

12-1-2021

AI, on the Law of the Elephant: Toward Understanding Artificial Intelligence

Emile Loza de Siles

Duquesne University School of Law

Follow this and additional works at: <https://digitalcommons.law.buffalo.edu/buffalolawreview>



Part of the [Law and Politics Commons](#), and the [Law and Society Commons](#)

Recommended Citation

Emile L. de Siles, *AI, on the Law of the Elephant: Toward Understanding Artificial Intelligence*, 69 Buff. L. Rev. 1389 (2021).

Available at: <https://digitalcommons.law.buffalo.edu/buffalolawreview/vol69/iss5/3>

This Article is brought to you for free and open access by the Law Journals at Digital Commons @ University at Buffalo School of Law. It has been accepted for inclusion in Buffalo Law Review by an authorized editor of Digital Commons @ University at Buffalo School of Law. For more information, please contact lawscholar@buffalo.edu.

Buffalo Law Review

VOLUME 69

DECEMBER 2021

NUMBER 5

AI, on the Law of the Elephant: Toward Understanding Artificial Intelligence¹

EMILE LOZA DE SILES†

“Reality is one, though wise ones speak of it variously.”

Rigveda, 1500–1200 B.C.E.²

1. Back when the internet was nascent, Judge Easterbrook asserted that there was no need for or wisdom in the specific development of internet law. In short, he said that we might as well create a “law of the horse” for all the sense that would make. Frank H. Easterbrook, *Cyberspace and the Law of the Horse*, 1996 U. CHI. LEGAL F. 207, 207 (1996). That horse proving too provocative to resist, Professor Lawrence Lessig wrote a presciently insightful response. See Lawrence Lessig, *The Law of the Horse: What Cyberlaw Might Teach*, 113 HARV. L. REV. 501 (1999). Professor Lessig’s views went on to carry the day, and their conversation inspired countless others, including this author.

2. *Rigveda* (1500–1200 B.C.E.), quoted in PAUL J. GRIFFITHS, AN APOLOGY FOR APOLOGETICS: A STUDY IN THE LOGIC OF INTERRELIGIOUS DIALOGUE 46 (1991); see *Rigveda, Hindu Literature*, BRITANNICA, <https://www.britannica.com/topic/Rigveda> (Mar. 12, 2020).

†Emile Loza de Siles is Visiting Assistant Professor at Howard University School of Law and Assistant Professor at Duquesne University School of Law. She founded Technology Law Group in 2003 and has provided technology and intellectual property legal services to Cisco, HP, Accenture, and other innovators. She is also Associate Professor (adjunct) of the University of Maryland Global Campus’s graduate cybersecurity program. Professor Loza de Siles has been nominated to serve on the inaugural National Artificial Intelligence Advisory Committee and on its Subcommittee on Artificial Intelligence and Law Enforcement. See Call for Nominations To Serve on the National Artificial Intelligence Advisory Committee and Call for Nominations To Serve on the

ABSTRACT

Machine learning and other artificial intelligence (AI) systems are changing our world in profound, exponentially rapid, and likely irreversible ways.³ Although AI may be harnessed for great good,⁴ it is capable of and is doing great harm at scale to people, communities, societies, and democratic institutions.⁵ The dearth of AI governance leaves unchecked AI's potentially existential risks.⁶ Whether

Subcommittee on Artificial Intelligence and Law Enforcement, 86 Fed. Reg. 50326-01 (Sept. 8, 2021). She chairs the Section on Minority Groups of the Association of American Law Schools. She is a member of the Institute of Electrical and Electronics Engineers (IEEE) and serves on its Artificial Intelligence and Autonomous Systems Policy Committee, P2863 Working Group on the Organizational Governance of Artificial Intelligence, and other AI working groups. Professor Loza de Siles holds a technology undergraduate degree, an MBA, a JD from The George Washington University, and cybersecurity strategy management graduate certificate from Georgetown, with a data science graduate certificate underway with Harvard.

Thanks to Steven Bender, Miguel Bordo, Sherley Cruz, César Cuauhtémoc García Hernández, Eric Goldman, Verónica Gonzales-Zamora, Pratheepan Gulasekaram, Jan Levine, Amy Lovell, Gary Marchant, Katherine Norton, Christine Pollack, Julia Tedjeske, and Kristy White. Special thanks to Don Simon, Shivam Rai, and Chetan Ganjihal for their technical reviews; my research assistants Danielle Mrjanovich, Kate Dumais, and Daniel Matesic; and the members of the *Buffalo Law Review*. Gratitude always to my family. *Animo!* Contact: lozae@duq.edu.

3. See *Artificial Intelligence for Europe*, at 1, COM (2018) 237 final (Apr. 25, 2018), [https://ec.europa.eu/transparency/documents-register/detail?ref=COM\(2018\)237&lang=en](https://ec.europa.eu/transparency/documents-register/detail?ref=COM(2018)237&lang=en).

4. See, e.g., *AI for Good Global Summit*, INT'L TELECOMM. UNION, <https://aiforgood.itu.int/> (last visited Dec. 7, 2021).

5. See generally Peter K. Yu, *The Algorithmic Divide and Equality in the Age of Artificial Intelligence*, 72 FLA. L. REV. 331, 343–61 (2020) (addressing AI's tremendous benefits and emerging problems, such as harmful algorithmic biases).

6. See Michael Guihot, Anne F. Matthew & Nicolas P. Suzor, *Nudging Robots: Innovative Solutions to Regulate Artificial Intelligence*, 20 VAND. J. ENT. & TECH. L. 385, 414–27 (2017) (pointing to regulatory deficiencies as AI technology rapidly advances). Few laws or regulations specifically address AI, although state and local authorities have begun to lead the way. See, e.g., City and County of S.F., Cal., Ordinance No. 103-19 (May 21, 2019), <https://sfbos.org/sites/default/files/o0103-19.pdf>; N.Y.C., N.Y., Local Law No. 49 (Jan. 11, 2018), <https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=3137815&GUID=437>

sounding urgent alarm or merely jumping on the bandwagon, law scholars, law students, and lawyers at bar are contributing volumes of AI policy and legislative proposals, commentaries, doctrinal theories, and calls to corporate and international organizations for ethical AI leadership.⁷ Unfortunately, erroneous, incomplete, and overly simplistic treatments of AI technology undermine the utility of a significant portion of that literature. Moreover, many of those treatments are piecemeal, and those gaps produce barriers to the proper legal understanding of AI.

Profound concerns exist about AI and the actual and potential crises of societal, democratic, and individual harm that it causes or may cause in future. On the whole, the legal community is not currently equal to the task of addressing those concerns, lacking sufficient AI knowledge and technological competence, despite ethical mandates for diligence and competence.⁸ As a result, law and policy debates and subsequent actions may be fundamentally flawed or produce devastating unintended consequences because they relied upon erroneous, uninformed, or misconceived understandings of AI technologies, inputs, and processes. Like the elephant in the ancient Jain parable, the wise ones may conceive of only a fraction of the AI creature and some more or less blindly.⁹

Now more than ever, lawyers need to be able to see around critically important corners. The general lack of understanding about AI technology robs the legal profession

A6A6D-62E1-47E2-9C42-461253F9C6D0.

7. See, e.g., Guihot et al., *supra* note 6, at 436–37; see also Emilie C. Schwartz, Note, *Human vs. Machine: A Framework of Responsibilities and Duties of Transnational Corporations for Respecting Human Rights in the Use of Artificial Intelligence*, 58 COLUM. J. TRANSNAT'L L. 232 (2019) (noting AI's harmful implications for international human rights).

8. See MODEL RULES OF PRO. CONDUCT r. 1.1, 1.3 (AM. BAR ASS'N 2020).

9. See Easterbrook, *supra* note 1, at 207 (“Beliefs lawyers hold about computers, and predictions they make about new technology, are highly likely to be false. This should make us hesitate to prescribe legal adaptations for cyberspace. The blind are not good trailblazers.”).

of that foresight. This state of affairs also raises significant ethical concerns. Worse, it undermines lawyers' power, authority, and legitimacy to bring forward truly valid, meaningful ideas and solutions to prevent AI from becoming humanity's apex predator.¹⁰

This Article responds with several descriptive and theoretical contributions. As to its descriptive contributions, it aims to correct and augment the record about AI, particularly machine learning and its underlying technologies and processes. It endeavors to present a concisely and accessibly stated foundational, but sufficiently comprehensive, single-source explanation. The Article draws extensively from the scientific and technical literatures and undertakes an important interdisciplinary¹¹ translational process by which to map the AI technical lexicon¹² to legal terms of art and constructions in patent and other cases.¹³ Because their understanding is foundational, the Article drills down on three principal AI inputs: data, including data curation; statistical models; and algorithms. It then engages in illustrative issue-spotting within these AI factual frames, sketching out some of the many legal implications associated with those vital understandings.

Toward its theoretical contributions, the Article presents two conceptual sortings of AI and introduces a systems- and process-engineering-inspired taxonomy of AI. First, it categorizes AI by the degree of human involvement in and, conversely, the degree of AI autonomy in AI-mediated decision-making. Second, it conceptualizes AI as being static or dynamic. Those distinctions are vital to AI's potential for

10. See the esoteric mind of *mi marido* for this terrifying vision of AI; see also *WAR OF THE WORLDS* (Paramount Pictures and DreamWorks Pictures 2005).

11. See generally Richard A. Posner, *The Decline of Law as an Autonomous Discipline: 1962–1987*, 100 HARV. L. REV. 761 (1987).

12. See ANTONIN SCALIA & BRYAN A. GARNER, *READING LAW: THE INTERPRETATION OF LEGAL TEXTS* 415 (2012) (“[O]ne of the chief functions of our courts is to act as an animated and authoritative dictionary.”).

13. See, e.g., *Markman v. Westview Instruments, Inc.*, 517 U.S. 370 (1996).

harm, meaningful accountability, and, ultimately, the proper prioritization of AI governance efforts. Third, the Article briefly introduces a taxonomy that conceptualizes AI as a human-machine enterprise made up of series of processes. By perceiving “the whole of the AI elephant,” the role of human decision-making and its limits may be understood, and the human-machine enterprise that is AI and its constituent processes may be deconstructed, comprehended, and framed for subsequent scholarship, doctrinal and procedural analyses, and law and policy developments. With these, the Article hopes to help inform and empower lawyers to improve the security, justness, and well-being of people in the increasingly algorithmic world.

CONTENTS

| | |
|--|------|
| ABSTRACT | 1390 |
| CONTENTS | 1394 |
| INTRODUCTION | 1395 |
| I. THE PROBLEM SPACE | 1400 |
| A. <i>Crossing the Lexicon: Lost in Translation</i> | 1401 |
| B. <i>Abstraction and Other Obfuscations</i> | 1405 |
| C. <i>Inadequate Adherence to Ethical Duties</i> | 1409 |
| II. THE GROUNDWORK | 1410 |
| A. <i>The Ethical Requirements: ABA</i> <i>Model Rules 1.1 and 1.3</i> | 1410 |
| 1. Ethical Duty of Technological Competence | 1411 |
| 2. Ethical Duty of Diligence..... | 1412 |
| B. <i>The Mental Model: AI as a Human-Machine</i> <i>Enterprise Comprised of Processes</i> | 1413 |
| III. THE ELEPHANT..... | 1416 |
| A. <i>What Is Artificial Intelligence?</i> | 1418 |
| B. <i>Legal Taxonomies for AI Decisional Context</i> <i>and System Mutability</i> | 1421 |
| 1. Decisional Use Context: Automated Decision System or Automated Decision Support System? | 1422 |
| 2. Operational Character: Static or Dynamic Artificial Intelligence? | 1424 |
| C. <i>What is Machine Learning?</i> | 1427 |
| 1. An Exemplar Machine Learning System in Operation | 1429 |
| 2. Models of Machine Learning | 1432 |
| IV. THE ELEPHANT AS PROCESS: AI INPUTS | 1436 |
| A. <i>Data, Big Data, and More and Different Data</i> . 1436 | |
| 1. What Are Data?..... | 1437 |
| 2. What Is “Big Data”?..... | 1439 |
| 3. What Is Data Curation?..... | 1443 |
| 4. Mapping Data to Artificial Intelligence | 1446 |
| B. <i>Statistical Models and AI Modeling Processes</i> . 1452 | |

| | | |
|-------|---|------|
| 2021] | <i>AI, ON THE LAW OF THE ELEPHANT</i> | 1395 |
| | 1. What is a Statistical Model?..... | 1453 |
| | 2. Model Building, Determination, Selection, and Finding | 1455 |
| | C. <i>Algorithms in Artificial Intelligence</i> | 1458 |
| | 1. What Is an Algorithm?..... | 1458 |
| | 2. Common Classes of Artificial Intelligence Algorithms..... | 1462 |
| V. | CONCLUDING FORWARD | 1469 |

INTRODUCTION

On an ancient day, a nomadic traveler, riding on the back of an elephant, arrived in a remote village. The massive and remarkable creature was an absolute mystery, something never before beheld by the village folk. Fearful and confused, they ran to the village leader. “Go,” he commanded, “bring our old ones. They will tell us what it is.” Presently, a group of wizened and blind villagers approached, led by the excited sighted who placed them before the beast and implored them to explain. The old blind women and men reached out their hands and thus encountered different parts of the animal. One of these found the fearsome point and heavy curve of a tusk, crying out, “Why, this is an enormous plow!” One, touching the writhing trunk, drew back and exclaimed, “It is a giant snake! Watch out lest it bite you!” Another felt one of the elephant’s big flapping ears, calling out, “No, it is a great fan!” Yet another old one exclaimed, “Of course! It is a mahogany tree!”, wrapping her arms around a massive leg. “No, it is a sturdy wall!,” another old one claimed, feeling his way down the elephant’s mountainous sides. “You’re all wrong!” the last cried, “It is a rope,” gripping the creature’s thick cable of tufted tail. In the midst of growing confusion and dissension, the nomad approached. “Why, good people! The reality is but one, although your wise ones speak of it variously. This reality is an elephant!”¹⁴

14. GRIFFITHS, *supra* note 2, at 46 (quoting *Rigveda* (1500–1200 B.C.E.)); see *Rigveda, Hindu Literature*, BRITANNICA, <https://www.britannica.com/topic>

This is the Jain parable of *andha-gaja-nyāya*.¹⁵ For the law and lawyers today, it is a metaphor for understanding the domain of artificial intelligence.

Anekāntavāda is the Jain doctrine of the multiplexity or many-sidedness of reality.¹⁶ Under the doctrine of *anekāntavāda*, reality, in addition to being many-sided, is in a constant and inevitable state of change.¹⁷ The Jain parable of the blind ones and the elephant illustrates that reality in its infinite nature is perceived and understood based upon differing predications, which, in turn, give rise to necessarily partial views.¹⁸ Where such partial views are unconditionally accepted, each view holder is weddedly blind to reality's other properties, those properties sitting outside the scope of that view. The proper comprehension of reality in its complexity requires a method of study and logical analyses that incorporates all viewpoints.¹⁹ Without this, a “superficial, deficient cognition” results by which view holders grasp at partial or scant data and underdeveloped notions.²⁰ The law can go seriously wrong when emanating from such cognition.

In the language of physics, reality under *anekāntavāda* reflects a high state of entropy, that is, a high degree of unpredictability and disorder.²¹ *Anekāntavāda* and high

/Rigveda (Mar. 12, 2020).

15. Transliterated from Sanskrit, *andha-gaya-nyāya* means “the maxim of the blind people and the elephant.” Piotr Balcerowicz, *Some Remarks on the Naya Method*, in *ESSAYS IN JAINA PHILOSOPHY AND RELIGION* 37, 40 (Piotr Balcerowicz ed., 1st Indian ed. 2003). I contribute my own version of this parable.

16. *See id.* at 37; *Anekāntavāda, Jainism*, BRITANNICA, <https://www.britannica.com/topic/aneantavada> (last visited Nov. 16, 2021).

17. *See Balcerowicz, supra* note 15, at 37.

18. *See id.* at 39–40.

19. *See id.* at 40–41.

20. *Id.*

21. *See* Gordon W.F. Drake, *Entropy*, BRITANNICA, <https://www.britannica.com/science/entropy-physics> (last visited June 2, 2021). Entropy is likewise a measure of disorder and, in machine learning, of a particular random variable's uncertainty. *See* MOHSSEN MOHAMMED, MUHAMMAD BADRUDDIN KHAN

entropy well describe the current state of law, legal practice at bar and bench, and legal scholarship and education regarding artificial intelligence and related technologies (collectively, AI).

In the many-sided, constantly changing reality of AI, the law is struggling. The technologies and their implications are complex and rapidly evolving. Compounding that complexity, too few lawyers understand AI and how it works or, at a more foundational level, its terminology. The terrible consequence is that lawyers' lack of sufficient AI knowledge and competency places people, communities, institutions, civil society, and even the rule of law itself at grave and perpetual risk through continued exposure to rapidly scaling, but virtually ungoverned AI.²²

A lack of effort by legal scholars and other legal writers is not the cause of this lack of AI knowledge and competency. The legal literature is replete and burgeoning with articles mentioning AI and related topics.²³ Indeed, the legal literature around AI exploded at least five years before the topic came to any arguably meaningful Congressional attention in 2018.²⁴ It is urgent that the AI legal scholarship

& EIHAB BASHIER MOHAMMED BASHIER, *MACHINE LEARNING: ALGORITHMS AND APPLICATIONS* 38 (2017).

22. *But see* Daniel L. Chen, *Machine Learning and the Rule of Law*, in *LAW AS DATA: COMPILATION, TEXT, AND THE FUTURE OF LEGAL ANALYSIS* 429, 433–41 (Michael A. Livermore & Daniel N. Rockmore eds., 2019) (arguing machine learning used to detect judicial bias, arbitrariness and variability may enhance rule of law and improve judicial education and decision-making).

23. For example, in 2019, there were 2,107 law review and bar journal articles about AI. Westlaw Queries, <https://1.next.westlaw.com/Search/Home.html> (searches using following criteria: “algo!,” “artificial intelligence,” “big data,” and “machine learning”) (on file with author). By contrast, in 2013, 814 such articles appeared. *Id.* In 2017, there were 1,881 such articles, 1,475, or some 78% of which, appeared in law review journals. *Id.*

24. *See* Comparative Analysis of Artificial Intelligence Search Results (Jan. 1, 2020) (based upon Westlaw Searches of Law Review and Journal Articles, <https://1.next.westlaw.com/Search/Home.html> (searches using following criteria: “artificial intelligence,” “machine learning,” “algorithm!,” and “big data”) and of Federal Congressional Record, <https://www.congress.gov/congressional-record> (searches using following criteria: “artificial intelligence,” “machine learning” of

should direct great energy toward rationalizing AI within legal constructs and to contribute guidance for those who will, and must, govern AI-mediated societies. Nothing short of the future of the rule of law and of humankind is at stake.

Contemporary explanations of AI in the legal literature, however, are often superficial²⁵ and, in some instances, even erroneous. Coverage of the technological foundations of AI phenomena is a patchwork with legal scholarly writings tending to focus on small pieces of the elephant as bases for the legal theories or policies therein proposed. Adopting a narrower and caveated scope of view is a common and important scholarship practice, but one that, here, confounds what law needs: a sufficiently detailed, comprehensive, and carefully expressed understanding of AI and related technologies.

