to be truly "beings" deserving of independent legal status. In the wake of IBM's Watson's win against its human Jeopardy opponents,[65] or Google's AlphaGo beating the

---

63. An insurance model for AI-enabled devices, dubbed the Turing registry, was notably proposed in Curtis E.A. Karnow. *Liability for Distributed Artificial Intelligences*, 11 BERKELEY TECH. L. J. 147, 193–94 (1996):

> [D]evelopers seeking coverage for an agent could submit it to a certification procedure, and if successful would be quoted a rate depending on the probable risks posed by the agent. That risk would be assessed along a spectrum of automation: the higher the intelligence, the higher the risk, and thus the higher the premium and vice versa. If third parties declined to deal with uncertified programs, the system would become self-fulfilling and self-policing. Sites should be sufficiently concerned to wish to deal only with certified agents. Programmers (or others with an interest in using, licensing or selling the agent) would in effect be required to secure a Turing certification, pay the premium and thereby secure protection for sites at which their agents are employed.

Id. Although this form of remedy might seem unconscionable to some, there is legal precedent in the United States. Such a system was put forth back in the days of slavery to account for the autonomous acts of slaves—concededly a discomforting comparison. *See* Jenny Bourne Wahl, *Legal Constraints on Slave Masters: The Problem of Social Cost*, 41 AM. J. LEGAL HIST. 1, 20 (1997) ("In some states, owners bore no liability for willful, malicious, intentional acts of slaves, just as masters did not pay for such acts committed by servants. The costs of these acts were thus spread widely over the slaveholding community.").

64. This could be modelled, for instance, on the International Oil Pollution Compensation Funds, created under the auspices of the International Maritime Organization pursuant to the 1992 International Convention on Civil Liability for Oil Pollution Damage and the 1992 International Convention on the Establishment of an International Fund for Compensation for Oil Pollution Damage. *See* *International Oil Compensation Funds*, IOPC FUNDS, https://www.iopcfunds.org/ [https://perma.cc/K4Z9-6KGH] (last visited Apr. 4, 2018). We are grateful to Michael Z. Snider, a current second-year William & Mary law student for this suggestion.

65. *See* Jason Hanna, *Computer Finishes Off Human Opponents on 'Jeopardy!'*, CNN (February 17, 2011, 5:50 AM), http://www.cnn.com/2011/TECH/innovation/02/16/jeopardy.watson/index.html [https://perma.cc/6TZJ-4ZX3].

world Go champion,[66] one could posit that computers can now be pro-
grammed to be as intelligent as humans—a related but clearly different
classification. However, this implies both that intelligence is no more
than the capacity to conduct probabilistic analysis and that intelli-
gence is perceived as the main criteria to establish legal capacity.
"Intelligence" is not enough for personhood, at least in most juris-
dictions. Rather, the test for capacity is that of reason; a person has to
be endowed with reason to be held civilly or criminally liable, to enter
into a contract, or to exercise other forms of legal autonomy.[67]

As Erich Fromm put it:

> Reason is man's faculty for *grasping* the world by thought, in
> contradiction to intelligence, which is man's ability to *manipulate*
> the world with the help of thought. Reason is man's instrument
> for arriving at the truth, intelligence is man's instrument for
> manipulating the world more successfully; the former is
> essentially human, the latter belongs to the animal part of man.[68]

Whether or not one agrees with Fromm's postulate, it remains un-
deniable that reason and intelligence are intrinsically linked and that
true "intelligence," for lack of a better word, is more than computing
capacities, no matter how sophisticated. A case in point: individuals
suffering from savant syndrome. Brought to public consciousness
through Dustin Hoffman's character in the 1988 film "Rain Man,"
savant syndrome "is a rare, but extraordinary, condition in which
persons with serious mental disabilities, including autistic disorder,
have some 'island of genius' which stands in marked, incongruous con-
trast to overall handicap."[69] Individuals afflicted with this condition will
often display impressive calculating abilities,[70] yet can still be con-
sidered legally incompetent.

---

66.   *See, e.g.,* Paul Mozur, *Google's AlphaGo Defeats Chinese Go Master in Win
       for AI,* N.Y. TIMES (May 23, 2017), https://www.nytimes.com/2017/05/
       23/business/google-deepmind-alphago-go-champion-defeat.html [http://
       perma.cc/A6R5-94V4].

67.   Arguably, this is inadequate as human beings are inherently sentient and
       self-aware as well. The issue of who or what is "human" is an old one, gaining
       importance in the field of "animal rights" as well. For a highly unusual
       variation, see EMILY BARTON, THE BOOK OF ESTHER 129–34 (2016), in
       which the protagonists debate whether artificial beings, "golems," must be
       both human and Jewish as they insist on Jewish prayer and performing
       traditional and sacred Jewish activities.

68.   ERICH FROMM, THE SANE SOCIETY 65 (1955).

69.   Darold A. Treffert, *The Savant Syndrome: An Extraordinary Condition.,*
       2009 PHIL. TRANSACTIONS OF THE ROYAL SOC'Y 1351, 1351.

