

Living in the Present, Anticipating the Future: Ascribing Liability for Artificial Intelligence

Aman Rehan and Hammad Ali Kalhoro*

Abstract:

For any legal system, determining how liability will be ascribed to a particular person is a difficult task. However, a recently popularised conundrum in legal literature considers the question of legal liability for artificially intelligent computer systems. With the advent of COVID-19, the adoption of new technologies is accelerating, and the role of AI in our lives is only going to increase. What is often overlooked is that such technologies are usually premised on the “deep learning” system, creating uncertainty in decision making, experience-based learning, and reactions to events. Considering the issue of ascribing liability for harms caused by AI, this paper scrutinises these shortcomings. It highlights how legal systems have the propensity to do more in the promulgation of industry-wide standards relating to AI products. A failure to do so may have catastrophic consequences with the rapid development of AI technology and the increasing reliance on it by humans.

Introduction

In 1842 Ada Lovelace,¹ an aspiring computer scientist, showcased an anomalistic combination of several abilities, including a proper appreciation of the scope, capability, and future of computer science and technology. She perceived human beings to have become too comfortable maintaining the traditional way of doing things – the “comfort” in question refers to the notion of stagnation of the development of humanity. Therefore, in a then inconsequential event, her penning down of the first algorithm for a computing engine would forever alter the way we would interact with each other, and more importantly with machines. Since then, the constantly evolving world of technology has created significant legal challenges which can easily be mistaken for “anomalies.”

The world of Artificial Intelligence (“AI”) is but one stream of this transformative technology expected to have an everlastingly influence on the world of robotics, transportation, manufacturing, cybersecurity, and even medicine.² The benefits are clear, however, the use of AI to complete tasks also involves an undertaking of a certain degree of risk for error. Given the sheer number of products and services that rely on AI, there will naturally be instances in which AI does not produce desirable results. While the majority of these failures will be benign,

* Aman Rehan and Hammad Ali Kalhoro are BA-LL.B (Hons.) graduates from LUMS. Aman is a lawyer based in Lahore, and Hammad is a management consultant based in Karachi.

¹ Eugene E. Kim and Betty A. Toole, ‘Ada & The First Computer’ (1999) *Scientific American* 76, 78 <http://www.cs.virginia.edu/~robins/Ada_and_the_First_Computer.pdf> accessed 13th November 2019.

² Yudong Zhang and others, ‘Artificial Intelligence and its Applications’ (2014) 10 *Mathematical Problems in Engineering* 1, 7 <https://www.researchgate.net/publication/261548333_Artificial_Intelligence_and_Its_Applications/link/00463537d95b102743000000/download> accessed 14th November 2019.

the law must adequately cover situations wherein the failure of AI can directly cause tangible harm to both people and property.³ This conundrum has led to increased advocacy for a re-evaluation of consumer liability laws around the globe.

Considering this perplexing legal challenge, this paper aims to explore the potential of product liability laws as an effective mechanism for addressing AI harms. The following legal questions will be explored in detail: firstly, are algorithms and products similar, and if so, what metric can be used to establish their similarity? Secondly, can certain algorithms be compartmentalised as “products” using the metric described? If so, what kind of liability regime will be applicable to them? Thirdly, do the existing legal instruments adequately protect AI consumers? Fourthly, what can be done to overcome the shortcomings of existing legal instruments? And finally, if liability can be ascribed to robots, should rights be granted to them as well? A systematic scrutinisation of these questions will help uncover the extent to which work needs to be done to protect consumers from the potential dangers of AI.

Definition of Artificial Intelligence

Before grappling with these questions, it is prudent to delve into a brief exposition of both the history and definition of AI, because AI as we know today is a product of historical developments rooted in religion, mythology, literature, and even pop-culture. Robert M. Geraci highlights the ways in which technologists have derived inspiration regarding AI from stories found in scriptures and popular culture: “to understand robots, we must understand how the history of religion and the history of science have twined around each other, quite often working towards the same ends and quite often influencing another’s methods and objectives.”⁴ The history of AI is commonly traced back to Charles Babbage and Ada Lovelace, who are deemed to have not only predicted the advent of AI but also put together designs of machines which were geared towards carrying out “intelligent tasks”.⁵ However, AI is not a child of the modern era; and the concept of intelligent beings being created from inanimate objects can be traced back to ancient texts.⁶ Along with scriptures, AI has also been explored in literature and the arts,⁷ as well as pop culture.⁸

³ Dr. Saleemi Amershi, ‘Embracing AI Failure’ (2009) *CSCW University of Texas* 1, 2 <<https://sites.utexas.edu/goodsystemscscw/files/2019/10/GoodSystemsCSCW2019WorkshopPapers.pdf>> accessed 2nd January 2020.

⁴ Robert M. Geraci, *Apocalyptic AI: Visions of Heaven in Robotics, Artificial Intelligence, and Virtual Reality* (Oxford University Press 2010) 147.

⁵ Christopher D. Green, Thomas Teo, and Marlene Shore (ed), *The Transformation of Psychology* (American Psychological Association Press 2001), 133; Ada Lovelace, ‘Notes by the Translator’ reprinted in R.A. Hyman (ed), *Science and Reform: Selected Works of Charles Babbage* (Cambridge University Press 1989) 268–310.

⁶ Chinese mythology and ancient Sumerian myths have alluded to the creation of mankind from “clay and blood” and while Chinese myths present humankind being made from “the yellow earth,” holy scriptures such as the Quran also allude to the creation of man from “a clot of congealed blood.” See T. Abusch, “Blood in Israel and Mesopotamia”, *Emanuel: Studies in the Hebrew Bible, the Septuagint, and the Dead Sea Scrolls in Honor of Emanuel Tov* (Brill 2003) 673; —, ‘Nuwa,’ <<http://www.newworldencyclopedia.org/entry/Nuwa>> accessed 3 Feb 2020; AI- Quran 96:2.

⁷ From Mary Shelly’s *Frankenstein*, in which the author warns about the human ambitions of creating intelligence, to Homer’s *Iliad*, in which a blacksmith had “servant maids” which he made from gold. See Jordan (tr), Homer, *The Iliad* (University of Oklahoma Press, 2008) 1, 352

⁸ Popular cinema has also advanced the advent of AI - this can be seen from the rather innocent “C-3PO” from

While religion and popular culture alike have provided insight into the development of AI, the myriad of representations and portrayals have led to misleading impressions in people's minds. However, legislation or regulation based on such impressions is not acceptable in any developed legal system. This principle is also expounded by legal theorist Lon L. Fuller, who defined eight formal requirements for a legal system to function in conjunction with a set of moral norms which allows humans the opportunity to not only engage with the law but also amend their actions accordingly. One of these requirements is that the citizens under a legal system must know of the standards which are applicable to them, implying that the laws should be comprehensible.⁹ Therefore, without a proper definition, the application of a regulatory mechanism to something as omnipotent, rapidly changing, and fluid as AI is a Herculean task. The definition to be used for this paper is the one proposed by Jacob Turner in his book *Robot Rules: Regulating Artificial Intelligence*: "Artificial Intelligence is the ability of a non-natural entity to make choices by an evaluative process."¹⁰

Within this definition, it is implied that the ability to make choices confers a certain level of autonomy, albeit not absolute autonomy. An artificially intelligent entity will be able to make an autonomous choice even if there is human input at any stage. As this paper focuses specifically on algorithms, this paper will follow Jack Balkin's classifications which treat both robots and algorithms as being part of the "algorithmic society".¹¹ In an "algorithmic society", societal organisation revolves around social and economic decision making by algorithms. The algorithms not only make the decisions but also carry them out in some cases. In this sense, robots and AI merely become a "special case of the Algorithmic society".¹² Additionally, the "algorithms" referred to in this paper are those which are computerised. These algorithms can cause damage without any physical embodiment (other than computer hardware) or human intervention.

