



SCHOOL OF LAW
TEXAS A&M UNIVERSITY

**Texas A&M Journal of Property
Law**

Volume 4

Number 5 Symposium Edition - Artificial Intelligence
and the Legal Profession

Article 2

2018

Transparency and Fairness in Machine Learning Applications

Jim Shook

Dell EMC, jimshook@gmail.com

Robyn Smith

Texas A&M University School of Law, robyn.smith1992@email.tamu.edu

Alex Antonio

Texas A&M University School of Law, anto8575@tamu.edu

Follow this and additional works at: <https://scholarship.law.tamu.edu/journal-of-property-law>



Part of the [Intellectual Property Law Commons](#)

Recommended Citation

Jim Shook, Robyn Smith & Alex Antonio, *Transparency and Fairness in Machine Learning Applications*, 4 Tex. A&M J. Prop. L. 443 (2018).

Available at: <https://scholarship.law.tamu.edu/journal-of-property-law/vol4/iss5/2>

This Symposia Article is brought to you for free and open access by Texas A&M Law Scholarship. It has been accepted for inclusion in Texas A&M Journal of Property Law by an authorized editor of Texas A&M Law Scholarship. For more information, please contact aretteen@law.tamu.edu.

TRANSPARENCY AND FAIRNESS IN MACHINE LEARNING APPLICATIONS

By Jim Shook, Robyn Smith, and Alex Antonio

Businesses and consumers increasingly use artificial intelligence (“AI”)—and specifically machine learning (“ML”) applications—in their daily work. ML is often used as a tool to help people perform their jobs more efficiently, but increasingly it is becoming a technology that may eventually replace humans in performing certain functions.¹ An AI recently beat humans in a reading comprehension test,² and there is an ongoing race to replace human drivers with self-driving cars and trucks. Tomorrow there is the potential for much more—as AI is even learning to build its own AI.³

As the use of AI technologies continues to expand, and especially as machines begin to act more autonomously with less human intervention, important questions arise about how we can best integrate this new technology into our society, particularly within our legal and compliance frameworks. The questions raised are different from those that we have already addressed with other technologies because AI is different. Most previous technologies functioned as a tool, operated by a person, and for legal purposes we could usually hold that person responsible for actions that resulted from using that tool. For example, an employee who used a computer to send a discriminatory or defamatory email could not have done so without the computer, but the employee would still be held responsible for creating the email.

While AI can function as merely a tool, it can also be designed to act after making its own decisions, and in the future, will act even more autonomously. As AI becomes more autonomous, it will be more difficult to determine who—or what—is making decisions and taking actions, and determining the basis and responsibility for those actions. These are the challenges that must be overcome to ensure AI’s integration for legal and compliance purposes.

TABLE OF CONTENTS

I. INTRODUCTION.....	444
II. A BRIEF OVERVIEW OF ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING.....	447
A. <i>Traditional Code</i>	447
B. <i>Machine Learning</i>	448
C. <i>Potential Legal and Compliance Problems</i>	449
III. AI USE AND THE LAW.....	450
A. <i>Lack of Transparency</i>	451
1. EU General Data Protection Regulation.....	451

1. Stephen Hawking, *This is the Most Dangerous Time for Our Planet*, Guardian (Dec. 1, 2016, 1:28 PM), <https://www.theguardian.com/commentisfree/2016/dec/01/stephen-hawking-dangerous-time-planet-inequality> [<https://perma.cc/4DW9-K5SB>].

2. Sherisse Pham, *Computers are Getting Better Than Humans at Reading*, CNN (Jan. 16, 2018, 4:16 AM), <http://money.cnn.com/2018/01/15/technology/reading-robot-alibaba-microsoft-stanford/index.html> [<https://perma.cc/T33M-QLA8>].

3. Tom Simonite, *AI Software Learns to Make AI Software*, MIT Tech. Rev. (Jan. 18, 2017), <https://www.technologyreview.com/s/603381/ai-software-learns-to-make-ai-software/> [<https://perma.cc/39TZ-DBKU>]. Cade Metz, *Building A.I. That Can Build A.I.*, N.Y. Times (Nov. 5, 2017), <https://www.nytimes.com/2017/11/05/technology/machine-learning-artificial-intelligence-ai.html> [<https://perma.cc/NR8V-EQEF>].

2. Litigation and Discovery or Disclosure 452
 3. Use of Prohibited Information..... 452
 4. Compliance 453
 B. *Lack of Repeatability and Consistency* 454
 1. Compliance 455
 2. Discovery and Disclosure 455
 C. *Autonomy* 456
 1. Contracts 456
 2. Intent and Consent 457
 IV. A WAY FORWARD 457
 A. *Legal and Regulatory Liability and Responsibility* .. 458
 1. AI as a Tool 458
 2. AI as an Independent Decision Maker 459
 V. AUTONOMOUS AI..... 459
 A. *AI as a Legal Entity* 461
 B. *AI in Discovery or Disclosure* 462
 1. Discovery and Investigations 462
 2. Witnesses and Representatives..... 463
 VI. CONCLUSIONS 463

I. INTRODUCTION

This Article focuses mainly on reinforcement-learning techniques using ML algorithms as a subset of the larger field of artificial intelligence. AI is “the broader concept of machines being able to carry out tasks in a way that we would consider ‘smart[;]’” while ML is “a current application of AI based around the idea that we should really just be able to give machines access to data and let them learn for themselves.”⁴

AI and ML are increasingly used in ways that have a direct impact on people in the real (physical) world:

- AI is used to make decisions on whether to extend credit to individuals with very little (if any) human involvement beyond setup and deployment;⁵
- AI is used to review and monitor communications and activities by employees in some workplaces, and can “flag” an individual’s conduct as anomalous or even bad;⁶

4. Bernard Marr, *What Is The Difference Between Artificial Intelligence And Machine Learning?*, Forbes (Dec. 6, 2016, 2:24 AM), <http://web.archive.org/web/20180206045712/https://www.forbes.com/sites/bernardmarr/2016/12/06/what-is-the-difference-between-artificial-intelligence-and-machine-learning/>.

5. *Machine Learning in Credit Scoring*, TIETO (Feb. 15, 2017), <https://www.tieto.com/news/machine-learning-in-credit-scoring> [<https://perma.cc/6MC2-DRRU>].