Absent that, readers of narrow and partial legal treatments of AI may ignore the finer nuances of and carefully constructed caveats applicable to AI and, consequently, may arrive at faulty cognitive understandings about AI.²⁶ This, in turn, undermines the logical bases and appropriate application of the AI law and policy constructs

'neural network!' or 'deep learning' or 'reinforcement learning,' "algorithm!," and "big data' or 'data /2 broker!' or 'aggregate!' or 'miner' or 'mining' or 'appender!'")) (on file with author). The latter search criteria were expanded in comparison to the former to produce more numerous results from the Congressional Record. Analyses of the Congressional Record results showed spiking instances of the appearance of the search terms starting in 2018. In the year with the highest number of results, and the latest examined year, however, there were still only 215 appearances of any of the search terms in the Congressional Record. *Id.* The yearly Congressional Record appearances totaled 34, 43, 54, 50, 53, 98, and 215 for 2013 through 2019, respectively. *Id.* The yearly law review and journal appearances totaled 430, 526, 513, 654, 901, 1,170, 1,291, 1,673, 2,273, and 2,837 for 2010 through 2019, respectively. *Id.*

25. This observation holds even when accounting for the necessary brevity of bar journal articles, which constituted about one quarter of the literature in a sample year.

26. "It is a capital mistake to theorise before one has data. Insensibly one begins to twist facts to suit theories, instead of theories to suit facts." SIR ARTHUR CONAN DOYLE, THE ADVENTURES OF SHERLOCK HOLMES 5 (1892) (*The Adventures of a Scandal in Bohemia*).

proposed in the writings. Worse yet, inferences, subsequently discovered to have been erroneously drawn, may be lain later at the feet of legal scholars and unjustly call into question the validity of their arguments and recommendations. The continued lack of a comprehensively laid out understanding of AI in the legal literature as support for other AI legal scholarship is deeply problematic. Forsooth, how does the elephant run—walk—survive with only three legs? Poorly, if at all.

In contrast, with such a companion and support as this Article aspires to be, clearer and more cohesive lines of inquiry can emerge within the AI legal scholarship. By doing so, the needs to repetitively explain the current technological underpinnings of AI and to expend valuable publication “real estate” diminish. This frees legal scholars to focus with even greater depth and intensity upon the theoretical and procedural questions at hand. The Article also hopes to expand the field of inquiry for legal scholars by offering a scaffolding upon which to build and consider more finely drawn legal implications of, for example, uses of convolutional neural networks and imputed data for international economic development, human rights, and sustainability²⁷ or recurrent neural networks used in medical imaging contexts.²⁸ The impact of legal scholarship to inform thinking, public policy, lawmaking, judicial and regulatory decision-making, and, in sum, the “long game” about AI and what it means for justice and law²⁹ will increase as readers better adopt and elaborate upon those

27. See, e.g., Carla Gomes et al., *Computational Sustainability: Computing for a Better World and a Sustainable Future*, 62 COMM'NS ACM, Sept. 2019, at 56, 57–58, <https://dl.acm.org/doi/10.1145/3339399>.

28. See generally Robert DiPietro & Gregory D. Hager, *Deep Learning: RNNs and LSTM*, in HANDBOOK OF MEDICAL IMAGE COMPUTING AND COMPUTER ASSISTED INTERVENTION 503 (S. Kevin Zhou et al. eds., 2020).

29. Robin West & Danielle Citron, *On Legal Scholarship*, ASS'N OF AM. L. SCHS. 14 (Aug. 2014), <https://www.aals.org/wp-content/uploads/2014/08/OnLegalScholarship-West-Citron.pdf>. See generally DiPietro & Hager, *supra* note 28.

teachings because they are built upon a strong foundational understanding of AI. They will no longer be those wise, but blind persons trying to comprehend the whole of the elephant by grasping whatever part of the gargantuan creature is at hand.³⁰ It will take time, however, to flesh out the beast.

Toward that, this Article offers the following contributions. First, it reads the language of AI technologies, inputs, and processes into the legal literature through an interdisciplinary translational map by which to relate the AI technical lexicon to legal terms of art and constructions in patent and other cases. Second, it sketches some exemplary legal implications associated with those vital understandings. Third, it offers an engineering and law-driven taxonomy for conceptualizing the whole of the AI elephant as a process to be deconstructed and analyzed. Because their understanding is fundamental to subsequent learnings and work regarding AI, the Article then drills down to analyze three principal inputs for artificial intelligence within that taxonomic model: data, statistical models, and algorithms.

I. THE PROBLEM SPACE

The AI problem space for lawyers consists of three interlinking problems. First, the most fundamental basis for comprehending AI is the language of its technical arts. The AI lexicon is confounding. Second, the constituent components of AI are equally confounding, each having its own deep complexity. Add to those further complexities in the diversity of AI applications, delivery models, and markets, for example. Automation bias, the use of abstraction as an analytic tool before AI has been unpacked from its black boxes, and other perspectives and analytical

30. See JOHN GODFREY SAXE, *The Blind Men and the Elephant*, in THE POEMS OF JOHN GODFREY SAXE 259, 259–61 (1873); see also David M. Zlotnick, *The Buddha's Parable and Legal Rhetoric*, 58 WASH. & LEE L. REV. 957, 958–61 & nn.5–9 (2001).

practices also create barriers. Third, the ethical rules that govern all attorneys unequivocally require competency, including technological competency in general and AI competency in particular, and the diligence by which to gain and maintain that competency. Absent independent enforcement of these duties by governments or private litigation, however, competency compliance remains a matter of self-enforcement, which seems a gravely inadequate approach in the face of all that AI portends for law and justice.

This Section discusses each of these three interlinking problems in turn and then closes with a summary of some of the AI dangers that threaten lives, communities, and civil society if those problems are not addressed soon.

A. *Crossing the Lexicon: Lost in the Translation*

When Julius Caesar dared to cross the Rubicon, he broke the *lex cornelia majestatis*, committing treason and embarking on an irrevocable course of action that sparked revolution, civil war, and, ultimately, his triumphant ascendancy as the Roman emperor.³¹

The unknown likewise confronts those who wish to cross into and read the AI lexicon. Terms of foreign arts lie ahead, terms that are highly technical, esoteric, and seemingly abstruse in the extreme. When the reading starts discussing feature vectors,³² the strengths and weaknesses of various statistical models, and zero-shot learning,³³ it can be

31. See David Luban, *On the Commander in Chief Power*, 81 S. CAL. L. REV. 477, 494 n.56 (2008); Emily Rodriguez et al., *Rubicon*, BRITANNICA, <https://www.britannica.com/place/Rubicon> (last visited Dec. 8, 2021); Arnold Joseph Toynbee, *Julius Caesar*, BRITANNICA, <https://www.britannica.com/biography/Julius-Caesar-Roman-ruler> (last visited Dec. 8, 2021); Ernesto Valgiglio, *Sulla*, BRITANNICA, <https://www.britannica.com/biography/Sulla#ref141208> (last visited Dec. 8, 2021).

32. See, e.g., ANDRIY BURKOV, *THE HUNDRED-PAGE MACHINE LEARNING BOOK* 1–4, 9–10 (2019) (ebook), <http://themlbook.com/>.

33. See, e.g., *id.* at 95–96.

intimidating. Add in mathematical notation and formulae, programming syntax, and model diagrams, and it can get downright fearsome.³⁴

Furthermore, the terms of art from cognitive science, statistics, computer science, and other AI-involved disciplines do not necessarily map to terms of art in law.³⁵ Some words may be so deeply ingrained and veneered in the legal arts that technical terms of art may be immediately imbued with legal connotations. Such “automated thinking” may engender error and blind readers to the term’s proper understanding within the AI context.³⁶

Bias is one such term: essential on one side, anathema on the other. For example, inductive bias, or “bias,” is “the set of all factors that influence hypothesis selection,”³⁷ and that bias is essential to the proper functioning of AI and, specifically, to machine learning.³⁸ Briefly, under a supervised machine learning model, the AI system is exposed

34. See, e.g., JONAS PETERS, DOMINIK JANZING & BERNHARD SCHÖLKOPF, ELEMENTS OF CAUSAL INFERENCE: FOUNDATIONS AND LEARNING ALGORITHMS 66, 85 (2017).

35. See, e.g., *State v. Torgerson*, 611 N.W.2d 182, 184 (N.D. 2000) (random selection under Federal Jury Selection and Service Act); *United States v. Rioux*, 97 F.3d 648, 655 (2d Cir. 1996) (statistical decision theory applied to jury selection).

36. JUDITH S. HURWITZ, MARCIA KAUFMAN & ADRIAN BOWLES, COGNITIVE COMPUTING AND BIG DATA ANALYTICS 13 (2015); see *id.* at 13–14.

37. PAUL E. UTGOFF, MACHINE LEARNING OF INDUCTIVE BIAS 5 (1986). The views on biases in machine learning have developed greatly since 1980 when bias, as a core theoretical concept, was introduced. See Thomas Hellström, Virginia Dignum & Suna Bensch, *Bias in Machine Learning - What Is It Good For?*, ARXIV 2 (Apr. 1, 2020), <https://arxiv.org/pdf/2004.00686.pdf> (citing TOM M. MITCHELL, THE NEED FOR BIASES IN LEARNING GENERALIZATIONS, RUTGERS UNIV. COMP. SCI. DEP’T, TECH. REP. NO. CBM-TR-117, at 1 (1980) (referring to “bias” as “any basis for choosing one generalization over another, other than strict consistency with the observed training instances” (emphasis omitted))). Today, there are many types of bias in machine learning, some useful and some problematic. See generally *id.* This exemplary discussion considers inductive bias, now a cornerstone of statistical learning theory. See *id.*

38. See UTGOFF, *supra* note 37, at 5–29.

to training sets.³⁹ From those exposures, the system is to learn its statistical modeling by correlating combinations of input data and meaningful parameters, or features, within such data sets to yield computational results that reflect the so-called “target concept,” such as the concept that criminal offenders present some comparatively higher risk of recidivism.⁴⁰

Another way to express this is that the machine learner, once in operation post-training, carries out inductive computational processes by which meaningful attributes, i.e., features, within the data are analyzed, and the degree to which those features fit the target concept is hypothesized.⁴¹ Thus, the machine formulates various hypotheses and uses bias to sort through them, iteratively seeking to optimize its selection of the best hypothesis from among the available candidates, the best one being assumed to reflect the target concept.⁴² It is essential⁴³ to supply the activation points at which the learning system’s decision-making fires to decide

39. *See id.* at 5.

40. *Id.* at 3; *see also id.* at 5; Memoona Khanum et al., *A Survey on Unsupervised Machine Learning Algorithms for Automation, Classification and Maintenance*, 119 INT’L J. COMPUT. APPLICATIONS 34, 34 (2015); Donald Firesmith, *Multicore Processing*, SOFTWARE ENG’G INST.: SEI BLOG (Aug. 21, 2017), https://insights.sei.cmu.edu/sei_blog/2017/08/multicore-processing.html.

41. *See* STUART RUSSELL & PETER NORVIG, *ARTIFICIAL INTELLIGENCE: A MODERN APPROACH* 785–87 (3d ed. 2010); UTGOFF, *supra* note 37, at 3–6.

42. *See* UTGOFF, *supra* note 37, at 5.

43. *See id.*

A program that learns concepts from examples is successful *only* when it has a bias that guides it to make a satisfactory selection from among the available hypotheses. Without the bias, the program has no basis for rejecting one hypothesis in favor of another. If the concept learner [referring to the machine learning system] is to make choices on a non-random basis, then *bias is necessary*.

Id. (emphases supplied and citation omitted); *see also* JIAWEI HAN, MICHELENE KAMBER & JIAN PEI, *DATA MINING: CONCEPTS AND TECHNIQUES* § 9.2.3, at 403–04 (3d ed. 2012) (discussing updating of biases and weights in machine learning); PEDRO DOMINGOS, *THE MASTER ALGORITHM: HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD* 78–79 (2015) (analyzing differences and interrelationships between bias and variance).

“yes” or “no” when iteratively comparing one hypothesis with another and choosing between them toward the best hypothesis. Bias is what informs those points of decision.⁴⁴

In simplest technological terms, bias equals good. For machine learning, bias is more than good. It is essential.

For the computational minds of the law, bias has precisely the opposite character. The ideals toward which law strives are fairness, justice, and equality.⁴⁵ Bias is antithetical to those concepts. Its meaning and connotations are closely tied to racism, sexism, other forms of illegal discrimination, and other illegitimate bases for decision-making.⁴⁶ For law, bias equals bad.⁴⁷ At scale and as systematized, iteratively propagated, and perpetuated in the form of discrimination arising from the design, use, and misuse of AI systems, bias is not merely bad.⁴⁸ Unchecked, it

44. See, e.g., MOHAMMED ET AL., *supra* note 21, at 90–93 (discussing optimal weights and bias for AI system’s decisional activation).

45. *But see* Paul Butler, *The System Is Working the Way It Is Supposed To: The Limits of Criminal Justice Reform*, 104 GEO. L.J. 1419 (2016).

46. See, e.g., *Price Waterhouse v. Hopkins*, 490 U.S. 228 (1989).

47. See, e.g., Jon M. Garon, *AI and the Labor Laws*, in THE LAW OF ARTIFICIAL INTELLIGENCE AND SMART MACHINES: UNDERSTANDING A.I. AND THE LEGAL IMPACT 75, 81–83, 87 (Theodore F. Claypoole ed., 2019); Nizan Geslevic Packin & Yifat Lev-Aretz, *Learning Algorithms and Discrimination*, in RESEARCH HANDBOOK ON THE LAW OF ARTIFICIAL INTELLIGENCE 88, 95–113 (Woodrow Barfield & Ugo Pagallo eds., 2018); MICHELLE ALEXANDER, THE NEW JIM CROW: MASS INCARCERATION IN THE AGE OF COLORBLINDNESS 186–87 (rev. ed. 2012); Anjanette H. Raymond, Emma Arrington Stone Young & Scott J. Shackelford, *Building a Better HAL 9000: Algorithmics, the Market, and the Need to Prevent the Engraining of Bias*, 15 NW. J. TECH. & INTELL. PROP. 215, 233 (2018).

48. See, e.g., AI NOW INST., LITIGATING ALGORITHMS: CHALLENGING GOVERNMENT USE OF ALGORITHMIC DECISION SYSTEMS 13 (2018) [hereinafter AI NOW LITIGATING ALGORITHMS], <https://ainowinstitute.org/litigatingalgorithms.pdf> (discussing successful challenge to “long-standing” use of violence risk assessment system in juvenile criminal matters by District of Columbia courts and many “significant concerns about embedded racial bias” therein). See also, among key case documents discussed in AI NOW LITIGATING ALGORITHMS, *supra*, at 13, and on file with the author, Motion to Exclude Results of the Violence Risk Assessment and all Related Testimony and/or Allocution Under FRE 702 and *Daubert v. Merrell Dow Pharmaceuticals*, *In re* T.K. (D.C. Super. Ct. Feb. 5, 2018) [hereinafter *SAVRY Motion*] (challenging use of Structured Assessment of

is catastrophic.⁴⁹

The journey in this Article starts with the AI lexicon, to be crossed like Hannibal on war elephants and on to triumph in Rome.⁵⁰

B. *Abstraction and Other Obfuscations*

The challenges of the AI lexicon, the first AI problem, are compounded by the use of abstraction. Abstraction has been an indispensable analytic and explanatory tool for law since time immemorial.⁵¹ Abstraction enables collections of details to be organized and distilled into principles, theories, and factors.⁵² Thus, the contours of what suffices as control over *ferae naturae* can be distilled from parsing the facts of various circumstances presented over time, as the court did in the classic property case of *Pierson v. Post*.⁵³ Further, courts choose and apply different levels of abstraction to decide what are and are not fundamental rights under the

Violence Risk in Youth, or SAVRY, violence risk predictive system against juvenile defendant).

49. See CATHY O'NEIL, *WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY* 3 (2016); Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 14–15 (2014) (discussing need to test credit-scoring systems for human bias).

50. See William Culican, *Hannibal: Carthaginian General [247-c.181 BC]*, BRITANNICA, <https://www.britannica.com/biography/Hannibal-Carthaginian-general-247-183-BC> (last visited Nov. 17, 2021). To Jan and Alejandro, forgive me. I know. Caesar crossed the Rubicon, and I have mixed historical metaphors.

51. See *Abstraction Definition*, SIMPLYPHILOSOPHY, <https://simplyphilosophy.org/study/abstraction-definition/> (last visited Dec. 8, 2021); Brent Cooper, *How to Humanize AI with Abstraction*, MEDIUM (July 26, 2017), <https://medium.com/the-abs-tract-organization/how-to-humanize-ai-with-abstraction-bd379036e67a>.

52. Abstraction is also an organizing approach within AI and information technology more broadly. See HURWITZ ET AL., *supra* note 36, at 251. For example, in cloud computing, the computing infrastructure is abstracted away from the user in an infrastructure-as-a-service, or IaaS, model. See *id.*

53. See, e.g., *Pierson v. Post*, 3 Cai. R. 175, 178 (N.Y. Sup. Ct. 1805); Angela Fernandez, *Fuzzy Rules and Clear Enough Standards: The Uses and Abuses of Pierson v. Post*, 63 U. TORONTO L.J. 97, 99–108 (2013).

Constitution, for example.⁵⁴

The problem with applying abstraction to AI law and policy analyses is that, in most instances, it is premature to do so. The facts, specifically the relevant facts, must always come first. At least two analytical steps must occur prior to an exercise in abstraction. First, the facts must be laid out. Second, the facts must be winnowed to separate out only those facts relevant to the legal questions at hand. As cautioned by Judge Easterbrook in the context of constitutional rights, the abstracted principles underlying those rights are, or should be, the products—not the progenitors—of the relevant facts, those being the constitutional text and its history.⁵⁵

This is a chicken and elephant problem. These days, scholars grapple mightily to abstract legal principles to apply to AI and to construct coherent doctrinal theories and governance rules for it. A thorough understanding of AI's technological bases and interrelated workings is necessary to better sequence the analyses for more completely informed abstractions. A clearer perception of those interrelations then allows issues to emerge and thus implicate the requisite law. The elephant must always come first.

It is also important to avoid mystifying AI and to be aware of automation bias and other human weaknesses as they may obfuscate the points at which abstraction should be applied.⁵⁶ Automation bias is the use of and reliance upon automation to replace one's own processes of discovery,

54. See Frank H. Easterbrook, *Abstraction and Authority*, 59 U. CHI. L. REV. 349, 359–71 (1992).

55. See *id.* at 362–63.

56. See John D. Lee & Katrina A. See, *Trust in Automation: Designing for Appropriate Reliance*, 46 HUM. FACTORS 50, 51 (2004) (“[M]isuse and disuse of automation may depend on certain feelings and attitudes of users, such as trust. This is particularly important as automation becomes more complex”); see also, e.g., William E. Foster & Andrew L. Lawson, *When to Praise the Machine: The Promise and Perils of Automated Transactional Drafting*, 69 S.C. L. REV. 597, 598 (2018) (“[R]eliance on software creates risks of undue deference to computer-generated outputs”).

research, and analysis.⁵⁷ The current difficulty, if not near impossibility, of assessing AI's functions, accuracy, and validity tends to drive the analytical treatment of AI into the category of a "credence good," one that is "consume[d] on faith."⁵⁸ This may lead to viewing and accepting AI systems as impenetrable black boxes,⁵⁹ which, along with other challenges outlined herein, may short-circuit the depth and precision of factual inquiries necessary to properly enable abstraction approaches to AI.⁶⁰

Artificial intelligence is a human institution⁶¹ and, like other human institutions, must be appropriately understood, employed, and controlled, but never accepted with unquestioning deference. Although AI systems may be used for great and abundant good,⁶² they are not "infallible

57. See Danielle Citron, *Technological Due Process*, 85 WASH. U. L. REV. 1249, 1262, 1271 (2008); *Heuristic*, OXFORD ENG. DICTIONARY (3d. ed. 2014) (Entry B2).

