

The Optimal Agent: The Future of Autonomous Vehicles & Liability Theory

Brian S. Haney¹

Table of Contents

I. Introduction.	2
II. AV Technology.	6
A. LIDAR.	7
B. DEEP LEARNING.	10
C. DEEP REINFORCEMENT LEARNING.	16
III. AV Liability.	22
A. LIABILITY THEORIES.	23
B. DESIGN DEFECT LIABILITY.	30
IV. Conclusion.	33
Appendix A. Summary of Notation.	35
Appendix B. Laws by Year.	36

I. INTRODUCTION

Car accidents result in 40,000 deaths and 2.4 million injuries each year in the United States.² Further, the World Health Organization estimates car crashes result in 1.3 million deaths globally each year.³ Driver error is the cause of the vast majority of these injuries and deaths.⁴ The development of Autonomous Vehicles (“AVs”) will drastically reduce the number of accidents, injuries, and deaths associated with car accidents over the next

¹ J.D. Notre Dame Law School, B.A. Washington & Jefferson College. Thanks to Tom Sweeney, Ed Gray, and Steve Kramer for the helpful guidance, suggestions, and discussion.

² See Kenneth S. Abraham & Robert L. Rabin, *Automated Vehicles and Manufacturer Responsibility for Accidents: A New Legal Regime for a New Era*, 105 VA. L. REV. 128, 129 (2019).

³ See LAWRENCE D. BURNS & CHRISTOPHER SHULGAN, *AUTONOMY* 8 (1st ed. 2018).

⁴ See Abraham & Rabin, *supra* note 2, at 129 (explaining driver error is the cause of the majority of car accidents).

few years.⁵ In fact, as AVs become more prevalent in everyday transportation, accident rates are estimated to decline by at least ninety percent.⁶

The quest to build AVs is reshaping the world.⁷ The quest began with the Defense Advanced Research Project Agency (“DARPA”) Grand Challenge in 2004.⁸ The Grand Challenge was the first contest that offered a cash prize for an AV that completed a test course.⁹ The first Grand Challenge competition took place on a course located in the Mojave Desert.¹⁰ However, the one-million-dollar prize went unclaimed as none of the contestants finished the course.¹¹ Carnegie Mellon’s Red Team designed an AV, a modified AM General Humvee, which traveled the furthest in the competition at 7.3 miles.¹² The next year, in the second Grand Challenge five AVs successfully completed a 132 mile off road course near the California-Nevada border.¹³ The fastest team to complete the course was the Stanford Racing Team, which claimed the two-million-dollar prize for completing the course in six hours and fifty-four minutes.¹⁴ Then, in 2007 DARPA hosted the Urban Challenge, which took place in Victorville, California.¹⁵ The course was made up of sixty miles of a staged city landscape and was required to be completed in six hours.¹⁶ Carnegie Mellon’s Tartan

⁵ See MAX TEGMARK, LIFE 3.0 BEING HUMAN IN THE AGE OF ARTIFICIAL INTELLIGENCE 99 (Alfred A. Knopf 2017) (describing that AVs will decrease the amount of injuries related to car crashes).

⁶ See *id.*

⁷ See generally BURNS *supra* note 3 (explaining that AVs are changing the area of transportation around the world).

⁸ See BURNS, *supra* note 3, at 11 (stating the DARPA challenge set in motion the race to build AVs).

⁹ See *The DARPA Grand Challenge: Ten Years Later*, DEF. ADVANCED RESEARCH PROJECTS AGENCY (Mar. 13, 2014), <https://www.darpa.mil/news-events/2014-03-13> (describing that DARPA was the first AV challenge to award a cash prize).

¹⁰ See *id.* (stating the first DARPA challenge took place in California).

¹¹ See BURNS, *supra* note 3, at 58 (explaining that none of the contestants finished the first DARPA race).

¹² See BURNS, *supra* note 3, at 57 (stating out of the competitors, Carnegie Mellon made it the closest to the finish line).

¹³ See DARPA, *supra* note 9 (describing that competitors were able to finish the DARPA course in the second year).

¹⁴ See *id.* (stating that Stanford Racing won the second DARPA race).

¹⁵ See *id.* (describing the urban challenge took place in a city instead of a desert).

¹⁶ See BURNS, *supra* note 3, at 87 (explaining the Urban Challenge had a set course and time limit).

Racing Team claimed victory and the two-million-dollar prize, finishing the race in four hours and ten minutes.¹⁷ Ultimately, the DARPA Grand Challenges and Urban Challenge lead to the development of breakthrough AV technologies that serve as a foundation for AV today.¹⁸

In its September 2016 policy statement, the National Highway Traffic Safety Administration (“NHTSA”) adopted the five-tiered levels of AV technology:¹⁹

- At Level 1, an automated system on the vehicle can sometimes assist the human driver conducting some parts of the driving task;
- At Level 2, an automated system on the vehicle can actually conduct some parts of the driving task, while the human continues to monitor the driving environment and performs the rest of the driving task;
- At Level 3, an automated system can both actually conduct some parts of the driving task and monitor the driving environment in some instances, but the human driver must be ready to take back control when the automated system requests;
- At Level 4, an automated system can conduct the driving task and monitor the driving environment, and the human need not take back control, but the automated system can operate only in certain environments and under certain conditions; and
- At Level 5, the automated system can perform all driving tasks, under all conditions that a human driver could perform them.²⁰

Some speculate the transition from current AV technology, Level 3, to a world where all cars are Levels 4 and 5 may take as long as thirty years.²¹ In short, these estimates are wrong. These projections rest on a common, but fallacious assumption that AV technology will continue to evolve at a linear rate.²²

However, AV is driven by Artificial Intelligence (“AI”)

¹⁷ See DARPA, *supra* note 9 (describing that Carnegie Mellon won the Urban Challenge).

¹⁸ See *generally* BURNS, *supra* note 3 (arguing that the AV challenges have impacted transportation technology).

¹⁹ See Abraham & Rabin, *supra* note 2, at 130 (stating that there is a ranking system for the levels of vehicle automation).

²⁰ See SOC’Y OF AUTO. ENGINEERS INT’L, SURFACE VEHICLE INFORMATION REPORT 2 (2014) (describing the taxonomy of AV systems).

²¹ See Abraham & Rabin, *supra* note 2, at 130 (explaining that automated vehicles have years to develop).

²² See RAY KURZWEIL, HOW TO CREATE A MIND: THE SECRET OF HUMAN THOUGHT REVEALED (2013) (arguing AV technology will develop at an exponential rate).

technology, which evolves at an accelerating rate.²³ Indeed, the Law of Accelerating Returns (“LOAR”), states that fundamental measures of information technology follow predictable and exponential trajectories.²⁴ Moore’s Law describes the LOARs application to the price and performance of computing and is commonly generalized to describe the observation that the power of information technology doubles every eighteen months.²⁵ The past fifty-three years have proven Moore’s Law correct.²⁶ As a result, a smartphone today has more computing power than all of NASA had in 1969, when Apollo 11 landed on the Moon.²⁷ Applied to AV, Moore’s law has led many researchers to believe that society is on the cusp of developing breakthrough technology that will allow all vehicles on the road to be Levels 4 and 5.²⁸

Yet, a recent piece of legal scholarship notes “[i]t will be a long time – probably several decades – before AVs are the only type of motor vehicle on the road.”²⁹ The reason for such hyper-conservative estimates in the kinetics of AV evolution is the anthropomorphism of technology.³⁰ Human anthropomorphism of technology refers to the ascription of human levels of intelligence to non-human entities.³¹ Indeed, AV technology is driven by non-human intelligence in the form of machine learning algorithms.³²

²³ See KURZWEIL, *supra* note 22.

²⁴ *Id.*

²⁵ See MARTINE ROTHBLATT, VIRTUALLY HUMAN: THE PROMISE AND THE PERIL OF DIGITAL IMMORTALITY 48 (1st ed. 2014).

²⁶ See KURZWEIL, *supra* note 22, at 251.

²⁷ See MICHIO KAKU, PHYSICS OF THE FUTURE: HOW SCIENCE WILL SHAPE HUMAN DESTINY AND OUR DAILY LIVES BY THE YEAR 2100 23 (2011).

²⁸ See *generally* NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES (2014); see also KURZWEIL, *supra* note 22 at 99; see *generally* Tegmark, *supra* note 5 at 99.

²⁹ Abraham & Rabin, *supra* note 2, at 6.

³⁰ See BOSTROM, *supra* note 28, at 85.

³¹ See Eliezer Yudkowsky, *Artificial Intelligence as a Positive and Negative Factor in Global Risk*, in GLOBAL CATASTROPHIC RISKS 308, 21 (Nick Bostrom & Milan M. Ćirković eds. 2008) <https://intelligence.org/files/AIPosNegFactor.pdf>.