The limitations which come with functional definitions, however, apply to any legislative effort. Hence, while it is important to define AI for conferring certainty into the law, it is also imperative to avoid precise boundaries and ossify the law. This is also logical given the rapid developments which are made in this field. In this paper, algorithms will be compartmentalised into the larger ambit of machine learning and adaptation, which occurs whenever a machine can alter its data, structure or program in a way that its performance in the future is expected to improve.¹³ The term "machine learning" was first defined by Arthur Samuel as computers being given the "ability to learn without being explicitly programmed."¹⁴

the Star Wars franchise to more complex conceptions of robots with a moral compass such as "Robocop" or "Terminator."

⁹ Lon L. Fuller, *The Morality of Law* (Yale University Press 1969).

¹⁰ Jacob Turner, *Robot Rules: Regulating Artificial Intelligence* (Palgrave Macmillan 2019) 27-33.

¹¹ Jack M. Balkin, 'The Three Laws of Robotics in the Age of Big Data' (2017) 78 *Ohio State University Law Journal*.

¹² *Ibid* 11.

¹³ Nils J. Nilsson, *Introduction to Machine Learning: An Early Draft of a Proposed Textbook* (Department of Computer Science, Stanford University 1998) 1 <<https://ai.stanford.edu/~nilsson/MLBOOK.pdf>> accessed 1 June 2018.

¹⁴ Andres Munoz, 'Machine Learning and Optimization' (2015) Courant Institute of Mathematical Sciences New

This categorisation results from AI being capable of “independent development” i.e. the ability to learn from data sets in a manner which is unforeseen by its designers.¹⁵

Relevance

Since this paper lies in an intersection of law and technology, it might be deemed to be too futuristic by some. Often, one is not even aware of the leaps being made in the field of technology. Indeed, it is common for companies to produce new technologies through upgrades and software patches; while these changes may be unnoticeable at first, they are cumulatively quite significant. An example of this is the changing user interface of social media platforms such as Facebook and Instagram. The tendency to ignore incremental changes may lead to undesirable yet avoidable consequences. McKinsey and Co., an international management consultancy company, has provided research which estimates that the technological revolution is “happening ten times faster and at 300 times the scale, or roughly 3000 times the impact.”¹⁶ Verily, urgency in this case is not only justified via the magnitude of change but also consolidated by a sharp increase in the number of aggrieved people around the globe.

The culmination of all the fears related to AI was the horrific death of Elaine Herzberg on 18 March 2018, which played a significant role in bringing AI technology to the forefront of both the local and international media. Herzberg, a 49-year-old resident of Arizona, was immediately declared dead after being struck by a Volvo SUV. The vehicle was said to have been cruising at a speed of 80 kph at night in Tempe. The horrifying incident was directly attributed to the AI lacking “the capability to classify an object as a pedestrian unless that object was near a crosswalk,” as was affirmed by the National Traffic Safety Board, or NTSB in Arizona¹⁷. As a direct consequence of this shortcoming, it could not correctly predict her path and concluded that it needed to brake just 1.3 seconds before it struck her as she wheeled her bicycle across the street a little before 10 p.m.¹⁸ For critics, the *laissez-faire* attitude adopted by the state of Arizona was particularly problematic. Many went as far as to question the rationale behind introducing such nascent technology to the state, specifically without giving much forethought to its potential dangers. A fate similar to Elaine’s was also suffered by Joshua Brown, a 40-year-old resident of Ohio, after he placed his newly purchased Tesla Model S in its self-driving “autopilot” mode. A malfunction of the AI at the heart of Tesla’s autopilot mode resulted in its failure to distinguish a large white 18-wheel truck from a trailer. Resultantly, the car attempted to drive at full speed under the trailer, amounting to the fatality.¹⁹ An example closer to home, within Pakistan, can be that of the machine-learning algorithm guiding the U.S.

York University 1 <https://cims.nyu.edu/~munoz/files/ml_optimization.pdf> accessed 1 June 2018.

¹⁵ Turner (n 10) 7.

¹⁶ Richard Dobbs, James Manyika, and Jonathan Woetzel, ‘No Ordinary Disruption: The Four Global Forces Breaking All the Trends’ (*McKinsey Global Institute*, 2015) <<https://www.mckinsey.com/mgi/no-ordinary-disruption#>> accessed 12 February 2020.

¹⁷ DeArman, ‘The Wild Wild West: A Case Study Of Self Driving Vehicle Testing In Arizona’ (2019) 61 *Arizona Law Review* 991.

¹⁸ *Ibid.*

¹⁹ Megan McArdle, ‘How safe are driverless cars? Unfortunately, it’s too soon to tell’ *The Washington Post* (20 March 2015) <https://www.washingtonpost.com/opinions/no-driverless-cars-arent-far-safer-than-human-drivers/2018/03/20/5dc77f42-2ba9-11e8-8ad6-fbc50284fce8_story.html> accessed 12 February 2020.

drone program. It is argued, in a report published by *Ars Technica*, that ‘SKYNET’ (the algorithm at the heart of the planes) may have wrongly targeted thousands of innocent civilians, leading to many unnecessary deaths.²⁰ It was also found that the Algorithm performed well strictly in terms of the outcomes it was trained for – with 0.008% of the targets being wrongly classified.²¹ However, if this data were viewed not as mere numbers, around 15000 innocent people were killed. All these cases highlight AI’s propensity to cause physical harm; however, such harms may not always be physical.

For instance, in late 2013, IBM teamed up with the Cancer Center of the University of Texas in the hope of developing a new “Oncology Expert Advisor” system. The first line of their launch press release stated the following: “MD Anderson is using the IBM Watson cognitive computing system for its mission to eradicate cancer.”²² Five years following the press release, a review of the internal IBM documents uncovered how their AI system was giving not only erroneous, but quite dangerous, cancer treatment advice. Ultimately, the entire venture failed to achieve IBM’s ambition, while simultaneously costing them \$62 million.²³ Thankfully, the AI system was trained on hypothetical patient data, resulting in only monetary loss rather than loss of life.