6. Oscar Williams-Grut, *Goldman Sachs and UBS monitor traders with a machine learning tool the US uses to find terrorists*, BUS. INSIDER (Sep. 17, 2016, 3:00 AM), <http://www.businessinsider.com/machine-learning-company-digital-reasoning-helps-goldman-sachs-ubs-and-more-spot-insider-trading-2016-9> [<https://perma.cc/NLP5-E5P5>].

- AI is increasingly being used as a “chatbot” to autonomously interact with customers, creating a savings of potentially \$8 billion by 2022;⁷
- Amazon uses ML in inventory optimization and product recommendations;⁸
- AI can review resumes and other job application information to determine whether an individual would be appropriate for a job, even making decisions on which candidates should not be approved for a next step in the hiring process;⁹ and
- a ML tool called COMPAS has been used to make recommendations on criminal sentencing. Today, these are only recommendations which are reviewed by a judge, but there are concerns about racial bias, and the impact and weight of the recommendation is largely unknown.¹⁰

With a little imagination, it is easy to imagine how these technologies could be used even more extensively and autonomously in the future:

- a ML “supervisor” that reviews communications, work product and other factors to determine if employees should be eligible for promotions, demoted for poor productivity or given raises or bonuses;¹¹
- an autonomous “chatbot” that communicates with a customer to answer questions about products or services, makes recommendations, completes sales transactions and otherwise handles activities that would traditionally be performed by a human;¹²
- an “eDiscovery Bot” which could understand a legal case and autonomously collect legally relevant electronically stored information based upon its own experience;¹³ and

7. Jennifer Kite-Powell, *Artificial Intelligence and Machine Learning Set to Change the Customer Experience*, FORBES (Nov. 11, 2017, 11:17 AM), <https://www.forbes.com/sites/jenniferhicks/2017/11/29/artificial-intelligence-and-machine-learning-set-to-change-the-customer-experience/#40389653437c> [<https://perma.cc/3E8X-ALK5>].

8. Erik Brynjolfsson & Andrew McAfee, *The Business of Artificial Intelligence*, HARV. BUS. REV., <https://hbr.org/cover-story/2017/07/the-business-of-artificial-intelligence> [<https://perma.cc/L4RU-AWVU>] (last visited Jan. 3, 2018).

9. Jennifer Alserver, *How AI is Changing Your Job Hunt*, FORTUNE (May 19, 2017), <http://fortune.com/2017/05/19/ai-changing-jobs-hiring-recruiting/> [<https://perma.cc/92BD-EM2Z>].

10. Aaron M. Bornstein, *Are Algorithms Building the New Infrastructure of Racism*, NAUTILUS (Dec. 21, 2017), <http://nautil.us/issue/55/trust/are-algorithms-building-the-new-infrastructure-of-racism> [<https://perma.cc/UEA5-8AWR>].

11. Ted Greenwald, *How AI is Transforming the Workplace*, WALL ST. J. (Mar. 10, 2017, 6:21 PM), <https://www.wsj.com/articles/how-ai-is-transforming-the-workplace-1489371060> [<https://perma.cc/6S6X-3P9Y>].

12. Michael Quoc, *11 Examples of Conversational Commerce and Chatbots*, CHATBOTS MAG. (May 31, 2016), <https://chatbotsmagazine.com/11-examples-of-conversational-commerce-57bb8783d332> [<https://perma.cc/7QBZ-UZMR>].

13. Jim Shook, *The E-Discovery Bots Are Coming: Machine Learning in Preservation and Collection*, LAW.COM (Feb. 7, 2017), <https://www.law.com/legaltechnews/almlID/1202778606612/> [<https://perma.cc/8KBU-8TYZ>].

- a fleet of AI vehicles to replace the trucking industry and the largest automobiles on the road are now driverless.¹⁴

ML used in these scenarios will be challenging for legal and compliance reasons. The decision-making process with ML is usually inscrutable, and even the creators may not understand how decisions are being made.¹⁵ Thus, when situations arise that require a review of the decision process, whether for legal, compliance, or regulatory purposes, there is little or no information for that review, and no person or thing who can answer typical questions.¹⁶

Traditionally, people handled these processes, and those individuals were responsible—legally and under compliance frameworks—for their activities. If problems needed to be resolved or investigated, people could be interviewed, their communications could be reviewed, they could be required to testify under oath or there were documents or witnesses to provide additional information. These information sources could be used to investigate any concerns or issues, and ultimately hold the right people accountable or exonerate them for acting appropriately.

The same processes are not available when ML capabilities are being used, at least not in their current forms. Decisions by ML applications are made generally on a “black box” basis, usually with a very limited or no understanding of how the decisions were made.¹⁷ There may not even be information logs or other data to help review the process of how decisions were made—just the inputs and the final output.¹⁸ As an example, a review of an ML hiring application may uncover information about a candidate, including name and address, school and training information, prior jobs and certifications, etc., and then a decision on whether to hire the candidate.¹⁹ What is missing is any information about the basis for that decision—why was an offer of employment made to one candidate and not to another?

This process is opaque than ever with traditional software, where source code could be reviewed to determine the actual algorithm or

14. Jonathan Vanian, *In 10 Years, Artificial Intelligence Will Transform Trucking, Says Otto Exec*, FORTUNE (Mar. 27, 2017), <http://fortune.com/2017/03/27/uber-otto-artificial-intelligence-truck-driving/> [<https://perma.cc/QV3W-T75H>].

15. Will Knight, *The Dark Secret at the Heart of AI*, MIT TECH. REV. (Apr. 11, 2017), <https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/> [<https://perma.cc/L9CN-8G7R>].

16. *Id.*

17. Colin Lewis & Dagmar Monett, *AI & Machine Learning Black Boxes: The Need for Transparency and Accountability*, KDNUGGETS (Apr. 2017), <https://www.kdnuggets.com/2017/04/ai-machine-learning-black-boxes-transparency-accountability.html> [<https://perma.cc/2XDD-CNWH>].

18. *Id.*

19. Thomas Claburn, *Mysterious Algorithms, Black-Box AI Recruiters are Binning our Résumés*, REG. (Oct. 17, 2016, 9:45 PM), https://www.theregister.co.uk/2016/10/17/hr_sw_for_job_applicants_lacking/ [<https://perma.cc/YNV8-YHZD>].

process of decision making.²⁰ Source code for a ML application mainly implements a learning function, which can be very similar (possibly even the same) for two programs with very different goals—because the algorithm is only used to learn from the data sets that are supplied for training. Thus, the only detailed information about how an ML application actually functions is hidden in the numerical weights and paths that have been assigned to hundreds, thousands or possibly even millions of nodes that make up the ML model—information that is opaque to humans.