³² See Manon Legrand, Deep Reinforcement Learning for Autonomous Vehicle Control among Human Drivers 24 (2017) https://ai.vub.ac.be/sites/default/files/thesis_legrand.pdf. See also Harry Surden, Mary-Anne Williams, *Technological Opacity, Predictability, and Self-Driving Cars*, 38 CARDOZO L. REV. 121, 147 (2016). See also Emily Berman, A Government of Laws and Not of Machines, 98 B.U. L. REV. 1277, 1278 (2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3098995. (Machine

And, while humans may consider a village idiot and Albert Einstein extreme ends of the intelligence spectrum, the difference between the two on a larger relative scale is actually *de minimis*.³³ Similarly, the difference between Level 3 and Level 5 AV intelligence systems is miniscule on a larger intelligence scale. Thus, the kinetics of AV advancement from current Level 3 systems to Level 5 systems will be faster than most expect.³⁴

Today, there are 212 million drivers and 252 million vehicles in the United States.³⁵ Collectively, these drivers and vehicles travel 3.2 trillion miles each year.³⁶ The emissions from these vehicles create one-fifth of the total greenhouse gas emissions created in the United States.³⁷ Now, society is on the verge of a new transportation era that will radically reshape a cleaner and more efficient world.³⁸ As this evolution takes place, the law must be prepared for the new challenges it will face. This paper takes an informatics-based approach to analyzing AV liability, adding to current legal scholarship in two ways. First, this is the first paper in legal scholarship to explore and explain deep reinforcement learning applications in AV technology. Second, this is the first paper to take an informatics-based approach to analyzing liability issues surrounding the evolution of AVs. This paper is divided into two parts. Part II explores and explains the state of the art in AV technology. Part III analyzes and assesses different frameworks for liability regimes in the AV era.

II. AV TECHNOLOGY

Most scholars and researchers analyze AV technology as consisting of two elements: perception and decision making.³⁹ Perception refers to a cars ability to perceive its environment and

learning is a strand of artificial intelligence that sits at the intersection of computer science, statistics, and mathematics, and it is changing the world.)

³³ Yudkowsky, *supra* note 31, at 21.

³⁴ BOSTROM, *supra* note 28, at 86.

³⁵ See BURNS *supra* note 3, at 2.

³⁶ *Id.*

³⁷ *Id.*

³⁸ See Abraham & Rabin, *supra* note 2, at 1.

³⁹ See The Information, *Advantages and Disadvantages of Legacy Car Manufacturers: Bryan Salesky of Argo AI*, YOUTUBE (September 14, 2017) [hereinafter *Salesky*]

<https://www.youtube.com/watch?v=0AZiOQer4aI>.

understand the meaning of the objects within that environment.⁴⁰ Technologies commonly used for AV perception are LIDAR, Deep Learning, and Convolutional Neural Networks (“CNNs”).⁴¹ Decision making refers to an AV’s ability to make decisions and appropriately interact with its environment.⁴² The most recent breakthrough in machine learning, deep reinforcement learning, is uniquely structured to allow for rapid advancements in AV decision making technology.⁴³ Deep reinforcement learning combines two traditional models of machine learning, deep learning and reinforcement learning, to allow algorithms to learn independently of humans.⁴⁴ This Part will explain these technologies and discuss their potential in the future of AV.

A. LIDAR

The most common tool for AV perception is a Light Detection and Ranging Device (“LIDAR”).⁴⁵ Indeed, LIDAR sensors are often contained in the dome-shaped structures on top of driverless cars.⁴⁶ LIDAR is a type of optical radar sensor.⁴⁷ All LIDAR systems consist of a transmitter and a receiver.⁴⁸ The transmitter includes a laser and a beam expander to set the outgoing beam divergence.⁴⁹ The receiver includes a telescope to collect backscattered signal, and appropriate optics to direct the return signal from the telescope to a detector, which records the signal.⁵⁰

⁴⁰ *Id.*

⁴¹ *See generally* Legrand *supra* note 32, at 2.

⁴² *See Salesky, supra* note 39.

⁴³ TEGMARK, *supra* note 5, at 85.

⁴⁴ *Id.*

⁴⁵ Jeff Hecht, *Lidar for Self-Driving Cars*, 29 OPTICS & PHOTONICS NEWS 26, 28 (Jan. 2018),

https://www.osapublishing.org/DirectPDFAccess/51790D97-AC5F-7CDF-A29E387D2461E9AC_380434/opn-29-1-26.pdf?da=1&id=380434&seq=0&mobile=no.

⁴⁶ *Id.* at 31.

⁴⁷ Matthew J. McGill, *LIDAR Remote Sensing*, NASA TECHNICAL REPS. SERVER (NTRS) (2002),

<https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/20020070892.pdf>.

⁴⁸ *Id.* at 47; *see, e.g.*, Devices and Methods for a Rotation LIDAR

Platform with a Shared Transmit/Receive Path, U.S. Patent No.13/971,606 (filed Aug. 20, 2013) (issued Sept. 16, 2014).

⁴⁹ McGill, *supra* note 47.

⁵⁰ *See Id.*

LIDAR sensors start by transmitting infrared light pulses.⁵¹ Then, the pulses travel to the nearest object and backscatter to the receiver.⁵² The time it takes for the pulse to travel to the object and return to the receiver is multiplied by the speed of light and divided by two:

$$\frac{tc}{2} = d$$

where t is travel time, c is the speed of light, and d is the distance between the LIDAR sensor and the object.⁵³ At its core, the input of a LIDAR system is backscattered laser light and the output is a point cloud that models an environment.⁵⁴

The two factors that enable LIDAR measurements are lasers with discrete pulses, and the constancy of the speed of light.⁵⁵ LIDAR uses discrete pulses to measure distances and the orientation of the lasers allows for the association of a three-dimensional position with each returning pulse.⁵⁶ The accuracy of these measurements is made possible by the constancy of the speed of light c because all light particles travel at 299792458 m/ps.⁵⁷ Thus, the time it takes for a laser light pulse to leave a transmitter and return to a receiver is multiplied by its speed, c and divided by two because the photon travels to the object and back.⁵⁸ A common quip is “LIDAR light only travels half the speed of normal light.”⁵⁹ Ultimately, each measurement of distance is recorded in a detector as a data point.⁶⁰

In the context of AVs, LIDAR sensors record each data point

⁵¹ See, e.g., ‘606 Patent fig.1.

⁵² McGill, *supra* note 47.

⁵³ *Id.*

⁵⁴ *Id.*

⁵⁵ *Id.*

⁵⁶ *Id.*

⁵⁷ J. W. Moffat, A Model of Varying Fine Structure Constant and Varying Speed of Light V2 5- 6 (Oct. 11, 2001) (unpublished research) (on file with Cornell University Library), <https://arxiv.org/abs/astro-ph/0109350>.

⁵⁸ McGill, *supra* note 47.

⁵⁹ *Id.*

⁶⁰ *Id.*; see also Brian S. Haney, *Applied Artificial Intelligence in Modern Warfare & National Security Policy*, 11 HASTINGS SCI. & TECH. L.J. ___ (2019)(Forthcoming)(See manuscript page 22, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3454204).

collected in a point cloud.⁶¹ A point cloud is a mapping of LIDAR data points representing a model of an AVs environment.⁶² AVs are commonly equipped with multiple LIDAR sensors that simultaneously collect data to build 3D models of AV environments.⁶³ This technology allows for AV to navigate through environments without input from a human driver.⁶⁴ And, LIDAR technology is rapidly evolving.⁶⁵ Indeed, as one scholar argues, “LIDAR has become the inevitable technology to provide accurate 3D data fast and reliably even in adverse measurement situations and harsh environments.”⁶⁶ Further, a recent piece of scholarship argues LIDAR technology is a necessary element of developing Level 4 and Level 5 AV systems.⁶⁷ As LIDAR technology evolves Geiger-mode LIDAR (“GmLIDAR”) presents enormous potential for commercial deployment in AVs Levels 4 and 5.⁶⁸

GmLIDAR utilizes an array of avalanche photo diodes (“APDs”).⁶⁹ APDs are highly sensitive semiconductor electronic devices that convert light to electricity using the photoelectric effect.⁷⁰ The photoelectric effect describes the emission of electrons when light shines on an object.⁷¹ In contrast to conventional LIDAR, which utilizes high energy pulses and low repetition rates, GmLIDAR uses low energy pulses at high repetition rates and measures what percentage of pulses register a return.⁷² The use of low energy pulses at high repetition rates has been successfully deployed in military settings to improve the accuracy and

⁶¹ See, e.g., ‘606 Patent.

⁶² *Id.*

⁶³ See Generally Xiangyu Yue, et. al., A LiDar Point Cloud Generator: From a Virtual World to Autonomous Driving (Mar. 31, 2018) (unpublished research) (on file with Cornell University Library), <https://arxiv.org/abs/1804.00103>.

⁶⁴ See generally *id.*

⁶⁵ See A. Ullrich & M. Pfennigbauer, *Linear LIDARR Versus Geiger-mode LIDAR: Impact on Data Properties and Data Quality*, 9832 LASER RADAR TECH. & APPLICATIONS XXI (2016), <http://www.rieglusa.com/pdf/linear-lidar-versus-geiger-mode-lidar.pdf>.

⁶⁶ *Id.*

⁶⁷ See Xiangyu Yue et. al., *supra* note 63, at 1.

⁶⁸ See generally A. Ullrich & M. Pfennigbauer, *supra* note 65.

⁶⁹ *Id.* at 4.