58. Philip M. Napoli, *What if More Speech Is No Longer the Solution? First Amendment Theory Meets Fake News and the Filter Bubble*, 70 FED. COMM'NS. L.J. 55, 80 (2018).

59. STEPHEN LUCCI & DANNY KOPEC, *ARTIFICIAL INTELLIGENCE IN THE 21ST CENTURY* 6 (2d ed. 2016); see also Frank Pasquale, *A Rule of Persons, Not Machines: The Limits of Legal Automation*, 87 GEO. WASH. L. REV. 1, 51 (2019) (warning against reductive "black box" treatment of human decisional values, control, and contributions). *But see generally* Cooper, *supra* note 51.

60. See Sarah Valentine, *Impoverished Algorithms: Misguided Governments, Flawed Technologies, and Social Control*, 46 FORDHAM URB. L.J. 364, 394–99 (2019); Vera Eidelman, *The First Amendment Case for Public Access to Secret Algorithms Used in Criminal Trials*, 34 GA. ST. U. L. REV. 915, 923–25 (2018).

61. *Accord* O BROTHER, WHERE ART THOU? (Universal 2000) (Sheriff Cooley, responding to protagonists' protests to their extrajudicial lynching after being granted pardon, states: "The law? The law is a human institution.").

62. See, e.g., *AI for Good Global Summit*, INT'L TELECOMM. UNION, <https://aiforgood.itu.int/> (last visited Dec. 7, 2021); Michael A. Livermore, Vladimir Eidelman & Brian Grom, *Computationally Assisted Regulatory Participation*, 93 NOTRE DAME L. REV. 977, 1007–23 (2018). See generally, e.g., Marta Poblet & Jonathan Kolieb, *Responding to Human Rights Abuses in the Digital Era: New Tools, Old Challenges*, 54 STAN. J. INT'L L. 259 (2018); Lois R. Lupica, Tobias A. Franklin & Sage M. Friedman, *The Apps for Justice Project: Employing Design Thinking to Narrow the Access to Justice Gap*, 44 FORDHAM URB. L.J. 1363 (2017); SWEETIE 2.0: USING ARTIFICIAL INTELLIGENCE TO FIGHT WEBCAM CHILD SEX TOURISM (Simone Van der Hof et al. eds., May 2019); Watson Health, *Bridging the data-to-study gap to solve Rare Disease research challenges*,

oracles.”⁶³ Otherwise, for example, the deferential treatment that the broadly powered administrative state now enjoys from the judiciary may cement agencies’ algorithmically-mediated decisions in place as unchallengeable, even where AI systems and their input data and models and their uses are inaccurate, grossly racialized, or deeply problematic.⁶⁴

The elephant must be demystified. Otherwise, these issues will obscure the distinctions upon which legal decisions turn,⁶⁵ enabling AI’s unilluminated black boxes to have profound legal consequences.⁶⁶ Otherwise, the danger

IBM (May 24, 2021), <https://www.ibm.com/blogs/watson-health/bridging-the-data-to-study-gap-to-solve-rare-disease-research-challenges/> (last visited Dec. 8, 2021).

63. See Eidelman, *supra* note 60, at 923; *Commonwealth v. Serge*, 896 A.2d 1170, 1174 n.1 (Pa. 2006) (contrasting “product of *neutral infallible artificial intelligence*” with demonstrative computer animation in jury trial (emphasis supplied)); Eric Wang, *What Does It Really Mean for an Algorithm to be Biased?*, THE GRADIENT (May 1, 2018), <https://thegradiant.pub/ai-bias/> (Some tend to “think algorithmic reasoning is always rational and objective, regardless of the situation. They might even believe that uncomfortable or undesirable results of the data simply reflect ‘politically incorrect’ truths in the data.”).

64. See *Zirkle Fruit Co. v. U.S. Dep’t of Lab.*, 442 F. Supp. 3d 1366, 1382 (E.D. Wash. 2020) (“[A]t most this renders the dataset imperfect.”); James A. Allen, *The Color of Algorithms: An Analysis and Proposed Research Agenda for Detering Algorithmic Redlining*, 46 FORDHAM URB. L.J. 219, 223 (2019); EXEC. ETHICS COMM’N OF STATE OF ILLINOIS, OEIG FINAL REP’T (REDACTED) 36–37, 39 (2017) (Eckerd Rapid Safety Feedback system used in child welfare system); OKLA. DEP’T OF HUMAN SVC’S, SFTP APPLICATION SPECIFICATIONS MINDSHARE HRDM TIII Program 1, 4, <https://www.muckrock.com/foi/oklahoma-248/oklahoma-department-of-human-services-eckerd-rapid-safety-feedback-74007/#file-818375> (last visited Nov. 16, 2021) (finding 70% error rate in child risk prediction using above-referenced Eckerd system). Thanks to my Artificial Intelligence and Social Justice students Diana Bruce, Chelsea Hill, and Jared Myers for these two cited Eckerd-related sources. See also generally Stephanie K. Glaberson, *Coding Over the Cracks: Predictive Analytics and Child Protection*, 46 FORDHAM URB. L.J. 307 (2019) (discussing Eckerd system).

65. See, e.g., *Neuromedical Sys., Inc. v. NeoPath, Inc.*, No. 96 Civ. 5245(JFK), 1998 WL 264845, at *6 (S.D.N.Y. May 26, 1998) (discussing plaintiff Neopath’s “fuzzy decision tree”-equipped system, and stating that “[NeoPath’s] neural network is a ‘black box’ where reasoning for the neural network’s decision cannot be traced through the neural network to explain its decision rationale”).

66. See, e.g., *People v. Super. Ct. (Dominguez)*, 239 Cal. Rptr. 3d 71, 76, 84 (Ct. App. 2018) (granting discovery relief in murder conspiracy case where petitioner sought “right to look inside the proverbial ‘black box’” of probabilistic

is that humanity will self-subjugate to Dark Law. This is a post-humanist construct that will rise, if unchecked, from AI's great capacity to approximate law and from the rapid rise of algorithmic government regimes, including as outsourced to private companies.⁶⁷ Otherwise, AI's code, big data, and statistical models become law,⁶⁸ and people its impotent, fungible objects, reduced to mere data production units.⁶⁹

C. *Inadequate Adherence to Ethical Duties*

As elaborated in the next Section, Rules of Professional Conduct (Rules) 1.1 and 1.3 establish the attorneys' ethical duties and standards of technological competence and diligence, respectively.⁷⁰ Few ethical complaints have been adjudicated as to alleged violations of competency and related diligence, however, and none involving AI seems to have yet appeared. There may be several underlying causes for this apparent dearth of ethical enforcement as to technological competency and associated diligence. Whatever the causes, the problem is that compliance with

DNA genotyping system).

67. See O'NEIL, *supra* note 49, at 30–31; Robert Brauneis & Ellen P. Goodman, *Algorithmic Transparency for the Smart City*, 20 YALE J.L. & TECH. 103, 114–18, 126–28 (2018) (predictive algorithms and big data analytics as a form of governance); Deirdre K. Mulligan & Kenneth A. Bamberger, *Saving Governance-By-Design*, 106 CALIF. L. REV. 697, 722 (2018); Christoph B. Graber, *Freedom and Affordances of the Net*, 10 WASH. U. JURIS. REV. 221, 239 (2018) (*de facto* regulatory power of online market dominators); Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 IND. L.J. 1401, 1421 (2017); Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147, 1149 (2017); Omer Tene & Jules Polonetsky, *Taming the Golem: Challenges of Ethical Algorithmic Decision-making*, 19 N.C. J.L. & TECH. 125, 146–60 (2017).

68. See LAWRENCE LESSIG, *CODE AND OTHER LAWS OF CYBERSPACE* 60 (1999).

69. See Emile Loza de Siles, *Slave.io*, Remarks at Biennial 2021 LatCrit Conference, University of Denver Sturm College of Law (Oct. 9, 2021) (transcript on file with author); Emile Loza de Siles, *New Directions in Law and Society Workshop*, Center for Justice, Law, and Societies at University of Massachusetts Amherst (Oct. 9, 2021) (transcript on file with author).

70. See MODEL RULES OF PRO. CONDUCT r. 1.1, 1.3 (AM. BAR ASS'N 2020).

these ethical duties remains a matter of self-enforcement. That is the elephant in the room.

II. THE GROUNDWORK

This Section lays a two-part foundation toward a more comprehensive view of artificial intelligence. First, it discusses the ethical requirements for all lawyers, including licensed lawyers who are law professors, to attain and maintain AI competency through diligence. Second, the Article sketches out a mental model for conceptualizing and better understanding AI. That model derives from a proposed taxonomy for comprehending AI as a process, a useful approach where legal analyses and recommendations must first deconstruct the AI process to properly analyze, better theorize, and govern it.

A. *The Ethical Requirements: ABA Model Rules 1.1 and 1.3*

Artificial intelligence is broadly relevant, irrespective of practice, doctrinal, clinical, or judicial focus.⁷¹ “Deliberate ignorance of technology is inexcusable.”⁷² Everyone in, or headed into, the legal profession has, or will have, ethical duties to diligently learn about AI and its underlying technologies to become sufficiently competent and stay that way.⁷³

71. See Emile Loza de Siles, *The Future Is Now: Top Ten Strategic Technology Trends & Competence in Practice*, Remarks at Idaho State Bar Ann. Meeting (July 13, 2017) (transcript on file with author); Shannon Brown, *Peeking Inside the Black Box: A Preliminary Survey of Technology Assisted Review (TAR) and Predictive Coding Algorithms for Ediscovery*, 21 SUFFOLK J. TRIAL & APP. ADVOC. 221, 227–29 (2016).

72. *James v. Nat'l Fin. LLC*, Case No. CV-8931-VCL, 2014 WL 6845560, at *12 (Del. Ch. Dec. 5, 2014) (quoting Judith L. Maute, *Facing 21st Century Realities*, 32 MISS. C. L. REV. 345, 369 (2013)).

73. See Wilson Ray Huhn, *A Proposed Code of Ethics of Law Educators*, 6 J. L. & RELIGION 25 (1988). *But see id.* at 29 (“Incompetence on the part of a law teacher is not necessarily unethical.”). Most law professors likely are licensed attorneys, but research did not reveal the professoriate’s licensure rate.

Rules 1.1 and 1.3 establish the ethical duties and standards of technological competence and diligence, respectively.⁷⁴ These two duties are tightly coupled. Sufficient competency is achieved through diligence.⁷⁵ Where disciplinary actions are brought on competency grounds under Rule 1.1, failures of diligence under Rule 1.3 are often also alleged with those failures almost always being the cause of the competency breaches.⁷⁶

1. Ethical Duty of Technological Competence

Rule 1.1 reads: “A lawyer shall provide competent representation to a client. Competent representation requires the legal knowledge, skill, thoroughness and preparation reasonably necessary for the representation.”⁷⁷ Each of the elements under Rule 1.1 must be achieved to a reasonableness standard, that is, reflecting “conduct of a reasonably prudent and competent lawyer.”⁷⁸

Necessary study to the level of reasonable preparation enables a lawyer to achieve an ethically appropriate level of AI competency.⁷⁹ The reasonableness standard is met as to AI-related matters when, at a minimum, the lawyer can sufficiently assess the circumstances and then accurately spot legal issues.⁸⁰ A cautionary note, however, is not to overly rely upon Comment 5 to Rule 1.1, where it indicates that the required level of attention and preparation turns, in part, on the client’s interests at stake.⁸¹ Indeed, the very core

74. See MODEL RULES OF PRO. CONDUCT r. 1.1, 1.3 (AM. BAR ASS’N 2020).

75. See *id.* at 1.1 cmt. 2.

76. See, e.g., *Atty. Grievance Comm’n v. Zuckerman*, 872 A.2d 693, 703 (Md. 2005) (subsequent history omitted); GEOFFREY C. HAZARD, W. WILLIAM HODES & PETER R. JARVIS, *THE LAW OF LAWYERING* 4–7 (4th ed., 2014) (citing *id.*).

77. MODEL RULES OF PRO. CONDUCT r. 1.3 (AM. BAR ASS’N 2020).

78. *Id.* r. 1.0(h).

79. See *id.* r. 1.1 cmts. 2, 4.

80. See *id.* r. 1.1 cmt. 2.

81. See *id.* r. 1.1 cmt. 5.

of a lawyer's ability to determine what client interests are at stake in an AI-related matter depends upon the lawyer having sufficient, indeed, perhaps "the most fundamental," legal skill: to spot the issues that may elucidate what those interests are.⁸² In addition, the AI-competent lawyer must be able to: accurately identify and analyze appropriate legal precedent; identify and evaluate the relevant facts and legal elements and then to apply those facts to the law; and appropriately draft legal documents.⁸³ As to whether the attorney's legal knowledge and skills related to AI are reasonable, the commentary to the Rule provides a non-exhaustive list of factors to consider.⁸⁴

2. Ethical Duty of Diligence

Rule 1.3 reads: "A lawyer shall act with reasonable diligence and promptness in representing a client."⁸⁵ Rule 1.3 requires the attorney's zealous advocacy of and dedication and commitment to the client's interests.⁸⁶ The ethical duties as to AI under Rule 1.3 are dependent upon ethical compliance with the Rule 1.1 duty of competency. If the level of AI competency is ethically insufficient, then zealous advocacy, dedication, and commitment are undermined, if not rendered impossible, and are not a matter merely left to the lawyer's exercise of professional discretion.⁸⁷ As a consequence, the Rule 1.3 duty of diligence likewise may be breached when there is an AI competency breach under Rule 1.1.

82. *Id.* r. 1.1 cmt. 2.

83. *See id.* r. 1.1 cmts. 2, 5.

84. *See id.* r. 1.1 cmt. 1 (listing lawyer's general experience, lawyer's training and experience in the relevant field, preparation and study lawyer can commit to the matter, relative complexity and specialized nature of the matter, and feasibility to consult or refer the matter to another lawyer established in the relevant field as factors).

85. *Id.* r. 1.3.

86. *See id.* r. 1.3 cmt. 1.

87. *See id.*

This Section has discussed the twin ethical duties of technological competence and diligence assigned to licensed attorneys, the satisfaction of which are constant requirements, as emphasized by the ABA's Commission on Ethics 20/20.⁸⁸ These duties are the *minimum* performance requirements incumbent upon lawyers as to AI. Fortunately, there are multiple strategies by which to meet these required standards,⁸⁹ and this Article aims to help toward that end.

B. *The Mental Model: AI as a Human-Machine Enterprise Comprised of Processes*

A world-renowned giant in quality engineering, Dr. W. Edwards Deming minced no words: "If you can't describe what you are doing as a process, you don't know what you're doing."⁹⁰ A process engineering discipline and its corresponding methodologies serve the purposes of breaking down the problem set, that is, the system, here of AI, for which understanding and mastery of control are sought, into iteratively discrete subunits in that overall process.⁹¹ This

88. See Ann M. Murphy, *Is It Safe? The Need for State Ethical Rules to Keep Pace with Technological Advances*, 81 *FORDHAM L. REV.* 1651, 1661 (2013) (quoting Debra Cassens Weiss, *Lawyers Have Duty to Stay Current on Technology's Risks and Benefits*, *New Model Ethics Comment Says*, *ABA J.* (Aug. 6, 2012, 7:46 PM), https://www.abajournal.com/news/article/lawyers_have_duty_to_stay_current_on_technologys_risks_and_benefits); ABA Comm'n on Ethics 20/20, Resolution 105A (2012).

89. See MODEL RULES OF PRO. CONDUCT r. 1.1 cmts. 2, 6 (AM. BAR ASS'N 2020); *James v. Nat'l Fin. LLC*, No. CV-8931-VCL, 2014 WL 6845560, at *12 (Del. Ch. Dec. 5, 2014) (quoting Judith L. Maute, *Facing 21st Century Realities*, 32 *MISS. C. L. REV.* 345, 369 (2013)). *But see also* *Mia. Bus. Servs., LLC v. Davis*, 299 P.3d 477, 487 (Okla. 2013) (commentary to professional conduct rules are persuasive interpretative tools, but not binding). As to competency strategies, see Emile Loza de Siles, *supra* note 71.

90. See Jane K. Winn, *Reports of a Blockchain Revolution in Trade Finance Are Greatly Exaggerated* 18 (Jan. 27, 2020) (unpublished draft) (<https://ssrn.com/abstract=3526521>) (quoting Deming); *see also* *W. Edwards Deming Quotes*, *GOODREADS*, https://www.goodreads.com/author/quotes/310261.W_Edwards_Deming (last visited Dec. 8, 2021).

91. Some opt for "lifecycle," rather than "process." See Eric Horvitz, Comm'r, Nat'l Sec. Comm'n on A.I. & Chief Sci. Officer, Microsoft Corp., Remarks at the

deconstruction enables careful analyses of each subunit as a separate inquiry; the mapping of how those subunits interact and the dependencies and interdependencies of those interactions; and a deeply and more accurately comprehensive understanding of the overall system-as-process.

Aligning with Dr. Deming's straight talk, this Article conceives a mental model of AI as it is: a human-machine enterprise, an enterprise in which humans and machines are engaged in common in a systematized set of processes to be understood in detail, engineered, optimized, examined, and controlled. In keeping with its aims to lay a descriptive and comprehensive foundation for the legal understanding of AI technologies, this Article leaves the enterprise aspects of the model for another work and concentrates on processes. It considers a systematized AI process to be documented in the legal lexicon; deconstructed into its component subunits, process sequences, interactions, and dependencies; and analyzed in detail with clarity, rigor, and according to the law, or the law to be, within applicable technological and use case contexts. As one of its principal advantages, the model serves vital translating, orienting, and mapping functions.⁹² As with any process, the AI-as-process model encompasses inputs, outputs, and interacting dependencies.⁹³

Nat'l Inst. of Standards and Tech. (NIST), Exploring AI Trustworthiness Kickoff Webinar (Aug. 6, 2020), <https://www.nist.gov/news-events/events/2020/08/exploring-ai-trustworthiness-workshop-series-kickoff-webinar>.

92. See W. EDWARDS DEMING, *THE NEW ECONOMICS FOR INDUSTRY, GOVERNMENT, EDUCATION* 63 (3d ed. 2018).

93. This process view of AI as a human-machine enterprise model integrates concepts from three disciplines: process engineering; strategic business management, particularly product marketing; and legal landscape modeling. The model's first part invokes process engineering principles and focuses on three inanimate categories of AI inputs: data, statistical models, and algorithms. There is a fourth and human-driven category of AI input encompassing AI design: market deployment; use cases, uses, disuses, and misuses.

The model's second part groups AI outputs into three categories: the computational results produced by AI system use, such as recidivism risk score; outcomes, such as a harsher prison sentence extended, in part, based upon that

The basic purpose of an AI system is to predict outputs based upon inputs.⁹⁴ Input data—more specifically, the predictors, or “features,”⁹⁵ within data sets—are the group of independent variables that impact upon one or more AI system outputs, and that are subject to dependencies within the process.⁹⁶ This Article focuses on the three principal types of inanimate inputs in an AI system: data, statistical models, and algorithms.

Harkening to the elephant in question, this process model offers to wise ones, well-sighted in the worlds of law and policy, but less well so in engineering, computer science, cognitive science, mathematics and statistics, and technology-driven business, to see AI as a system and,

score; and impacts as the logical, if unanticipated, consequences of AI system use, such as the racial biases in some recidivism risk predictive systems that, in turn, result in disproportionate numbers of people of color being given such harsher sentences. *See, e.g.*, *State v. Loomis*, 881 N.W.2d 749, 770–71 (Wis. 2016); Katherine Freeman, Recent Development, *Algorithmic Injustice: How the Wisconsin Supreme Court Failed to Protect Due Process Rights in State v. Loomis*, 18 N.C. J.L. & TECH. 75 (2016).