II. A BRIEF OVERVIEW OF ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

ML applications are not just a different type of software but are significantly different from how applications have traditionally been written in programming languages.²¹ To fully understand the new challenges ML presents, it must be understood that ML “represents a fundamentally different approach to creating software: The machine learns from examples, rather than being explicitly programmed for a particular outcome.”²²

A. Traditional Code

In a traditional computer program, an application is written in computer code, usually called a language, to perform a certain task. The program includes instructions, or code, that implement rules or algorithms to perform different processes or logical operations, which may change depending on the inputs and will produce different results.

For example, if a credit company developed a program to evaluate an applicant’s ability to repay a loan, it might rate applicants based upon a total “score” earned from work and credit history, salary history, the value of a home, bank and credit balances and other information related to the capability to repay. One rule might be to give an applicant points based upon the salary earned in the previous year. In pseudo-code,²³ instructions to implement this rule might be:

```
If Applicant.Salary > 200,000 then score+=10
Elsif Applicant.Salary > 100,000 then score+=7
Elsif Applicant.Salary > 50,000 then score+=4
```

20. Brynjolfsson & McAfee, *supra* note 8.

21. *Id.*

22. *Id.*

23. “Pseudocode is an informal high-level description of the operating principle of a computer program or other algorithm. It uses the structural conventions of a normal programming language, but is intended for human reading rather than machine reading.” *Pseudocode*, WIKIPEDIA, <https://en.wikipedia.org/wiki/Pseudocode> [<https://perma.cc/G32R-BUEP>] (last visited Mar. 19, 2018).

In natural language, this implements a rule that an applicant with a salary greater than \$200,000 would receive ten additional points; a salary greater than \$100,000 (and up to \$200,000) would receive seven additional points; and one with a salary greater than \$50,000 (and up to \$100,000) would receive four additional points. Applicants with salaries less than \$50,000 would not receive any additional points.

The importance of this example is that the rules are explicit in the source code—and must be so that the computer can execute them. Programmers can readily review the source code to understand how the program works, and to make changes or updates.

B. *Machine Learning*

The workings of a ML application, built upon a neural network, are more complex:

You can't just look inside a deep neural network to see how it works. A network's reasoning is embedded in the behavior of thousands of simulated neurons, arranged into dozens or even hundreds of intricately interconnected layers. The neurons in the first layer each receive an input, like the intensity of a pixel in an image, and then perform a calculation before outputting a new signal. These outputs are fed, in a complex web, to the neurons in the next layer, and so on, until an overall output is produced. Plus, there is a process known as back-propagation that tweaks the calculations of individual neurons in a way that lets the network learn to produce a desired output.²⁴

ML takes a very different approach than traditional computer code. The code for a ML application is mainly the code for its learning algorithm.²⁵ This algorithm is used to “teach” the application using data sets that are provided to it.²⁶ In fact, ML code may not vary much between applications that are completely unrelated; the code for a ML application to recommend criminal sentences could be very similar to code which evaluates credit applications.²⁷ The main difference between applications applying the same ML algorithm is the data used to learn, and after the learning process has been completed, the underlying structure of neurons and weights.

Returning to the credit example, a ML application would start with a learning algorithm and data sets consisting of historical loan and repayment information. This data would include the same inputs that we described before in the source code example (e.g., salary, history,

24. Knight, *supra* note 15.

25. See David Lhr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 688–90 (2017).

26. *Id.*

27. Catherine Dong, *The Evolution of Machine Learning*, TECHCRUNCH (Aug. 8, 2017), <https://techcrunch.com/2017/08/08/the-evolution-of-machine-learning/> [https://perma.cc/648S-CCUZ].

etc.), additional information related to the applicant (e.g., home address, zip code), and the outcomes from past decisions made for those applicants such as repayment history, defaults, etc. All that information is fed to the ML application, which then applies its algorithm to learn—determining which combination of inputs are important to a final score and weighing them (and adjusting those weights) as it continues through the learning process. After training is complete, the state of the neural network is evaluated with a separate data set, and once the results are within an acceptable test range, it can be used.

C. *Potential Legal and Compliance Problems*

People cannot normally understand what is happening in the ML decision-making process—there is no algorithm to follow showing how input information is being used. For example, even with dozens of inputs, depending on the training data, it is possible that a neural network could “learn”—on its own—to deny a pool of credit applicants based solely upon gender or race. Worse, if this did occur it would be difficult to identify. In fact, “subtle biases in training data can utterly change a neural network’s personality—an issue that will become increasingly significant as AI increasingly comes to rely upon big data.”²⁸

A real-world example shows this problem in action. A ML application trained to assist on diagnosing patients with pneumonia “learned” from its data that asthma sufferers were at low risk for death from contracting pneumonia, and could safely be sent home for further care. Later, people evaluating the outcomes from the application determined that data fed to the algorithm had not accounted for the fact that the asthma sufferers had better outcomes because they had been immediately sent to intensive care.²⁹

Taking a closer look at the differences between traditional code and ML further highlights the difficulty with identifying the problem. Assume an application is being developed to evaluate candidates for a job in a large organization. Traditionally, this position has been filled mainly by males—thus the data set used for training consists mostly of males. Across a large data set, candidate qualifications should be distributed across a bell curve—a few will have performed exceptionally, a few poorly, and the bulk within a standard deviation or two of average. If there are only a handful of female candidates represented in the data set, and they are rated exceptionally well, the system could learn to tie the input of female gender to better performance and

28. Ilan Moscovitz, *Artificial Intelligence’s “Holy Grail” Victory*, MOTLEY FOOL (Dec. 31, 2017, 1:22 PM), <https://www.fool.com/investing/2017/12/31/artificial-intelligences-holy-grail-victory.aspx> [https://perma.cc/L7QA-YYKX].

29. Anastassia Fedyk, *How to Tell if Machine Learning Can Solve Your Business Problem*, HARV. BUS. REV. (Nov. 25, 2016), <https://hbr.org/2016/11/how-to-tell-if-machine-learning-can-solve-your-business-problem> [https://perma.cc/2SNF-SHUB].

males to average or poor performance. Due to the potentially biased training data, the application could learn to discount or discard all other data and simply select candidates based upon their gender.