Another product that proves the potential for non-physical harm through AI reliance is that of the new Apple iPhone X.²⁴ A well marketed feature of the new phone was its “Face ID” technology which allows its owner to unlock their phone by simply showing their face to the front camera. Apple described this mechanism as being 10 times more secure than the traditional fingerprint mechanism. One year after the release of the phone, hackers successfully attempted to utilise 3D printed masks as a loophole to the system. A Vietnam-based security firm, Bkav, affirmed these claims and further stipulated that at a mere cost of \$200, people could access the personal data of anyone who relied on the Face ID technology.²⁵ The work of Bkav provides a fascinating glimpse into the shortcomings of AI. More importantly, it shows

²⁰ Christian Grothoff and J.M Proup, ‘The NSA’s SKYNET program may be killing thousands of innocent people’ (*Ars Technica*, 16 February 2016) <<https://amp.theguardian.com/science/the-lay-scientist/2016/feb/18/has-a-rampaging-ai-algorithm-really-killed-thousands-in-pakistan>> accessed 13 September 2021.

²¹ Martin Robbins, ‘Has a rampaging AI algorithm really killed thousands in Pakistan?’ (*The Guardian*, 18 February 2016) <<https://amp.theguardian.com/science/the-lay-scientist/2016/feb/18/has-a-rampaging-ai-algorithm-really-killed-thousands-in-pakistan>> accessed 13 September 2021.

²² MD Anderson News Release, ‘MD Anderson Taps IBM Watson to Power “Moon Shots” Mission’ (*MD Anderson Cancer Centre*, 18 October 2013) <<https://www.mdanderson.org/newsroom/md-anderson--ibm-watson-work-together-to-fight-cancer.h00-158833590.html?fbclid=IwAR0FxQEG4txoo5ldUuqeNeFTv6mt9cGVqSkE7tL0-VVxHTSRNdbIV9QpAuM>> accessed 8 August 2021.

²³ Mathew Herper, ‘MD Anderson Benches IBM Watson In Setback For Artificial Intelligence In Medicine’ (*Forbes*, 19 February 2017) <<https://www.forbes.com/sites/matthewherper/2017/02/19/md-anderson-benches-ibm-watson-in-setback-for-artificial-intelligence-in-medicine/?sh=78cb92163774&fbclid=IwAR2tQq2YYFR6POVWe1qJ7Fta2TmtqHYLYNvKJeJIX0FYDLT4tnjMqim2Bu>> accessed 8 August 2020.

²⁴ Garofalo, Rimmer and Van Hamme, ‘Fishy Faces: Crafting Adversarial Images to Poison Face Authentication’ (2014) *KU Leuven* 4 <https://www.usenix.org/sites/default/files/conference/protected-files/woot18_slides_garofalo.pdf> accessed 2 March 2020.

²⁵ Webster, Kwon, Clarizio, ‘Anthony & Scheirer, Visual Psychophysics for Making Face Recognition Algorithms More Explainable’ (2018) *Arxiv Cornell Tech* 6 <<https://arxiv.org/pdf/1803.07140.pdf>> accessed 2 March 2020.

that the rise of technology coincides with an increase in our reliance on algorithms to regulate our daily lives. The resultant risk to privacy and data security is a consequence that can be linked directly or indirectly to AI, as data-dependency is a fundamental characteristic of algorithms.

The legal implications become further pronounced when one delves into contracts involving AI. Members of the public enter into contractual arrangements daily, through a tap on their smart phone. Ideally, such an arrangement should involve both parties being fully aware of the obligations which bind them. In reality, however, mobile app users generally gloss over the “Terms and Conditions” or the “End User License Agreement” before clicking the “accept” box. Such quasi-hidden contracts are a feature of many of the free utilities which users enjoy- from mapping services to photo-editing applications. A significant manifestation of the use of data acquired through these quasi-hidden contracts occurred in 2016 when Cambridge Analytica, a data- analysis firm, used the psychological profiles of millions of American Facebook users for the Trump campaign in the US elections.²⁶ It is clear, therefore, that more must be done to determine the important legal questions raised at the helm of ascribing liability, especially when we fail algorithms or when algorithms fail us.

Are Algorithms and Products Similar?

In ancient Rome, there was debate on whether liability could be ascribed to a horse, which was characterised as a “semi-intelligent entity.”²⁷ Although there was a view that the horse should pay for its actions, the more popular view was to extend the liability to its human owner. The US Judge Frank Easterbrook elaborated on this example while opposing the idea of a separate regime for cyber law, stating that doing so is as futile as asking for a “Law of the Horse.”²⁸ Instead, he advocated for general rules to be studied in order to approach specialised areas of the law—otherwise, “the Law of the Horse is doomed to be shallow and to miss unifying principles.”²⁹ Keeping this principle in mind, this paper will approach the idea of creating a product liability regime for algorithms by extrapolating from already—established legal principles.

It is undoubted that the positive benefits of AI are immense: they can eliminate human error by making decisions which are more consistent, efficient, objective, and reliable. However, as mentioned before, even AI is susceptible to mistakes; in the event of an AI error, aggrieved humans will seek compensation and turn to liability regimes already in place. The questions of attribution which arise at this point include how the fault in the algorithm should be organised? Who should be held liable in the event of an AI error? And the type of approach

²⁶ Nicholas Confessor, ‘Cambridge Analytica and Facebook: The Scandal and the Fallout So Far’ (*The New York Times*, 4th April 2020) <<https://www.nytimes.com/2018/04/04/us/politics/cambridge-analytica-scandal-fallout.html>> accessed 8 April 2021.

²⁷ D.I.C. Ashton-Cross, ‘Liability in Roman Law for Damage Caused by Animals’ (1953) 11 (3) *The Cambridge Law Journal* 395–403.

²⁸ Frank H. Easterbrook, ‘Cyberspace and the Law of the Horse’ (1996) *University of Chicago Legal Forum* 207–215, 207.

²⁹ *Ibid.*

to be taken towards relief and remedy. To approach these questions, one must create a comparison between algorithms and products, as the existing product liability framework needs to accommodate the advances being made in technology.

Firstly, it is pertinent to define a “product.” A product is simply defined as “something that is made to be sold, usually something that is produced by an industrial process.”³⁰ While this definition does not immediately clarify the distinction between a “product” and an “algorithm”, as an algorithm can be made for sale but algorithms have a specific quality which distinguish them from the typical washing machine or television: an inherent decision-making process. For instance, algorithms have the unique ability to not only perform complex actions and take intricate decisions, but they do it at a level which goes beyond computations—an example being the e-commerce industry and the predicted omnipotence of algorithmic agents which will eventually bypass most human decisions.³¹

Algorithms can also make decisions of a moral character, i.e., making choices which would be considered as moral or immoral if made by a human. Germany has the unique distinction of introducing a set of ethical guidelines which must be followed by autonomous vehicles. For example, the “Ethical Rules for Automated and Connected Vehicular Traffic” include that the “protection of individuals takes precedence over all utilitarian considerations.”³² Another instance of this was when a medical algorithm was found to prefer white patients over black patients.³³ The algorithm was aimed at predicting which patients would benefit more from extra care-giving. Even though the algorithm itself was not intended to be racist i.e., the way it categorised data did not factor in a patient’s race, yet it had prioritised patients in terms of how much the person chosen would cost the healthcare system in the future. Costs incurred by black patients were around \$1800 less than white patients with the same chronic conditions.³⁴ It should be noted that costs incurred by an individual is not a race-neutral metric as it depends on, among many other things the person’s capabilities to afford healthcare and the healthcare facilities available. As a result, the algorithm scored both white patients and black patients as having an equal risk of health problems in the future, even though black patients had many more health problems. In instances such as this, one may conclude that the same laws which apply to human moral choices should also apply to algorithms carrying out tasks of a moral character. However, the decision making of the algorithm was again based on the information, which was being provided to it, so there was a degree of human input as well. This is where AI takes a departure from the traditional confines of the product liability regime.