⁷⁰ *Id.*

⁷¹ *Id.*

⁷² Point of Beginning, *What is Geiger Mode LiDAR?*, YOUTUBE (May 26, 2017), <https://www.youtube.com/watch?v=XbDpnOOnfOw&t=113s>.

specificity of LIDAR measurements.⁷³ Now, GmLIDAR is making its way into the commercial sphere and being deployed to develop AV perception technology.⁷⁴ However, to make sense of LIDAR data, AV researchers have turned to machine learning to allow AV to understand and make sense of 3D LIDAR point clouds.⁷⁵

B. DEEP LEARNING

Deep learning is a common type of machine learning that is used for computer vision and object classification.⁷⁶ Deep learning allows machines to learn with architectures inspired by the biological neocortex.⁷⁷ The human brain is composed of processing units called neurons.⁷⁸ Each neuron is connected to other neurons through connections called synapses.⁷⁹ A biological neuron includes multiple dendrites, receivers of various electrical impulses from other neurons, which are gathered in the cell body of a neuron.⁸⁰ If the neuron's cell body has collected enough electrical energy to exceed a threshold amount, the neuron transmits an electrical charge to other neurons in the brain through synapses.⁸¹ This transfer of information in the biological brain provides the foundation for the way in which modern neural networks operate.⁸²

⁷³ *Id.*

⁷⁴ *Id.*

⁷⁵ Bin Yang et. al., PIXOR: Real-time 3D Object Detection from Point Clouds (Mar. 2, 2019) (unpublished research paper, Department of Computer Science, University of Toronto) at 7652, http://openaccess.thecvf.com/content_cvpr_2018/papers/Yang_PIXOR_Real-Time_3D_CVPR_2018_paper.pdf.

⁷⁶ Gary Marcus, Deep Learning: A Critical Appraisal (Jan. 8, 2018) (unpublished research paper) at 3,

<https://arxiv.org/ftp/arxiv/papers/1801/1801.00631.pdf>. *See also* Olga Russakovsky, et al., Best of both worlds: human-machine collaboration for object annotation (2015),

<https://ieeexplore.ieee.org/document/7298824>. (Introducing a model that integrates multiple computer vision models with multiple sources of human input in a Markov Decision Process.)

⁷⁷ *See generally* KURZWEIL, *supra* note 22.

⁷⁸ ETHEM ALAPAYDIN, MACHINE LEARNING: THE NEW AI 86 (2016).

⁷⁹ *Id.*

⁸⁰ SEBASTIAN RASCHKA & VAHID MIRJALILI, PYTHON MACHINE LEARNING: MACHINE LEARNING AND DEEP LEARNING WITH PYTHON, SCIKIT-LEARN, AND TENSORFLOW 18 (2d. ed. 2017).

⁸¹ *Id.*

⁸² ALAPAYDIN, *supra* note 78.

Artificial neurons are logic gates inspired by the biological neuron.⁸³ Both artificial and biological neurons receive input from various sources and map input information to an output value.⁸⁴ In an artificial neural network (“ANN”) the output value is associated with some abstraction.⁸⁵ For example, whether an object is a tumbleweed or a human may be abstracted from point cloud data. An ANN is a group of interconnected artificial neurons that are able to influence each other’s behavior.⁸⁶ The single layer perceptron can be extended to a multi-layer perceptron, which is an ANN composed of two or more layers of connected perceptron units.⁸⁷ It is important to note that ANNs are capable of generalization.⁸⁸ In other words, a multi-layer perceptron model is a “universal approximator”, which is an algorithm that “can approximate any function with desired accuracy given enough neurons.”⁸⁹

A deep neural network (“DNN”) is a neural network with multiple hidden layers.⁹⁰ The below illustration is a simple model of a DNN:⁹¹

⁸³ KURZWEIL, *supra* note 22, at 38.

⁸⁴ See JOHN D. KELLEHER & BRENDEN TIERNEY, DATA SCIENCE 131 (2018).

⁸⁵ TEGMARK, *supra* note 5, at 72.

⁸⁶ *Id.*


⁸⁷ ALAPAYDIN, *supra* note 78, at 100 (explaining a perceptron is an algorithm that automatically learns the weight coefficients of artificial neurons.)

⁸⁸ *Id.* at 99.

⁸⁹ ALAPAYDIN, *supra* note 78, at 99.

⁹⁰ TEGMARK, *supra* note 5, at 76.

⁹¹ See KELLEHER & TIERNEY, *supra* note 84 at 132 (discussing a model based on illustration at the following citation); see also Brian S. Haney, *The Perils & Promises of Artificial General Intelligence*, 45 J. LEGIS. 151, 160 (2018).



DNNs learn through a training process using large data sets.⁹² In a DNN, the neurons are connected by weight coefficients modeling the strength of synapses in the biological brain.⁹³ The training process allows the weight coefficients to adjust, so the output or prediction of the DNN is accurate.⁹⁴ After a DNN is trained, new data is fed through the network to make predictions.⁹⁵

Deep learning is based on “learning of feature extractions in each layer.”⁹⁶ In a DNN each neuron represents a hidden layer and defines a complex feature of the model.⁹⁷ Hidden layers are layers in between the input and output layer of a DNN.⁹⁸ Integrating hidden layers in an ANN allows the network to account for more abstraction.⁹⁹ Indeed, each layer of hidden units acts as a feature extractor by providing analysis of slightly more complicated features.¹⁰⁰ Feature extraction is a method of dimensionality reduction, allowing raw inputs to be converted to output so data

⁹² ALPAYDIN, *supra* note 87, at 88.

⁹³ *Id.* at 88.

⁹⁴ *Id.* at 89.

⁹⁵ See KELLEHER & TIERNEY, *supra* note 84, at 127; see also Harry Surden, *Machine Learning and Law*, 87 Wash. L. Rev. 87, 90 (2014).

⁹⁶ Legrand, *supra* note 32, at 26.

⁹⁷ *Id.*

⁹⁸ KURZWEIL, *supra* note 22, at 132.

⁹⁹ ALPAYDIN, *supra* note 78, at 88.

¹⁰⁰ KELLEHER & TIERNEY, *supra* note 84, at 127.

scientists are able to observe previously latent features in data.¹⁰¹ The later hidden units extract hidden features by combining the previous features in a larger space.¹⁰² The output layer observes the whole network to produce a final prediction.¹⁰³ In other words, DNNs learn more complicated functions of their initial input when a hidden layer combines the values of the preceding layer.¹⁰⁴ Additionally, DNNs have been proven to be excellent for making predictions in several contexts, one of which is computer vision.¹⁰⁵

DNNs are used in conjunction with LIDAR point cloud data to label and classify objects so AVs are able to make sense of their environment.¹⁰⁶ And, point cloud labeling is a critical element of computer vision and object recognition for AVs.¹⁰⁷ Recent progress in deep learning techniques demonstrate point cloud data may be directly combined with DNNs and trained together.¹⁰⁸ Indeed, a convolutional neural network (“CNN”) is a type of DNN used in this context.¹⁰⁹

CNNs are modeled based upon the biological visual cortex.¹¹⁰ The biological visual cortex is composed of receptive fields made up of cells that are sensitive to small sub-regions of the visual field.¹¹¹ In an artificial visual cortex, the response of a neuron to a

¹⁰¹ ALPAYDIN, *supra* note 78, at 102.

¹⁰² *Id.*

¹⁰³ KURZWEIL, *supra* note 22, at 132.

¹⁰⁴ ALPAYDIN, *supra* note 78, at 104.

¹⁰⁵ See generally KEVIN D. ASHLEY, ARTIFICIAL INTELLIGENCE AND LEGAL ANALYTICS 107-26 (2017) (explaining the different types of ways computational models of legal reasoning can be used to predict legal outcomes).

¹⁰⁶ Jing Huang & Suyu You, *Point Cloud Labeling using 3D Convolutional Neural Network*, 23RD INT’L CONF. ON PATTERN RECOGNITION (2016),

http://www.cvlabs.net/projects/autonomous_vision_survey/literature/Huang2016ICPR.pdf.

¹⁰⁷ *Id.*

¹⁰⁸ Daniel Maturana & Sebastian Scherer, *3D Convolutional Neural Networks for Landing Zone Detection from LiDAR*, 3472 IEEE INT’L CONF. ON ROBOTICS & AUTOMATION (2015), http://wcl.cs.rpi.edu/pilots/library/papers/SafeLanding/3D_Convolutional_Neural_Networks_LIDAR.pdf.

¹⁰⁹ Axel Bender & Elías Marel Þorsteinsson, *Object Classification using 3D Convolutional Neural Networks*, at 11 (2016) (unpublished Master’s thesis in Systems, Control and Mechatronics, Chalmers University of Technology) (on file with the Department of Energy and Environment).

¹¹⁰ Legrand, *supra* note 32, at 22.

¹¹¹ *Id.* at 22-23.

stimulus in its receptive field is modeled with a mathematical convolutional operation.¹¹² Convolution is a form of mathematical operation with two matrices: an input matrix and a kernel, also called a filter.¹¹³ A kernel is a small square matrix that is applied each element of the input matrix.¹¹⁴ Applied to AVs, the input of a CNN is often 3D LIDAR data and the output is a mapping of an AV's environment that allows an AV to identify objects in its environment and to ascribe meaning to those objects through relativistic associations.¹¹⁵ The function of the CNN is in essence a classification task, where the CNN classifies objects based upon their similarity.¹¹⁶ The model learns through both supervised and unsupervised learning processes and is designed to capture human intuition in visual classifications.¹¹⁷ CNNs are specifically used in AV vision because of their ability to map the locality of data.¹¹⁸ And, CNNs have shown state-of-the-art performance in computer vision tasks.¹¹⁹

A CNN contains at least one convolution layer, a layer whose parameters are learnable kernels.¹²⁰ Each kernel is convolved across an input matrix and the resulting output is called a feature map.¹²¹ The full output of the layers is obtained by stacking all of the feature maps to create dimensionality.¹²² In contrast to some DNNs, the weight coefficients in a CNN are not all connected.¹²³

¹¹² *Id.* at 23.