If the organization used a standard software package to make a similar decision, the encoding of an algorithm to evaluate candidates based upon gender would be evident in the source code. Depending on the language, the source code could look like this

```
If $Gender == "Male" then HiringScore:= 0
```

While this code would not be evident to any outsiders (because source code is not generally available for review), it would be easy to locate for anyone who knew to look.

ML, however, has no parallel which would enable anyone to identify this issue. The only difference between a "fair" and "unfair" neural network would be in the weights assigned to the different neurons, and perhaps in the overall structure of the neural network. The complexity of these networks, particularly deep learning networks, means that there is simply no straightforward way to review the weights and determine how the data inputs are being used to make a final decision.

Overall, the implications for this opaqueness of ML can be important for legal and compliance purposes because:

- There is often no way to explain how a decision was reached;
- there is little or no code to inspect which can assist in that understanding;
- the ML learns on its own, meaning that it selects and weighs the factors that are used in making its decisions. This could be problematic if it uses factors such as race, gender or socioeconomic status, which might be illegal, unethical or against company policy; and
- with decisions made by a machine instead of a person, there is no way to instill a sense of ethics, code of conduct or to incorporate other compliance tools.

These types of problems will be very difficult to uncover and could work improperly for years before being discovered.

III. AI USE AND THE LAW

The opaque and sometimes mysterious way ML functions, and its ability for more autonomous use, creates categories of legal issues and concerns requiring a proper approach to resolve in a fair and efficient manner.

It is difficult to predict how AI may challenge existing legal and regulatory frameworks, in part due to the lack of regulations acknowledging AI and the lack of case law applying existing regulations to this

technology.³⁰ With the exception of a few regulations on drone use, the United States has yet to determine which agency or institution is best equipped to govern the use of AI, as it now has a presence in most industries.³¹ Understanding if and how AI must comply with current regulations touches a variety of legal fields from basic tort liability to employment discrimination and basic constitutional protections. Yet at a high level, the difficulties will be caused by certain characteristics of AI: the lack of transparency; the lack of repeatability and consistency; and autonomy.

A. *Lack of Transparency*

The methods by which ML applications are trained and deliver results makes it difficult to obtain an explanation about how a specific result was computed. In the credit score example above, an applicant may receive a score of X, but there is no information regarding how individual factors applied to yield that score. Another candidate with slightly different input characteristics—say a different zip code or salary information—might produce the same or a very different result. Regardless, there is no way to query the application to determine the reasons for the score, or even for the difference between two scores.

This opacity creates problems with transparency. Without transparency – in this context, the ability to explain how a result was achieved—it is difficult to trust the underlying system and to know whether appropriate, legally and ethically proper, and accurate criteria are being used in the decision-making process. Opacity creates problems with several existing laws and obligations.

1. EU General Data Protection Regulation

In some situations, the ability to explain a decision may be explicitly required. Many commentators believe that the upcoming General Data Protection Regulation (“GDPR”),³² requires organizations to provide detailed information about automated decision-making—the so-called “right to explanation.”³³ This requirement includes “mean-

30. Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 378 (2016).

31. *Id.* at 356.

32. See generally Regulation 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC 2018 O.J. (L 119).

33. Article 15 EU GDPR “Right of Access by the Data Subject”, PRIVACY PLAN, <http://www.privacy-regulation.eu/en/15.htm> (last visited Feb. 24, 2018) (stating the data subject shall have the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed, and, where that is the case, access to the personal data and the following information: . . . (h) the existence of automated decision-making, including profiling, referred to in Article 22(1) and (4) and, at least in those cases, meaningful information about the logic

ingful information about the logic involved” in the decision.³⁴ Because the regulation is not yet enforceable there are no decisions to help guide the level of detail required. But if the requirement is more than specifying how an application was trained, it will be difficult for organizations to meet these requirements with most ML applications.

2. Litigation and Discovery or Disclosure

Even when transparency is not legally mandated as a part of using an ML application, if a dispute arises, the lack of transparency will create difficulties in the litigation process.³⁵

In the U.S. litigation process, parties are entitled to information about the matters involved in a controversy:

(1) Scope in General. Unless otherwise limited by court order, the scope of discovery is as follows: Parties may obtain discovery regarding any nonprivileged matter that is relevant to any party’s claim or defense and proportional to the needs of the case, considering the importance of the issues at stake in the action, the amount in controversy, the parties’ relative access to relevant information, the parties’ resources, the importance of the discovery in resolving the issues, and whether the burden or expense of the proposed discovery outweighs its likely benefit. Information within this scope of discovery need not be admissible in evidence to be discoverable.³⁶

A party must preserve evidence, including electronically stored information (“ESI”), which is relevant to the issues of the case. Depending on other factors, such as relevance, proportionality and privilege, they may also have to produce this ESI.

Discovery is problematic when it involves ML. As previously noted, unlike with a witness who can testify or be deposed, there may not be any logs, code or other information that can help to explain a decision made by an ML application. There will be few categories of ESI to preserve, and a party may have to preserve and produce the ML application itself.³⁷

3. Use of Prohibited Information

Many federal and state laws, particularly in the employment area, prohibit discrimination based upon specific characteristics or prohibited classes. If ML is used by organizations within their hiring, recruit-

involved, as well as the significance and the envisaged consequences of such processing for the data subject).

34. *Id.*

35. Although the U.S. has the broadest discovery requirements, this issue may create problems in other countries such as disclosure in the United Kingdom.