³⁰ ‘Product meaning in The Cambridge English Dictionary’ (*Dictionary.cambridge.org*, 2020) <<https://dictionary.cambridge.org/dictionary/english/product>> accessed 5 April 2020.

³¹ Michal S. Gal and Niva Elkin-Koren, ‘Algorithmic Consumers’ (2017) 30 (2) *Harvard Journal of Law & Technology*, 310-311.

³² Ethics Commission at the German Ministry of Transport and Digital Infrastructure, *Automated and Connected Driving* (Report 2017) <<https://www.bmvi.de/SharedDocs/EN/Documents/G/ethic-commissionreport.pdf?blob=publicationFile>> accessed 11 December 2019.

³³ Carolyn Y. Johnson, ‘Racial Bias in a medical algorithm favors white patients over sicker black patients’ (*The Washington Post*, 24 October 2019) <<https://www.washingtonpost.com/health/2019/10/24/racial-bias-medical-algorithm-favors-white-patients-over-sicker-black-patients/>> accessed 5 April 2020.

³⁴ *Ibid.*

Algorithms also differ from products in the sense that these are capable of learning from datasets, even in manners not perceived by their manufacturers. While this point was amply underlined by the medical algorithm mentioned above, another example from daily life is Instagram. Being a social networking site, Instagram allows users to upload pictures and videos, using algorithms which learn preferences of users, filters out spam, and carries out targeted advertising.³⁵ It also contains an in-built text analytics algorithm called DeepText which not only understands the context of language with human-like accuracy, but also helps in combatting cyberbullying and harassment.³⁶ The ability to adapt and improve an AI system in manners not “predetermined by its designer”³⁷ has implications when it comes to ascribing liability: harm caused by a product may be traced back to the manufacturer, but legal concepts may be challenged if the resultant algorithm does not operate in a way intended by the manufacturer. Foreseeability is one of these legal concepts, as demonstrated, there is a key difference between products and algorithms: the latter involves less human foreseeability in its use.

In order to determine a conclusive metric for differentiating algorithms from products, it is prudent to further categorise machine learning into “supervised”, “unsupervised”, and “reinforcement” learning. These categorisations may be used to determine the level of autonomy which an algorithm has. While the terms “autonomous decision-maker” and “autonomous algorithm” are used to a great extent—and often interchangeably—they differ in meaning.³⁸ On one hand, autonomy can refer to whether an algorithm has the required authorisation to perform a specific task, without human input or permission.³⁹ On the other hand, in a different context, autonomy could signify a characteristic of the algorithm itself i.e. its ability to “teach” itself certain tasks or “understand” its actions and their implications.⁴⁰ In essence, the level of autonomy depends on the type of algorithms i.e. whether its learning is supervised, unsupervised, or reinforced.

While there are several ways to categorise autonomy, this article will now delve into the algorithm’s ability to “self-learn” and carry out tasks not foreseen by its programmer or manufacturer.

Autonomy and the Type of Algorithms

Within the context of this paper, a discussion of autonomy and algorithm types is important as autonomy remains one of the core differentials between AI and a product. The autonomous nature of AI makes it impossible for a manufacturer to envisage all potential actions carried

³⁵ Bernard Marr, ‘The Amazing Ways Instagram Uses Big Data And Artificial Intelligence’ (*Forbes*, 16 March 2020) <<https://www.forbes.com/sites/bernardmarr/2018/03/16/the-amazing-ways-instagram-uses-big-data-and-artificial-intelligence/#359411265ca6>> accessed 5 April 2020.

³⁶ *Ibid.*

³⁷ ‘The Amazing Ways Instagram Uses Big Data And Artificial Intelligence’ (*Forbes*, 2020) accessed 1 April 2020; See also Pei Wang, *Rigid Flexibility: The Logic of Intelligence* (New York: Springer 2006).

³⁸ Thomas B. Sheridan and William L. Verplank, *Human and Computer Control of Underseateleoperators* (Defense Technical Information Service 1978) <<http://www.dtic.mil/dtic/tr/fulltext/u2/a057655.pdf>> 1-3.

³⁹ *Ibid.*

⁴⁰ *Ibid* 1.

out by the AI. The three types of algorithms help us identify which AI products have a higher propensity to be autonomous in the future, and in turn are more distinct than products.

In Supervised Learning, the algorithm is trained with data, such as a “training set,” and is used to derive “good” predictors for a required value.⁴¹ In such algorithms, it is not sufficient to merely provide feedback that the system was erroneous; rather, specific messages which highlight the error are required for proper functioning. The feedback allows the system to hypothesise ways to categorise data which may be unlabelled in the future—data which is also updated based on the feedback the algorithm is provided.⁴² While there is some level of human input involved, which may allow one to ascribe liability easily, it should be noted that the hypotheses regarding the data as well as the improvements made with each feedback turn the algorithm into a version which was not programmed by its manufacturers.

In Unsupervised Systems, the algorithm is not trained with data but carries out the task of deciphering patterns in the information that may lead to the correct answer for a particular example.⁴³ The degree of autonomy enjoyed by unsupervised system is greater than supervised systems. The Chief Scientist of Uber, Zoubin Ghahramani, has described unsupervised learning as “finding patterns in the data above and beyond what would be considered pure unstructured noise.”⁴⁴ However, both these systems involve development to a stage which was not pre-programmed at the time of manufacture.

In Reinforced Learning, the algorithm is not pre-programmed to take specific actions; it has to map out situations and actions through machine learning in order to yield the maximum reward. Essentially, it tries different options until it achieves a certain goal because it is not taught the process to achieve a certain goal.⁴⁵ Reinforcement Learning has been particularly successful in games such as chess, which was shown by the program AlphaGo. The CEO of DeepMind has described this program as neither a human, nor a program, but “almost alien.”⁴⁶ Along with games, recent research has shown the possibilities of reinforcement learning in the field of medicine as well.⁴⁷ This also sets algorithms apart from products, as algorithms may reach a point whereby, they can function without human input.

⁴¹ Andrew Ng, ‘CS229 Lecture Notes: Supervised Learning’ (2018) *Studylib*, <<https://studylib.net/doc/14126957/cs229-lecture-notes-supervised-learning-andrew-ng>> accessed 1 January 2020.

⁴² Amir Gandomi and Murtaza Haider, ‘Beyond the Hype: Big Data Concepts, Methods and Analytics’ (2015) 35 (2) *International Journal of Information Management* 137, 144.

⁴³ Avigdor Gal, ‘It’s A Feature, Not A Bug: On Learning Algorithms and What They Teach Us’ (2017) Organisation for Economic Co-operation and Development <<https://one.oecd.org/document/DAF/COMP/WD%282017%2950/en/pdf>> accessed 29th January 2020; Harry Surden, ‘Machine Learning and Law’ (2014) 89 (1) *Washington Law Review* 87.