The model’s third part identifies four categories of significant associations, or “dependencies,” that exist between, interoperate with, and impact upon AI inputs and outputs. *See Process Dependency Analysis Technique*, PROJECT MGMT. INST., <https://www.projectmanagement.com/process/popup.cfm?ID=23931> (last visited Dec. 8, 2021). Those categories are: (1) AI subjects, meaning principally, but not exclusively, the individual human persons who are exposed to a particular AI system and as whom the system is to operate; (2) AI market participants, those being the public and private organizations that design, develop, and deploy AI systems; (3) AI end users who operate a given AI system or otherwise procure outputs from its use; and (4) AI markets, conceived broadly to encompass the applications, use cases, industry sectors, and other market categories as to which AI systems are directed. *Id.*

94. *See* TREVOR HASTIE, ROBERT TIBSHIRANI & JEROME FRIEDMAN, *THE ELEMENTS OF STATISTICAL LEARNING: DATA MINING, INFERENCE, AND PREDICTION* 9–10 (2d ed. 2009).

95. *See* *Trs. of Columbia Univ. v. Symantec Corp.*, 811 F.3d 1359, 1366 n.3 (Fed. Cir. 2016) (appeal, in part, from patent claim construction by trial court defining “feature” as “a property or attribute of data which may take on a set of values”); *accord, e.g.*, Mohammed Osman & Edward Imwinkelried, *Facial Recognition Systems*, 50 CRIM. L. BULL. 695, 713 (2014) (likening feature extraction in facial recognition systems to identifying points of comparison in latent fingerprinting techniques).

96. HASTIE ET AL., *supra* note 94, at 9.

moreover, the human-machine enterprise that it is. It provides the sight to see that the trunk is connected to the giant head affixed to which are great fan-like ears and, in turn, to a mountainous back, tall-walled sides, tree trunk legs, and thick ropey tail.

The parable and the Jain doctrine expressed therein teach that “all viewpoints with no exception are false views when strictly related to their respective spheres . . . ; however, when understood as mutually dependent, they become viewpoints conducive to truth.”⁹⁷ Therefore, toward a truth-conducive end, this Article builds toward an interdisciplinary, that is, many-sided, mental map by which to conceptualize the whole of AI and its intersectionalities with and relevancies to the law.

Before that map becomes interpretable, however, a common and informed language of AI is needed, a set of defining terms that form cornerstones for the discussion. Next, this Article turns to the lexiconic morass in an attempt to sort and read the technical language of AI into a collective legal understanding, starting with AI, its types, and the inputs that engender it.

III. THE ELEPHANT

Some eighty years ago, Alan Turing laid down his systems of logic⁹⁸ that, in turn and with other ideas, gave birth to today’s “thinking machines.”⁹⁹ During that span of time, those in computer science and other AI disciplines have developed robust terminologies, theories, and

97. Balcerowicz, *supra* note 15, at 41 (citation and internal punctuation omitted).

98. See Alan M. Turing, *Systems of Logic Based on Ordinals*, in ALAN TURING’S SYSTEMS OF LOGIC: THE PRINCETON THESIS 31–140 (Andrew W. Appel, ed. 2012) (Turing’s 1939 doctoral thesis).

99. Alan M. Turing, *Computing Machinery and Intelligence*, 49 MIND 433, 436 (1950) (“thinking machine”). Turing’s publication was the “first serious, scholarly treatment of the concept of artificial intelligence.” MURRAY SHANAHAN, THE TECHNOLOGICAL SINGULARITY 1 n.1, 233 n.1 (2015).

understandings of their creations.¹⁰⁰ In law, however, AI-related work has only recently begun in earnest. The law has neither achieved consensus as to the meanings of “artificial intelligence” and related terms nor methodically and comprehensively catalogued, mapped, or adopted the scientific, technological, and mathematical terms of the AI arts into the legal vernacular.

The Article’s next two Sections contribute to the understanding of AI terminology and present contexts for the application of those terms and their understanding in the law. All large and arduous efforts are daunting, and this one is no less so. Never fear. There is but one answer to the question, “How do you eat an elephant?!” The response is always “One bite at a time.”

This Section presents that feast in three courses. First, it maps out the meaning of “artificial intelligence” and rationalizes the multiplicity of potentially confusing ways in which the term is used. Second, it offers some straightforward, but powerful taxonomic tools with which to categorize types of AI in legally relevant ways. Third, it drills down on “dynamic artificial intelligence,” arguably the most legally impactful type of artificial intelligence, to parse out important machine learning models. In this last part, the Article also starts helping to build neural pathways toward some exotic types of dynamic AI as may wander otherwise unrecognized into the law’s villages.

100. See ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT, *ARTIFICIAL INTELLIGENCE IN SOCIETY* ch. 1 (2019) (ebook), https://www.oecd-ilibrary.org/sites/eedfee77-en/1/2/1/index.html?itemId=/content/publication/eedfee77-en&_csp_=5c39a73676a331d76fa56f36ff0d4aca&itemIGO=oecd&itemContentType=book (last visited Nov. 7, 2021) (providing “A Short History of Artificial Intelligence”).

A. *What Is Artificial Intelligence?*

The definition of AI is varied¹⁰¹ and the subject of much thought.¹⁰² This next discussion briefly unpacks some useful levels of meaning for “artificial intelligence” like a set of Russian nesting elephants to facilitate a clearer discernment and discussion when analyzing AI.

As a discipline, AI is an interdisciplinary branch of computer science that deals with models and data processing systems for the performance, emulation, or recreation of functions that earlier have been associated with human intelligence, such as reasoning, learning, and self-improvement.¹⁰³ As a capability, the masters who dreamed up the concept of AI had ideas as to what AI is. Turing considered artificial intelligence to be the ability of digital computers to imitate humans in typed conversational exchanges to such a remarkable degree as to be indistinguishable from human conversants.¹⁰⁴ Fast forward

101. See NAT'L INST. OF STANDARDS & TECH., U.S. DEP'T OF COM., U.S. LEADERSHIP IN AI: PLAN OUTLINES PRIORITIES FOR FEDERAL AGENCY ENGAGEMENT IN AI STANDARDS DEVELOPMENT 7 (Aug. 9, 2019), https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf.

102. See, e.g., *AI Researcher: Stop Calling Everything “Artificial Intelligence.”* MIND MATTERS NEWS (Apr. 7, 2021), <https://mindmatters.ai/2021/04/ai-researcher-stop-calling-everything-artificial-intelligence/>; SOFIA SAMOILI ET AL., EUR. COMM'N JOINT RSCH. CTR., AI WATCH: DEFINING ARTIFICIAL INTELLIGENCE 7–9, 15–85 (2020), <https://publications.jrc.ec.europa.eu/repository/handle/JRC118163>; Selmer Bringsjord & Naveen Sundar Govindarajulu, *Artificial Intelligence*, STAN. ENCYCLOPEDIA OF PHIL. ARCHIVE (Edward N. Zalta ed., 2020) <https://plato.stanford.edu/archives/sum2020/entries/artificial-intelligence/> (last visited Nov. 7, 2021) (especially *What Exactly Is AI?* in Section 2).

103. See Artificial Intelligence, LEXICO, https://www.lexico.com/definition/artificial_intelligence (last visited Nov. 7, 2021); HIGH-LEVEL EXPERT GROUP ON A.I., EUR. COMM'N, A DEFINITION OF AI: MAIN CAPABILITIES AND SCIENTIFIC DISCIPLINES 7 (Dec. 18, 2018), https://ec.europa.eu/futurium/en/system/files/ged/ai_hleg_definition_of_ai_18_december_1.pdf.

104. See Turing, *Computing Machinery and Intelligence*, *supra* note 99, at 433. John McCarthy and his colleagues in the landmark 1955 Dartmouth AI summer project had similar thoughts of defining AI by its human imitative capacity. See JOHN MCCARTHY, MARVIN L. MINSKY, NATHANIEL ROCHESTER & CLAUDE E. SHANNON, A PROPOSAL FOR THE DARTMOUTH SUMMER RESEARCH PROJECT ON

and the International Organization for Standardization, known as ISO, and International Electrotechnical Commission, known as IEC, have harmonized their definition of artificial intelligence as the “capability of a functional unit to perform functions that are generally associated with human intelligence such as reasoning and learning.”¹⁰⁵

Collectively, an artificial intelligence system may be called simply an artificial intelligence or, singularly or plurally, “AI.”¹⁰⁶ An AI is a computational engine that is comprised of software, firmware, or hardware or a combination thereof; includes one or more databases or access to such; and, in very simple terms, runs on a computer.¹⁰⁷ In reality, AI are instantiated in a dazzling

ARTIFICIAL INTELLIGENCE 11 (Aug. 31, 1955) [hereinafter DARTMOUTH PROJECT], <http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf> (characterizing project focus as “making a machine behave in ways that would be called intelligent if a human were so behaving”).

105. INT’L ORG. FOR STANDARDIZATION & INT’L ELECTROTECHNICAL COMM., ISO/IEC 2382:2015, INFORMATION TECHNOLOGY – VOCABULARY (2015) [hereinafter ISO/IEC 2382:2015], <https://www.iso.org/standard/63598.html>.

106. AIs are frequently anthropomorphized and gendered as “female.” See generally Emile Loza de Siles, *AI, on the Law of Being: “Feminine” Imagery in Humanoid Robots, Evolving Law as to What Constitutes a Human* (Duquesne Univ. Sch. of L. Rsch., Paper No. 2020-12, 2020), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3658667; James Veitch, *Siri vs Alexa*, YOUTUBE (Aug. 4, 2017), https://www.youtube.com/watch?v=f_1dhKsELzs; see also Kate Darling, *Extending Legal Protection to Social Robots: The Effects of Anthropomorphism, Empathy, and Violent Behavior Towards Robotic Objects*, in ROBOT LAW 213–31 (Ryan Calo, A. Michael Froomkin & Ian Kerr eds., 2016) [hereinafter ROBOT LAW].

107. NIST has collaboratively defined AI systems “to comprise software and/or hardware that can learn to solve complex problems, make predictions or undertake tasks that require human-like sensing (such as vision, speech, and touch), perception, cognition, planning, learning, communication, or physical action.” NAT’L INST. OF STANDARDS & TECH., U.S. DEP’T OF COM., U.S. LEADERSHIP IN AI: PLAN OUTLINES PRIORITIES FOR FEDERAL AGENCY ENGAGEMENT IN AI STANDARDS DEVELOPMENT 7–8 (Aug. 9, 2019), https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf. This definition, however, focuses only on machine learning and is thus narrower than “artificial intelligence,” which also encompasses non-learning AI systems.

diversity of form factors,¹⁰⁸ across multi-core and multiple other processing configurations, and globally and even extra-globally distributed locations.¹⁰⁹ The form factors,

108. It may run on sensors inside a robot, an implanted medical device, or a self-driving vehicle, for example. *See, e.g.*, Alexandros Gazis, Evangelos Ioannou & Eleftheria Katsiri, *Examining the Sensors that Enable Self-driving Vehicles*, 39 IEEE POTENTIALS, Jan./Feb. 2020, at 46. AI may run on the tiniest microprocessors to newly huge ones. *See, e.g.*, Stacey Higginbotham, *Machine Learning on the Edge*, 57 IEEE SPECTRUM, Jan. 2020, at 20 (tiny machine learning); Samuel K. Moore, *Huge Chip Smashes Deep-Learning's Speed Barrier*, 57 IEEE SPECTRUM, Jan. 2020, at 24, 24–25, 27 (expectations that Cerebra's new dinner plate-sized chip, i.e., more than 50 times larger than any commercially available chip, will enable training of deep learning neural network systems to occur within hours, not weeks).

Form factors in which AI may be increasingly embodied include, for example, handheld microcontrollers. *See, e.g.*, Limor Fried, *Making Machine Learning Arduino Compatible: A Gaming Handheld that Runs Neural Networks*, 56 IEEE SPECTRUM, Aug. 2019, at 14. Furthermore, the component parts of an AI may be distributed, including in locations across the globe and beyond. *See, e.g.*, Mina Mitry, *Routers in Space*, 57 IEEE SPECTRUM, Feb. 2020, at 39. Finally, AI systems may be combinatorial, such as multilingual virtual “assistants” bringing together AI and cloud-mediated speech capabilities to serve those transiting German train stations and airports. *See AI Operator Standing By*, 56 IEEE SPECTRUM, Aug. 2019, at 13.

109. For example, AI may run on processors that are accessed through online cloud computing or under fog or edge computing models. *See, e.g.*, *In re Intel Corp. Sec. Litig.*, No. 18-CV-00507-YGR, 2019 WL 1427660, at *16–17 (N.D. Cal. Mar. 29, 2019) (cloud computing); NAT'L INST. OF STANDARDS & TECH. BIG DATA PUB. WORKING GRP., U.S. DEP'T OF COM., NIST BIG DATA INTEROPERABILITY FRAMEWORK: VOLUME 1, DEFINITIONS § 4.3.2, at 19 (2018) [hereinafter NIST BIG DATA DEFINITIONS], <https://doi.org/10.6028/NIST.SP.1500-1r1> (hybrid clouds); Yuan Ai, Mugen Peng & Kecheng Zhang, *Edge Computing Technologies for Internet of Things: A Primer*, 4 DIGITAL COMM'NS & NETWORKS 77, 78 (2018); *AI Operation Standing By*, *supra* note 108, at 13–20 (fog computing).

Increasingly, AI operates within cyberphysical systems, which consist of interacting physical and digital components connected by closed networks or the Internet. *See* Stefano Zanero, *Cyber-Physical Systems*, 50 COMPUTER 14, 14 (2017) (publication of IEEE COMPUT. SOC'Y); Ahmad-Reza Sadeghi, Christian Wachsmann & Michael Waidner, *Security and Privacy Challenges in Industrial Internet of Things*, in PROCEEDINGS OF THE 52ND ANNUAL DESIGN AUTOMATION CONFERENCE 1 (2015), <https://dl.acm.org/doi/pdf/10.1145/2744769.2747942>; NIST BIG DATA DEFINITIONS, *supra*, § 4.3.3, at 20; *see also, e.g.*, Emile Loza de Siles, *Google Glass: Wearable Technology for a Better Life for Persons with Autism Spectrum Disorder & Other Medical Conditions* (July 13, 2018) (Harvard eportolio) (on file with author); *Cyber Physical Systems Security*, U.S. DEP'T OF HOMELAND SEC., <https://www.dhs.gov/science-and-technology/cpssec> (last visited

configurations, and loci of an AI may present numerous jurisdictional and other legal issues.¹¹⁰

B. *Legal Taxonomies for AI Decisional Context and System Mutability*

This framework discussion introduces two organizing perspectives among the *anekāntavāda* that is AI-as-elephant. To think about and analyze AI, it is useful to categorize the subject system as follows: (1) by the decisional context of its use as an automated decision system versus an automated decision support system; and (2) the mutability of the system's operational character as comparatively static or dynamic. Note that these two ways of categorizing AI systems of categorization are not mutually exclusive. For example, an AI system under scrutiny could be used as an automated decision support tool and likewise be relatively immutable, that is, operationally static. The following sketches out helpful ways to frame the elephant as we read and think further about AI and the law.

Dec. 8, 2021).

AI operates within sensors and other devices connected via the Internet of Things (IoT) or the Industrial Internet of Things (IIoT). See NIST BIG DATA DEFINITIONS, *supra*, § 4.3.3, at 20; Emile Loza de Siles, *Cybersecurity Law & Emerging Technologies: The Federal Trade Commission, Reasonable Security Measures, and IoT*, IEEE FUTURE DIRECTIONS: TECH. POL'Y & ETHICS (May 2017), <https://cmte.ieee.org/futuredirections/tech-policy-ethics/may-2017/cybersecurity-law-and-emerging-technologies-part-1/> (text accompanying nn.29–34); Sadeghi et al., *supra*, at 1.

110. See JONAH FORCE HILL & MATTHEW NOYCE, NEW AMERICA, RETHINKING DATA, GEOGRAPHY, AND JURISDICTION 2 (Feb. 2018); see, e.g., Bernard Marr, *What Is the Artificial Intelligence of Things? When AI Meets IoT*, FORBES (DEC. 20, 2019, 12:22AM), <https://www.forbes.com/sites/bernardmarr/2019/12/20/what-is-the-artificial-intelligence-of-things-when-ai-meets-iot/?sh=107f85d7b1fd>; NAT'L INST. OF STANDARDS & TECH., *Consumer Cybersecurity Labeling for IoT Devices: A Q&A with NIST's Katerina Megas*, TAKING MEASURE BLOG (Oct. 21, 2021), <https://www.nist.gov/blogs/taking-measure/consumer-cybersecurity-labeling-iot-devices-qa-nists-katerina-megas>; Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 1004–05 (2014) (“This trend of firms initiating the interaction with the consumer will only accelerate as our thermometers, appliances, glasses, watches, and other artifacts become networked into an ‘Internet of Things.’”).

1. Decisional Use Context: Automated Decision System or Automated Decision Support System?

AI systems fall into one of two decisional use contexts: automated decision systems (ADS) or automated decision support systems (ADSS). The distinction between these two categories is in the presence and involvement of human mediation in the decision toward which the computational power of the AI system is directed. No direct human mediation exists in the former, but some degree of human mediation exists in the latter. Within the ADSS decisional context, there likely is a range of the degrees of human mediation or a sliding scale running from minimally human-mediated at one end to fully human-mediated at the other.

In the former category, ADS are used to computationally produce a result, and that result, when applied against some predefined threshold, produces the final decision or *de facto* final decision. An online credit application system is an example.¹¹¹ Briefly, the credit applicant provides the required information. Within seconds, the AI-equipped system operates, accessing the individual's credit score and history, along with the applicant-provided information and other undisclosed dark data from social media and other organizations.¹¹² The system computes and returns a result in the binary form of credit approval or disapproval. No human acts or is likely to act to mediate the computationally driven decision. The AI system is an automated decision system, and its computed result equates to the credit decision and is, respectively, final or *de facto* final.

An ADSS, by contrast, does not produce an autonomous decision, but rather computes results that are presented to human decisionmakers who, in turn, make the subject

111. See generally FRANK PASQUALE, *THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION* (2015).

112. See Janine S. Hiller, *Fairness in the Eyes of the Beholder: AI, Fairness, and Alternative Credit Scoring*, 123 W. VA. L. REV. 907, 923–24 (2021); Matthew Bruckner, *The Promise and Perils of Algorithmic Lenders' Use of Big Data*, 93 CHI.-KENT L. REV. 3, 12–17 (2018).

decision. For example, judges use recidivism or violence risk predictive systems in sentencing and other decisions.¹¹³ Because the ADSS result is mediated by humans and, presumably, their appropriate discretion, judgment, contextual understanding, and intuition,¹¹⁴ human decision-making is to be informed and facilitated by the AI system's result, but not autonomously supplanted by it, as in the case of ADS.

The distinctions between AI system autonomy and gradations of human mediation in resultant decision-making present a helpful framework for exploring AI risk and the assignment of liability and other legal responsibilities within the human-machine enterprise.¹¹⁵ Among other impacts to consider when applying this ADS-ADSS taxonomy is whether the actual use of the system aligns with its intended and designed use. For example, the use by courts of an AI system for sentencing decisions when the system was intended and designed only to inform probation decisions¹¹⁶

113. See, e.g., *State v. Loomis*, 881 N.W.2d 749, 753 (Wis. 2016) (Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) system used in sentencing decision); see also, e.g., *Henderson v. Stensberg*, No. 18-cv-555-jdp, 2021 WL 1221249, at *1 (W.D. Wis. Mar. 26, 2021) (COMPAS used by parole commission).