70.   *Id.*

So how does this relate to AI? Like individuals suffering from savant syndrome, AI-enabled devices have great computing capacities, but lack in overall reason. This was brought to light, for example, by Tay's racist Twitter rants.[71] The chatbot was intelligent enough to generate coherent tweets, but lacked the reason to understand the insensitive nature, to put it lightly, of its postings. As Mireille Hildebrandt put it:

> It seems to me that artificial intelligence in itself does not qualify as [reasonable], even if some kind of consciousness would emerge. Animals have consciousness but we do not consider them fit to be subjected to legal punishment, because we have no indication that they can reflect on their actions as their own actions. Their consciousness is an awareness of the environment, without the concomitant awareness of this awareness which is typical of the human sense of self. Helmuth Plessner actually took this to be the crucial difference between humans and non-human life forms: the self-consciousness of the human person creates a distance between the self, the world and the self itself, condemning humans to what he called indirect directness, natural artificiality and a utopian position. To be sensitive to censure, rather than mere discipline, a subject needs to be conscious of its self, allowing the kind of reflection that can lead to contestation or repentance in the case of a criminal charge.[72]

The animal comparison is interesting as it seems to be a position shared by both legal scholars and AI experts. For example, Yoshua Bengio, one of the foremost international experts on machine learning, has stated on more than one occasion that the intelligence of most AI-enabled devices is comparable to that of a frog.[73] As frogs do not have legal personhood, logic would dictate that AI-enabled devices, for the very reasons described in the quote above, should not either. This would imply that if we reject placing AI within the inanimate property

---

71. Sophie Kleeman, *Here Are the Microsoft Twitter Bot's Craziest Racist Rants*, GIZMODO (March 24, 2016, 11:43 AM), https://gizmodo.com/here-are-the-microsoft-twitter-bot-s-craziest-racist-ra-1766820160 [https://perma.cc/HS97-JNDD].

72. Mireille Hildebrandt, *Ambient Intelligence, Criminal Liability and Democracy*, 2 CRIM. L. & PHIL. 163, 178 (2007).

73. *See* Lucie Luneau, *Retour Sur l'Intelligence Artificielle en 10 ans et 10 points*, ACS (May 15, 2004), http://www.acs.qc.ca/actualite/192-retour-sur-lintelligence-artificielle-en-10-ans-et-10-points.html [https://perma.cc/6F3Q-D7XF]. The author states [translation]: "The current learning ability of a computer is about that of a frog. 'To worry that a computer surpasses us would be as if ancient Egyptians worried about the pollution that will be created by the traffic of spaceships on Mars.'"

category, classic tort law applicable to animals[74] would best suit current advances in AI. Therefore, in the absence of known misconduct formerly committed by an AI entity, only a clearly "dangerous" AI-enabled device would dictate liability, or rather, the level of care that its manufacturer, programmer, or owner should take in its development.

However, if we are to believe IT pioneers like Bill Gates, Elon Musk, and Steve Wozniak,[75] the AI-animal metaphor could be short-lived, as AI is becoming increasingly powerful and could eventually reach a level of ability or consciousness equal to that of humans. If this is the case, does the position attributing legal personhood to AI become the only solution? Some authors such as Lawrence B. Solum have held this position for years.[76] As Solum puts it, refusing legal personhood to AI "is akin to American slave owners saying that slaves could not have constitutional rights simply because they were not white or simply because it was not in the interests of whites to give them rights."[77] Although we disagree with this premise, which in our view understates the true effects of slavery on the African-American community to this day,[78] slavery laws—when stripped from their historical, societal, and moral contexts—do offer interesting insight on how more advanced AI could be approached from the standpoint of liability.

As explained in *Wright v. Weatherly*,[79] in some states "a master was liable for every [slave's] trespass, whether the act be done when in the master's service, or not, and whether with or without the master's knowledge."[80] Putting aside the obvious ethical and legal repulsion to

---

74. *See generally* Behrens v. Bertram Mills Circus, Ltd. [1957] 2 QB 1, 11 (Eng.). The acts of wild animals give rise to strict liability. Others, especially domestic animals impose tort liability only if harm is foreseeable.

75. Pater Holley, *Apple Co-Founder on Artificial Intelligence: 'The Future Is Scary and Very Bad for People'*, WASH. POST (March 24, 2015), https://www.washingtonpost.com/news/the-switch/wp/2015/03/24/apple-co-founder-on-artificial-intelligence-the-future-is-scary-and-very-bad-for-people/?utm_term=.5b9ced4fdbff [http://perma.cc/Y6J8-47HV].

76. *See* Solum, *supra* note 33, at 1261; *see also* Hubbard, *supra* note 33, at 434; Susan W. Brenner, *Humans and Humans+: Technological Enhancement and Criminal Responsibility*, 19 B.U. J. SCI. & TECH. L. 215, 244–48 (2013).