¹¹³ *Id.*

¹¹⁴ *Id.*

¹¹⁵ See Prabhu, *Understanding of Convolutional Neural Network (CNN)--Deep Learning*, MEDIUM (Mar. 4, 2018),

<https://medium.com/@RaghavPrabhu/understanding-of-convolutional-neural-network-cnn-deep-learning-99760835f148> (explaining how inputs and outputs function); Victor Vaquero et al., *Deconvolutional Networks for Point-Cloud Vehicle Detection and Tracking in Driving Scenarios*, ARXIV (Aug. 23, 2018), <https://arxiv.org/pdf/1808.07935.pdf> (explaining that to increase robustness, redundancy in AVs is often tackled by using other sensors, such as 3D lidar range-scanners).

¹¹⁶ See Legrand, *supra* note 32, at 23.

¹¹⁷ See *Id.* at 7-9.

¹¹⁸ Bender & Porsteinsson, *supra* note 109, at 11.

¹¹⁹ Zetao Chen, Convolutional Neural Network-based Place Recognition, 1 (2014) arXiv:1411.1509.

¹²⁰ Damien Matti et al., Combining LiDAR Space Clustering and Convolutional Neural Networks for Pedestrian Detection 3 (Oct. 17, 2017) arXiv: 1710.0160.

¹²¹ See Legrand, *supra* note 32, at 24.

¹²² *Id.*

¹²³ ALPAYDIN, *supra* note 78, at 101.

Instead, a window is defined over a smaller input space and the units are connected to a small subset of the inputs.¹²⁴ In other words, the kernel is centered over a subset of the input matrix and then multiplied for the purpose of feature abstraction.¹²⁵

For example, one study explained how 3D CNNs are being used for pedestrian detection.¹²⁶ According to the study, there are two methods for solving the problem of pedestrian detection for AVs.¹²⁷ The first uses LIDAR sensors and focuses on creating a map of agents in motion around the vehicle.¹²⁸ The second, consists of applying computer vision algorithms to captured images.¹²⁹ The study argues “[w]ith the recent development of [DNNs] for image classification, current state-of-the-art performance is achieved by CNNs.”¹³⁰ The algorithm used in the study “is built upon the idea of clustering the 3-D point cloud of the LiDAR.”¹³¹ Clustering is the grouping of input objects based on their analogous properties.¹³² Indeed, the model used in the study used a CNN to perform a clustering analysis for the classification of LIDAR data points.¹³³ The algorithm takes raw data as input and the CNN generates potential candidates for pedestrians in the AVs line of vision.¹³⁴

A second study divided the computer vision problem into two tasks: categorized detection and general obstacle detection.¹³⁵ In

¹²⁴ *Id.* See also Serena Yeung, et al., A computer vision system for deep learning-based detection of patient mobilization activities in the ICU, *npj Digit. Med.* 2, 11 (2019). <https://doi.org/10.1038/s41746-019-0087-z>. (Introducing an algorithm for detection of mobility activity occurrence.)

¹²⁵ See Legrand, *supra* note 32, at 23.

¹²⁶ See generally Damien Matti et al., *supra* note 120, at 1 (proposing a new pedestrian detector for autonomous vehicles).

¹²⁷ *Id.*

¹²⁸ *Id.*

¹²⁹ *Id.* See also Olga Russakovsky, et al., Object-Centric Spatial Pooling for Image Classification (2012), <http://ai.stanford.edu/~olga/papers/eccv12-OCP.pdf>. (The goal for image recognition is two-fold: deciding what objects are in an image, classification and where these objects are in the image, localization.)

¹³⁰ *Id.*

¹³¹ *Id.* at 1-2.

¹³² ALAPAYDIN, *supra* note 78, at 111-13.

¹³³ Matti et al., *supra* note 120, at 1-2.

¹³⁴ *Id.* at 2.

¹³⁵ Noa Garnett et. al., *Real-time Category-Based and General Obstacle Detection for Autonomous Driving*, GENERAL MOTORS R&D 198, 198 (2017), http://openaccess.thecvf.com/content_ICCV_2017_workshops/papers/w3/Garnett_Real-Time_Category-Based_and_ICCV_2017_paper.pdf.

categorized detection, “a bounding box and class is marked for each object belonging to a set of predefined classes.”¹³⁶ “However, objects outside the pre-defined class set are not detected.”¹³⁷ In general obstacle detection, “[t]he task is to identify in each image column the row position of the nearest *roughly vertical* obstacle.”¹³⁸ In the method used in the study, vertical objects were defined as any object higher than a curb.¹³⁹ Ultimately, the study established a method for performing both types of obstacle detection in a single network capable of running in real-time.¹⁴⁰ The method allows machines to develop real time classifications of 3D LIDAR point cloud data and camera images for both object detection and general obstacle identification.¹⁴¹ This method has resulted in a rapid development in the ability of automated systems to understand the world in which they operate.¹⁴² The path to Levels 4 and 5 becomes more clear when these systems are combined with cutting edge machine learning algorithms for sequential decision making.

C. DEEP REINFORCEMENT LEARNING

AV decision making capabilities are rapidly evolving due to the promises of reinforcement learning.¹⁴³ Reinforcement learning is a machine learning technique for solving optimization problems.¹⁴⁴ Formally, reinforcement learning is described through an agent-environment interaction, with a Markov Decision Process (“MDP”).¹⁴⁵ The model below describes the agent-environment

¹³⁶ *Id.*

¹³⁷ *Id.*

¹³⁸ *Id.*

¹³⁹ *Id.*

¹⁴⁰ *Id.* at 199.

¹⁴¹ Garnett et al., *supra* note 135, at 198-99.


¹⁴² *See generally* Gregory Kahn et al., *Uncertainty-Aware Reinforcement Learning for Collision Avoidance* (Feb. 3, 2017) <https://arxiv.org/pdf/1702.01182.pdf> (showing how automated systems can navigate the environment around them).

¹⁴³ *Id.*

¹⁴⁴ RICHARD S. SUTTON & ANDREW G. BARTO, REINFORCEMENT LEARNING: AN INTRODUCTION 341 (MIT Press, 2d ed., 2018). *See also* Leslie Pack Kaelbling, *Learning in Embedded Systems* (1990), <https://apps.dtic.mil/dtic/tr/fulltext/u2/a323936.pdf>.

¹⁴⁵ Alex Kendall et. al., *Learning to Drive in a Day* (Sep. 11, 2018), <https://arxiv.org/abs/1807.00412>.

interaction in an MDP:¹⁴⁶



The model is designed to capture human intuition in sequential decision making. The environment is made up of states for each point in time in which the environment exists.¹⁴⁷ The agent's actions in each state determine the probabilistic evolution of the environment.¹⁴⁸ In the context of AVs, the agent is the car and the environment is the world in which it drives.¹⁴⁹

Initially, the agent is presented with a state of the environment, which includes several possible actions.¹⁵⁰ Then, the agent takes an action in the state and advances to the next state of the

¹⁴⁶ See SUTTON & BARTO, *supra* note 144, at 48 (illustrating the agent-environment interaction in a MDP). See also Leslie Pack Kaelbling, et al., Reinforcement Learning: A Survey, *J. OF ARTIFICIAL INTELLIGENCE RESEARCH* 237, 248-249 (1996). <http://www.cse.msu.edu/~cse841/papers/kaelbling.pdf>. (Discussing the Markov Decision Process.)

¹⁴⁷ Serena Yeung, et al., End-to-end Learning of Action Detection from Frame Glimpses in Videos (2016), https://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Yeung_End-To-End_Learning_of_CVPR_2016_paper.html. (Introducing a fully end-to-end approach for action detection in videos that learns to directly predict the temporal bounds of actions.) See also Fei-Fei Li et al., *Lecture 14 | Deep Reinforcement Learning*, YOUTUBE (Aug. 11, 2017), <https://www.youtube.com/watch?v=lvoHnicueoE>.

¹⁴⁸ MYKEL J. KOCHENDERFER, *DECISION MAKING UNDER UNCERTAINTY: THEORY AND APPLICATION* 77 (MIT Press, 2015).

¹⁴⁹ Kendall et al., *supra* note 145.

¹⁵⁰ SUTTON & BARTO, *supra* note 144, at 47-49.

environment, where a reward is returned.¹⁵¹ The agent chooses which action to take based upon the agent's policy.¹⁵² A policy is the way in which an agent makes decisions, or chooses actions within a state.¹⁵³ For example, a person with a high amount of integrity may have a policy that routinely guides their decision-making to choose to do the right thing when faced with an ethical dilemma. Similarly, a greedy person may have a policy that routinely guides their decision-making to choose the action that returns the most money. The goal for the policy is to allow the agent to advance through the environment so as to maximize a reward for the entire environment.¹⁵⁴ In the context of an AV, a reward may be associated with collision avoidance, driver safety, efficiency, or a combination of different variables.¹⁵⁵ In short, reward is a measure of short-term gain, however value is a measure of long-term reward.