114. Government decision-making may overly rely on or misapply the results of ADSS. See, e.g., Erin Collins, *Punishing Risk*, 107 GEO. L.J. 57, 65–66, 69–70, 85–104 (2018) (discussing “off-label” use of AI, or use of AI system for other than its intended and designed purpose); *Loomis*, 881 N.W.2d at 769–70 (discussing use of COMPAS risk recidivism system in sentencing, and stating that “COMPAS was not developed for use at sentencing”). But see Order, *In re T.K.* (D.C. Super. Ct. Mar. 15, 2018) [hereinafter *SAVRY Order*] (Granting, in Part, Respondent's Motion to Exclude Results of the Violence Risk Assessment and all Related Testimony and/or Allocation Under FRE 702 and *Daubert v. Merrell Dow Pharmaceuticals*).

115. Compare, e.g., SAMIR CHOPRA & LAURENCE F. WHITE, A LEGAL THEORY FOR AUTONOMOUS ARTIFICIAL AGENTS 119–51 (2011), with Curtis E.A. Karnow, *The Application of Traditional Tort Theory in Embodied Machine Intelligence*, in ROBOT LAW, supra note 106, at 51–77; see also Reux Stearns et al., *Panel 2: Accountability for the Actions of Robots*, 41 SEATTLE U. L. REV. 1101, 1101–03, 1108 (2018) (remarks of Howard Jay Chizeck, debating proposed liability taxonomies where robots or human-robot collaborations result in kinetic action).

116. See Cary Coglianese & Lavi M. Ben Dor, *AI in Adjudication and*

would require adjustment along the autonomous versus human-mediated decisional spectrum. These and other considerations could further inform this decisional use context model toward helping to discern, analyze, and make informed conclusions about the AI-relevant facts and legal issues.¹¹⁷

2. Operational Character: Static or Dynamic Artificial Intelligence?

AI systems also may be classified by the mutability of their operational character as being relatively static or dynamic AI.¹¹⁸ A static AI has its operations fixed in accordance with its software code or other enabling structure. The operation of a static AI system has no internal dynamism by which the system modifies its statistical modeling or other ways of reaching its results in real time, or “on the fly,” as the vernacular goes. Although they may and should be updated, static AI systems do not “learn” as dynamic AI systems do.

The Structured Assessment of Violence Risk in Youth AI system (SAVRY) is one such static AI. The SAVRY product is an AI-based violence risk prediction tool.¹¹⁹ With SAVRY,

Administration, 11–12 n.47, BROOK. L. REV. (forthcoming), https://scholarship.law.upenn.edu/cgi/viewcontent.cgi?article=3120&context=faculty_scholarship.

117. See, e.g., Ryan Abbott, *The Reasonable Computer: Disrupting the Paradigm of Tort Liability*, 86 GEO. WASH. L. REV. 1, 16–44 (2018) (considering tort law and attribution of causation and liability after first known death in autonomous vehicle crash in 2016). See generally, e.g., Yu, *supra* note 5.

118. Some categorize AI as being “rules-based,” “coded,” or “scripted,” on one hand, meaning the subject AI system’s operation is by rote processing in accordance with the applicable software code and “data-driven” or “learning” on the other hand, meaning the system is machine learning-based. See, e.g., Casandra Laskowski, Tech. & Empirical Servs. Libr., Univ. of Ariz. James E. Rogers Coll. of L., Remarks at the Am. Ass’n of L. Schs. AI Fundamentals for Faculty Webinar (July 15, 2020), https://ncculaw.zoom.us/webinar/register/WN_UbsWdoHcTe21MLp5JvwCKQ (on-demand webinar, see especially slides 5–6).

119. See AI NOW LITIGATING ALGORITHMS, *supra* note 48, at 13 (“Studies on Violence Risk in Youth”). As to this AI, the user or the consumer of the system’s output does not have access or visibility to computation, reference data,

the user, a probation officer or psychologist, for example, interviews the subject youth to glean information with which to complete a SAVRY rating form.¹²⁰ Through that form, the SAVRY user distills and captures violence risk predictive and protective factors as being low, moderate, or high levels or present or absent.¹²¹ The user then reports out his or her estimate for the youth's risk of violent behavior.¹²² The system's underlying input data, computational models, weighting of factors, and so on are fixed and are not updated during its operation.¹²³

By contrast, a dynamic AI system may be mutable in two aspects. First, it learns its modeling and thus function through mechanisms not explicitly programmed for it in code.¹²⁴ Second, the data upon which its modeling is based or

statistical models, and algorithms. Rather, the distribution model for this AI system occurs via a manual, rating forms, and training. *See, e.g.*, JOHN S. RYALS, JR., JEFFERSON PARISH DEPT OF JUV. SERVS., 2013 SCREENING & ASSESSMENT MANUAL 13 paras. 1 & 2(a), 14 para. 3(c) (2013), <https://jefferson-parish-government.azureedge.net/documents/departments/juvenile-services/juvenile-justice-reform-publications/ScreeningAssessmentManual-2013-09-01.pdf>; *SAVRY – Structured Assessment of Violence Risk in Youth*, ANN ARBOR PUBLISHERS, https://www.annarbor.co.uk/index.php?main_page=index&cPath=416_419_189 (last visited Dec. 8, 2021).

120. *See, e.g.*, Ryals, *supra* note 119, at 13 para. 1, 29–48, apps. 1–5; *SAVRY Motion*, *supra* note 48, at 2, para. 3; Child Guidance Clinic, District of Columbia Courts, *Clinical Staff*, <https://web.archive.org/web/20210809020140/https://www.dccourts.gov/services/juvenile-matters/child-guidance> (last visited Dec. 3, 2021) (Dr. Woodland).

121. *See SAVRY Motion*, *supra* note 48, at Ex. 2 (SAVRY criteria summary sheet).

122. *See SAVRY Order*, *supra* note 114, at 3. Under this Article's taxonomy, the SAVRY system is an ADSS.

123. The SAVRY system appears to be infrequently updated, although it recently migrated to an online platform. *See* ARK. DEPT OF YOUTH SERV.'S, SAVRY 2.0 RELEASE NOTES (June 8, 2020), https://www.arcourts.gov/sites/default/files/SAVRY_2.0_DYS_Release%20Notes.pdf; *see also* Randy Borum, Patrick Bartel, & Adelle Forth, *Structured Assessment of Violence Risk in Youth*, PAR, INC. (2019) (system apparently originating 2006).

124. INT'L. ORG. FOR STANDARDIZATION, ISO 20252:2019, MARKET, OPINION AND SOCIAL RESEARCH, INCLUDING INSIGHTS AND DATA ANALYTICS – VOCABULARY AND SERVICE REQUIREMENTS § 3 [hereinafter ISO 20252:2019], <https://www.iso.org/obp/ui/#iso:std:iso:20252:ed-3:v1:en> (last visited Dec. 8, 2021) (Entry No. 3.52).

other internal aspects of a dynamic AI system may change its computational trajectory on the fly.¹²⁵ The AI system iteratively operates in a state of dynamic flux, always, for example, toward its selection of the best outcome under its data and other conditions. Machine learning systems are dynamic AI under this taxonomy, with self-programming systems being perhaps the most dynamic of all.¹²⁶

One analytical utility of this comparative taxonomy of AIs is that the proof of facts, determinations as to appropriate governance, the conceptualization of proposed theories of liability, and other legal analyses generally may be more straightforwardly applied to static AI but more complex in application to dynamic AI. In addition, the scale and rapidity of informational injury¹²⁷ and other harms that may result from these two operational types of AI may be markedly distinguishable. As a general principal, the greater the mutability of an AI system, the greater the risk of potential harms that could arise from its use. Thus, the use of dynamic AI presents a greater potential for and more rapidly produced harms. In an earlier version, this Article proposed that, this static-dynamic dichotomy could supply an appropriate risk perspective from which to prioritize the formulation and deployment of informed AI legislation in the

This is not to suggest that machine learning systems learn autonomously without human decision-making in the design and development process. *See, e.g.*, Chanin Nantasenamat, *How to Build a Machine Learning Model*, YOUTUBE (Dec. 23, 2020), <https://www.youtube.com/watch?v=NRnaMCNOK7Y>.

125. *See, e.g.*, Hiller, *supra* note 112, at 908–09, 923–24, 927–32.

126. *See, e.g.*, *Pure Predictive, Inc. v. H2O.AI, Inc.*, No. 17-cv-03049-WHO, 2017 WL 3721480, at *1–2, (N.D. Cal. Aug. 29, 2017) (describing invention claimed in “Predictive Analytics Factory” patent, see U.S. Patent No. 8,880,446 (issued Nov. 4, 2014) (subsequent history omitted); *Human Labeling*, GOOGLE CLOUD, <https://cloud.google.com/vision/automl/docs/human-labeling> (last visited Dec. 8, 2021) (Google’s AutoML self-programming artificial intelligence).

127. *See* FTC INFORMATIONAL INJURY WORKSHOP: BE AND BCP STAFF PERSPECTIVE, U.S. FED. TRADE COMM’N (Oct. 19, 2018), https://www.ftc.gov/system/files/documents/reports/ftc-informational-injury-workshop-be-bcp-staff-perspective/informational_injury_workshop_staff_report_-_oct_2018_0.pdf.

vast landscape now populated by rogue AI elephants.¹²⁸ Since-proposed AI legislation is proceeding along similar risk perspectives.¹²⁹

C. *What Is Machine Learning?*

Machine learning is a subset of AI that may be categorized as operationally dynamic under the second sorting framework posited *supra* in Section III.B.2. This Section defines and describes machine learning generally calling upon the work of leading technological organizations and institutions. The Article then illustrates how a machine learning system operates by reference to an exemplar patent. Next, it discusses how computational machines learn and, with that information onboarded, the two principal models of machine learning: supervised and unsupervised learning, models that are less cabined than their adjectives suggest. The Section closes by sketching a spectrum model as a way of thinking about and balancing the respective roles, responsibilities, and other legal touchstones where machine learning models rest within direct human control, machine control, or within the control of the blended human-machine enterprise.

Information theory, competing schools of algorithm design, and other disciplines each create deep and nuanced differences in the varied designs and instantiations of

128. The August 2020 submission draft of this Article suggested that a risk matrix model may integrate the ADS-ADSS and the static-dynamic taxonomies, being diagrammed with the lowest risk quadrant occupied by static ADSS and the highest risk quadrant by dynamic ADS.

129. See, e.g., Thomas Burri & Fredrik von Bothmer, *The New EU Legislation on Artificial Intelligence: A Primer* (Apr. 21, 2021) (unpublished manuscript at 2–3), <https://ssrn.com/abstract=3831424> (discussing bans, with some exceptions, on real-time facial recognition and other high-risk AI proposed in *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonized Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts*, COM (2021) 206 final (Apr. 21, 2021)); see also Automated Decision Systems Accountability Act, A.B. 13, Reg. Sess. § 12115(b)(1), (3) (Cal. 2021) (high-risk ADS applications).

machine learning systems.¹³⁰ As a result, there are numerous types of machine learning systems, which, in turn, may employ differing models of learning even within type.¹³¹ This Article avoids those elephant weeds for now.

Machine learning has been synonymized as “automatic learning,” which is defined as the “process by which a functional unit improves its performance by acquiring new knowledge or skills, or by reorganizing existing knowledge or skills.”¹³² Continuing, a machine learning system has the ability “to automatically learn and improve from experience,”¹³³ that experience being iterative exposures to data, new or updated training datasets, and its own results over time “without being explicitly programmed” for that learning and improvement.¹³⁴

The National Institute of Standards and Technology (NIST) describes this automatic learning capacity as a machine learning system’s ability to “distill[] meaning” through its exposure to data.¹³⁵ Through this distillation, the

130. See, e.g., Laura Martignon, *Information Theory*, in INT’L ENCYC. OF THE SOC. & BEHAV. SCIS. 7476, 7476–80 (2001); ED FINN, WHAT ALGORITHMS WANT: IMAGINATION IN THE AGE OF COMPUTING 15–56 (2018); NILS J. NILSSON, THE QUEST FOR ARTIFICIAL INTELLIGENCE: A HISTORY OF IDEAS AND ACHIEVEMENTS 398–425, 515–35 (2010); see also *infra* text accompanying notes 140, 253.

131. The field of AI is overrun with the word “learning” used with varying modifiers. See, e.g., NILSSON, *supra* note 130, at 398–425 (categories of machine learning); Mariusz Bojarski et al., *End to End Learning for Self-driving Cars*, ARXIV 2 (Apr. 26, 2016), <https://arxiv.org/pdf/1604.07316.pdf>; ISO/IEC 2382:2015, *supra* note 105 (Entry No. 2123770, genetic learning); HASTIE ET AL., *supra* note 94, at 605, 622–24 (ensemble learning); NAT’L SCI. & TECH. COUNCIL, THE NATIONAL ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT STRATEGIC PLAN 24 (2016), https://www.nitrd.gov/pubs/national_ai_rd_strategic_plan.pdf (active learning); LUCCI & KOPEC, *supra* note 59, at 300 (reinforcement learning).

132. INTL. ORG. FOR STANDARDIZATION & INTL. ELECTROTECHNICAL COMM.’N, ISO/IEC 2382:2015, INFORMATION TECHNOLOGY – VOCABULARY § 1, <https://www.iso.org/obp/ui#iso:std:iso-iec:2382:ed-1:v1:en> (last visited Dec. 8, 2021) (Entry No. 2123789).

133. ISO 20252:2019, *supra* note 124 (Entry No. 3.52).

134. *Id.*; see Communications with Shivam Rai, Data and AI Lead, Cloudreach (Aug. 11, 2020) (on file with author) (discussing training datasets).

135. NAT’L INST. OF STANDARDS & TECH., U.S. DEP’T OF COM., U.S. LEADERSHIP

machine learning system accumulates and refines its “understanding” of its decisional domain. It distills the correlations and other patterns among features within its input data and the weights to which it assigns to those patterns and features and, therefore, their significances within that domain. Using the technical vernacular, the system is trained or is self-training and thereby “learns” one or more of its functions.¹³⁶

Capacities for degrees of real or apparent autonomy are inherent in machine learning systems’ design and operation. Such autonomy arises through the systems’ iterative application of inductive bias and other decisional logics seeking to optimize the correlations underlying their computational results. Speech recognition, spam detection systems,¹³⁷ internet searches that automatically formulate and offer to complete search query commands and other predictive text systems, and financial technology, or FinTech, predictive systems¹³⁸ exemplify such machine learning systems.

1. An Exemplar Machine Learning System in Operation

Studying patents for machine learning and other artificial intelligence inventions is a useful way to get a sense of how such systems operate. For example, U.S. Patent No. 5,361,201 (the “201 patent”) claimed a machine learning

IN AI: PLAN OUTLINES PRIORITIES FOR FEDERAL AGENCY ENGAGEMENT IN AI STANDARDS DEVELOPMENT 7 (Aug. 9, 2019), https://www.nist.gov/system/files/documents/2019/08/10/ai_standards_fedengagement_plan_9aug2019.pdf.

136. See LUCCHI & KOPEC, *supra* note 59, at 300.

137. See INT’L. ORG. FOR STANDARDIZATION, ISO 19731:2017, DIGITAL ANALYTICS AND WEB ANALYSES FOR PURPOSES OF MARKET, OPINION AND SOCIAL RESEARCH – VOCABULARY AND SERVICE REQUIREMENTS, <https://www.iso.org/obp/ui/#iso:std:iso:19731:ed-1:v1:en> (last visited Dec. 8, 2021) (Entry No. 3.22); ISO 20252:2019, *supra* note 124 (Entry No. 3.52).

138. See Bruckner, *supra* note 112, at 7–11, 16; see also, e.g., Willem Van de Wiele, *European FinTech: New Rules on the Way*, 37 BANKING & FIN. SERVS. POL’Y REP., May 2018, at 16, 19, 21–23 (Financial Stability Board’s study of AI implications for financial stability).

invention using predictive modeling to automate real property appraisals.¹³⁹ As with all U.S. patents, the patent must disclose the preferred embodiment of the claimed invention.¹⁴⁰ The claims of the '201 patent disclose that the subject machine learning system best operates through the following excerpted sequence of steps:

1. Training data are collected, including in a layered collection process by which first one category of data, e.g., geographic data, are collected followed by other categories, e.g., property valuation data within a selected geographic area;¹⁴¹
2. Iterative sub steps are performed to: (a) apply live input data to a predictive model, thereby generating output data; (b) apply one or more quality or other ranking measurements to those output data; and (c) adjust, or “tune,” the model’s operation in response to the ranking results;¹⁴²
3. A predictive model is developed based upon the training data;
4. The selected predictive model is stored;
5. Live data for use in producing the desired AI system outcome are obtained with each live data record being comprised of multiple data elements, e.g.,

139. Real Est. Appraisal Using Predictive Modeling, U.S. Patent No. 5,361,201 (issued Nov. 1, 1994) [hereinafter '201 patent]; *see id.* at claims 4, 8, 13, 16, 20 (“neural network”); *see id.* at abstract, summary (“neural network” or other predictive models); Corelogic Info. Sols., Inc. v. Fiserv, Inc., No. 2:10-CV-132-RSP, 2012 WL 4355394, at *10 (E.D. Tex. Sept. 21, 2012) (construing “training data” in '201 patent).

140. *See* 35 U.S.C. § 112(a); *see also, e.g.*, Rambus Inc. v. Rea, 731 F.3d 1248, 1253 (Fed. Cir. 2013).

141. *See* '201 patent, *supra* note 139, at claim 9.2.

142. Stacey Ronaghan, *Toward Demystifying Model Training & Tuning*, TOWARDS DATA SCI. (Oct. 28, 2019), <https://towardsdatascience.com/demystifying-model-training-tuning-f4e6b46e7307>; *see* '201 patent, *supra* note 139, at claim 7. More precisely, model tuning is the optimization of the model’s hyper-parameters or “hyper-parameter optimization.” Ronaghan, *supra*.

property street address, parcel number, valuation date;¹⁴³

6. An intermediate result, or “signal,” indicative of that desired outcome is generated by applying the live data to the stored predictive model;¹⁴⁴ and
7. For each data element within the live data, a value, associated coding, and a sub-intermediate result are generated by which to quantify and denote the relative contribution of each element to the intermediate result;¹⁴⁵
8. One or more error models, which may be an error range or lower and upper percentile error values, for example, is developed from the training data;¹⁴⁶
9. The selected error model(s) is(are) stored; and
10. An error result, or, again, “signal,” in the form of the associated error model(s) is generated for the intermediate result by applying the live data to the stored error model(s).¹⁴⁷

Because patent law dictates that inventions be novel, neither the ‘201 patent nor the AI process steps outlined here are definitive for how all machine learning systems work. Nevertheless, the operative steps disclosed in this patent serve as a useful exemplar by which to illuminate those workings generally.

143. See ‘201 patent, *supra* note 139, at claim 14.

144. *E.g., id.* at claim 1.

145. See *id.* at claim 14.

146. See *id.* at claims 1, 3.

147. In addition to the sources cited interstitially, see *id.* at claims 1, 3, 7, 9, 12, 14.

2. Models of Machine Learning

A machine learning system learns a function by an iterative process of feedback and decisional enhancement.¹⁴⁸ Two common models of machine learning are supervised learning and unsupervised learning. As their monikers respectively suggest, these models involve humans or they do not, or only scantily do, involve humans in the system's learning processes.