77. Solum, *supra* note 33, at 1261.

78. *E.g.*, MICHELLE ALEXANDER, THE NEW JIM CROW: MASS INCARCERATION IN THE AGE OF COLORBLINDNESS 13 (2010).

79. 15 Tenn. (7 Yer.) 367, 378 (1835).

80. *Id. quoted in* Jacob I. Corre, *Thinking Property at Memphis: An Application of Watson*, 68 CHI.-KENT L. REV. 1373, 1376 (1993); *see also* Anthony R. Chase, *Race, Culture, and Contract Law: From the Cottonfield to the Courtroom*, 28 CONN. L. REV. 1, 29 (1995)("[I]f the slave was acting as a tradesman or carrier, the courts held the master liable for the slave's trespass or negligence since the master in such a situation invited the public to have

one person owning another, this outcome makes sense from a purely compensatory standpoint as slaves had no means to offer financial redress to their victims. The same logic could apply to AI as computers have no property, while their owners, manufacturers and programmers do.

C. Liability in the Near Future

In the immediate future, it seems clear that AI technology will be regarded as property. It is unlikely that an AI device would be held civilly or criminally liable for harm done by it. Rather, the primary issue likely will be the classic one of civil liability under tort law. The owner or operator will be liable for injury caused by its property whether "intelligent" or not. Product liability and negligence[81] will be the primary causes of action. Although the law may be clear in concept, it may be very difficult to apply in practice given the IoT and impossibility of tracing the sources of data relied upon by an algorithm. Imagine a dam failure caused by an AI control system reliant on thousands of sensors supplied by multiple vendors, data supplied from independent third parties, many of which are derived from other AI devices, with decision-making shared with other non-owned AI devices. The *Restatement (Third) of Torts—Product Liability* § 5 declares:

> One engaged in the business of selling or otherwise distributing product components who sells or distributes a component is subject to liability for harm to persons or property caused by a product into which the component is integrated if:
>
> the component is defective in itself, as defined in this Chapter, and the defect causes the harm; or
>
> (b)(1) the seller or distributor of the component substantially participates in the integration of the component into the design of the product; and
>
> > (2) the integration of the component causes the product to be defective, as defined in this Chapter; and
> >
> > (3) the defect in the product causes the harm.[82]

---

confidence in the slave's ability . . . ."); Wahl, *supra* note 63, at 19 ("Entrusting one's slave was a double-edged sword, however, because owners were often responsible for injuries caused by their slaves, much as masters can be liable for the actions of their servants.").

81. Negligence here obviously includes also medical and other forms of professional malpractice.

82. RESTATEMENT (THIRD) OF TORTS: PROD. LIAB. § 5 (AM. LAW. INST. 1998). In Quebec, the law is simpler; a manufacturer is responsible for a

In one sense, this will permit the classic tort suit: sue everyone. Pragmatically, however, given a sufficiently large enough harm, IoT distributed AI blame may be so large as to defy legal resolution. Assuming an adequate duty of care, we may be unable to prove factual fault or, if we can, proximate cause.[83] Put differently, the harm caused may have been unforeseeable from the perspective of a specific component manufacturer.

And, to be fair, we might add the contractual issue. Is a predictive AI a "good" or a "service" under Article 2 of the Uniform Commercial Code, given the differences that classification can yield? When is a smart contract a "contract," an alternative to the classic contract,[84] or a device?

In short, the AI age starts with traditional legal concepts increasingly applied to new and previously unforeseen circumstances impelling legal change. This has happened before, of course, but the AI age will not only be immense in scope it will also proceed incredibly quickly. Our legal systems tend to be reactive and not proactive, especially when we cannot predict what the future will be like. One author writes that in 1880 experts charged with predicting what New York City would look like a hundred years later reported that it would be destroyed. The manure that would be generated by the more than six million horses needed by the city's people would make it uninhabitable.[85] The modern internal combustion engine and the automobiles it produced was unpredictable. Predicting the evolution of AI and its related technologies may be equally unsuccessful.

Of course, there is an inherent risk that, if we wait, contemporary liability rules, which in the United States are designed to not only compensate injured victims but also to deter wrong doing, will stifle AI

"safety defect." *See* Civil Code of Québec, S.Q. 1991, c 64, art 1468 (Can.). What does this mean in a distributed causation environment?

83. Perhaps breathing new life into the landmark torts case *Palsgraf v. Long Island R.R. Co.*, 162 N.E. 99 (1928).

84. There is discussion in the literature of smart contracts being considered as an alternative to legally enforceable contracts or as somewhat analogous to the way in which letters of credit operate. *See, e.g.*, Stephen McJohn & Ian McJohn, *The Commercial Law of Bitcoin and Blockchain Transactions* 17 (Suffolk Univ. Law Sch., Legal Studies Research Paper Series, Paper No. 16-13, 2016), https://ssrn.com/abstract=2874463 [https://perma.cc/M923-8ATU].