A value-function defines the value of being in a state s and following a policy π until the final state of the environment, the terminal state.¹⁵⁶ The terminal state concludes the episode, which is all of the states in an environment.¹⁵⁷ The expected value of executing a policy π given state s is denoted $V^\pi(s)$.¹⁵⁸ In the context of an MDP, the value function V^π is equal to the expected sum of the discounted future rewards for executing policy π :¹⁵⁹

$$V^\pi(s) = E[R(s_0) + \gamma R(s_1) + \dots | s_0 = s, \pi(s)]$$

The expected future rewards are discounted with a discount

¹⁵¹ KOCHENDERFER, *supra* note 148, at 77; *see also* SUTTON & BARTO, *supra* note 144, at 38-39. *See also* Katerina Fragkiadaki, et al., Grouping-Based Low-Rank Trajectory Completion and 3D Reconstruction (2014), <https://www.cs.cmu.edu/~katel/nrsfm.html>. (Exploring how an agent can be equipped with an internal model of the dynamics of the external world, and how it can use this model to plan novel actions by running multiple internal simulations.)

¹⁵² SUTTON & BARTO, *supra* note 144, at 37.

¹⁵³ KOCHENDERFER, *supra* note 148, at 79.

¹⁵⁴ SUTTON & BARTO, *supra* note 144, at 50.

¹⁵⁵ *See* Ahmad El Sallab, et al., *Deep Reinforcement Learning Framework for Autonomous Driving*, IS&T ELECTRONIC IMAGING, AUTONOMOUS VEHICLES AND MACHINES, (Apr. 8, 2017) <https://arxiv.org/pdf/1704.02532.pdf>.

¹⁵⁶ *See generally* Li, *supra* note 147.

¹⁵⁷ *See generally id.*

¹⁵⁸ KOCHENDERFER, *supra* note 148, at 80.

¹⁵⁹ El Sallab et al, *supra* note 155.

factor γ .¹⁶⁰ The discount factor is typically defined: $0 < \gamma < 1$.¹⁶¹ This allows the value of rewards in the present to be more valuable than rewards in the future.¹⁶² The optimal policy $\pi^*(s)$ is defined as the policy that maximizes the expected value relative to other policies.¹⁶³ The objective of the MDP model is to find the optimal policy.¹⁶⁴

$$\pi^*(s) = \underset{\pi}{arg \max} V^{\pi}(s)$$

The problem of finding the optimal policy for a given MDP is commonly solved with Q-learning.¹⁶⁵

Q-learning solves the problem of discovering the optimal policy for an agent by maximizing a Q-value function: $Q(s, a)$.¹⁶⁶ A Q-value function describes the value of a state-action pair.¹⁶⁷ Indeed, the goal of a Q-learning algorithm is to discover the optimal Q-value function Q^* for any state-action pair.¹⁶⁸ The Bellman equation expresses the relationship between the value of a state and the values of its subsequent state(s).¹⁶⁹ The algorithm continues perpetually until the convergence of the Q-value function.¹⁷⁰ The convergence of the Q-value function represents Q^* and satisfies the Bellman Equation, defined:¹⁷¹

$$Q^*(s, a) = \mathbb{E}_{s' \sim \epsilon} \left[r + \gamma \max_{a'} Q^*(s', a') | s, a \right]$$

An agent's optimal policy π^* corresponds to taking the action in each state defined by Q^* .¹⁷²

However, one issue that arises is that the value of $Q(s, a)$ must

¹⁶⁰ SUTTON & BARTO, *supra* note 144, at 55.

¹⁶¹ *Id.*

¹⁶² KOCHENDERFER, *supra* note 148, at 78.

¹⁶³ *Id.* at 80.

¹⁶⁴ El Sallab et al., *supra* note 155.

¹⁶⁵ See SUTTON & BARTO, *supra* note 144, at 131.

¹⁶⁶ See *id.* at 130.

¹⁶⁷ See Volodymyr Mnih, Koray Kavukcuoglu, *Methods and Apparatus for Reinforcement Learning*, GOOGLE, INC., Appl. No: 14/097,862, 1 (Dec. 5, 2013) <https://patents.google.com/patent/US20150100530A1/en>.

¹⁶⁸ See generally Li, *supra* note 147.

¹⁶⁹ SUTTON & BARTO, *supra* note 144, at 59.

¹⁷⁰ El Sallab et al., *supra* note 155, at 70, 72.

¹⁷¹ '862, Patent at [5].

¹⁷² MAXIM LAPAN, DEEP REINFORCEMENT LEARNING HANDS-ON, 102 (2018).

be computed for every state-action pair, which may be computationally infeasible.¹⁷³ For example, computing the value of every state-action pair, where the raw input is pixels in an Atari game would require tremendous computational power.¹⁷⁴ Or in the context of AVs, computing the value of every state-action pair where the input is LIDAR data, would also be infeasible.¹⁷⁵ One solution is to use a function approximator to estimate the Q-value function:¹⁷⁶

$$Q(s, a; \emptyset) \approx (s, a)$$

Here, \emptyset represents the function's parameters.¹⁷⁷ And, if \emptyset is determined by a DNN, the algorithm is a deep reinforcement learning algorithm called a Deep Q-Network ("DQN").¹⁷⁸

A DQN is a deep learning model that combines a Deep Convolutional Neural Network ("DCNN") with a Q-learning algorithm.¹⁷⁹ The DQN uses a technique called *experience replay* to create and maintain a buffer of the algorithm's past experiences for training a neural network.¹⁸⁰ An experience consists of an observed state-action pair, the immediate reward obtained, and the next state observed.¹⁸¹ An agent's experience at a time step t is denoted e_t and is a tuple (s_t, a_t, r_t, s_{t+1}) consisting of the current state s_t , the chosen action a_t , the reward r_t , and the next state s_{t+1} .¹⁸² The experiences for all the time steps are stored in memory, and used to train the DCNN.¹⁸³ The DCNN's output corresponds to one valid action because the DCNN serves as an approximator for the Q-value function.¹⁸⁴ Thus, after a feedforward pass of that

¹⁷³ Fei-Fei Li et al., *supra* note 147.

¹⁷⁴ LAPAN, *supra* note 172 at 125.

¹⁷⁵ Lei, Xioyun, *Dynamic Path Planning of Unknown Environment Based on Deep Reinforcement Learning*, J. OF ROBOTICS, Vol. 2018, August 6, 2018, at 2.

¹⁷⁶ '862, Patent at [5].

¹⁷⁷ *Id.*

¹⁷⁸ Fei-Fei Li et al., *supra* note 147.

¹⁷⁹ Legrand, *supra* note 32.

¹⁸⁰ Hado van Hasselt et al., *Deep Reinforcement Learning with Q-Learning*, Google DeepMind, 2095.

¹⁸¹ Volodymyr Mnih, *Human-Level Control Through Deep Reinforcement Learning*, Vol. 518 Nature 529, 529 (February 26, 2015) <https://www.nature.com/articles/nature14236>.

¹⁸² See Legrand, *supra* note 32 at 27 (explaining a tuple is a data storage format similar to a list or an array).

¹⁸³ Hasselt et al., *supra* note 180, at 2095.

¹⁸⁴ Legrand, *supra* note 182, at 27.

network, the outputs are the estimated Q-values of the state-action pair.¹⁸⁵ This allows the algorithm to generalize from collected data of past experiences.¹⁸⁶ Indeed, according to MIT Professor, Max Tegmark, “. . . deep reinforcement learning is a completely general technique.”¹⁸⁷

Deep reinforcement learning has been deployed in multiple successful AV applications.¹⁸⁸ While these applications have yet to be developed to scale, deep reinforcement learning presents enormous promise for the future of AV.¹⁸⁹ Indeed, in a recent study, a vehicle was trained with a deep reinforcement learning algorithm to drive around a test track in only twenty minutes.¹⁹⁰ The study explained deep reinforcement learning provides a useful framework for AVs because it creates corrective mechanisms to improve learned autonomous behavior.¹⁹¹ Further, the study noted previous approaches to AVs faced complex engineering challenges due to the many “specific independently engineered components” such as state estimation, mapping, and planning.¹⁹² Instead, the study strived for a more simple approach through a reinforcement learning framework.¹⁹³ Ultimately, the study was able to “show that an off-the-shelf reinforcement learning algorithm with no task-specific adaption is capable of solving the MDP” through the use of deep deterministic policy gradients (“DDPG”).¹⁹⁴ DDPG is a method of deep reinforcement learning that uses two function approximators, and is commonly referred to as an actor-critic model.¹⁹⁵ The actor estimates the optimal Q-value to solve the Bellman Equation.¹⁹⁶ The critic attempts to facilitate that process by minimizing the actor’s error.¹⁹⁷

Further, another study demonstrated the capabilities deep

185 Mnih & Kavukcuoglu, *supra* note 176 at 529.

186 KOCHENDERFER, *supra* note 148, at 124 (MIT Press, 2015).

187 TEGMARK, *supra* note 5 at 85.

188 Kendall et al., *supra* note 145, at 2; *See also* Kahn et al., *supra* note 142.

189 *See* Kendal et al., *supra* note 145, at 2.