⁴⁴ Margaret Boden, *AI: Its Nature and Future* (Oxford University Press 2016) 47; See also Zoubin Ghahramani, ‘Unsupervised Learning’ (2004) Gatsby Computational Neuroscience Unit 3.

⁴⁵ Richard S. Sutton and Andrew G. Barto, ‘Reinforcement Learning: An Introduction’ (1998) 1 (1) Massachusetts Institute of Technology Press 4.

⁴⁶ Will Knight, ‘Alpha Zero’s “Alien” Chess Shows the Power, and the Peculiarity, of AI’ (*MIT Technology Review*, 8 December 2017) <<https://www.technologyreview.com/2017/12/08/147199/alpha-zeros-alien-chess-shows-the-power-and-the-peculiarity-of-ai/>> accessed 13 February 2020.

⁴⁷ Anders Jonsson, ‘Deep Reinforcement Learning In Medicine’ (2018) *Karger Journals* 21.

Liability Regimes

Before ascribing a liability regime, it is pertinent to first delve into the different liability regimes which may be applicable to the law on AI and algorithms. Legal systems are mostly two tiered: with civil law and criminal law. AI in general, and algorithms, can lead to challenges in both these regimes. Civil law, also referred to as private law, essentially governs the legal relationship between private parties, and is used to either create, remove, or alter rights. Civil law liability arising from tort or contract may not have effects which are as harsh as those arising from criminal liability.

Criminal law, on the other hand, is mostly enforced by the state and can be invoked even if the criminals have not agreed to be bound by them. To designate an act as a crime is society's way of denouncing conduct in the harshest way possible. Ergo, the burden of proof required to prove someone guilty is higher in criminal law as compared to private law.

Civil Law Liability Regimes

When it comes to private law, there are basically two sources which may relate to the ways in which algorithms may be governed: obligations through contract and obligations arising out of civil wrongs.⁴⁸ Within civil wrongs, there are a number of categories which may provide a liability regime. These are negligence, strict and product liability, and vicarious liability.

The application of these regimes to AI is problematic for several reasons. Firstly, upon examining key legal questions relating to the tort of negligence, one may arrive at the conclusion that the duty of care will not always fall on the owner of the AI. Rather, it can extend to the designer of the AI or an intermediary party who may have taught, trained, or added to it. This complexity of tracing liability across the supply chain can result in inconsistent application of the law. Secondly, the central concern in negligence cases is whether the defendant was acting in the same way an ordinary and reasonable person would act in a similar situation. A problem arises when this notion is being applied to humans relying on an algorithm or algorithms themselves. One option could be to deduce what the user of the algorithm or the reasonable designer of the AI might have done if faced with the same circumstances.⁴⁹ For instance, to avoid instances such as the death of Elaine Herzberg,⁵⁰ it may be reasonable to design a car in such a way that it enters a fully autonomous mode only when there is a relatively clear motorway, rather than in a crowded street.⁵¹ This solution, however, runs into problems

⁴⁸ Lord Justice Jackson, 'Concurrent Liability: Where Have Things Gone Wrong?' (Lecture to the Technology & Construction Bar Association and the Society of Construction Law in 2014) <<https://www.judiciary.uk/wp-content/uploads/2014/10/tecbarpaper.pdf>> accessed 23 April 2020.

⁴⁹ Ryan Abbot, 'The Reasonable Computer: Disrupting the Paradigm of Tort Liability' (2017) 86(1) *The George Washington Law Review* 101, 138–139.

⁵⁰ 'Self-Driving Uber In Fatal Crash Had Safety Flaws' *BBC News* (6 November 2019) <<https://www.bbc.com/news/business-50312340>> accessed 28 April 2020.

⁵¹ F. Patrick Hubbard, 'Sophisticated Robots: Balancing Liability, Regulation, and Innovation' (2015) 66. *Florida Law Review* 1803, 1861–1862.

in situations where there is no human input in any functions of the AI, which raises the question of whom the liability can be imposed upon.

Similarly, applying strict liability may lead to certain drawbacks for the technology industry. For the victim, the advantages of strict liability are obvious: it does not require them to prove causation between the harm caused and the loss suffered by the victim.⁵² This liability regime only expects the victim to prove that the risk posed by the technology surfaced by causing them harm. It should be noted, however, that strict liability alone would result in an increased risk of liability of those in the technology industry or those who benefit from the technology.⁵³ To counterbalance this effect, restrictions and liability caps may be used. However, such caps are justified with the view that the risk becomes insurable, given that strict liability statutes usually prescribe insurance for liability risks. Naturally, such a regime is deemed to have a negative effect on the advancement of technology, as manufacturers and companies may see strict liability as a deterrent to promote technological research, which in the 21st century is an important economic and social goal for many countries across the globe.

Criminal Law Liability Regime

In addition to civil law liability regimes, instruments within the ambit of criminal law have also been used to play an increasingly relevant role in the context of AI. The notion of exclusively utilising criminal liability for AI entities is challenging for many reasons. For example, in situations wherein an AI entity is successfully incarcerated for one year, how may the implementation of such a sentence manifest itself? This conundrum is extenuated in cases wherein the AI software is not part of something physical (such as a robot or a machine), which essentially makes it impossible for an arrest to take place. Similarly, in more critical cases involving sentences of capital punishment, the lack of a physical body to arrest and incarcerate may make such liability impractical.⁵⁴ These issues are not just restricted to physical sentences, but also extend to monetary punishments, particularly fines. Most sentenced AI entities will lack the abilities to manage their own finances, such as own a bank account, thus making the notion of fining an AI entity unrealistic.⁵⁵ These challenges greatly undermine the inherent foundational aims of imposing criminal liability in the first place: retribution, deterrence, rehabilitation, and incapacitation. Therefore, imposing liability based on a criminal liability system may prove to be counterintuitive in terms of limiting harms that may arise from the failures of AI.

⁵² Ibid.

⁵³ Expert Group on Liability and New Technologies – New Technology Formation, ‘Liability For Artificial Intelligence And Other Emerging Digital Technologies’ (*European Union*, 2019) <<https://ec.europa.eu/transparency/regexpert/index.cfm?do=groupDetail.groupMeetingDoc&docid=36608>> accessed 10 April 2020.

⁵⁴ Aleš Završnik, ‘Criminal justice, Artificial Intelligence Systems, and Human rights’ (2020) SpringerLink <<https://link.springer.com/article/10.1007/s12027-020-00602-0>> accessed 28 May 2020.

⁵⁵ Francesca Lagioia and Giovanni Sartor, ‘AI Systems Under Criminal Law: a Legal Analysis and a Regulatory Perspective’ (2019) Springer Link <<https://link.springer.com/article/10.1007/s13347-019-00362-x#Abs1>> accessed 29 April 2020

Ascribing Liability

In light of these challenges and the discussion above, utilising a product liability regime, which on its own, entails an intricate mix of both contract and tort law, seems most fitting for AI.⁵⁶ Product liability deals with establishing liability in the event that a product causes harm. The party deemed responsible for the harm caused can either be the producer of the product or the intermediate suppliers as well.⁵⁷ The defect in a product is given more importance than the fault of an individual. For Product liability laws to apply to algorithms, harm caused by any AI can be redressed if the affected party brings a claim against the producer or any supplier at any stage of the supply chain.