85. JEFF STIBELL, BREAKPOINT: WHY THE WEB WILL IMPLODE, SEARCH WILL BE OBSOLETE, AND EVERYTHING ELSE YOU NEED TO KNOW ABOUT TECHNOLOGY IS IN YOUR BRAIN 23–24 (2013) (citing STEVEN D. LEVITT & STEPHEN J. DUBNER, SUPERFREAKONOMICS: GLOBAL COOLING, PATRIOTIC PROSTITUTES, AND WHY SUICIDE BOMBERS SHOULD BUY LIFE INSURANCE 8–10 (2009)).

innovation while we ponder how best to change them. If nothing else, that poses our last matter: how can AI flourish while staying within the confines of a society of rights?

III. HOW CAN AI FLOURISH WHILE STAYING WITHIN THE CONFINES OF A SOCIETY OF RIGHTS?

In his 1999 article, *The Law of the Horse: What Cyberspace Might Teach Us*,[86] Lawrence Lessig asked a series of questions on how those in the legal community should address the regulation of cyberspace. As Lessig put it:

> [L]aw faces a choice—whether to regulate to change this architectural feature, or to leave cyberspace alone and disable this collective or individual goal. Should the law change in response to these differences? Or should the law try to change the features of cyberspace, to make them conform to the law? And if the latter, then what constraints should there be on the law's effort to change cyberspace's "nature"? What principles should govern the law's mucking about with this space? Or, again, how should law *regulate*?[87]

To borrow a famous quote attributed to Spanish philosopher George Santayana, "Those who cannot remember the past are condemned to repeat it."[88] In this sense, we should study how the law of cyberspace came to be, as Lessig's interrogations remain extremely relevant when addressing AI.

As Lessig posited, technology—in his example, cyberspace; in our case, AI—can be regulated in different manners. The classic route for regulation, of course, remains legislation. However, the author continues by suggesting that "Code" could also be the key to regulating technology. These are therefore the two avenues this Article will now broach as they pertain to AI.

A. Changing Laws to Address AI Innovations

As legal professionals, our initial reaction when faced with technologies we do not quite understand is often to take the legislative route and draft a legal framework destined to control the use and spread of these technologies. AI has not escaped this trend as many states have

---

86. Lawrence Lessig, *Commentary: The Law of the Horse: What Cyberspace Might Teach Us*, 113 HARV. L. REV. 501 (1999).

87. *Id.* at 505.

88. Matthew Caleb Flamm, *George Santayana (1863–1952)*, INTERNET ENCYCLOPEDIA PHIL., https://www.iep.utm.edu/santayan/ [https://perma.cc/QJ9S-HNPT] (last visited Apr. 3, 2018).

already adopted legislation aimed at curtailing the use of AI in certain fields.[89] In fact, some authors are even predicting the drafting of a Uniform Artificial Intelligence Act by the end of the decade.[90] Even famous businessman Elon Musk has implored legislators to act quickly in regulating AI.[91]

Unfortunately, to borrow a few lines from Justice Easterbrook's famous "Law of the Horse" speech—the very speech that inspired Lawrence Lessig to publish his aforementioned article on the same topic—"Beliefs lawyers hold about computers, and predictions they make about new technology, are highly likely to be false. This should make us hesitate to prescribe legal adaptations for cyberspace. The blind are not good trailblazers."[92] Although Justice Easterbrook's general thesis can be, and was,[93] disputed, history has proven him right when it comes to trying to predict and legislate on technological change.[94] In fact, his statement can already be verified in one field of AI, that of self-driving cars.

To this day, twenty-one states have adopted legislation regarding self-driving cars, and more are expected to follow suit.[95] Even the US government is currently working on a bill to regulate the use of autonomous vehicles.[96] As this technology is still in its infancy, the drafters of these bills have taken to predict the future, and some of their predictions have already proven to be problematic. For example, in the

---

89. This is the case, for example regarding driverless cars. *See* Statistics Can., *supra* note 6; *see also* To Provide for Information on Highly Automated Driving Systems to be Made Available to Prospective Buyers, H.R. 3388, 115th Cong. (1st Sess. 2017).

90. WEAVER, *supra* note 4, at 61.

91. Ali Breland, *Elon Musk: We Need to Regulate AI Before 'It's Too Late'*, HILL (July 17, 2017), http://thehill.com/policy/technology/342345-elon-musk-we-need-to-regulate-ai-before-its-too-late [https://perma.cc/U85X-9ENP].

92. Frank H. Easterbrook, *Cyberspace and the Law of the Horse*, 1996 U. CHI. LEGAL F. 207, 207 (1996).

93. *See generally* Lessig, *supra* note 86.

94. For example, the Utah Digital Signatures Act, UTAH CODE ANN. §§ 46-3-101 to 46-3-504 (West 1995), which was adopted in 1995, was later repealed by 2006 Utah Laws, c. 21 § 13 (West), notably because it required the use of a specific technology that proved ill-chosen.