190 Kendall et al., *supra* note 145, at 6.

191 *Id.* at 1.

192 *Id.*

193 *Id.*

194 *Id.*

195 *Id.*

196 Kendall et al., *supra* note 145, at 1.

197 *Id.* at 3.

reinforcement learning provides for collision avoidance systems.¹⁹⁸ The study explains “[d]eveloping reinforcement learning algorithms that reason about perception and control in unknown environments, understand uncertainty, and explore safely is crucial to deploying” Level 4 and 5 AV models.¹⁹⁹ The study demonstrated a method that “effectively minimizes dangerous collisions in an obstacle avoidance task” for an autonomous drone.²⁰⁰ The algorithm learned to choose to proceed cautiously in unfamiliar environments, while increasing velocity in settings where it was familiar.²⁰¹ Both studies demonstrate the potential for deep reinforcement learning in future AV technology and in the commercial deployment of Levels 4 and 5 AVs.²⁰² However, legislators and regulators will be tasked with developing legal frameworks that allocate liability in the event that an AV fails.²⁰³

III. AV LIABILITY

Suppose an AV is driving along a narrow mountain pass.²⁰⁴ As the AV comes around a sharp bend, a group of four hikers walk across the road.²⁰⁵ Based upon all of the information available to the AV, the AV calculates there is no time to stop.²⁰⁶ The only way to save the hikers is to veer off of the road over the mountain side, which would kill the AV’s lone passenger.²⁰⁷ So, should the AV be programmed to save the four hikers, or the solo driver?

This thought experiment is useful in understanding the philosophical dilemmas that must be solved by AV software developers, legislators, and courts.²⁰⁸ From a technical perspective,

¹⁹⁸ Kahn, *supra* note 142, at 1.

¹⁹⁹ *Id.*

²⁰⁰ *Id.*

²⁰¹ *Id.*

²⁰² See generally William J. Tronsor, *The Omnipotent Programmer: An Ethical and Legal Analysis of Autonomous Cars*, 15 RUTGERS L.J. 213, 224 (2018) (discussing increasing use of autonomous vehicles).

²⁰³ See generally *id.* at 263 (suggesting that government intervention is required to hold vehicle manufacturers liable for accidents involving their autonomous vehicles).

²⁰⁴ *Id.* at 230.

²⁰⁵ *Id.*

²⁰⁶ *Id.*

²⁰⁷ *Id.*

²⁰⁸ *Id.* at 229-30. See also Veronica Root, *More Meaningful Ethics*, U. CHI. L. REV. Online, 9 (2019),

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3474344. (“What one person finds to be unethical may be considered entirely appropriate by another individual. Similarly, what might seem like improper

the answer to the question would be a matter of shaping an agent's reward function and policy with respect to whether to save a maximum number of lives, or the driver. Rapid technical developments in AV technology create a number of liability issues for insurers, manufacturers, and drivers.²⁰⁹ This Part will focus on two of those issues. The first issue this Part will discuss is the allocation of liability for AV caused accidents. The second issue this Part will discuss is the design standards for AV decision making functions.

A. LIABILITY THEORIES

State governments have begun to acknowledge the potential benefits of AVs, “and many states have responded with legislation aimed at facilitating the testing of [AV] technology.”²¹⁰ The trend started in Nevada in 2011, when the state legislature enacted a legislative scheme authorizing the testing of autonomous vehicles.²¹¹ Washington D.C. followed, passing legislation broadly authorizing the operation of autonomous vehicles on public roads within its jurisdiction.²¹² In 2015, Tennessee also enacted legislation preventing local governments from prohibiting the use of AV technology.²¹³

The following chart illustrates a complete picture of AV law in the United States through 2018:²¹⁴

States with AV Laws

behavior in the United States may be considered normal business practice in another part of the world.”)

²⁰⁹ Tronsor, *supra* note 202, at 253-54.

²¹⁰ Rustin Diehl & Matthew I. Thue, *Autonomous Vehicle Testing Legislation: A Review of Best Practices from States on The Cutting Edge*,

21 J. TECH. L. & POL'Y 197, 207 (2017) (discussing legislation to facilitate autonomous vehicle testing). (Hereinafter “Diehl”).

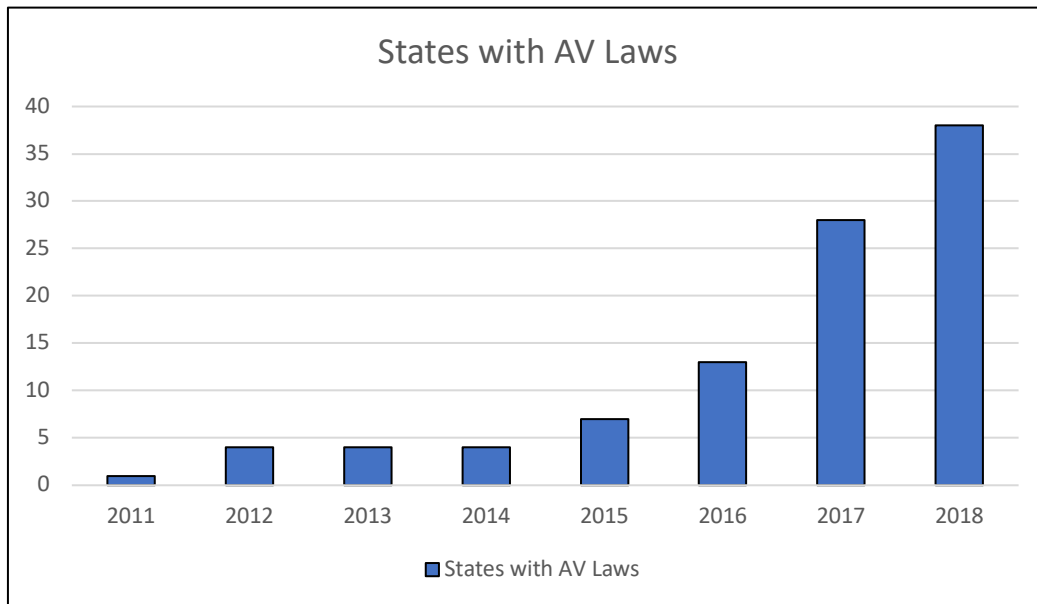
²¹¹ H.R. 511, 76th Gen. Assemb., Reg. Sess. (Nev. 2011).

²¹² H.R. 19-0931, “Autonomous Vehicle Act of 2012,” 19th Council Period (D.C. 2012).

²¹³ H.R. 598, 109th Gen. Assemb., Reg. Sess. (Tenn. 2015).

²¹⁴ Brian S. Haney, et al., *Products Liability Advisory Council, Eckert, Seamans, Cherin, & Mellot, Autonomous Vehicle Legislative Survey*, ECKERT SEAMANS, 4 (2018),

https://www.eckertseamans.com/app/uploads/PLAC_Eckert_Seamans_AV_Survey-1.pdf. (Hereinafter “Haney”).



Indeed, thirty-seven states have enacted laws relating to the regulation of AVs.²¹⁸

Interestingly, all jurisdictions that have successfully enacted legislation authorizing the operation of AVs have also outlined certain conditions that must be met to test an AV on public roads.²¹⁹ For example, California law outlines the requirements for authorized testing, and then expressly prohibits the general operation of AVs on public roads until after a manufacturer has successfully applied for permission to release the technology.²²⁰ For the most part, the existing state legislation focuses on testing only and does not contemplate AV use by the general public.²²¹ Interestingly, the NHTSA's basic liability and insurance guidelines, primarily defer liability decisions to the states.²²² However, as the following chart illustrates, most state statutes are

²¹⁸ *Id.* at 5.

²¹⁹ Diehl, *supra* note 210, at 212.


²²⁰ See CAL. CODE REGS. tit. 13, §§ 227.00, 227.04 (West, Westlaw through Register 2019, No. 37); see also CAL. VEH. CODE § 38750 (West, Westlaw through Ch. 333 of 2019 Reg. Sess.).

²²¹ Diehl, *supra* note 210, at 214.

²²² See Diehl, *supra* note 210, at 219.

relatively silent with respect to liability for AV accidents.²²³

States with AV Laws Allocating Liability



Liability allocation for AV collisions and defects is an important topic in state legislatures across the country.²²⁴

State laws allocating liability are generally narrow in scope.²²⁵ Indeed, most liability clauses have the effect of insulating auto manufactures from liability if a third party converts a conventional vehicle to an AV causing an accident.²²⁶ Additionally, some states require AV testers to purchase liability insurance as

²²³ See, e.g., H.B. 1754 91st Gen. Assemb., Reg. Sess. (Ark. 2017); cf. Haney, *supra* note 214, at 46.