Under supervised learning models, the “learner,” i.e., the machine learning system being trained, is exposed to paired sets of input data, or “variables,” and output data, or “outcomes.”¹⁴⁹ Generally, but not always, humans pre-interpret and label these as input and output.¹⁵⁰ Stated sparsely, these human trainers instruct the learner that “if input is A, then outcome is B.” Through this intensive process,¹⁵¹ the system through its algorithmic and statistically grounded logic “learns” the mapping function by which, for a given set of input data and features within those data, its corresponding output is to be returned.¹⁵² The

148. See LUCCHI & KOPEC, *supra* note 59, at 300.

149. See Saghar Sukla, *Regression and Classification: Supervised Machine Learning*, GEEKSFORGEEKS, <https://www.geeksforgeeks.org/regression-classification-supervised-machine-learning/> (June 1, 2021); *Types of Artificial Intelligence Algorithms You Should Know (A Complete Guide)*, UPGRAD BLOG (Nov. 13, 2019) [hereinafter *AI Algorithms Guide*], <https://www.upgrad.com/blog/types-of-artificial-intelligence-algorithms/> (variables).

150. See *Human Labeling*, GOOGLE CLOUD, <https://cloud.google.com/vision/automl/docs/human-labeling> (last visited Nov. 16, 2020); Josh Taylor, *No Labels? No Problem!*, TOWARDS DATA SCI. (Mar. 5, 2020), <https://towardsdatascience.com/no-labels-no-problem-30024984681d>.

151. For example, the Bank of New York Mellon Corporation created an artificial intelligence analytics system for its millions of legal contracts across fifty-four entities. See Roman Regelman, *How We are “Digitizing This Very Bank” at BNY Mellon*, BNY MELLON: NEWSROOM (Oct. 17, 2019), <https://www.bnymellon.com/us/en/newsroom/news/expert-voices/true-digital-artificial-intelligence-plus-human-intelligence.jsp>; Communication with Kyle Johnson, Vice President, BNY Mellon (Sept. 19, 2019) (on file with author). Some 150 attorneys worked to train the system and validate its feature extractions. See *id.*

152. See Sukla, *supra* note 149; *AI Algorithms Guide*, *supra* note 149 (variables).

machine learns to connect the two concepts, i.e., input and outcome, including by induction: the principle for and means of “establishing [] universal statements by a consideration of particular cases falling under them.”¹⁵³ Restated and contextualized, the machine learning system computes and “learns” the general rule(s) or finds patterns for correlating inputs to outcomes as the result of repeated exposures to associations between those elements.¹⁵⁴

Once a supervised machine learning system has been trained, the system’s performance is evaluated against previously unexposed input data, or so-called “testing data.” Testing of that performance should occur and be satisfactorily completed and validated prior to the live deployment of the system.¹⁵⁵

Under unsupervised models of learning, machine learning systems computationally derive, that is, “learn” or “find,” optimized correlative statistical models by being exposed to vast quantities of unlabeled or otherwise uncategorized input data.¹⁵⁶ Humans are not or are only minimally involved in these systems’ learning. As they also are with supervised learning, the resultant models are highly complex and characterized by potentially millions of features within the input data.¹⁵⁷ For example, a machine

153. JOHN PATRICK DAY, *INDUCTIVE PROBABILITY* 3 (1961) (discussing Aristotelian origins of inductive reasoning).

154. See RUSSELL & NORVIG, *supra* note 41, at 6; Marvin L. Minsky, *Proposal for Research by M. L. Minsky*, in DARTMOUTH PROJECT, *supra* note 104, at 6 (Minsky’s pioneering machine learning proposal); see also Richa Bhatia, *Understanding the Difference Between Symbolic AI & Non Symbolic AI*, ANALYTICS INDIAN MAG. (Dec. 27, 2017), <https://analyticsindiamag.com/understanding-difference-symbolic-ai-non-symbolic-ai/> (discussing these AI design theories as relevant to whether machine learning occurs pursuant to learning rules or finding patterns).

155. See Nantasenamat, *supra* note 124.

156. See Khanum et al., *supra* note 40, at 34; Ocean Tomo, LLC v. PatentRatings, LLC, 375 F. Supp. 3d 915, 956 (N.D. Ill. 2019) (summarizing expert testimony as to unsupervised machine learning).

157. See Khanum et al., *supra* note 40, at 34–35 (“massively parallel” resources and methods); Firesmith, *supra* note 40.

learning system in a DARPA-originated autonomous vehicle project was able to detect, on its own, important road outlines and other features where there was no human involvement in its training beyond a mere 100 hours of steering the vehicle on roadways.¹⁵⁸

For machine learning systems, there is an iterative or periodic process by which the system is “tuned.”¹⁵⁹ These tuning adjustments drive toward the selection of an optimized model for the subject system’s operations. The twin objectives of the tuning process are to minimize error rates while simultaneously narrowing the range of the types of error types that occur during the system’s operation.¹⁶⁰ Machine learning tools are being developed to expedite and otherwise improve the creation and operation of other machine learning systems. For example, auto-machine learning, or Auto-ML, a new subdiscipline, is focused on generating auto-tuning systems in which machine learning tools are used to evaluate thousands of potential models in tandem and then to select and optimize the selected model.¹⁶¹

As with other aspects of artificial intelligence, access to tuning information can result in contentious discovery proceedings.¹⁶² Particularly where Auto-ML is employed,

158. Bojarski, *supra* note 131, at 1. DARPA is the acronym for the Defense Advanced Research Projects Agency.

159. See MIT Lab. for Info. & Decision Sys’s, *Auto-Tuning Data Science: New Research Streamlines Machine Learning*, MIT NEWS (Dec. 19, 2017) [hereinafter MIT, *Auto-Tuning*], <http://news.mit.edu/2017/auto-tuning-data-science-new-research-streamlines-machine-learning-1219>.

160. See Brauneis & Goodman, *supra* note 67, at 119–21.

161. See, e.g., MIT, *Auto-Tuning*, *supra* note 159.

162. See, e.g., Def.’s 702 Motion to Strike Aaron DeShaw, Esq., & Mark Romano at 2–3, *Schreiner v. Allstate Fire and Casualty Insurance Co.*, No. 2014-cv-31147, 2015 WL 9901600 (D. Colo. Oct. 13, 2015) (tuning information for Colossus, Allstate’s AI system for claims evaluation and valuation); see also Dawn R. Bonnett, Note, *Use of Colossus to Measure the General Damages of a Personal Injury Claim Demonstrates Good Faith Claims Handling*, 53 CLEV. ST. L. REV. 107, 110–14 (2005) (describing Colossus system).

other challenges will be to interpret the information, even if such information were disclosed.

It is helpful to think of supervised and unsupervised models of machine learning as sitting at opposite ends of a spectrum that reflects the comparative degrees of human involvement, if any, in training these systems. Human-only mediated training that results in the system's learning sits at one end, and machine-only mediated learning, that is, "self-" training, at the other.¹⁶³

Consider a scenario in which an individual alleges reputational damage caused by a social media platform's machine learning-generated reputational score.¹⁶⁴ Among other considerations, the analysis of whether and to what extent the platform purveyor bears liability under a given tort theory should consider whether the learning model was mediated by humans and, if so, to what comparative degree vis-à-vis machine mediation. If the learning model were executed with machine-only mediation and, moreover, if the resultant model finding were unknown, uninterpretable, or unexplained, the liability analysis would require a potentially different theoretical trajectory. Agency theory, for example, provides a useful construct for evaluating liability caused by an autonomous system, such as in an unsupervised machine learning context.¹⁶⁵

163. See, e.g., Gomes et al., *supra* note 27, at 58, 60 fig. 4 (semi-supervised learning model).

164. See, e.g., Citron & Pasquale, *supra* note 49, at 24–27; Abbey Stemler, *Feedback Loop Failure: Implications for the Self-Regulation of the Sharing Economy*, 18 MINN. J.L. SCI. & TECH. 673, 712 (2017).

165. See generally CHOPRA & WHITE, *supra* note 115. "Agency" is a term of art in law and in computer science and AI. See LUCCI & KOPEC, *supra* note 59, at 300.

IV. THE ELEPHANT AS PROCESS: AI INPUTS

The Article earlier sketched out a mental map for the legal conceptualization of AI as a process-based human-machine enterprise. This Section focuses on the first part of that AI taxonomy dealing with the inanimate inputs to AI systems. As to the first AI input, it discusses data generally, some key types of data, the critical processes by which data curation occurs, and how databases and data systems are designed. This Section also discusses data as used in machine learning. It then turns to the second input: statistics, statistical models, and model-related activities. As the third input, the Article discusses algorithms, what they are, and the three most common types of machine learning algorithms.

A. *Data, Big Data, and More and Different Data*

Data are the lifeblood of AI, AI systems, and AI-mediated processes. As a foundational note, all data are backward-looking,¹⁶⁶ that is, they were captured in or synthesized from data captured in the past. Any future-looking analyses computed upon historical data are necessarily predictive and based upon probabilities, rather than causal relationships or determinative facts.

AI input data range from AI subject- or end user-supplied data to metadata to synthesized data, for example.¹⁶⁷ All these implicate corresponding data sourcing,

166. See Caryn Devins et al., *The Law and Big Data*, 27 CORNELL J.L. & PUB. POL'Y 357, 360 (2017).

167. See *Williams v. Sprint/United Mgmt. Co.*, 230 F.R.D. 640, 646 (D. Kan. 2005) (Metadata are “information about a particular data set which describes how, when and by whom it was collected, created, accessed, or modified and how it is formatted (including data demographics such as size, location, storage requirements and media information”) (quoting SEDONA CONF., THE SEDONA GUIDELINES: BEST PRACTICE GUIDELINES & COMMENTARY FOR MANAGING INFORMATION & RECORDS IN THE ELECTRONIC AGE app'x F (2005)); Don Libes, David J. Lechevalier & Sanjay Jain, *Issues in Synthetic Data Generation for Advanced Manufacturing* 1 (Dec. 11, 2017), <https://tsapps.nist.gov/publication>

preparation, and other data curation processes. Data may be collected, such as by questionnaires or online forms, for immediate and one-time algorithmic use.¹⁶⁸ They also may be sourced at scale for propagating use through what Apple CEO Tim Cook has warningly called the “data industrial complex.”¹⁶⁹

1. What Are Data?

In its simplest construction, “data” means “information.”¹⁷⁰ That synonymization, however, is insufficient to supply the context and perspectives needed for

/get_pdf.cfm?pub_id=921398 (“[D]ata analytics applications can use synthetic data to test that training algorithms perform adequately. Factories can also use the data to experiment with proposed changes.”) (2017 IEEE Big Data Conference, Boston, Mass., Dec. 11-14, 2017). Regarding synthetic data, see *infra* text accompanying notes 227–32.

168. See, e.g., *supra* text accompanying nn.119–23 (regarding SAVRY rating form).

169. Natasha Lomas, *Apple’s Tim Cook Makes Blistering Attack on the ‘Data Industrial Complex’*, TECHCRUNCH (Oct. 24, 2018), <https://techcrunch.com/2018/10/24/apples-tim-cook-makes-blistering-attack-on-the-data-industrial-complex/> (quoting Tim Cook, CEO, Apple, Inc.); see also FED. TRADE COMM’N, DATA BROKERS: A CALL FOR TRANSPARENCY AND ACCOUNTABILITY 19 (2014) [hereinafter FTC, DATA BROKERS], <https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf>; Giridhari Venkatadri et al., *Privacy Risks with Facebook’s PII-Based Targeting: Auditing a Data Broker’s Advertising Interface*, 2018 IEEE SYMP. ON SEC. AND PRIV. 89 (2018), <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8418598> (discussing data brokers’ revenue model based on aggregating information about individuals from numerous public and private sources).

170. *Pinpoint Inc. v. Amazon.com*, No. 03-CV-4954, 2004 WL 5681471, at *5 (N.D. Ill. Sept. 1, 2004); see also, e.g., *Kroll Ontrack, LLC v. Comm’r of Revenue*, 931 N.W.2d 371, 375 (Minn. 2019) (data are “detailed information of any kind”); *Keezer v. Spickard*, 493 N.W.2d 614, 617–18 (Minn. Ct. App. 1992) (requiring recordation to meet “government data” definition); *Servais v. Port of Bellingham*, 904 P.2d 1124, 1130 (Wash. 1995) (en banc) (rejecting lower court’s construction of “research data” as “scientific facts” and holding same as “body of facts and information collected for a specific purpose and derived from close, careful study, or from scholarly or scientific investigation or inquiry”); *RGIS Inventory Specialists v. Palmer*, 544 S.E.2d 79, 87 (W. Va. 2001) (tax exemption “data” definition).

AI.¹⁷¹ A better working definition for data might be representational, symbolic, or abstract information, which, once recorded, may be digitally transmitted¹⁷² or transformed into other such information.¹⁷³

Data are organized within relational or, less frequently, hierarchical databases.¹⁷⁴ That said, there is a hierarchical structure to the way data are stored in those databases.¹⁷⁵ For present purposes, a data element is the smallest informational item.¹⁷⁶ “Doe” might be the information contained in one such data element, surname. Data elements about one particular person or transaction, say, are collected into a data record.¹⁷⁷ Thus, data elements for date of birth, first name, and surname, for example, may be gathered into a data record about the identity of a person, John Doe. Continuing up the hierarchy, a dataset is a collection of personal identification records for multiple people.¹⁷⁸ This dataset aggregated or otherwise gathered together with

171. Context is key for determining the meaning of “data,” which may have “such a wide range of meanings, in different contexts, that reliance on a specific dictionary definition is not much help in answering the questions before us.” *RGIS Inventory Specialists*, 544 S.E.2d at 85.

172. *See* Philips Elecs. N. Am. Corp. v. Contec Corp., 312 F. Supp. 2d 592, 601 (D. Del. 2004).

173. *See* *RGIS Inventory Specialists*, 544 S.E.2d at 85–86; *see also* Skinner v. State, 956 S.W.2d 532, 540 (Tex. Crim. App. 1997) (employing a dictionary definition to “data” as “factual information (as measurements or statistics) used as a basis for reasoning, discussion, or calculation.” (internal citation omitted)).

174. *See generally* Lithmee, *What Is the Difference Between Relational and Hierarchical Database*, PEDIAA (Oct. 30, 2018), <https://pediaa.com/what-is-the-difference-between-relational-and-hierarchical-database/>.

175. Communication with Donald L. Simon, Professor, Mathematics & Comput. Sci. Dep’t, Duquesne Univ. (Aug. 14, 2020) (on file with author).

176. *See* NAT’L INST. OF STANDARDS & TECH., BIG DATA INTEROPERABILITY FRAMEWORK: VOLUME 2, TAXONOMIES § 3.1, at 26–27, 26 fig. 10 (2018) [hereinafter NIST BIG DATA TAXONOMIES], <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1500-2r2.pdf>; *see also* PersonalWeb Tech., LLC v. NEC Corp. of Am., No. 6:11-CV-655, 2013 WL 4015332, at *4 (E.D. Tex. Aug. 5, 2013) (*Markman* order construing “data item” as “sequence of bits”).

177. *See* NIST BIG DATA TAXONOMIES, *supra* note 176, § 3.2, at 27.

178. *See id.* § 3.3, at 28.

other datasets, for example, about the credit experiences of those people, their healthcare treatments and visits, and others.¹⁷⁹

2. What Is “Big Data”?

Data are big—astronomically big. Every year, 175 zettabytes of data are produced, consumed, and stored.¹⁸⁰ Check that number: 175 followed by 21 zeroes, i.e., 175,000,000,000,000,000,000,000 bytes of data every year.¹⁸¹ For perspective, 175 zettabytes occupies many multiples of millions of times more storage than would all the documents held by the U.S. Library of Congress,¹⁸² the largest library in the world with more than 170 million items in its collections.¹⁸³ The John Deere Company processes more data than Twitter.¹⁸⁴ The agricultural giant, one of the largest users of cloud computing in the world, gathers up to 15 million measurement readings *per second* from a global network of some 130,000 machines.¹⁸⁵ That is big data and the rocket fuel driving the percussive explosion of AI around the world.

Big data describes the vast “deluge of data in today’s networked, digitized, sensor-laden, and information-driven world.”¹⁸⁶ Many provably hard problems confound traditional analytical methods and resources. Provably hard

179. *See id.* § 3.4, at 28.

180. José M.F. Moura, *IEEE President’s Column: An IEEE for the Next Technological Revolution*, IEEE SPECTRUM: THE INSTITUTE (Dec. 2, 2019), <https://spectrum.ieee.org/ieee-presidents-column-an-ieee-for-the-next-technological-revolution> (IEEE President and CEO).

181. *See id.*

182. *See id.*

183. *See Fascinating Facts*, U.S. LIBR. OF CONG., <https://www.loc.gov/about/fascinating-facts/> (last visited Dec. 8, 2021).

184. *See* Tekla S. Perry, *John Deere’s Quest to Solve Agriculture’s Deep-Learning Problems*, 57 IEEE SPECTRUM, Feb. 2020, at 4.

185. *Id.*

186. NIST BIG DATA DEFINITIONS, *supra* note 109, § 1.1, at 2.

problems are those that are solvable in theory, but not as a practical matter because their scope and complexity outstrip computational methods and resources. Big data, along with advances in computational performance and design, represent the potential to such previously practicably unanswerable questions, such as, presciently in view of the COVID-19 crisis, “How can a potential pandemic reliably be detected early enough to intervene?”¹⁸⁷

The quantity, or volume, of big data is only one characteristic of these extensive datasets, however. An alliterative list summarizes big data’s principal characteristics: volume, variety, velocity, variability, veracity, validity, volatility, and value.¹⁸⁸ These

187. *Id.*

188. **Volume** refers to the size of the subject dataset. *See id.* § 2, at 6; *id.* § 3.2.1, at 10.

Variety refers to the characteristic of the data within the subject datasets as arising from multiple database sources and being of varying types, data structures, domains, and other characteristics. *See id.* § 2, at 6; *id.* § 3.2.3, at 10; *id.* § 3.3.1, at 11.

Velocity refers to the rates at which the data flow into and within the computational systems. *See id.* § 2, at 6; *id.* § 3.2.2, at 10. *Compare id.*, with PETE GUERRA ET AL., BOOZ ALLEN HAMILTON, THE FIELD GUIDE TO DATA SCIENCE 55 (2d ed., 2015) [hereinafter BAH, DATA SCIENCE], <http://www.boozallen.com/insights/2015/12/data-science-field-guide-second-edition> (using “data rate” as synonym for data velocity, but more broadly scoped to include rate at which data are created, gathered, and processed, data rate as significant influencer).

Variability refers to changes within the subject datasets that, in turn, impact upon the applications using those datasets. Those changes may include, for example, variability in data velocity, their format or structure, semantics, or quality. *See* NIST BIG DATA DEFINITIONS, *supra* note 109, § 2, at 6; *id.* § 3.2.4, at 10.

Veracity refers to the truthfulness, accuracy, and completeness of the data. *See id.* § 2, at 6; *id.* § 5.4.1, at 26; BAH, DATA SCIENCE, *supra*, at 94 (discussing data veracity and subjective and other means to ascertain same).

Validity means the appropriateness of the subject data for its intended purpose. *See* NIST BIG DATA DEFINITIONS, *supra* note 109, § 2, at 6; *id.* § 5.4.2, at 26. *But cf. id.* § 5.1, at 22–23 (data mining or knowledge discovery as uses beyond prospectively intended data analytics purpose).