There are certain advantages of the product liability regime. Firstly, a sense of certainty is attached to this regime in identifying the party to be held responsible; the aggrieved party will not have to seek out different parties in the supply chain and ask for their relative contribution to determine their relative fault. Instead, upon locating the supplier or producer of the algorithm, the party can claim the entire amount from them. The burden of proof will lie on the relevant producer or supplier, who may deflect liability to other parties if necessary. In contrast to a fault-based liability regime, a strict liability regime would not entail the courts determining the level of duty of care accrued in the process of manufacturing and selling AI, as this is a difficult exercise keeping in view the heterogeneous nature of AI. Moreover, strict product liability also encourages developers of algorithms to ensure that the products containing them have control and safety mechanisms intact. An example of this was the announcement made by Volvo that it would assume complete liability for the actions of its autonomous vehicles.⁵⁸ This placed pressure on its competitors to meet the same standards to ensure that self-driving cars become safe to use in everyday circumstances. Additionally, even if an algorithm develops and acts in unforeseeable ways, the producer or designer of the algorithm itself will be looked upon as the person best equipped to control and understand the associated risks.⁵⁹ An example of this are the prompt and sophisticated measures taken by Google in the wake of an accident caused by one of its self-driving cars. Google took cognizance of the causes of the accident, stated the ways in which the scenario was similar to normal interactions and expectations between human drivers, and also took responsibility by improving its software further.⁶⁰

Considering these advantages, the product liability approach makes sense as opposed to strict liability for a multitude of reasons. Firstly, within the broad ambit of both contract and tort law, there are various theories of liability that can be asserted. These include breach of

⁵⁶ John Villasenor, 'Products liability law as a way to address AI harms' (*Brookings*, 31 Oct 2019) <<https://www.brookings.edu/research/products-liability-law-as-a-way-to-address-ai-harms/>> accessed 20 February 2020.

⁵⁷ Horst Eidenmüller, 'The Rise of Robots and the Law of Humans' (2017) 8 Oxford Legal Studies Research Paper No. 27/2017.

⁵⁸ Kirsten Korosec, 'Volvo CEO: We Will Accept All Liability When Our Cars Are In Autonomous Mode' (*Fortune*, 8 Oct 2015) <<https://fortune.com/2015/10/07/volvo-liability-self-driving-cars/>> accessed 11 January 2021.

⁵⁹ *Ibid.*

⁶⁰ Jon Fingas, Google self-driving car crashes into a bus (update: statement), (*Engadget*, 29 February 2016) <<https://www.engadget.com/2016-02-29-google-self-driving-car-accident.html>> accessed 11 January 2021.

warranty, misrepresentation, negligence, design defects, failure to warn, manufacturing defects and more.⁶¹ As mentioned before, the majority of AI is mostly comprised of decision-assistance tools, and it makes sense to turn to negligence law in case the usage of such a tool results in harm. Therefore, to ensure a maximum coverage of a multitude of claims, it is more fitting to impose a product liability system. Secondly, applying product liability laws will resultantly force the courts to fall back on the reasonableness standard, which in turn should ensure a greater access to justice while bringing down trial costs.⁶² The reasonableness standard is ideal as it involves adopting a holistic mechanism of scrutiny when coming to a decision. In instances of product liability, the courts will therein be able to look at factors including but not limited to the actual harm caused, the circumstances surrounding the harm and the decision-making process adopted by both the parties which in turn, should lead to fairer decisions. Thirdly, a relatively lenient reasonableness standard will not come at the cost of computer innovation and a reduction in the usage of machines. This is important for a number of reasons, as innovation is an important aim for many countries into the future. For instance, the UAE in its vision for 2030 highlights innovation as an important aim for its foreseeable future. It has taken many steps, such as setting up special economic zones to promote start-up's, launching accelerators, such as the Ghaddan 21 and offering subsidies, support, and funding to innovative companies. For countries like this, any legislative instrument governing machines cannot hamper innovation, otherwise they will be disincentivised to adopt it. Lastly, in lieu of deterrence, a rule of no-fault liability might not be as effective as the reasonableness standard.⁶³ For instance, within the parameters of a no-fault liability regime, "normal risks" of using technology and machines could be actively excluded from meriting compensation. Therefore, not many organisations will be discouraged from adopting unsafe practices. In a regime falling back on the foundations of the reasonable standard, "normal risks" would not exist. Conversely, every judgment will be premised on the factors listed above (decision making, harm, etc.) making it a better fit for ascribing liabilities to algorithms and their creators.

Consequently, the utilisation of product liability laws will prove to be a viable solution to the question of ascribing liability posed at the onset of this paper. In this light, the compatibility conundrum must also be scrutinised. Thankfully, products liability has been one of the most dynamic fields of law since the middle of the 20th century. This is in part due to the new technologies that have emerged over this period, leading courts to consider a continuing series of initially novel products liability questions; courts have generally proven quite capable of addressing these questions. There are several tactics that can be utilised to smoothen the process of adaption of this liability regime. Primarily, inquiries into AI-based systems and their faults must be informed by the rationale that alleged harms are made by the

⁶¹ Ruben Graaf, 'Concurrent Claims in Contract and Tort: A Comparative Perspective' (2019) Research Gate 713 <https://www.researchgate.net/publication/319701592_Concurrent_Claims_in_Contract_and_Tort_A_Comparative_Perspective> accessed 12th February 2020.

⁶² John W. Ely et al., 'Determining the Standard of Care in Medical Malpractice: The Physician's Perspective' (2002) 37 Wake Forest Law Review 861, 864-865, 869-873.

⁶³ Alan Marco & Casey Salvietti, 'What Does Tort Law Deter? Precaution and Activity Levels in No-Fault Automobile Insurance' (2019) Research Gate <https://www.researchgate.net/publication/228141345_What_Does_Tort_Law_Deter_Precaution_and_Activity_Levels_in_No-Fault_Automobile_Insurance> accessed 01 March 2020.

intelligence software; however, their decisions can be traced to choices made by companies, programmers and users. Liability will thus accordingly need to be placed if harm is caused.

The three classifications of algorithms discussed in the section above are pertinent to this discussion i.e., Supervised Learning, Unsupervised Learning, and Reinforced Learning.⁶⁴ By utilising these three classifications, the level of autonomy of the algorithm can be determined, and the liability of the companies/manufacturers can subsequently be ascertained. For instance, the level of human input required in supervised systems is much greater than that required in unsupervised systems. Reinforcement systems, on the other hand, have no human input whatsoever. These differentiators are pertinent to the determination of liability. For the courts, the case-by-case determinations of liability for the specific algorithm can be made by utilising expert testimony of industry specialists.