36. FED. R. CIV. P. 26(b)(1).

37. The algorithm and machine learning application could be confidential and/or proprietary. But this should not stop production when it is critical to the issues in a case, and the concerns can be covered by a confidentiality order under FED. R. CIV. P. 26(c)(1)(g).

ing, retention and other employee-related workflows, it may be impossible for those organizations to know whether improper information is being considered in violation of certain laws. For example:

- Title VII of The Civil Rights Act of 1964 prohibits discrimination based upon race, color, religion, sex or national origin.³⁸
- The Equal Employment Opportunity Act of 1972 protects against employment discrimination based on an individual's race, color, religion, sex, or national origin.³⁹
- The Age Discrimination in Employment Act of 1967 extended protected class status to individuals barring discrimination based on age.⁴⁰

If an ML application has access to prohibited data during training, it may be impossible to prove that such information was not used and did not play any role in a decision. But even excluding such data from a training set may not be a sufficient safeguard: data about other activities or traits can serve as a proxy for prohibited information. For example, shopping patterns can be used to identify age, gender and socio-economic status.⁴¹

However, the Wisconsin Supreme Court held that due process is not violated by the opaqueness of the algorithms used by ML and could be used in criminal sentencing cases to determine if an offender would likely offend again,⁴² even if the algorithm may have taken race into account.⁴³ The supreme court refused a writ of certiorari to appeal the Wisconsin decision and additional concerns related to due process remain unanswered.⁴⁴

4. Compliance

While not strictly a legal requirement, the concept that a decision is fair may be necessary for corporate values and compliance, marketing, shareholders, and other concerns, and can be difficult to establish when there is no transparency. For example, the concept of fairness in

38. See Civil Rights Act of 1964, 42 U.S.C. § 2000e-2 (1964).

39. Equal Employment Opportunity Act of 1972, Pub. L. No. 92-261, 86 Stat. 103.

40. Age Discrimination in Employment Act of 1967, Pub. L. No. 90-202, 81 Stat. 602.

41. Finale Doshi-Velez et al, *Accountability of AI Under the Law: The Role of Explanation*, BERKMAN KLEIN CTR. FOR INTERNET & SOC'Y (Nov. 3, 2017), available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3064761.

42. The limitations on this decision, such as a determination that the algorithm was being used as one of many inputs and the decision ultimately left to the judge, will be explored more fully in section D(a)(i), *infra*.

43. Garry G. Mathiason et al., *Litter on Legal Compliance Solutions for the Transformation of the Workplace Through Robotic Artificial Intelligence, and Automation*, 2017 SZ002 A.L.I.-A.B.A. 2823 § 1.2.

44. Michelle Liu, *Supreme Court refuses to hear Wisconsin predictive crime assessment case*, MILWAUKEE J. SENTINEL (June 26, 2017, 10:47 AM), <https://www.jsonline.com/story/news/crime/2017/06/26/supreme-court-refuses-hear-wisconsin-predictive-crime-assessment-case/428240001/> [<https://perma.cc/Q3KY-9RXA>].

a hiring context would likely imply that a decision regarding a candidate was: (a) well-founded, for example that applicants selected for jobs truly are qualified, and those that are rejected lacked necessary skills, experience, or other characteristics; and (b) not based upon factors or biases that while not prohibited, probably should not be considered (such as sexual orientation).⁴⁵ Similarly, a shopping recommendation bot might violate corporate values by making recommendations based upon ethnicity or economic status, or just information considered to be private, even if those recommendations made algorithmic sense.⁴⁶

Interestingly, these non-legal requirements can become legal issues if ML undermines specific representations made by an organization. Many organizations make broad statements on their websites and elsewhere about how they conduct business. If their actions run contrary to their representations, the FTC could view this practice as unfair under its broad Section 5 powers.⁴⁷ For example, in 2011 the FTC investigated Google for allegedly preferring its own products in website search results. Although some of the results could have been preferred by programmers, at least some of the results were driven by ML algorithms. These algorithms could have been trained with data skewing in favor of Google, which would have been more difficult to ascertain than traditional code showing Google favoritism. The FTC ultimately took no action against Google.

B. *Lack of Repeatability and Consistency*

Another inherent and useful capability of ML is that it can continue to learn and improve, using data from the feedback loop of its decisions. For example, if a candidate selected by the hiring application turns out to be a poor choice, that data can be used to further train the application. If the same candidate reapplies six months later, the ap-

45. In one recent study, a machine learning application could predict sexual orientation based solely on photographs. At its best, the ML was 91% accurate, compared to 61% accuracy for humans. Bernard Marr, *The AI That Predicts Your Sexual Orientation Simply By Looking At Your Face*, Forbes (Sep. 28, 2017, 12:28 AM), <https://www.forbes.com/sites/bernardmarr/2017/09/28/the-ai-that-predicts-your-sexual-orientation-simply-by-looking-at-your-face/#67a4f9c43456> [https://perma.cc/XP3T-LV2G].

46. In a well-publicized matter, Target determined that a high school student was pregnant based upon her shopping habits, and sent her coupons for pre-natal vitamins. The fact that Target knew about and used the pregnancy for marketing, before she had told her family, caused a public relations problem. See e.g., Charles Duhigg, *How Companies Learn Your Secrets*, N.Y. TIMES MAG. (Feb. 16, 2012), <http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html> [https://perma.cc/U4WA-XY5R].

47. “The basic consumer protection statute enforced by the [FTC] is Section 5(a) of the FTC Act, which provides that ‘unfair or deceptive acts or practices in or affecting commerce . . . are . . . declared unlawful.’” *A Brief Overview of the Federal Trade Commission’s Investigative and Law Enforcement Authority*, FED. TRADE COMM’N, <https://www.ftc.gov/about-ftc/what-we-do/enforcement-authority> [https://perma.cc/B8JN-VAFK] (last updated July 2008).

plication may rate him substantially different than the first time because it has learned that certain qualifications may not be as important to success as originally calculated.

However, this constant improvement process creates problems with repeatability and consistency. While every organization has an interest in improving its applications, there can be problems when a decision made today will be reversed tomorrow after an update to the algorithm.

1. Compliance

Mainly, the lack of consistency may present problems with compliance and corporate oversight and culture. For example, an organization might be working to instill a corporate culture that places the customer “first.” In deploying a chatbot to improve the speed and efficiency of customer service, this core value would be difficult to implement. Further, even with a successful initial implementation, the ML’s ability to learn to become more efficient may help it to realize that certain customer characteristics, or specific issues, can seldom be resolved. In its quest for efficiency, it could learn to spend less time on these issues, leaving the customer unsatisfied.

Many issues in this area may not be manifested until ML is being used in a more autonomous manner. For example, an organization which values its long-term relationships with suppliers—or perhaps prefers a specific partner characteristic such as MBE/WBE businesses—could lose these preferences as an ML learns that other characteristics such as price or shipping times are more valuable to achieving its established goals.