95. *See Autonomous Vehicles Self-Driving Vehicles Enacted Legislation*, NAT'L CONF. STATE LEGISLATURES (Jan. 2, 2018), http://www.ncsl.org/research/transportation/autonomous-vehicles-self-driving-vehicles-enacted-legislation.aspx [https://perma.cc/GP5W-ZSFR].

96. *See* To Provide for Information on Highly Automated Driving Systems to be Made Available to Prospective Buyers, H.R. 3388, 115th Cong. (1st Sess. 2017).

District of Columbia, an autonomous vehicle must have "a driver seated in the control seat of the vehicle while in operation who is prepared to take control of the autonomous vehicle at any moment."[97] This obviously limits how self-driving cars could be used and designed. For example, under these rules, GM's recently announced autonomous cars without steering wheels or pedals[98] will never be able to drive on D.C. roads. It also means that driverless taxi services[99] will not be able to establish themselves in the Capitol. This might be exactly what the drafters of the Automated Vehicle Act of 2012 had in mind, or it could simply be that, six years ago, they could not fathom that strides in AI would make it possible to have fully automated vehicles. Whichever the case may be, this demonstrates that the technology is evolving in a manner that is incompatible with what the drafters of these laws had in mind.

Of course, getting back to Justice Easterbrook's statement, this is not to say that we shouldn't legislate on AI, smart contracts, or the Internet of Things, or wait until we have understood all there is to know about these technologies—something that could take centuries[100]—before adopting further AI-related legislation. History does teach us, however, that we should be careful in drafting said laws.[101] To quote iconic French jurist Jean Carbonnier, "one should always tremble when legislating."[102] However, how should the current legal framework be adapted—through the modification of current laws, or the adoption of new legislation—to take into account AI innovations?

---

97. D.C. CODE § 50-2352(2) (2013); *see also* WEAVER, *supra* note 4, at 56.

98. Alex Davies, *GM Will Launch Robocars Without Steering Wheels Next Year*, WIRED (Jan. 12, 2018), https://www.wired.com/story/gm-cruise-self-driving-car-launch-2019/ [https://perma.cc/PRK8-9GDF].

99. Timothy J. Seppala, *Waymo's Driverless Taxi Service Will Open to the Public Soon*, ENGADGET (July 11, 2017), https://www.engadget.com/2017/11/07/waymo-autnomous-taxi-phoenix/ [https://perma.cc/FC99-2MFE].

100. In WALTER J. ONG, ORALITY AND LITERACY (30th Anniversary ed. 2012), the author explains that it was only with the advent of the Internet that we came to fully understand how paper, as a technology, had truly impacted our lives. In fact, the author argues that technology can only truly be understood in hindsight, i.e. when it has been replaced by another. *Id.* at 2–3.

101. On this issue, see Roger Brownsword, *So What Does the World Need Now? Reflections on Regulating Technologies*, *in*, REGULATING TECHNOLOGIES— LEGAL FUTURES, REGULATORY FRAMES AND TECHNOLOGICAL FIXES 23 (Roger Brownsword & Karen Yeung eds., 2008).

102. The original quote reads: "Ne légiférez qu'en tremblant." OLIVIER ABEL, PAUL RICOEUR, JACQUES ELLUL, JEAN CARBONNIER, PIERRE CHAUNU: DIALOGUES 75 (2012).

In order to answer this specific question, we suggest, as noted above, to study recent history and how the law of the Internet has evolved. For all of the discussions about "Internet sovereignty"[103] and how "cyberspace law is different," very few laws were ultimately adopted to strictly address Internet-related issues, The Digital Millennium Copyright Act,[104] and Communications Decency Act[105] being the main exceptions to this rule. In most other Internet-related issues, current legislation and common law rules were tweaked or simply applied as is. Keeping this in mind, one could argue that the same should be true for AI.

For example, in a recent New-York Times Op-Ed, Oren Etzioni proposed three rules that he believes should apply to A.I:

∞ an AI system must be subject to the full gamut of laws that apply to its human operator;

∞ an AI system must clearly disclose that it is not human; and

∞ an AI system cannot retain or disclose confidential information without explicit approval from the source of that information.[106]

These rules, which are more of a tip of the hat to Isaac Asimov's aforementioned three laws of robotics than directives aimed at state legislators, do somewhat support the argument that current laws should apply to AI.[107] The problem is, which ones, and how should they be adapted?

---

103. John Perry Barlow, *A Declaration of the Independence of Cyberspace*, ELECTRONIC FRONTIER FOUND. (1996), https://www.eff.org/cyberspace-independence [https://perma.cc/37QR-HV9S].

104. 17 U.S.C. § 512 (2012).

105. 47 U.S.C. § 230 (2012).

106. Oren Etzioni, *How to Regulate Artificial Intelligence*, N.Y. TIMES (September 1, 2017), https://www.nytimes.com/2017/09/01/opinion/artificial-intelligence-regulations-rules.html [https://perma.cc/C247-KMEW].