Another approach to improve adaptability is the development of risk utility tests⁶⁵ in relation to AI.⁶⁶ These tests have actively been employed in AI liability lawsuits to ascertain whether alleged defects in design could have been avoided “through the use of an alternative solution that would not have impaired the utility of the product or unnecessarily increased its cost.”⁶⁷ However, the mechanism of application will need to take into account not only the human-designed portions of an algorithm, but the post-sale design decisions and substitutes available to the system in light of it automatically updating itself as well. Additionally, it has been discussed that all three types of algorithms on the autonomy scale may lead to a stage of development which was not conceived by its manufacturers, which must also be considered.

It must be recognised that it will take many years to develop a substantial body of case law and statutory law specific to the intersection of AI and product liability; the judiciary will not be consistent in its decisions in each case. However, over time, adoption in lieu of the intricacies of AI will be considered by product liability legislation, particularly in terms of emerging technologies. One way to streamline this process is through the utilisation of law reform agencies and voluntary frameworks. The American Law Institute (ALI), for example, is a respected organisation that produces “scholarly work to clarify, modernise, and otherwise improve the law.”⁶⁸ If the ALI or a similar organisation were to develop and publish model principles of law and/or legislation specific to AI products liability, this could help promote greater certainty, predictability, and uniformity in state-level approaches to AI law.

Should Robots Have Rights?

⁶⁴ See discussion under ‘Autonomy and the type of Algorithms’.

⁶⁵ William Beatty, ‘The Illinois Supreme Court Examines the Risk-Utility Test in Design Defect Cases’ (*Johnson & Bell*, 2011) <<http://johnsonandbell.com/alerts-blog/product-liability/the-illinois-supreme-court-examines-risk-utility-test-in-design-defect-cases-2/>> accessed 2nd March 2020.

⁶⁶ Sunghyo Kim, ‘Crashed Software: assessing Product Liability for Software Defects in Automated Vehicles’ (2019) 16 *Duke Law & Technology Review* 315 <<https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=1322&context=dltr>> accessed 1st March 2020.

⁶⁷ John Villasenor, ‘Products Liability and Driverless Cars: Issues and Guiding Principles for Legislation’ *Brookings Institution* 9 <https://www.brookings.edu/wp-content/uploads/2016/06/Products_Liability_and_Driverless_Cars.pdf> accessed 8 August 2021.

⁶⁸ ‘About ALI’ (*American Law Institute*, 2021) <<https://www.ali.org/about-ali/>> accessed 19 January 2021.

So far, this paper has delved into ascribing liability to AI by developing a liability regime which builds on established legal principles. However, if it is conceded that there are different types of AI with varying degrees of autonomy, then should the varying degree of liability associated with a robot's decision making be accompanied with rights as well? This question, which may seem bizarre at first, has been brought up at many instances, as rights and liabilities are often conceptualised as co-existing concepts. In 2015, Victor Collins was found dead in the hot tub of James Bates. James Bates was charged with murder and his Amazon Echo, a home speaker device which incorporated an AI virtual assistant, was indeed the "key witness" to the alleged crime. While the Arkansas police asked for a divulsion of data from the period relevant to the murder, it was in 2017 that Amazon argued that the human voice commands and the device's responses are capable of protection under the US First Amendment. While this argument was not agreed with, it raised important questions as to whether AI has a right to protection of its speech.⁶⁹ Another example is a robot called "Random Darknet Shopper," that purchased ecstasy and a fake Hungarian passport on the dark web. This robot was part of an art installation in Switzerland. It should be noted that it was the robot, not the artist or another human, that was arrested by the St. Gallen police for the unlawful transactions. While the Swiss authorities took cognizance of the artistic value of the robot, the occurrence opened up a debate on the measures to be taken if a robot does cause harm, and whether such liability should also be accompanied by rights being accrued to robots.⁷⁰

While ascribing liability is a key component of protecting consumers from AI harm, the standalone imposition of liability under an effective regime may raise questions about a state's moral duty towards new technology and AI. All in all, it might lead one to ponder whether robots can and should have rights.⁷¹ These questions stem from the debate in the European Union Parliament in 2017, where concrete recommendations were made to the Commission on Civil Law Rules on Robotics. Section 59(f) laid out the notion of corporate personhood as a model of robot rights:⁷²

Creating a specific legal status for robots in the long run, so that at least the most sophisticated autonomous robots could be established as having the status of electronic persons responsible for making good any damage they may cause, and possibly applying electronic personality to cases where robots make autonomous decisions or

⁶⁹ *State of Arkansas v James A. Bates* CR-2016-370-2; Rich McCormick, 'Amazon Gives up Fight for Alexa's First Amendment Rights After Defendant Hands Over Data' (*The Verge*, 7 March 2017) <<https://www.theverge.com/2017/3/7/14839684/amazon-alexa-firstamendment-case>> accessed 23 May 2020.

⁷⁰ Daniel Rivero, 'That Robot Who Bought Ecstasy And A Fake Passport Online Is Finally Out Of Prison' (*Splinter*, 17 April 2015) <<https://splinternews.com/that-robot-who-bought-ecstasy-and-a-fake-passport-online-1793847213>> accessed 11 January 2021.

⁷¹ David J. Gunkel, 'The Other Question: Can and Should Robots Have Rights?' (2018) Researchgate 1-2 <https://www.researchgate.net/publication/320463916_The_other_question_can_and_should_robots_have_rights> accessed 7 May 2020.

⁷² Nathalie Nevejans, 'European Civil Law Rules In Robotics: A Study For The JURI Committee' (2016) European Union <[https://www.europarl.europa.eu/RegData/etudes/STUD/2016/571379/IPOL_STU\(2016\)571379_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2016/571379/IPOL_STU(2016)571379_EN.pdf)> accessed 3 May 2020.

otherwise interact with third parties independently.⁷³

On the surface, this idea may seem inherently problematic for the establishment of a liability regime, because it gives manufacturers a way to escape responsibility for defects that can be directly attributed to them. However, this notion of various entities being characterised within the ambit of ‘legal personhood’ is not as recent as one might assume: for example, the seminal case of *Santa Clara County v Southern Pacific Railroad Co.*⁷⁴ expanded the ambit of the Fourteenth Amendment to the US Constitution to corporations and established the base for personhood to such entities as well.⁷⁵ Indeed, corporations are some of the most common and oldest examples of non-human entities who have been granted legal personhood. It should be noted, however, that it is an abstraction which “has no mind of its own any more than it has a body of its own.”⁷⁶ While it can be said that corporations are capable of performing actions independent of their directors, owners, and employees, in reality, it is humans who take decisions on the company’s behalf. Otto von Gierke, a legal scholar from the nineteenth century, argued that companies are real “group-persons” and cannot be categorised as mere fictions.⁷⁷ This argument can account for the decision-making processes of companies, which, barring sole proprietorships, may not comprise of opinions of a single person but rather a the collective will of the company, which may be expressed by procedures such as board meetings. Considering the human input involved in companies, it is difficult to make a case for AI personhood based on the same logic.