2. Discovery and Disclosure

The challenges related to repeatability and consistency can also create legal issues. As detailed in Section D(A)(2), above, parties in litigation may have a duty to preserve evidence, including ESI, that is relevant to the issues in a case. This duty of preservation generally triggers upon reasonable anticipation of litigation (subject to concerns about proportionality, relevance, privilege, etc.).⁴⁸

In the context of a continuously learning application, legal hold preservation could prove to be difficult and fraught with risk. Without pre-established procedures, an update to an ML could be irreversible, and there would be no way to re-run the application with the same

48. The obligation to preserve electronically stored information requires reasonable and good faith efforts to retain information that is expected to be relevant to claims or defenses in reasonably anticipated or pending litigation. However, it is unreasonable to expect parties to take every conceivable step or disproportionate steps to preserve each instance of relevant electronically stored information. *The Sedona Principles, Third Edition: Best Practices, Recommendations & Principles for Addressing Electronic Document Production*, 19 SEDONA CONF. J. 1, Principle 5, 38–40 (2018).

parameters to achieve the same result as that which triggered the litigation. In fact, in many cases it would be likely that the application could have already updated itself before litigation was anticipated, with the ability to review and evaluate the decision-making process gone forever. Further, if the ML is not generating logs or leaving behind other information explaining its decisions or interim calculations, there will be little information to add to the discovery process.

C. *Autonomy*

An ML's ability to act more autonomously—more like an independent human—can also produce a multitude of legal issues.

1. Contracts

There are many advantages to enabling an AI to enter into a contract, for both the seller and buyer. Commercial benefits include quicker transactions, more sophisticated decision making, reduction in clerical errors, and cost cutting.⁴⁹ However, both the seller employing the AI and the buyer interacting with the AI must ensure that the transaction is legally binding.⁵⁰

Thus, as ML applications act more autonomously, the ability to enter into agreements⁵¹—and the valid scope of those agreements—will likely become a more significant legal issue. At least one commentator determined that an AI cannot enter into an agreement because it is not a natural person, but agency, apparent authority, or even an AI serving as a “device” of an organization can suffice to make the agreement binding.⁵² Today, much of online commerce is driven without human interaction on the seller side and is normally subject to express terms and conditions provided by the selling site.

However, potential legal issues will grow as AI becomes more involved in negotiating actual terms and when there are AIs on both sides of the agreement negotiating those terms. Can there be a meeting of the minds if there are no human minds involved in creating the terms and conditions of a contract? Should human operators of AI be bound by contracts that they take no direct part in negotiating—and what if the negotiations are outside the intended scope? Questions like these and other policy determinations will need to be answered

49. Tom Allen & Robin Widdison, *Can Computers Make Contracts?*, 9 HARV. J.L. & TECH. 25, 29 (1996).

50. *Id.*

51. The classical formation of a contract requires: (1) two or more parties to the contract; (2) the two parties are in agreement; (3) the parties must intend to create a legal relation that promises each side to perform to the contract because of the contract; (4) the promises of each party must be supported by consideration. *Id.* at 30, citing CHARLES FRIED, *CONTRACT AS PROMISE: A THEORY OF CONTRACTUAL OBLIGATION* (1981).

52. *Id.* Or the AI could be a “device” of the company. See *State Farm Mut. Auto. Ins. Co. v. Bockhorst*, 453 F.2d 533 (10th Cir. 1972).

when determining if contract law needs to be modified when dealing with future AI.

2. Intent and Consent

Some laws require a finding of intent that may be difficult to conceptualize when a human is not involved. For example, the common law claim of tortious interference with contract requires, among other elements, knowledge of a pre-existing agreement and a finding of intent to act.⁵³ Similarly, many states require a human-based standard for an award of punitive damages. For example, in Arizona punitive damages require an “evil mind”: an intent to injure or wrongful conduct motivated by spite or ill will.⁵⁴

Similarly, an autonomous AI might also collect and store locally recorded data, either in efforts for providing auditability or in continuing to learn and improve. Regardless, this activity could invoke state and federal wiretapping laws.⁵⁵ Most current regulations require consent from parties to be videotaped or recorded and individuals near certain types of AI may be unaware that they are being recorded. This activity comes with stiff legal fines for the party that uses such recordings, but an AI employer may be unaware that this information is being taken in or used.⁵⁶ Further, if the AI has acted autonomously in making the recording, it may be difficult to determine whether and how existing laws apply (since classically there may be only one actual “party” to the conversation and no other human party triggers, hears, or uses the recording) and how and when consent requirements could be met. These standards of intent and other human-focused requirements lack clarity in their applicability to machine-based actors.

IV. A WAY FORWARD

AI presents many novel legal and compliance challenges. Some suggest that the problems should be solved technologically—that AI can be designed to explain its decisions, and that logs and audit information can help with the explanation process.⁵⁷ Those improvements could happen soon or in the future. However, it is likely that not all applications will have these capabilities, and that we will create new forms of AI that advance quickly and lack such capabilities. In addi-

53. Tortious interference with contract requires proof of four elements: (1) that a valid contract existed, (2) that defendant had knowledge of the contract, (3) that defendant acted intentionally and improperly, and (4) that plaintiff was injured by the defendant’s actions. *Intentional Interference with Contractual Relations*, LEGAL INFO. INST., https://www.law.cornell.edu/wex/intentional_interference_with_contractual_relations [<https://perma.cc/MU2D-8344>] (last visited Jan. 3, 2018).

54. *E.g.*, *Linthicum v. Nationwide Ins. Co.*, 723 P.2d 675, 679 (Ariz. 1986).

55. Mathiason, *supra* note 43.

56. *See id.*

57. *E.g.*, Doshi-Velez, *supra* note 41, at 12.

tion, there are other issues previously described, such as intent, that are difficult to interpret in an AI context.

Rather than rely on technological advancement, it is simpler and more direct to look at well-developed law and legal theories to determine which can apply directly to the challenges presented by AI. There are two distinct areas of concern: (1) assessing legal, regulatory and other liability or responsibility for AI; and (2) determining how to handle evidentiary and witness concerns when an AI holds the information necessary for those processes.

A. *Legal and Regulatory Liability and Responsibility*

An organization deploying any form of AI should generally be responsible for the actions and decisions resulting from its use of AI under the existing legal structure.⁵⁸ However, because AI can be applied on so many different levels, and those capabilities are growing, a continuum may be required for any analysis to determine responsibility. On one end of that continuum AI is being applied as a basic tool, while at the other a fully autonomous AI is making its own decisions, entering into agreements, and taking other independent actions.