107. This follows the general rule that "[w]hen AI agents act autonomously, we expect them to behave according to the formal and informal norms to which we hold our fellow humans. As fundamental social ordering forces, law and ethics therefore both inform and adjudge the behavior of AI systems. The dominant research needs involve both understanding the ethical, legal, and social implications of AI, as well as developing methods for AI design that align with ethical, legal, and social principles." NAT'L SCI. & TECH. COUNCIL, THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN 26 (2016), https://obamawhite

According to the National Science and Technology Council, the answer to this question lies in classic risk analysis:[108]

> [T[he approach to regulation of AI-enabled products to protect public safety should be informed by assessment of the aspects of risk that the addition of AI may reduce, alongside the aspects of risk that it may increase. If a risk falls within the bounds of an existing regulatory regime, moreover, the policy discussion should start by considering whether the existing regulations already adequately address the risk, or whether they need to be adapted to the addition of AI.[109]

Although risk analysis is a process that has been used by lawmakers for years, it remains more prevalent in other fields. For example, risk analysis is at the very core of cybersecurity, i.e. the degree to which information technology is safe from unwanted external interference. Over the years, numerous conceptual frameworks were developed to structure risk analysis as it relates to cybersecurity.[110] Although all have valid tenets, we are partial to Bruce Schneier's simplified five-step process:

1) What assets are you trying to protect?

2) What are the risks to these assets?

3) How well does the security solution mitigate those risks?

4) What other risks does the security solution cause?

5) What costs and trade-offs does the security solution impose?[111]

To answer these questions, one must first understand the concept of risk. Risk is usually defined as the probability that a threat can exploit a vulnerability in the system before the proper safeguards are put

---

house.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NS TC/national_ai_rd_strategic_plan.pdf [https://perma.cc/M43S-XP2S].

108. On the general topic of risk analysis, see DAVID VOSE, RISK ANALYSIS (3rd ed., 2008).

109. NAT'L SCI. AND TECH. COUNCIL, PREPARING FOR THE FUTURE OF ARTIFICIAL INTELLIGENCE (2016), https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf [https://perma.cc/56T2-FP33].

110. On this topic, see NICOLAS VERMEYS, RESPONSABILITÉ CIVILE ET SÉCURITÉ INFORMATIONNELLE (2010).

111. BRUCE SCHNEIER, BEYOND RISK: THINKING SENSIBLY ABOUT SECURITY IN AN UNCERTAIN WORLD 14–15 (2003).

into place.[112] Threats obviously include events such as surreptitious ex-
ternal "hacking" from a network or the Internet, but they also include
such intrusions as an employee sitting down at a friend's computer over
the lunch hour and making improper use of it.

Because AI-enabled devices frequently use data from the Internet
or implement their algorithms via the Internet, AI functions are espe-
cially vulnerable to cybersecurity threats. In July 2017, for example,
Forbes reported that "*Criminals Hacked a Fish Tank to Steal Data
from a Casino.*"[113] The fish tank was connected to the Internet to permit
remote monitoring of water conditions, and the thieves used that con-
nection as the route into the casino's computers.[114]

Getting back to applying risk analysis to AI from a legislative
standpoint, if we adapt Schneier's five-step process to legislative analy-
sis regarding AI, the process could be imagined as follows:

1) What rights are you trying to protect?

2) What are the risks that AI poses to these rights?

3) How well does current legislation mitigate those risks?

4) What risks would the application of current legislation to AI
cause?

5) What costs and trade-offs does current legislation impose?

Looking at these steps, the main issue remains that of identifying
the risks associated with the use of AI under step 2. Only then will we
be able to establish whether current legislation can sufficiently mitigate
those risks under step 3.[115] As for steps 4 and 5, they are mostly linked
to the risk of current legislation stifling innovation. Getting back to the
Internet parallel, the "notice and takedown" doctrine[116] was created for

---

112. IRA WINKLER, ZEN AND THE ART OF INFORMATION SECURITY 26–27
(2007).

113. Lee Mathews, *Criminals Hacked a Fish Tank to Steal Data from a Casino*,
FORBES (July 27, 2017), https://www.forbes.com/sites/leemathews/2017/
07/27/criminals-hacked-a-fish-tank-to-steal-data-from-a-casino/#1547e65c3
2b9 [https://perma.cc/PDF7-4XQQ].

114. *Id.*

115. One could argue that insofar as code must periodically be updated to protect
against accidental and intentional risk—and that most computer users fail
to update their systems, we should deter this negligent behavior by
legislating liability at least for those whose failure to install security upgrades
harms others or their property.

116. *See* ORG. FOR ECON. COOPERATION AND DEV., THE ROLE OF INTERNET
INTERMEDIARIES IN ADVANCING PUBLIC POLICY OBJECTIVES 144 (2011)
("Legal frameworks such as the European Union E-Commerce Directive

that very reason. It was meant to ease liability constraints that existing legislation put on Internet service providers.