Volatility means the degree to which the data structures tend to change over time. *See id.* § 2, at 6; *id.* § 5.4.3, at 26.

characteristics do or may have legal significance, as recognized by the courts.¹⁸⁹ Scholars likewise are increasingly investigating these characteristics of big data, and its impacts generally.¹⁹⁰

Due to some or all of its “V” characteristics and their interactions,¹⁹¹ big data far outstrips the capabilities and

Value signifies the economic, social, or other wealth represented by or resident within the subject dataset. *See id.* § 2, at 6; *id.* § 5.4.5, at 27. Note that not all of the foregoing characteristics may or must be present in all datasets that are considered “big data.” For example, data operating with the Internet of Things might be of relatively small volume, but the velocity with which the data transmission and processing must occur could nevertheless qualify the small subject dataset as big data. *See id.* § 3.2.2, at 10; *id.* § 4.3.3, at 20. Although the V-alliterative mental model is useful, other characteristics long associated with data analytics, e.g., completeness, comprehensiveness, and others, continue to be relevant in big data applications and likely also for legal inquiry. *See id.* § 5.4.7, at 28.

189. *See, e.g.*, *LSSI Data Corp. v. Comcast Phone, LLC*, 696 F.3d 1114, 1117 (11th Cir. 2012) (“This customer data has a number of uses, so it is valuable.”).

190. *See, e.g.*, Frank Pasquale & Danielle Keats Citron, *Promoting Innovation While Preventing Discrimination: Policy Goals for the Scored Society*, 89 WASH. L. REV. 1413, 1417–18 (2014); Karen Levy & Solon Barocas, *Designing Against Discrimination in Online Markets*, 32 BERKELEY TECH. L.J. 1183, 1223–28 (2018); John Frank Weaver, *Artificial Intelligence and Governing the Life Cycle of Personal Data*, 24 RICH. J.L. & TECH., no. 4, 2018, at 1, 2–18; *see also* Margaret Hu, *From the National Surveillance State to the Cybersurveillance State*, 13 ANN. REV. L. & SOC. SCI. 161, 162–63 (2017); Mary Madden et al., *Privacy, Poverty, and Big Data: A Matrix of Vulnerabilities for Poor Americans*, 95 WASH. U. L. REV. 53 (2017); Yoni Har Carmel & Tammey Harel Ben-Shahar, *Reshaping Ability Grouping Through Big Data*, 20 VAND. J. ENT. & TECH. L. 87, 109 (2017); Timothy M. Snyder, Note, *You’re Fired! A Case for Agency Moderation of Machine Data in the Employment Context*, 24 GEO. MASON L. REV. 243, 254–56 (2016); GLENN J. VOELZ, U.S. ARMY WAR COLL., *THE RISE OF iWAR: IDENTITY, INFORMATION, AND THE INDIVIDUALIZATION OF MODERN WARFARE* 89, 109–20 (2015), <https://ntrl.ntis.gov/NTRL/dashboard/searchResults/titleDetail/ADA624745.xhtml> (government military and surveillance use of AI with data from non-contact mass collection and compilation of facial images and behavioral and other biometrics). *See generally* Shlomit Yanisky-Ravid & Sean K. Hallisey, “Equality and Privacy by Design”: A New Model of Artificial Intelligence Data Transparency via Auditing, Certification, and Safe Harbor Regimes, 46 FORDHAM URB. L.J. 428 (2019) (concise helpful treatment of data and its role in AI and machine learning).

191. Compare NIST BIG DATA DEFINITIONS, *supra* note 109, § 3.1, at 8 (scalability requirements as driven by only four V-named characteristics), with *id.* § 3.1, at 9, and *id.* § 3.2, at 10 (all V-named characteristics as potential scalability requirement drivers).

capacities of traditional computational approaches.¹⁹² Instead, they necessitate scalable architectures that can efficiently and cost-effectively store, manipulate, and analyze these data.¹⁹³ It is this extensiveness and scale that distinguish big data from “small data,” the latter being datasets that are sufficiently circumscribed or presented through visualizations so that humans can understand and evaluate them.¹⁹⁴

Data elements within any dataset possess another characteristic, that being the presence and some varying degrees of complexity between those elements.¹⁹⁵ These interchanging complexities gave rise to a large international, multi-stakeholder effort convened by the NIST to establish a reference architecture for software, extensive supporting materials, and other tools to better facilitate the use and transmission of data agnostically across computing

192. See NIST BIG DATA TAXONOMIES, *supra* note 176, § 1.1, at 7.

193. See NIST BIG DATA DEFINITIONS, *supra* note 109, § 2, at 6; see also *Computer Science: Architecture & Organization*, BRITANNICA, <https://www.britannica.com/science/computer-science/Architecture-and-organization> (last visited Nov. 16, 2021); ANTHONY SNEED & MANUEL FRADINHO OLIVEIRA, NETWORKED GRAPHICS: BUILDING NETWORKS GAMES AND VIRTUAL ENVIRONMENTS 393–458 (2010) (scalability chapter); NIST BIG DATA DEFINITIONS, *supra* note 109, § 3.1, at 9; *id.* § 4.3.1, at 18 (scalability, including horizontal and vertical scalability); *id.* § 2, at 6 (latency); *id.* § 4.3.1, at 18 (high performance computing, including massively parallel processing).

194. See NIST BIG DATA DEFINITIONS, *supra* note 109, § 2, at 7 (“small data”); *id.* § 5.4.4, at 27. The visualization of complex weather data for visualization on a map is an example of small data. See The Weather Company, *Round-the-clock Accurate Weather Reports Help VTV Keep Citizens Informed*, IBM, <https://www.ibm.com/weather/industries/broadcast-media> (last visited Dec. 3, 2021) (*The Weather Company Advantage* video). “Small data” is also a shorthand term denoting small data sets, a usage not to be confused with the data visualization and human cognition concepts discussed here. See, e.g., Karen Hao, *A Radical New Technique Lets AI Learn with Practically No Data*, MIT TECH REV. (Oct. 16, 2020), <https://www.technologyreview.com/2020/10/16/1010566/ai-machine-learning-with-tiny-data/>.

195. See NIST BIG DATA DEFINITIONS, *supra* note 109, § 5.4.7, at 28; see also, e.g., Jamie Pamela Rasmussen, *Horseless Carriages with Buggy-whip Holders: The Failure of Legal Citation Reform in the 1990s*, 110 LAW LIBR. J. 221, 222, 229 (2018) (court opinions increasingly incorporating images, tables, diagrams, video files, and other non-textual content).

platforms, these being useful to understanding data in AI terms.¹⁹⁶

3. What Is Data Curation?

Big data brings with it dynamic, widely sourced, structured, and unstructured data of varying quality, formats, provenance, and dates of capture,¹⁹⁷ and repeated data changes over time.¹⁹⁸ Data curation creates order from this chaos. Without the rigor of data curation and its painstaking organizing and transformative processes, data would remain an “archipelago of information” with its great potential utility and value locked inside and difficult, if not impossible, to reach.¹⁹⁹ For these and other reasons, data curation can be distinguished, but not decoupled from the AI process.

Many AI data curation processes are human-intensive with multiple decision points throughout.²⁰⁰ Automated data curation tools, themselves machine learning-based, are also increasingly being developed and deployed to achieve greater efficiencies and economies of scale than can be achieved with human-only data curation.²⁰¹ Thus, data curation is becoming, like AI, a blended human-machine enterprise.

Views differ as to the scope of activities and processes

196. See NIST Final ‘Big Data’ Framework Will Help Make Sense of Our Data-Drenched Age, NAT’L INST. OF STANDARDS & TECH. (Oct. 29, 2019), <https://www.nist.gov/news-events/news/2019/10/nist-final-big-data-framework-will-help-make-sense-our-data-drenched-age>.

197. See A.M. Turing Award Laureate Michael Stonebraker, Lecture on Tackling the Challenges of Big Data 9–18 (2019) (on file with author) (lecture transcript from 2019 graduate data science course).

198. See Michael Dumiak, *Data Project Aims to Organize Scientific Records*, 57 IEEE SPECTRUM, Mar. 2020, at 9.

199. *Id.* at 10; MIT, *Auto-Tuning*, *supra* note 159 (complex series).

200. See Francesca Rossi, IBM A.I. Glob. Ethics Leader, Remarks at the Nat’l Inst. of Standards and Tech., Exploring AI Trustworthiness Kickoff Webinar (Aug. 6, 2020), <https://www.nist.gov/news-events/events/2020/08/exploring-ai-trustworthiness-workshop-series-kickoff-webinar>.

201. See, e.g., MIT, *Auto-Tuning*, *supra* note 159.

within the data curation wheelhouse.²⁰² This Section introduces data curation from a holistic framing reference. Minimally stated, data curation encompasses work involving the “four A’s of data”: architecture, including database design; acquisition and other pre-analysis data preparations; analysis; and archiving.²⁰³

A University of Edinburgh’s Digital Curation Centre collaboration, however, fields a more detailed data curation lifecycle model.²⁰⁴ Contextualized here for AI, the model’s sequential activities are to: (1) conceptualize the data system, including its architecture design and data modeling; (2) discover and source the data²⁰⁵ and create associated metadata²⁰⁶ and other attributional data; (3) appraise, including as to the data’s virtuousness, provenance, accuracy, and other quality and potentially compliance

202. Professor Stonebraker and colleagues consider data curation processes to encompass from data sourcing through final pre-analysis data preparations. See Michael Stonebraker et al., *Data Curation at Scale: The Data Tamer*, Remarks at the 6th Biennial Conf. on Innovative Data Sys. Rsch. 1, in Asilomar, Cal. (Jan. 6–9, 2013); see also Michelle Cheatham & Catia Pesquita, *Semantic Data Integration*, in HANDBOOK ON BIG DATA TECHNOLOGIES 263, 264 (Sherif Sakr & Albert Zomaya eds., 2017) (ebook), <https://daseelab.cs.ksu.edu/sites/default/files/semantic-data-integration.pdf>.

Others, however, use data curation to collectively refer to four groups of processes going beyond analysis and through to the point of archival, as follows: (1) work regarding database design and other data architecture; (2) data acquisition and cleaning and other processing preparatory to (3) data analysis; and (4) data archival, or “the four A’s of data,” i.e., architecture, acquisition, analysis, and archiving. JEFFREY S. SALTZ & JEFFREY M. STANTON, AN INTRODUCTION TO DATA SCIENCE 2 (2018).

203. SALTZ & STANTON, *supra* note 202, at 2.

204. See Sarah Higgins, *The DCC Curation Lifecycle Model*, 3 INT’L J. DATA CURATION 134, 135–36 (2008); see also UNIV. OF EDINBURGH DIGIT. CURATION CTR., *Curation Lifecycle Model*, <https://www.dcc.ac.uk/guidance/curation-lifecycle-model> (last visited Dec. 8, 2021) (excellent model visualization); UNIV. OF EDINBURGH DIGIT. CURATION CTR., *Introduction to Curation*, <https://www.dcc.ac.uk/guidance/briefing-papers/introduction-curation> (last visited Dec. 8, 2021).

205. See *Sandvig v. Sessions*, 315 F. Supp. 3d 1, 9 (D.D.C. 2018) (housing discrimination allegedly caused by automated algorithmic decision-making).

206. See, e.g., *Williams v. Sprint/United Mgmt. Co.*, 230 F.R.D. 640, 646–47 (D. Kan. 2005) (detailed discussion of metadata).

indicia,²⁰⁷ and select the data; (4) clean, including to deduplicate,²⁰⁸ complete, and otherwise transform the data and, from those, create new data; (5) deploy the data in distributed stores across multicore parallel processing and other computational resources for greater efficiency, privacy, and security, for example;²⁰⁹ (6) ingest the data, meaning to expose the subject AI system to and perform computations upon the data; (7) preserve the data, including as required under public records acts and for other archival purposes; and (8) access, use, and reuse data for downstream AI applications, for example.²¹⁰

These essential processes impact the quality, validity, and trustworthiness of AI systems' functioning and output and are sure to raise significant legal considerations.²¹¹

207. *See generally, e.g.*, Madden et al., *supra* note 190 (examining and warning that ubiquitous data collection and aggregation are expanding discrimination against the poor and increased use of thus-enabled predictive analytics is disproportionately eroding their privacy); Andrew Sellars, *Twenty Years of Web Scraping and the Computer Fraud and Abuse Act*, 24 B.U. J. SCI. & TECH. L. 372 (2018) (web scraping as automated data harvesting); Bruckner, *supra* note 112, at 15 (quoting ZestFinance CEO, "All data is [sic] credit data").

208. *See* Perry, *supra* note 184, at 4.

Much of this information [i.e., data processed by John Deere] is so-called *dirty data* that doesn't share the same format or structure because it's coming from some 100 other companies that have access to the John Deere platform, in addition to the wide variety of John Deere machines. Those companies add data about weather conditions, aerial imagery, soil analyses, and so on. As a result, *Deere has had to make tremendous investments in back-end data cleanup.*

Id. (emphases supplied).

209. *See* Higgins, *supra* note 204, at 136 fig. 1, 138 (data store action).

210. *See id.* at 138; Stonebraker et al., *supra* note 202, at 1.

211. *See, e.g.*, Christine L. Borgman, *Open Data, Grey Data, and Stewardship: Universities at the Privacy Frontier*, 33 BERKELEY TECH. L.J. 365, 408–09 (2018) (discussing rarity of researchers with legal expertise and short- and long-term data curation resources).

4. Mapping Data to Artificial Intelligence

Building upon this focus on data, this Section now returns to the AI-as-process model and presents some major data categories and perspectives mapping those categories to that model. To somewhat stem the flood of data terms running out of the elephant's ears, this Section offers clarity around some of those terms. First, it discusses data structure, or lack thereof, examining what are structured, unstructured, and semi-structured data. Second, it presents another view of data, exploring three categories of data as operative subjects of AI: real data, derived data, and synthetic data, with a nod also to imputed data. Third, it considers machine learning categories of testing and training data.

a. Categories of Data Structure

First, unprocessed, or uncurated, data are called “raw data.”²¹² These data may be collected from various sources and exist and are stored in various formats, but they need to be processed before they can be of analytical value.²¹³ As discussed, data curation imposes order and quality requirements upon the dataset(s) to be amalgamated and used by an AI system. In big data characteristics terms, *supra* Section IV.A.2, raw data exhibit a high degree of variety, if not also volatility and variability.

Second, data are “structured” when the information in those data is clearly organized and easily searchable in distinct fields, such as in Microsoft Excel spreadsheet cells or in database or data table fields, and when those fields have an express meaning, such as “eye color,” that is numeric, ordinal, or otherwise categorical, such as “brown.”²¹⁴

212. See NIST BIG DATA DEFINITIONS, *supra* note 109, at 6.

213. See *id.*

214. See BAH, DATA SCIENCE, *supra* note 188, at 55; Christine Taylor, *Structured vs. Unstructured Data*, DATAMATION (May 21, 2021), <https://www.datamation.com/big-data/structured-vs-unstructured-data.html>.

Unstructured data, however, are not initially organized in such a clear and distinct manner. Viewed as data in natural English language, and a Microsoft Word document format, for example, the text of this Article would be categorized as unstructured, as likewise are audio recordings, podcasts, photographs, and videos.²¹⁵ To provide a sense of comparative scale, about eighty percent of an organization's data may be unstructured.²¹⁶

The categorization of data as structured or unstructured indicate the degree and complexity of preprocessing work that must be done before pertinent features contained or represented therein those data can be identified and made available, or “extracted,” for analysis.²¹⁷ Generally, unstructured data require comparatively more preprocessing work than do structured data, and this suggests a greater potential for the introduction of errors into the data.

A third category of semi-structured data lies between the two polar ends of these structural characterizations of data. Semi-structured data are similar to unstructured data because they likewise are not organized in distinct fields in data tables. They are rendered similar to structured data, however, by categorizing keywords, or “tags” or “labels,” or other informational markers placed in association with those data by human or algorithmic annotators.²¹⁸

215. See BAH, DATA SCIENCE, *supra* note 188, at 55.

216. See Michael Chen, *Structured vs. Unstructured Data*, ORACLE: BIG DATA BLOG (Oct. 9, 2019), <https://blogs.oracle.com/bigdata/structured-vs-unstructured-data>.

217. See BAH, DATA SCIENCE, *supra* note 188, at 55.

218. “Tags” are metadata that are attached to the subject data to describe the information within those data. See INFO. SHARING & ACCESS INTERAGENCY POLY COMM. ET AL., PRIORITY OBJECTIVE 3: DATA TAGGING FUNCTIONAL REQUIREMENTS 1, 4, 7–8 (2014), <https://www.dni.gov/files/ISE/documents/DocumentLibrary/PO3-Data-Tagging-Functional-Requirements.pdf>; BAH, DATA SCIENCE, *supra* note 188, at 55; see also *Amazon Introduces a New Way to Label Data for Machine Learning with MTurk*, AMAZON: MECHANICAL TURK (Dec. 13, 2018), <https://blog.mturk.com/aws-introduces-a-new-way-to-label-data-for-machine-learning-with->

The structural category(ies) of data upon which AI operates may have legal implications. For example, if an AI system uses semi-structured or unstructured data as input, then the accuracy and other characteristics of its computational results are dependent upon the accuracy and other quality measures of the associated tag placements or other data preprocessing. The persons and processes used to tag data may raise other legal issues.²¹⁹ Constitutionality, evidentiary sufficiency, discriminatory bias, and compliance with reasonableness standards are among the other legal issues that may arise from the data used or the structures existing as to or imposed upon those data.²²⁰

b. Types of Data for Computation

This Section identifies three types of data used in AI computations: real, derived, and synthetic data. Real data means those data that are created by actual, as opposed to computational, events.²²¹ Wind turbine actuator and sensor data, analyzed by machine learning to detect faults that may signal maintenance needs, are examples of real data.²²²

Derived data are also created and used. Albeit somewhat circularly, a United Nations commission defines a derived

mturk-2f9c19866a98.

219. See, e.g., Dhruv Mehrotra, *Horror Stories from Inside Amazon's Mechanical Turk*, GIZMODO (Jan. 28, 2020, 10:00 AM), <https://gizmodo.com/horror-stories-from-inside-amazons-mechanical-turk-1840878041>; Moshe Z. Marvit, *Amazon & Mechanical Turk: How Crowdworkers (the Low-wage Virtual Labor) Became the Ghosts in the Digital Machine*, EUR. SOLIDAIRE SANS FRONTIÈRES (Feb. 4, 2014), <http://europe-solidaire.org/spip.php?article31067>.

220. See, e.g., U.K. INFO. COMM'R'S OFF., *BIG DATA, ARTIFICIAL INTELLIGENCE, MACHINE LEARNING AND DATA PROTECTION* 44 (2017), <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf> (“[E]ven when the raw data used in the analysis is [sic] recorded accurately, there may be issues as to how representative the dataset is and whether the analytics contain hidden bias.”).

221. See Libes et al., *supra* note 167, at 1. Real data also may be called “live data,” *id.*, or “actual data,” FTC, *DATA BROKERS*, *supra* note 169, at 19.

222. See, e.g., Magda Ruiz et al., *Wind Turbine Fault Detection and Classification by Means of Image Texture Analysis*, 107 *MECH. SYS'S & SIGNAL PROCESSING* 149, 149–50 (2018).

datum as a “data element derived from other data elements using a mathematical, logical, or other type of transformation, e.g.[.] arithmetic formula, composition, aggregation, etc.”²²³ By whatever means, derived data are those that are inferred based upon the real data.²²⁴ As an example, your author’s surname, i.e., the real data, is of Spanish origin. Therefore, the derived data that might be inferred from this real data is that I am Latina/X/e, which is, in my case, a correct inference. Inferred data may be factually incorrect, however, such as if I had acquired my surname by marriage and were actually of Lithuanian heritage. Inferred data’s role in AI systems is important to illuminate and question. For example, in the significantly segregated landscape that is the United States, zip codes are strongly correlated with race, and their use in AI systems may constitute illegal proxies for race in AI-mediated decision-making.²²⁵

Data also may be synthesized.²²⁶ Also called artificial, virtual, imputed, simulated, or generated data,²²⁷ synthetic data may be generated *de novo* where, for example, real data are in short supply and a larger corpus of training data is

223. See UNITED NATIONS STAT. COMM’N & ECON. COMM’N FOR EUR., TERMINOLOGY ON STATISTICAL METADATA 1, 11 (2000), <http://www.unece.org/fileadmin/DAM/stats/publications/53metadaterminology.pdf>.