1. AI as a Tool

In many situations, especially today, AI is acting merely as a tool to provide humans with additional insights, efficiencies, or other capabilities. A person is wielding that tool and is ultimately responsible for any decisions, which are merely being informed by the tool.⁵⁹ Examples abound, but these are situations where AI is perhaps reviewing data and providing recommended scores for credit, job applicants, and sentencing. The key factor here is that these are mere starting points that are viewed and vetted by humans, and are simply factors in an overall decision-making process.

In the COMPAS sentencing algorithm matter, the Wisconsin Supreme Court held that due process is not violated by the opaqueness of the algorithms used by ML and could be used in criminal sentencing cases to determine if an offender would likely offend again, even if the algorithm may have taken race into account.⁶⁰ However, the court noted concerns and provided limits on the use of the algorithm, including that “risk scores may not be used as the determinative factor”—essentially noting that COMPAS is permitted as a mere tool for the judge to use in sentencing.⁶¹ The US Supreme Court refused a writ

58. David C. Vladeck, *Machines Without Principals: Liability Rules and Artificial Intelligence*, 89 WASH. L. REV. 117, 121 (2014).

59. *Id.* at 120–21.

60. *See* State v. Loomis, 881 N.W.2d 749 (Wis. 2016); Mathiason, *supra* note 43.

61. *Loomis*, 881 N.W.2d at 768.

of certiorari to appeal the Wisconsin decision and additional concerns related to due process remain unanswered.⁶²

Situations where AI is a tool are the most direct to handle from a legal perspective—they should not require any changes to existing legal approaches. An individual using the tool remains responsible for any actions taken, just as with an employee who is using a computer or a spreadsheet to help perform their day-to-day activities.

2. AI As An Independent Decision Maker

The next point to examine on the continuum is when AI is acting as an independent decision maker in an overall process that includes humans. This differs from the tool scenario because the AI is making decisions, and the basis or reasoning for the decisions are typically unknown to others using or relying upon those decisions.

Problems will arise when another member of an organization makes a decision or takes action based upon the AI, and that individual is not responsible for reviewing or questioning the AI decision. For example, an AI may cull an employee applicant pool to a small group for the next step, meaning that it made the decision to remove all other applicants from the process.

This is a more difficult scenario than the tool situation previously described. However, for policy and social reasons, it seems clear that an organization must remain responsible for any decisions made by its AI. If a poorly designed AI selects only male candidates from a pool, creating a gender discrimination liability issue, the organization should remain responsible for those actions under standard legal theories. The fact that a ML application—selected and used by the organization—is making decisions instead of a (human) employee should not change the analysis.

V. AUTONOMOUS AI

An AI making decisions and undertaking further activity based on those decisions requires a different approach, and common legal frameworks will be more difficult to apply. An example is an online credit application that is processed from start to finish without human intervention. Today, there are generally humans involved at some level of a review process, at least for larger and more sophisticated loans such as mortgages and commercial loans or lines of credit. But as confidence and faith grow in AI, and business models transform, this may no longer be the case.

For legal responsibility and liability, a starting point is to treat an autonomous AI as an employee. The doctrine of *respondeat superior* holds an employer liable “for torts committed by employees while act-

62. Liu, *supra* note 44.

ing within the scope of their employment.”⁶³ Since the AI is acting on behalf of the employer, replacing activities normally performed by an employee, the employer should similarly be liable for the results of the AI’s activities. Although current AI technology can usually be recognized as a tool of an organization, as previously described, there may be times in which AI is better classified as an agent of a person or an organization and *respondeat superior* is better suited in determining liability.⁶⁴

However, even *respondeat superior* is merely a starting point and there are situations that present more difficult scenarios. What if the AI works very differently than expected—should the employing organization be liable for all actions? Microsoft released an AI bot in 2016 that was to learn how to “tweet” —send messages on the Twitter platform.⁶⁵ Within a day, the bot had learned some bad and apparently racist tendencies. The result was not Microsoft’s intent, and it shut down the bot. But if the bot had caused actionable legal harm, such as slandering an individual, should Microsoft have been responsible? Microsoft did not intend this type of activity and it does not appear to be within the course and scope of the bot’s intended “employment.” Should Microsoft be permitted to argue course and scope as a defense to a legal claim; or should Microsoft, which was solely responsible for the underlying code and the decision to unleash this bot on the public, be prohibited from arguing that defense? After all, Microsoft chose the learning algorithms and the environment for the data set to enable learning.

Compliance can present similar challenges for autonomous AI. Most companies have a code of conduct that employees must follow. What if the actions of an AI violate that code? Since the AI probably has not pledged to follow the code, should it be disciplined as an employee—and what form would that discipline take? Perhaps some consideration might be given to disciplining the employees responsible for deploying and monitoring the AI. However, that strategy could be overly punitive and might preclude people from actively working on such projects. One way to ensure compliance measures are met by AI is for corporations to only utilize AI that is able to explain the actions that it takes.⁶⁶

63. RESTATEMENT (THIRD) OF AGENCY § 2.04 (AM. L. INST. 2006).

64. Joel Espelien, *The Brave New World of Robot Law*, L. PRAC. TODAY (Jan. 14, 2016), <http://www.lawpracticetoday.org/article/the-brave-new-world-of-robot-law/> [https://perma.cc/BTR9-WKZS].

65. Rob Price, *Microsoft is Deleting Its AI Chatbot’s Incredibly Racist Tweets*, BUS. INSIDER (Mar. 24, 2016, 7:31 AM), <http://www.businessinsider.com/microsoft-deletes-racist-genocidal-tweets-from-ai-chatbot-tay-2016-3> [https://perma.cc/ASM8-EC6E].

66. Guruduth Banavar, *Learning to Trust Artificial Intelligence Systems: Accountability, Compliance and Ethics in the Age of Smart Machines*, IBM 4 (2016), www.research.ibm.com/software/IBMResearch/multimedia/AIEthics_Whitepaper.pdf [https://perma.cc/3D8N-VLNB].