This model is voluntarily imperfect as it starts from the postulate that legislation is the only way to curtail the risks caused by AI. However, there are other, sometimes more successful, ways to arrive at this same end through the use of code.

B. Coding Legal Constructs and Barriers into Algorithms

As Lawrence Lessig put it:

> In real space, we recognize how laws regulate—through constitutions, statutes, and other legal codes. In cyberspace we must understand how a different "code" regulates—how the software and hardware (i.e., the "code" of cyberspace) that make cyberspace what it is also regulate cyberspace as it is. As William Mitchell puts it, this code is cyberspace's "law." "Lex Informatica," as Joel Reidenberg first put it, or better, "code is law."[117]

This statement, which was made regarding cyberspace, holds as true with respect to AI-enabled devices. Computers exist to perform given functions. These functions are programmed in by their programmers— who serve somewhat as legislators as they can force a device to act in a certain manner or forbid it from doing so. Isaac Asimov's afore-mentioned three laws of robotics serve this point. The reason, according to Asimov's fictional universe, a robot:

∞ May not injure a human being or, through inaction, allow a human being to come to harm;

∞ Must obey orders given it by human beings except where such orders would conflict with the First Law; and

∞ Must protect its own existence as long as such protection does not conflict with the First or Second Law,

is because its programming does not allow it to go against these "laws." In this sense, code could be used to ensure compliance with current legislation. For example, autonomous vehicles can be programmed to obey the speed limit, making speeding violations a thing of the past. Of course, as we discussed earlier in this article, there exists an issue of

---

(ECD) or the US Digital Millennium Copyright Act DMCA create a safe harbor from liability for copyright infringement for various Internet actors when they meet certain conditions. One of the common elements of these regimes is that intermediaries must respond when they receive notice of an alleged infringement from the rights holder or his or her representative, by expeditiously removing the alleged infringing content.").

117. LAWRENCE LESSIG, CODE: VERSION 2.0, 5 (2006).

liability when there is a flaw in the code. But, what if the code is not flawed, yet the AI stops obeying the "laws" it was pre-programmed to obey? What if AI acted in a way that is *inexplicable* by reference to the code and is, in fact, incompatible with it? After all, anyone who has read an Isaac Asimov novel or seen a movie based on his books knows that robots will ultimately break the three aforementioned laws if only by being confronted by unforeseen circumstances. This is where Lessig's teachings stop being useful when discussing code as a means of controlling AI; it is also the major issue we are confronted with not only from a legal standpoint, but from a societal one as well.

As eluded to in the first section of this Article, AI-enabled devices are dependent on data. The more data they have access to and are trained with, the better they can predict an outcome or address a given situation. In this sense, "data analytics," which can be envisioned as the sophisticated and complex analysis by computer of enormous amounts of data, called "big data," is really at the heart of the current boom in AI. Given appropriate data, a computer's algorithms will produce a given result. Inadequate and flawed data will produce erroneous results.[118] A good example of this is Amazon's ability to suggest books a customer might like—if you do not train the algorithm properly, for example if you buy a book for a niece or nephew without indicating to the platform to disregard said purchase when making its suggestions, it will most probably offer poor recommendations going forward.

Data comes in many forms. Like a human being, a classic movie robot listens, sees, and often "feels." The robot takes in the raw observational data and then pursuant to its algorithms interprets the data and decides what, if anything, to do as a result. A computer that places stock trades without human intervention is doing the very same thing but using different data obtained in a different form.

This is where an issue could arise. In the case of advanced machine learning, i.e. deep learning, devices will eventually outgrow their initial coding and use new sets of data to produce an outcome. This implies that the calculation that led to said outcome is unknown to the consumer of the generated answer—which is often the case as the algorithm is protected by trade secrets[119]—or worse, unknown to the programmers

---

118. For example, incomplete or erroneous data could cause bias within the AI algorithm. *See* NAT'L SCI. AND TECH. COUNCIL, *supra* note 109, at 30 ("In the criminal justice system, some of the biggest concerns with Big Data are the lack of data and the lack of quality data. AI needs good data. If the data is incomplete or biased, AI can exacerbate problems of bias. It is important that anyone using AI in the criminal justice context is aware of the limitations of current data.").

119. As one author puts it: "algorithmic opacity is a largely intentional form of self-protection by corporations intent on maintaining their trade secrets and

because the AI-enabled device has acted upon data that they are unaware of or, unbeknownst to them, it has created its own algorithms to "solve problems."[120] In other words, "[a]s Machine Learning algorithms get smarter, they are also becoming more incomprehensible."[121] In this sense:

> [T]he fact that Machine Learning algorithms can act in ways unforeseen by their designer raises issues about the 'autonomy,' 'decision-making,' and 'responsibility' capacities of AI. When something goes wrong, as it inevitably does, it can be a daunting task discovering the behavior that caused an event that is locked away inside a black box where discoverability is virtually impossible.[122]

This opacity issue[123] is the one that seems most daunting for lawmakers. Although legislation can always be passed to make a protected line of coding available for analysis in case of an accident,[124] how does one identify how an algorithm produces an erroneous result when even its programmers cannot explain how this result was attained?