²²⁴ *Id.*

²²⁵ *Id.*

²²⁶ *Id.*

a prerequisite to testing.²²⁷ One argument is the inconsistencies among state legislatures across the country are the result of a general confusion about how to effectively run a government in the digital age.²²⁸ It may be further argued, the variance in language and terminology alone display legislatures' cluelessness towards AV technology from a design perspective. Thus, establishing clear language and clear liability standards for AV technology would be helpful to clarify current laws.²²⁹

Various liability regimes have been suggested for AVs.²³⁰ For example, one piece of scholarship argues once roadway traffic consists of twenty-five percent AVs, a reshaping of the current liability system's approach to deterrence and compensation will be required.²³¹ This piece argues such a reshaping would be centered around the concept of Manufacturer Enterprise Responsibility ("MER").²³² The MER approach entails auto manufacturer responsibility for all injuries arising out of AV operation.²³³ According to the piece, the proposal for a MER system departs significantly from the current tort approach.²³⁴ According to the authors of this piece, a new approach to tort law will need to be developed to deal with the long and uneven transition to automated technology and impose substantial, but appropriate, financial responsibility on AV manufacturers for any accidents.²³⁵ However, the MER approach is flawed for two main reasons. First,

²²⁷ *Id.*

²²⁸ See generally Austin Brown, et al., *Federal, State, and Local Governance of Automated Vehicles*, UC DAVIS INST. OF TRANSP. STUDIES & POL'Y INST. FOR ENERGY, ENV'T, AND THE ECON. (2018). See also Veronica Root, *Coordinating Compliance Incentives*, 102 CORNELL L. REV. 1003, 1029 (2017),

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2867048.

(Discussing regulatory agencies deficiencies in information and coordination.)

²²⁹ Diehl, *supra* note 210, at 212.

²³⁰ Diehl, *supra* note 210, at 218-19. See also Veronica Root, *The Compliance Process*, 94 IND. L.J. 203 (2019),

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3151893

(Introducing analytical frameworks for effective compliance programs.)

²³¹ Abraham & Rabin, *supra* note 2, at 5.

²³² *Id.* at 21.

²³³ *Id.* at 26.

²³⁴ *Id.* at 12. See also Maryam Jamshidi, *How the War on Terror is Transforming Private U.S. Law*, 96 WASH. U. L. REV. 559 (2018),

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3226331.

(Discussing international tort law and liability theory.)

²³⁵ *Id.* at 129.

the MER approach acts as a significant disincentive and impediment to innovation because it deters manufacturers from deploying AV.²³⁶ Second, the MER approach incorrectly assumes there must be significant departure from the current tort system.²³⁷

A second suggested approach to developing a liability theory for AVs centers around the concept of robot rights.²³⁸ Indeed, some scholars argue intelligent agents should be granted rights in a way analogous to human rights.²³⁹ Under this model an AV would be able to develop its own profit and be sued for liability if it breached a standard of care and caused injury.²⁴⁰ Proponents suggest conferring personhood on the vehicle similar to human Uber drivers.²⁴¹ The argument follows, victims of defects may bring claims against the vehicle without deterring manufacturers from producing more AVs if each AV is treated as a separate business entity.²⁴² However, a critical flaw in this approach is that if an individual AV may be sued, individual AVs would also need to retain revenue and profits, disincentivizing manufacturers from producing AVs. Ultimately, the regulatory scheme that should be adopted for AV liability issues should accomplish the task of reducing the number of automobile deaths on the road by promoting innovation, while fairly allocating cost to victims injured by AVs.²⁴³ In developing such a framework, it is important to keep in mind the quest to build level 5 AVs is a global competition.²⁴⁴

The law of agency comes from the Latin maxim “*Qui facit per alium, facit per se*,” which translates to, “he who acts through

²³⁶ *Id.* at 154. See also Michael Gallagher II, *The Coupon Quandary: Restructuring Incentives in CAFA Coupon Settlements*, 91 NOTRE DAME L. REV. 2091, 2091 (2016). (Discussing attorney’s fees as incentive in multi-million-dollar litigation.)

²³⁷ Abraham & Rabin, *supra* note 2, at 128-29.

²³⁸ Rothblatt, *supra* note 25, at 215.

²³⁹ *Id.*

²⁴⁰ Sunghyo Kim, *Crashed Software: Assessing Product Liability for Software Defects in Automated Vehicles*, 16 DUKE L. & TECH. REV. 300, 316 (2018).

²⁴¹ *Id.* at 315.

²⁴² *Id.*

²⁴³ Tronser, *supra* note 202, at 269.

²⁴⁴ See generally Legrand *supra* note 32 (Unpublished Master’s dissertation, Universite Libre de Bruxelles) (on file with author); see also Alex Kendall, et al., *supra* note 145.

another is deemed in law to have acted himself.”²⁴⁵ In the context of AVs, it is fitting the law of agency be applied to liability regimes. Indeed, the algorithms that determine the decisions of an AV are in fact agents acting within an MDP.²⁴⁶ These agents act on behalf of the owner of the vehicle.²⁴⁷ This is true whether the vehicle is personally owned for commuting to work, or if the vehicle is owned by a large company as part of a fleet of vehicles designed to return a profit.²⁴⁸ In both circumstances, the AV may be viewed as an agent of the owner.²⁴⁹

AVs sense their environment, plan what to do, and act based upon defined software architectures.²⁵⁰ The law recognizes computers as instrumentalities of the person using them.²⁵¹ Thus, agency theory is one answer to the question of who should be liable for AV accidents.²⁵² Indeed, holding owners of autonomous vehicles liable for their vehicles would accomplish the goal of reducing the number of automobile deaths on the road, while fairly allocating the cost to victims injured by autonomous vehicles.²⁵³ In the words of Oliver Wendell Holmes, “[t]he safest way to secure care is to throw the risk upon the person who decides what precautions shall be taken.”²⁵⁴

B. DESIGN DEFECT LIABILITY

Another important question that will arise as the quest to create the driverless car continues is how to define a design standard for AVs. Products liability is the main way consumers hold manufacturers liable for defects in a product.²⁵⁵ One of three products liability causes of action for a claimant is a “design defect”

²⁴⁵ Samir Chopra & Laurence F. White, *A Legal Theory for Autonomous Artificial Agents*, 18 (UNIV. OF MICH. PRESS 2011).

²⁴⁶ *Id.* at 137.

²⁴⁷ Tronser, *supra* note 202.

²⁴⁸ *Id.* at 268.

²⁴⁹ *Id.*

²⁵⁰ *See* Tronser, *supra* note 202, at 268.

²⁵¹ *Id.* at 269.

²⁵² *See* Joel Espelien, *The Brave New World of Robot Law*, LAW PRACTICE TODAY (Jan. 14, 2016), <http://www.lawpracticetoday.org/article/the-brave-new-world-of-robot-law/>.

²⁵³ Tronser, *supra* note 202, at 269.

²⁵⁴ OLIVER WENDELL HOLMES, JR., *THE COMMON LAW*, 117 (1881).

²⁵⁵ Tronser, *supra* note 202, at 254.

claim.²⁵⁶ A *design defect* claim may be brought when a product meets the manufacturer's design specifications, but the product's design itself creates an unreasonable risk.²⁵⁷ One scholar argues, the strongest claim consumers will have against manufacturers for AV accidents is a design defect claim.²⁵⁸

A design defect occurs when "foreseeable risks of harm posed by the product could have been reduced or avoided by the adoption of a reasonable alternative design"²⁵⁹ In other words, when a consumer asserts a design defect, they claim the product meets the manufacturer's design specifications, but argue the design itself creates an unreasonable risk of harm.²⁶⁰ The test for judging whether a design is defective is the risk-utility balancing test, which asks "whether a reasonable alternative design would, at reasonable cost, have reduced the foreseeable risk of harm posed by the product and, if so, whether the omission of the alternative design . . . rendered the product not reasonably safe."²⁶¹ To prevail for a design defect claim a plaintiff must prove that a reasonable alternative was, or reasonably could have been available at the time of sale of the product.²⁶² Courts use a reasonable person standard when comparing the product's design to an alternative design.²⁶³

However in the context of AVs, plaintiffs may need a highly specialized expert to develop a safer algorithm the vehicle could have been programmed with, and then they would need to testify as to how the algorithm would have prevented the accident.²⁶⁴ Currently, if a plaintiff brings a claim alleging a cruise control system was defectively designed, they need to prove a reasonable alternative design existed, and had that design been in use, the accident would not have occurred.²⁶⁵ The structure of what a

²⁵⁶ RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2 (AM. LAW INST. 1998).

²⁵⁷ *See id.* at cmt. d.

²⁵⁸ Tronser, *supra* note 202, at 257.

²⁵⁹ RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB. § 2(B) (AM. LAW INST. 1998).

²⁶⁰ *Id.*

²⁶¹ *Id.* at cmt. d.

²⁶² Tronser, *supra* note 202, at 257.

²⁶³ *Id.*

²⁶⁴ *See* Jeffrey Gurney, *Sue My Car Not Me: Prods. Liab. & Accidents Involving Autonomous Vehicles*, 13 U. ILL. J.L. TECH. & POL'Y 247, 263 (2013).