A final area of interest in autonomous applications is with establishing intent as an element of a cause of action or to establish punitive damages. If an AI application is inscrutable, then without evidence that project members intended a result—which is only circumstantial evidence of intent by the AI—there is no evidence to shed any light on the question of intent. This seems to be an unfair position for an aggrieved party, and would wrongly reward organizations that do not move forward with efforts to have AI explain its actions. Perhaps a better solution in determining intent is to shift the burden of proof to the party deploying the AI. This would encourage deploying parties to generate logs and other information from the AI. Even without this data, a deploying organization could be permitted to use design and other deployment information to establish a lack of intent, which should help to prevent unjust results. Regardless, shifting the burden has the beneficial effect of forcing deploying parties to think more about their logging and auditing processes, and to be more diligent in their design and deployment.

A. *AI As A Legal Entity*

Some surmise that in the future, AI could function autonomously as its own legal entity.⁶⁷ With smart contracts, crypto currencies such as Bitcoin and other processes that do not require actual human or physical intervention, an AI could, for example, act as an intermediary to buy and resell goods, create and sell digital content, provide business advice and number crunching capabilities.

Clearly similar scenarios are not anticipated by our current laws or legal frameworks. What if the AI is sued by another party and it has not been set up to receive service of process, respond to regulatory inquiries or hire counsel? It is easy to envision a situation where a default judgment is entered against the AI for failure to appear. However, this may not be the most equitable decision—and executing against the judgment could prove difficult; there is no way for an AI to respond to a debtor's examination and even garnishments may be impossible if cryptocurrencies are involved. Ultimately, the AI could be run out of business if its servers are located and shut down, or if it operates in a third-party cloud environment and fails to pay its bills.

However, applying existing laws seems to be the most sensible route. If an AI is unleashed as an autonomous entity, then the AI should suffer any pains that result from a failure to follow local laws and regulations. This is akin to a situation where a (human) individual operates a business and fails to respond to legal process or regulatory inquiries—the only difference being that the human is making this de-

67. Matt Scherer, *Is AI Personhood Already Possible Under Current Laws?*, L. & AI (May 14, 2017), <http://www.lawandai.com/2017/05/14/is-ai-personhood-already-possible-under-current-u-s-laws-dont-count-on-it-part-one/> [https://perma.cc/V6CN-HSAR].

cision by choice and the AI is unaware or incapable. In situations where a significant amount of money is involved or there are other important considerations, it is possible that the AI entity might be placed into a form of receivership where advanced technical capabilities could be used to operate and wind down the business under well-established rules that are already in place.

B. *AI in Discovery or Disclosure*

A second difficult issue is determining how an organization can be properly represented during discovery or a regulatory inquiry or investigation when an AI is at issue.

1. Discovery and Investigations

One problem with discovery when an AI is involved is that there may be little or no useful information available during the process. However, discovery generally does not obligate parties to create documents that do not exist.⁶⁸ Thus, if an AI is not creating logs and does not have other capabilities to explain its process, there may be little or no evidence regarding its activities.

However, there are other options. The ML application itself—including the neural network architecture and weights—may need to be preserved and produced as ESI. The requesting party could potentially gain information about the operation of the ML by changing various input points and reviewing the results.⁶⁹

Alternatively, the requesting party might issue a Rule 30(b)(6)⁷⁰ deposition request mandating production of a witness who can testify to the operation of the ML application. Whether this produces useful information will be tested on a case-by-case basis, although the testimony ordinarily is binding on the organization (in many jurisdictions, the proffering organization can still contradict or impeach that testimony).⁷¹ A more difficult case is when such information is truly not reasonably available to the organization, as is likely to occur in many AI matters regarding the process of how the AI made a decision. If

68. See e.g., *Mervyn v. Atlas Van Lines, Inc.*, No. 13 C 3587, 2015 WL 3649136 (N.D. Ill. Oct. 23, 2015).

69. See e.g., *Doshi-Velez*, *supra* note 41 (discussing how empirical evidence can be used in interpreting AI decision making).

70. “In its notice or subpoena, a party may name as the deponent a public or private corporation, a partnership, an association, a governmental agency, or other entity and must describe with reasonable particularity the matters for examination. The named organization must then designate one or more officers, directors, or managing agents, or designate other persons who consent to testify on its behalf; and it may set out the matters on which each person designated will testify” FED. R. CIV. P. 30 (explaining Depositions by Oral Examination).

71. See generally *ABA Section of Litigation Federal Practice Task Force Report On Rule 30(b)(6) Depositions of Organizations*, 17 (2015), https://www.americanbar.org/content/dam/aba/administrative/litigation/materials/2016_sac/written_materials/2_30b6_aba_report.authcheckdam.pdf [<https://perma.cc/D67Q-L7XC>].

the lack of information is perceived as a failure to properly prepare, sanctions could be ordered, or an organization could be barred from presenting other evidence at trial.⁷²

2. Witnesses and Representatives

Both in the present and the foreseeable future, an AI will not be able to testify or respond directly to deposition or investigative questioning. In situations where the AI has acted more autonomously, this may mean that there is very little or no evidence—other than the final result—regarding the issues in question. This seems to be an unjust result and one that could wrongly encourage organizations to minimize information generated by their ML applications.

In situations where there is little or no information to explain actions related to the ML application, the best course may be to handle the investigation similar to a Rule 30(b)(6) deposition, as previously discussed. The organization should be required to select a person or persons with knowledge about the issues related to the AI and provide testimony (sworn where appropriate) on those issues. The organization should then be bound by this testimony and unable to provide contrary information on the issues. This should not be a punishment, but instead designed to encourage more transparency by ML-deploying parties and to place the onus upon them where their ML creates a legal or regulatory issue.

VI. CONCLUSIONS

ML and artificial intelligence are rapidly growing in usage, and will generate novel legal and regulatory issues and concerns as their use becomes more prevalent. From a technology perspective, many are calling for and working on capabilities that will explain how these technologies are making decisions, much as we have traditionally been able to ask a human whose role they are overtaking. However, there are no guarantees that these capabilities will be developed or accepted. It is therefore prudent to consider how existing legal and regulatory frameworks can be applied to these technologies, based upon their level of autonomy and societal interest in balancing their efficiencies with the need for legal certainty.

72. *Id.*; See also *QBE Ins. Corp. v. Jorda Enters.*, 277 F.R.D. 676, 688 (S.D. Fla. 2012) (“We don’t know” answer on 30(b)(6) topics may be binding and preclude further testimony on that topic at trial).