This finding has led to political pressure "to have some form of explanation for any AI-based determination."[125] But this does not simply imply the need to have transcripts, for example, of the case law a robot lawyer has consulted to arrive to its decision, it further implies the need to comprehend the whole technological ecosystem in which a given AI-enabled device resides. This brings the analysis back to this Article's initial thesis: the main risk of the increasing reliance on AI, and, therefore, the most difficult obstacle for regulators, does not reside in the technology itself, but rather in the interaction between AI-

---

competitive advantage." *See* Jenna Burrell, *How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms*, BIG DATA & SOCIETY, Jan.–June 2016, at 1, 3.

120. Will Knight, *The Dark Secret at the Heart of AI*, MIT TECH. REV. (Apr. 11, 2017), https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/ [https://perma.cc/29NZ-RH9J].

121. Colin Lewis and Dagmar Monett, *AI & Machine Learning Black Boxes: The Need for Transparency and Accountability*, KDNUGGETS, https://www.kdnuggets.com/2017/04/ai-machine-learning-black-boxes-transparency-accountability.html [https://perma.cc/TR6G-GTV7] (last visited Feb. 4, 2018).

122. *Id.*

123. *See* Burrell, *supra* note 119, at 1.

124. It should be pointed out, however, that "A call for code 'audits' (where this means reading the code) and the employment of 'auditors' may underestimate what this would entail as far as the number of hours required to untangle the logic of the code within a complicated software system." *See id.* at 4–5.

125. *See* NAT'L SCI. AND TECH. COUNCIL, *supra* note 109, at 31.

enabled devices—the sharing of information between databases erro-
neously believed to be in silos. If a computer hacker wants to corrupt
an AI-enabled device, the hacker can certainly erase data or tamper
with its coding, but a much more insidious means to the end is to add
invalid and unverified data to the device's database and let it learn
itself into chaos.[126]

CONCLUSION

When one lives in such rapidly changing times, it is difficult to gain
sufficient perspective to grasp how myriad changes interface with each
other. Often it seems to be enough to just hold on to some form of se-
cure support. The goal in sharing the preliminary views on how the
increasing use of AI, the IoT, smart contracts, and other technologies
discussed in this Article will affect the law is to emphasize that they
exist in a technological ecosystem. Although it is important to consider
the legal implications of each new technology on its own, our view is
that more attention should be given to the risks that are posed when
new technologies, especially those that are AI-enabled, are inter-
connected and interact with each other in ways that can at least seem
to be unfathomable.

Understanding AI and its related technologies can be difficult.
Consider, however, their impact on daily human life. Although many
are reasonably concerned about technological unemployment, consider
what smart contracts may do to legal practice. No doubt we will have
standard "boilerplate" to be tailored to a specific transaction, but what
of sophisticated custom work? What will lawyers need to know about
the technological ecosystem, and what skills will they, or perhaps their
AI assistants, associates, or partners, be able to do to produce, inspect,
and enforce a smart contract?

Our increasing reliance on AI, which certainly has its useful and
justifiable ends, is in no means a challenge for the legal system. This
Article has sketched, for example, some of the difficulties of establishing
tortious liability for AI-enabled unlawful acts. It has posited that the
legal system may respond in any number of ways, from relying on classic
tort law regimes of negligence and strict liability, to considering AI

---

126. Curtin E.A. Karnow makes a similar point: "Indeed, one way to "poison" a
robot is to interfere with its on-the-job training as it seeks to make patterns
from instances in the environment by substituting in misleading training
data—that is, faking the environment." *See* Curtis E.A. Karnow, *The
Application of Traditional Tort Theory to Embodied Machine Intelligence*,
in CALO ET AL., supra note 4, at 51, 60. The Author notably refers to an
article by Alex Armstrong, *Poison Attacks Against Machine Learning*, I
PROGRAMMER (July 19, 2012), http://www.i-programmer.info/news/105-
artificial-intelligence/4526-poison-attacks-against-machine-learning.html
[https://perma.cc/R6FE-MBSE].

devices "semi-autonomous beings," to even granting them independent legal personhood. None of these options are obvious. Each will require the legal system and its practitioners to engage in understanding and conceptualizing what AI is; how best the legal system can control it, if at all, by way of rules, code, or both; and the extent to which our legal system needs to adapt the growing complexity of the technological ecosystem.

Surely, members of the legal professions must engage with these topics now lest we be entirely unprepared when faced with immediate need for legal advice, legislation or rule-making, or case resolution. There is an ancient Chinese saying, "May you live in interesting times." That we are doing so, we can say with certainty. We would do well, however, to recognize that that saying is usually said to be a curse. Let us work proactively to ensure that, legally at least, AI may prove a blessing and not a curse.