Considering the discussion above, it is still unclear whether robots should be granted rights. Rights may vary depending upon the liability regime that is established. However, wherein rights are granted, they may be contingent on the realisation of a future where robots may exhibit further functional similarities to humans, meriting a change in legal standards. As of now, violence against machines is not seen as a criminal wrongdoing. Legal systems throughout the globe offer no rights to robots despite them becoming more advanced and being developed with higher levels of AI. In an attempt to remedy this, some have suggested that the right for a robot to not be shut down against its will and the right to not have its source code manipulated against its will should form part of a set of rights for robots in the future.⁷⁸ It is futile to offer such summations of potential rights a robot could be given, especially wherein the technology in question has not yet evolved to its fullest potential.⁷⁹ This waiting period is the first obstacle towards protection, especially when such rights should be universal.

Similarly, another issue with granting rights to robots is formulating them in the first

⁷³ Ibid.

⁷⁴ *Santa Clara County v Southern Pacific Railroad Company* 118 U.S. 394 6 S. Ct. 1132; 1886 U.S. LEXIS 1942.

⁷⁵ Kurt Marko, ‘Robot rights - a Legal Necessity or Ethical Absurdity?’ (*Diginomica*, 03 January 2019) <<https://diginomica.com/robot-rights-a-legal-necessity-or-ethical-absurdity>> accessed 1 May 2020.

⁷⁶ *Lennard’s Carrying Co Ltd v Asiatic Petroleum Co Ltd* [1915] AC 705, 713.

⁷⁷ David Nicholls, *The Pluralist State* (St Martin’s Press in association with St Antony’s College 1994) 56.

⁷⁸ Mark Fishel, ‘Why Robots Should Be Given Rights’ (*Good Audience*, 17 September 2018) <<https://blog.goodaudience.com/5-reasons-why-robots-should-have-rights-4e62e8698571>> accessed 12 May 2020.

⁷⁹ Kenneth Kernaghan, ‘The Rights and Wrongs of Robotics: Ethics and Robots in Public Organizations’ (2014) Wiley Online Library <<https://onlinelibrary.wiley.com/doi/abs/10.1111/capa.12093>> accessed 23rd April 2020

place. While certain machines have the propensity to “think” rationally in the twenty first century, the notion of rationality for a machine will vastly differ from the human conception of rationality. Nevertheless, these two conceptions of rationality are slowly merging due to the advent of deep learning and its popularisation.⁸⁰ Resultantly, machines are input with statistics, situations, and moral principles from which the machine distinguishes between “right” and “wrong”. This interdependence means that machines still have a long way to go before they can be independently rational and therefore require the legal protection of rights.

Lastly, the parallel between animals and machines, especially in the context of rights, poses a relevant and interesting obstacle. One might argue that machines do not deserve rights protection over animals. Indeed, the discourse on animal rights has only recently gained increased momentum.⁸¹ From a utilitarian perspective, however, it is pertinent to provide a certain set of rights within the short term to AI entities and algorithms. It may not be desirable in the long-term to keep AI entities devoid of rights; thus, certain work must be done to provide a specific set of rights to AI entities.

Therefore, it remains reasonable to state that robot rights are neither a moral absurdity nor a legal urgency. It must be noted, however, that no matter how similar the treatment of robots may be to humans nowadays, there are many years from when robots may be capable of actions forcing us to confront issues as to their rights. Verily, as of now, the Section 56 approach to AI rights might make sense, namely via establishing laws of accountability and damage mitigation structures (like insurance) that reflect the differences between autonomous, adaptive, “intelligent” robots, and the algorithms that power them, and traditional machines.⁸² However, we must make sure that this approach is complemented through legal instruments that outline ownership of any intellectual property that such machines might create in their normal functioning that may be explicitly distinct from the underlying algorithms controlling them. In a few years, more heed can be given to future protections as the technology behind AI progresses to the extent where it is seamlessly integrated into every aspect of human life and is thus subjective to extensive liability. Currently, it remains more morally pertinent to focus on the protection of historically non-human exploited groups, such as animals and plants.

Conclusion

The paper uses several differentiating factors to conclude that algorithms indeed differ from products. Some of the most prominent differences is the AI’s ability to make decisions of a

⁸⁰ Stamatis Karnouskos, ‘The Interplay of Law, Robots and Society, in an Artificial Intelligence Era’ (2017) Digitala Vetenskapliga Arkivet 10-11 <<https://www.diva-portal.org/smash/get/diva2:1238452/FULLTEXT01.pdf>> accessed 3 May 2020.

⁸¹ Gary Francione, ‘Animal Rights Theory and Utilitarianism: Relative Normative Guidance’(1997) Lewis & Clark Law School <<https://digitalcommons.calpoly.edu/cgi/viewcontent.cgi?referer=https://www.google.com/&httpsredir=1&article=1054&context=bts>> accessed 2 May 2020.

⁸² Mathias Risse, ‘Human Rights and Artificial Intelligence: An Urgently Needed Agenda’ (2018) Harvard Kennedy School Working Paper No. RWP18-015 <https://carrcenter.hks.harvard.edu/files/cchr/files/humanrightsai_designed.pdf> accessed 2nd May 2020.

moral character, as well as the ability to learn from a data set in a manner which could not be anticipated by its manufacturers. However, the fact that there are certain factors which differentiate products from algorithms does not mean that AI should have a different legal regime altogether. Rather, the existing legal framework of product liability law- which contains a mix of both tort and contract law, would be most feasible in addressing the legal questions posed by AI. The compatibility conundrum between existing product liability laws and AI can hence be solved when the “autonomy classifications” proposed in this paper can be used to determine the extent to which liability can be traced to the manufacturer/company in case a harm occurs. Lastly, this paper argues that a system recognising the rights of robots is not possible soon—as humankind has a long way to go before robots make completely autonomous decisions with no human input.

AI can make decisions without human input and is characterised by a great degree of autonomy. More specifically, AI is different from products because the manufacturer may not have envisioned a potential action carried out by AI. This happens due to machine learning and the potential for AI to morph into something completely different than what it was at its conception. The paper displayed this by highlighting three kinds of algorithms: supervised, unsupervised, and reinforced.

A product liability regime needs to be enforced; however, it should be adapted to the novel nature of AI. Two reasons were highlighted for this: the first being the technological leaps being taken in this field, and the growing influence of AI in our lives. Indeed, we have seen an upsurge of digital solutions during 2020 itself due to the advent of COVID-19, and our interaction with AI has increased manifold. The second reason is that if the product liability framework does not advance and a holistic framework is not developed, there will be haphazard regulation and conflicting legislation. In this regard, the best practices of the EU may be instructive. Naturally, a multilateral framework will be required to address such an all-encompassing technological phenomenon which knows no borders.

This paper grappled with the question of imposing liability in the event that an algorithm causes harm and has attempted to propose a system of ascribing liability through an expansion of the existing product liability framework, rather than introducing a different area of law altogether. Additionally, this paper delved into possibility of granting robots rights akin to human beings, concluding that it may not be a legal necessity facing us today. The concept of AI, being mentioned by scriptures thousands of years ago, may not be as visibly frightening as the creature in Frankenstein, nor as threatening as the Terminator. However, it is capable of racial discrimination, breach of privacy, and fatal accidents. The liability framework hence needs to account for the potential undesirable actions of AI, because at this juncture of history, it is a concept that is continuously advancing and evolving.