²⁶⁵ *See* F. Patrick Hubbard, "Sophisticated Robots": *Balancing*

plaintiff needs to prove to make out a case for a design defect based on an algorithm may be similar to the one used by plaintiffs attempting to claim their cruise control was defectively designed.²⁶⁶

A second issue plaintiffs may face is when courts evaluate the reasonableness of a design alternative, they look at the overall safety of both designs in general.²⁶⁷ It is not enough to prove the alternative design would have reduced or prevented harm to the plaintiff if the alternative would have produced dangers of equal or greater magnitude.²⁶⁸ As largely uniform software becomes pervasive, the concept of reasonable alternative design is likely to become increasingly indeterminate.²⁶⁹ Indeed, the “confounding effect of technological innovation is perhaps most evident as [AV] software incorporates machine learning[.]”²⁷⁰As one recent piece argues, the “algorithm-based differences . . . would impose overwhelming stress on the premises of conventional analysis.”²⁷¹

One solution to the issue of how to design an AV’s software to maximize efficiency with respect to safety is to include multiple designs within an AV. This solution addresses the design defect and the agency problem. Different designs may be associated with different driving modes. For example, a safe passenger mode would have a reward structure designed to maximize passenger safety. Under this model, an AV’s policy would be designed to maximize a reward correlated with safe driving. This mode would be useful in scenarios where passenger safety is a more important aspect of the journey. Another example would be emergency mode. Under, emergency mode, the AV would be designed to maximize the arrival at a destination in a minimal amount of time. This mode would be useful in scenarios where speed is the most important aspect of the journey – such as an ambulance carrying a patient in critical condition to the hospital.

However, regardless of which mode is selected, the reward structure for an AV would be within a certain set of designed parameters. In other words, a combination of different driving

Liability, Regulation, and Innovation, 66 FLA. L. REV. 1803, 1851 (2015).

²⁶⁶ See Tronser, *supra* note 202, at 258.

²⁶⁷ *Id.*

²⁶⁸ *Id.*

²⁶⁹ Abraham & Rabin, *supra* note 2.

²⁷⁰ *Id.* at 18.

²⁷¹ *Id.* at 19.

metrics, such as efficiency and safety, would be combined. Indeed, one way to visualize this possibility is that a spectrum may be laid out for the way in which a car may travel. On one end, the AV would use a policy to maximize a reward structure based more on speed, and on the other end, the AV would use a policy to maximize a reward structure focused more on safety. The relationship between these two variables is inverse.²⁷² Allowing the driver to decide how the AV will operate within defined parameters maximizes individual control. This alternative allows for the AV to be viewed as an agent of the operator and not the manufacturer.

Under this model, manufacturers may insulate themselves from liability in certain circumstances. For example, while a different design safety mode is available, an occupant may elect to use an emergency mode to get to work faster and the AV may crash as a result.²⁷³ In such a circumstance manufacturers may use this type of setting to insulate themselves from liability by leaving the decision of risk allocation between speed and safety to the driver.

III. CONCLUSION

At the time of his death in 1988, Nobel Prize winning physicist Richard Feynman's blackboard contained the words: "What I cannot create, I do not understand."²⁷⁴ To this point, levels 4 and 5 have evaded human creation and understanding.²⁷⁵ However, the purpose of this essay is to demonstrate the difference between where AV technology is today and a world where all cars are levels 4 and 5 is much narrower than most expect.²⁷⁶ As a result, it is

²⁷² See generally SUTTON & BARTO, *supra* note 144, at 42-43 (demonstrating AI reward structure).

²⁷³ See Tronser, *supra* note 202 at 213 ("The law does not recognize computers as legal entities, but instead as instrumentalities of the person using them.").

²⁷⁴ Michael Way, *What I Cannot Create, I do not Understand*, 2941 J. CELL SCI., 2941 (2017), <http://jcs.biologists.org/content/joces/130/18/2941.full.pdf>.

²⁷⁵ See Abraham & Rabin, *supra* note 2, at 3-4 (explaining that Level 4 and 5 AVs may not reach the market until the mid-2020s); see also TODD LITMAN., AUTONOMOUS VEHICLE IMPLEMENTATION PREDICTIONS: IMPLICATIONS FOR TRANSPORT PLANNING, 4, 14, (Victoria Transp. Pol'y Inst., 2019), <https://www.vtpi.org/avip.pdf> (discussing reasons for slow progress on Level 4 and 5 AV development).

²⁷⁶ See Abraham & Rabin, *supra* note 2, at 3-4 (discussing current expectations of differences between AV levels).

critical states immediately develop liability regimes and design standards for AV technology so that when level 4 and 5 AVs are deployed, courts can turn to statutes rather than common law rules in reaching decisions. This will inevitably improve the efficiency and consistency with which courts analyze legal issues arising from AVs.

In addition to deep reinforcement learning algorithms, fully autonomous vehicles require two things: (1) a natural language processing (NLP) classifier for learning to read street signs; and (2) an object classifier to distinguish objects in the field of vision.²⁷⁷ Combining these three elements will allow for the development of fully autonomous vehicles.²⁷⁸ Another interesting consideration is for states to ban human drivers. By removing human drivers, many of the infrastructure needs associated with driving such as street signs could be removed and traffic – a product of human decision error – may be totally eliminated.²⁷⁹ Without human drivers on the road, fewer accidents will occur, further allowing AV technology to rapidly improve toward a safer world.²⁸⁰

The technology behind AVs' rapid past and future evolution, LIDAR, deep learning, and deep reinforcement learning. Two critical AV issues for lawmakers to consider are: (1) how should liability be allocated for AV caused accidents; and (2) how should design defect standards be set for AV machine learning algorithms? In sum, while the future of AVs remains uncertain and will not happen on its own, it is critical that legislators take action to resolve liability issues surrounding AV, today.²⁸¹

APPENDIX A. SUMMARY OF NOTATION

²⁷⁷ See generally Legrand, *supra* note 32, at 23 (explaining natural language processing); see also MARCUS, *supra* note 76, at 5 (explaining object recognition).

²⁷⁸ See generally MARCUS, *supra* note 76, at 5-7 (suggesting that deep reinforcement learning requires the extra data from NLP classifiers and object classifiers to work effectively in AVs).

²⁷⁹ See generally Abraham & Rabin, *supra* note 2, at 5-9 (discussing likelihood that movement to higher level AVs will decrease accidents and spur changes in road infrastructure).

²⁸⁰ See generally TEGMARK, *supra* note 5, at 99 (discussing likelihood that further development of AV technology will lessen likelihood of car accidents).

²⁸¹ See PETER THIEL & BLAKE MASTERS, ZERO TO ONE: NOTES ON STARTUPS, OR HOW TO BUILD THE FUTURE, 195 (Currency, 2014) (commenting on public responsibility to improve public issues).

Notation	Meaning ²⁸²
$q^*(s, a)$	Value of taking action a in state s under the optimal policy.
γ	Discount-rate parameter
$\mathbb{E}[x]$	Expectation of random variable.
$\operatorname{argmax}_a f(a)$	A value of a at which $f(a)$ takes its maximal value.
R or r	Reward.
s_t	State at time t .
π	Policy.
π^*	Optimal policy.
$v_\pi(s)$	Value of state s under policy π (expected return).

APPENDIX B. LAWS BY YEAR²⁸³

Year	State	Law
2011	Nevada	Assembly Bill No. 511
2012	California District of Columbia	Senate Bill No. 1298 Bill No. 19-0931
2013	Nevada	Senate Bill No. 313
2014	N/A	N/A
2015	Arizona North Dakota Tennessee Utah	Executive Order No. 2015-09 House Bill No. 1065 Senate Bill No. 598 House Bill No. 373

²⁸² See SUTTON & BARTO, *supra* note 144, at xv-xvii (showing notation for variables in analysis).

²⁸³ See Haney, *supra* note 214, at 9-34 (explaining active laws relating to AV technology).

2016	Alabama Florida Florida Louisiana Massachusetts Michigan Michigan Michigan Michigan Pennsylvania Utah	SJR No. 81 House Bill No. 7027 House Bill No. 7061 House Bill No. 1143 Executive Order No. 572 Senate Bill No. 995 Senate Bill No. 996 Senate Bill No. 997 Senate Bill No. 998 Senate Bill No. 1267 House Bill No. 280
2017	Arkansas Arkansas California Colorado Connecticut Delaware Georgia Hawaii Illinois Montana Nevada New York North Carolina North Carolina North Dakota South Carolina Tennessee Texas Texas Vermont Washington Wisconsin	House Bill No. 1754 House Bill No. 1754 Assembly Bill No. 669 Senate Bill No. 213 Senate Bill No. 260 Executive Order No. 14 Senate Bill No. 219 Executive Order No. 17-07 House Bill No. 0791 Joint Resolution No. 40 Assembly Bill No. 69 Senate Bill No. 2005 House Bill No. 469 House Bill No. 716 House Bill No. 1202 House Bill No. 3289 Senate Bill No. 151 House Bill No. 1791 Senate Bill No. 2205 House Bill No. 494 Executive Order No. 17-02 Executive Order No. 245

2018	Alabama Arizona Arizona California District of Columbia Hawaii Idaho Illinois Indiana Kentucky Maine Maine Minnesota Mississippi Nebraska Ohio Ohio Oregon Oregon Pennsylvania Utah Washington Wisconsin	Senate Bill No. 125 Executive Order No. 2018-04 Executive Order No. 2018-09 Assembly Bill No. 87 Bill No. 22-0901 House Bill No. 2253 Executive Order No. 2018-01 Executive Order No. 2018-13 House Bill No. 1290 Senate Bill No. 116 Executive Order No. 2018-001 H.P. 1204 – L.D. 1724 Executive Order No. 18-04 House Bill No. 1343 Legislative Bill No. 989 Executive Order No. 2018-01K Executive Order No. 2018-04K House Bill No. 4059 House Bill No. 4063 House Bill No. 1958 Senate Bill No. 56 House Bill No. 2970 Senate Bill No. 695
------	--	---