224. See FTC, DATA BROKERS, *supra* note 169, at 19; Paulina Gueorguieva, *Declared or Inferred Data and What It Means For Marketers?*, ADSQUARE (Mar. 28, 2017), <https://www.adsquare.com/declared-or-inferred-data-and-what-it-means-for-marketers/> [<https://web.archive.org/web/20210414165335/https://www.adsquare.com/declared-or-inferred-data-and-what-it-means-for-marketers/>].

225. See Oliver Rollins et al., *Proxies for Race: A Catalogue*, PRICE LAB FOR DIGIT. HUMANITIES, <https://pricelab.sas.upenn.edu/projects/proxies-race-catalogue> (last visited Dec. 8, 2021).

226. See Libes et al., *supra* note 167, at 1; Cem Dilmegani, *The Ultimate Guide to Synthetic Data in 2020*, AI MULTIPLE, <https://research.aimultiple.com/synthetic-data/> (Aug. 9, 2021). In addition, data are created to fill in gaps in real data. These new synthesized data are created by imputing them from the existing data. See, e.g., HASTIE ET AL., *supra* note 94, at 332–33.

227. See Libes et al., *supra* note 167, at 1.

needed.²²⁸ One method by which this occurs is by “seeding” synthetic data from a real data sample to augment an inadequate supply of real data to, for example, use AI to help diagnose rare diseases.²²⁹ In other instances, a dataset may combine synthetic and real data.²³⁰ Some synthetic data may be produced as intermediate outputs for other AI processes or to lay down additional layers of privacy protection for the data subjects.²³¹

Legal questions exist as to these categories of data used in AI systems. For example, a patentability challenge may turn, in part, upon whether a claim in the patent disclosed the training use of real or synthetic data.²³² Evidentiary issues also likely will arise as to the appropriateness of reliance upon synthetic data. There may be claims challenging the reasonableness of inferences made to generate derived data or whether discrimination occurs where inferences may be tied to racial stereotypes. Invasions of privacy may result when marketing analytics systems

228. See Alexandre Gonfalonieri, *Do You Need Synthetic Data for Your AI Project?*, TOWARDS DATA SCI. (Oct. 21, 2019), <https://towardsdatascience.com/do-you-need-synthetic-data-for-your-ai-project-e7ecc2072d6b>; Evan Nisselson, *Deep Learning with Synthetic Data Will Democratize the Tech Industry*, TECHCRUNCH (May 11, 2018, 2:11 PM), <https://techcrunch.com/2018/05/11/deep-learning-with-synthetic-data-will-democratize-the-tech-industry/>.

229. See SIMSON L. GARFINKEL ET AL., NAT'L INST. OF STANDARDS & TECH., U.S. DEP'T COMM., DE-IDENTIFICATION OF PERSONAL INFORMATION 52 (2015), <https://nvlpubs.nist.gov/nistpubs/ir/2015/NIST.IR.8053.pdf> (glossary entry for “synthetic data generation”); Computer Security Resource Center, *Synthetic Data Generation*, NAT'L INST. OF STANDARDS & TECH., https://csrc.nist.gov/glossary/term/synthetic_data_generation (last visited Dec. 8, 2021); see, e.g., Richard J. Chen et al., *Synthetic Data in Machine Learning for Medicine and Healthcare*, 5 NATURE BIOMEDICAL ENG'G 493–97 (2021), <https://www.nature.com/articles/s41551-021-00751-8.pdf>; Watson Health Perspectives, *Bridging the Data-to-study Gap to Solve Rare Disease Research Challenges*, IBM (May 24, 2021), <https://www.ibm.com/blogs/watson-health/bridging-the-data-to-study-gap-to-solve-rare-disease-research-challenges/>.

230. See, e.g., Gonfalonieri, *supra* note 228.

231. See Steven M. Bellovin et al., *Privacy and Synthetic Datasets*, 22 STAN. TECH. L. REV. 1, 30 (2019); GARFINKEL, *supra* note 229, § 2.2, at 6–8.

232. See Mercedes-Benz USA, LLC v. Am. Vehicular Scis. LLC, No. IPR2014-00647, 2014 WL 5462676, at *5–28 (P.T.A.B. Oct. 23, 2014).

infer a teenager's not-yet-revealed pregnancy based upon her purchases.²³³

c. Data Categories in Machine Learning

The prior Section discussed types of data ingested by AI systems to produce computational results. Before machine learning systems are deployed, data may be grouped into “training” and “testing” data categories during the creation of those systems, specifically during the design and development of supervised learning systems. It is important to understand how training and testing data are assembled and used during those and other stages of AI processes.

Recall the supervised model of machine learning discussed *supra* in Section III.C.2. Under that model, the learner is exposed to labelled input-output pairs so that the correlations between the two may be discovered and its statistical model accordingly built for the system's later use in the wild. When placed into operation, the system will be exposed to previously unseen or unknown inputs and, based upon its modeling, be expected to produce results that reflect the modeled correlation. These initial input-output pairs constitute the training data.²³⁴

Once the machine learning system has been trained, its computational decision-making is tested to determine whether, when exposed to previously unseen input data, it will return the proper predictive output based upon its earlier training and the system's intended purpose.²³⁵ One landmine to avoid, and one for lawyers and their experts to detect, is the use of training data as testing data. That poor practice predictably results in a falsely more favorable

233. See Kashmir Hill, *How Target Figured Out a Teen Girl Was Pregnant Before her Father Did*, FORBES (Feb. 16, 2012, 11:02 AM), <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#45c65ad66668>.

234. See LUCCHI & KOPEC, *supra* note 59, at 298.

235. See Sukla, *supra* note 149.

estimate of the system's performance.²³⁶ Unfortunately, this may occur to the peril of the courts, other AI users, and, most perniciously and harmfully, the people who are the AI's data subjects.²³⁷

B. *Statistical Models and AI Modeling Processes*

Courts have long analyzed and opined on matters involving the rarified art of statistics.²³⁸ For example, the courts have parsed litigants' proffered patent claim constructions attempting to distinguish unsupervised machine learning from regression modeling²³⁹ and predictive modeling from statistical modeling in a supervised learning context.²⁴⁰ Disparate impact discrimination cases and criminal cases involving DNA evidence are other examples of the courts' extensive experience in splitting such elephant

236. This practice is even more suspect than a vendor's self and close third-party validations of its AI-enabled DNA software purported to be able to probabilistically predict the identities of contributors to mixed DNA samples. *See generally* State v. Pickett, No. A-4207-19T4, 2021 WL 357765 (N.J. Super. Ct. App. Div. Feb. 3, 2021) (TrueAllele Casework system).

237. *See, e.g.*, Bruckner, *supra* note 112, at 25–26.

238. *See, e.g.*, Utah v. Evans, 182 F. Supp. 2d 1165, 1175–77 (D. Utah 2001) (distinguishing statistical sampling from imputed data for census purposes), *aff'd*, 536 U.S. 452 (2002); Cooper v. Univ. of Tex., 482 F. Supp. 187, 196–98 (N.D. Tex. 1979) (analyzing statistical sampling and standard deviations under chi-square test in sex discrimination case), *aff'd*, 648 F.2d 1039 (5th Cir. 1981); *see also, e.g.*, Castaneda v. Partida, 430 U.S. 482, 494–99, (1977) (examining “statistical disparities” over time as to Mexican Americans’ representation in general population (79%) and grand juries (39%)); Brown v. Bd. of Educ., 347 U.S. 483, 494 n.11 (1954) (subsequent history omitted) (citing, in part, KENNETH B. CLARK, EFFECT OF PREJUDICE AND DISCRIMINATION ON PERSONALITY DEVELOPMENT (1950)); Kenneth B. Clark & Mamie P. Clark, *Racial Identity and Preference in Negro Children*, in READINGS IN SOC. PSYCH. 169–78 (Theodore M. Newcomb & Eugene L. Hartly eds., 1947) (reporting statistical, other quantitative, and qualitative results of dolls test used in *Brown v. Bd. of Educ.* litigation). The confidential 1950 report by Kenneth Clark is not readily available, but this 1947 study presents the statistical results of the famous dolls test.

239. *See* Ocean Tomo, LLC v. PatentRatings, LLC, 375 F. Supp. 3d 915, 956 (N.D. Ill. 2019).

240. *See* Corelogic Info. Sols., Inc. v. Fiserv, Inc., No. 2:10-CV-132-RSP, 2012 WL 4355394, at *3–6 (E.D. Tex. Sept. 21, 2012).

hairs. Legal writers have also started examining statistics in AI contexts.²⁴¹

Here, the Article describes what statistical models are and how such models relate to AI in general and machine learning in particular. It reviews the processes associated with AI statistical modeling. The Section next sketches an example of how judicial review of agencies' decisions regarding statistical modeling and AI design might require additional inquiry to justify the courts' deference to those decisions. It concludes with a foreshadowing of an even greater complexity driving the need for explainable AI.

1. What Is a Statistical Model?

The paramount and permanent feature of statistical models is that “all models are approximations. Essentially, *all models are wrong*, but some are useful. However, the approximate nature of the model must always be borne in mind.”²⁴² Because all models are wrong, lawyers “must be alert to what is importantly wrong. It is inappropriate to be concerned about mice when there are tigers abroad.”²⁴³ Fix that in an elephant's memory.²⁴⁴

In AI, statistics operate within statistical models that are used to reflect relationships between features within the input data and the output of the system. Numerous types

241. See generally, e.g., Brown, *supra* note 71 (describing, in rare detail, statistical approaches incorporated within design of AI-based document review and predictive coding systems); Coglianese & Lehr, *supra* note 67, at 1156–60; Cassandra Jones Havard, “*On the Take*”: *The Black Box of Credit Scoring and Mortgage Discrimination*, 20 B.U. PUB. INT. L.J. 241, 262–63 (2011) (touching upon statistical underpinnings of machine learning).

242. GEORGE E. P. BOX & NORMAN R. DRAPER, *EMPIRICAL MODEL-BUILDING AND RESPONSE SURFACES* 424 (1987) (emphasis supplied).

243. George E. P. Box, *Science and Statistics*, 71 J. AM. STAT. ASS'N 791, 792 (1976). Dr. Box had “one of the greatest statistical minds of the 20th century.” Ronald Wasserstein, *George Box: A Model Statistician*, 7 SIGNIFICANCE Sept. 2010 at 134, 134, <https://rss.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1740-9713.2010.00442.x>.

244. See James Ritchie, *Fact or Fiction?: Elephants Never Forget*, SCI. AM. (Jan. 12, 2009), <https://www.scientificamerican.com/article/elephants-never-forget/>.

and variants of statistical models exist, and these models are linked to and embodied in corresponding AI algorithms. For this reason, this Article discusses those linked topics of algorithms expressing statistical models in the next Section. Here, the focus is on what a statistical model is and to briefly introduce the steps of model building, determination, model selection, and model finding.

Recall this Article's operational categorization of AI as static, that is, non-learning AI, and dynamic AI, that is, generally, machine learning. Statistical models underlie both categories of AI.²⁴⁵ A statistical model, or "model," is a set of mathematical functions that approximately express and, potentially, predict the relationship(s) of interest between features within the input data and the output.²⁴⁶

Where the prediction, for example, of criminal recidivism risk is the objective for the AI's use, the statistical model is called a "predictive model" or "inferential model."²⁴⁷ Briefly, predictive models work by accounting for the uncertainty and randomness associated with the model's observations of the data features' characteristics or behavior vis-à-vis the output and then inferentially extending those observations to further postulate the model's description of that characteristic or behavior.²⁴⁸ In machine learning, the predictive model iteratively operates upon the input data and aims toward increasingly better optimization of that

245. See, e.g., Richard Cook, *Statistical Modeling and Machine Learning Coexist, Not Compete*, CMSWIRE (July 24, 2019), <https://www.cmswire.com/digital-experience/statistical-modeling-and-machine-learning-coexist-not-compete/>.

246. See HAN ET AL., *supra* note 43, § 1.5.1, at 23 (defining statistical model as collection of "mathematical functions that describe the characteristics or behavior of the objects in a target class in terms of random variables and their associated probability distributions."); HASTIE ET AL., *supra* note 94, at 27–28; *What Is Statistical Modeling?*, XLSTAT, <https://help.xlstat.com/s/article/what-is-statistical-modeling> (last visited Dec. 8, 2021).

247. See, e.g., HASTIE ET AL., *supra* note 94, at 333; HAN ET AL., *supra* note 43, § 1.5.1, at 24.

248. See HAN ET AL., *supra* note 43, § 1.5.1, at 24.

model.

2. Model Building, Determination, Selection, and Finding

Model building, determining, and selection are an integrated and iterative process to developing a probabilistic model that best describes mathematically the relationship between the independent and dependent variables within the system, these being its input and output, respectively, either of which may be qualitative or quantitative.²⁴⁹ Among the major considerations during this process are to identify the proper mathematical, or perhaps geometric, form of the relationship between the two and the selection of which independent variables to include and the weights, or indicators of significance to the task at hand, to assign to those variables.²⁵⁰ The foregoing process is called “tuning,” and model selection is its endpoint.²⁵¹ “Model finding” is a term typically used to describe tuning as carried out autonomously, or largely so, by machine learning systems.

The building, determination, and selection of a contextually optimized model occurs for all artificial intelligence systems during the final steps of system design. In carrying out this work, data scientists and others on system design teams must identify and choose, potentially from among thousands, the modeling technique that is expected to yield optimal performance in the subject system and optimal results for the particular problem or use case to which the system is directed.²⁵² The decisions made during tuning are highly deliberated.²⁵³ Depending upon the numbers and complexity of potential models to be

249. See David R. Anderson, Dennis J. Sweeney & Thomas A. Williams, *Model Building*, BRITANNICA, <https://www.britannica.com/science/statistics/Residual-analysis#ref60719> (last visited Dec. 8, 2021).

250. See *id.*

251. For a discussion of tuning, see *supra* text accompanying notes 159-62.

252. See MIT, *Auto-Tuning*, *supra* note 159.

253. See Brauneis & Goodman, *supra* note 67, at 120; Rossi, *supra* note 200.

comparatively evaluated, the decisions may themselves be facilitated and expedited by the use of machine learning tools.²⁵⁴

In such cases, transparency, the reasonableness of reliance, and other legal questions harken back to human-machine constructs discussed earlier in this Article. New legal theories or new ideas for applying existing doctrine to these unprecedented human-machine collaborations are needed. For example, where government agencies employ AI and other algorithmic means, their decisions as to the choices of statistical model and datasets upon which those models operate are subject to judicial review, but with significant deference to the agencies' discretion and particularly so for "agency modeling of complex phenomena."²⁵⁵ Even an agency's reliance upon imperfect datasets or statistical models will not necessarily result in the overturning of an agency decision as arbitrary and capricious.²⁵⁶ Given that AI use can propagate error and harm at scale, and likely irreversibly, the deference to an agency decision in the face of imperfections in the chosen data inputs, weights assigned to features within those inputs, or statistical model seems a premature end of the analysis.²⁵⁷ As seen in enacted and proposed legislation, the inclusion of an assessment of the degree and impact of those imperfections seems a more justifiable approach.²⁵⁸

254. See MIT, *Auto-Tuning*, *supra* note 159.

255. *Zirkle Fruit Co. v. U.S. Dep't of Lab.*, 442 F. Supp. 3d 1366, 1380 (E.D. Wash 2020).

256. See *id.* at 1379. Agencies act arbitrarily and capriciously when they rely upon "a report or study without ascertaining the accuracy of the data contained in the study or the methodology used to collect the data," thus rendering their findings in reliance thereupon "unsupported by substantial evidence." *Id.* (citation and internal punctuation omitted). If challenged, agencies must mount a "complete analytic defense," including explanations of the methodology and assumptions underlying their formulation and selection of statistical models. *Id.* at 1380.

257. See Coglianese & Lehr, *supra* note 67, at 1147, 1183–84.

258. See N.Y.C., N.Y., Local Law No. 49 § 3(e) (Jan. 11, 2018),

The need for new or newly applied legal theories is even greater where these model-related activities occur in the contexts of unsupervised or self-programming machine learning. There, the system designers may and likely do not have complete insight and understanding of the model or how the system thereby functions to generate its outputs.²⁵⁹ The growing impetus toward the development of explainable AI, or so-called “XAI,” is in response to the complexity and opacity of these unsupervised and self-programming machine learning scenarios.²⁶⁰ In all cases, however, the building, determination, selection, and finding of AI statistical models are potentially high-impact activities and ones requiring careful legal understanding and scrutiny.²⁶¹ Having highlighted statistical modeling and related AI processes, the Article turns to the third input in the AI-as-process model to consider algorithms and their embodiments of AI-relevant statistical models.

<https://legistar.council.nyc.gov/LegislationDetail.aspx?ID=3137815&GUID=437A6A6D-62E1-47E2-9C42-461253F9C6D0>; Algorithmic Accountability Act of 2019, S. 1108, 116th Cong. § 3 (2019); Algorithmic Accountability Act of 2019, H.R. 2231, 116th Cong. § 3 (2019) (enacted). *See generally supra* note 129 (proposed European Union AI legislation and California Assembly Bill No. 13).

259. *See, e.g.*, Stephen C. Kleene, *Representation of Events in Nerve Nets and Finite Automata*, in *AUTOMATA STUDIES* 3, 4 (Claude E. Shannon & John McCarthy eds., 1956) (“Having set up such a model, the next step is to seek a thorough understanding of the model itself.”).

260. Explainable artificial intelligence, or XAI, is a technological aspiration by which an AI system, in theory, will be able to document or demonstrate in human-understandable ways how it functioned; the inputs it used or created or both; the “reasons” that it functioned in certain ways; and how it derived its results. *See* DEF. ADVANCED RSCH. PROJECTS AGENCY, BROAD AGENCY ANNOUNCEMENT: EXPLAINABLE ARTIFICIAL INTELLIGENCE (XAI) 5–6, 9 (2016), <https://www.darpa.mil/attachments/DARPA-BAA-16-53.pdf>; Yavar Bathaee, *The Artificial Intelligence Black Box and the Failure of Intent and Causation*, 31 *HARVARD J.L. & TECH.* 889, 913–14 (2018); Weston Kowert, Note, *The Foreseeability of Human-Artificial Intelligence Interactions*, 96 *TEX. L. REV.* 181 (Nov. 2017). *But see* INDEP. HIGH-LEVEL EXPERT GROUP ON A.I., EUR. COMM’N, ETHICS GUIDELINES FOR TRUSTWORTHY AI § 2.2, at 13 (2019), <https://op.europa.eu/en/publication-detail/-/publication/d3988569-0434-11ea-8c1f-01aa75ed71a1> (stating XAI may not be achievable or operationalizable).

261. *See supra* text accompanying note 61.

C. *Algorithms in Artificial Intelligence*

Humans have created and used algorithms for millennia, dating back to those written on Babylonian cuneiform clay tablets in 1800–1600 B.C.E.²⁶² The word “algorithm” and the study of algorithms within a mathematical discipline dates to the ninth century and the writings of Muḥammad ibn Mūsā al-Khwārizmī, a Persian scholar and mathematician, credited for introducing Arabic numbers and algebra into Europe.²⁶³ In simple homage to those lofty beginnings, this Section explains what an algorithm is within an AI context and then goes on to identify and briefly describe some types of algorithms most commonly used in machine learning AI.

1. What Is an Algorithm?

The concept of an algorithm is not difficult to understand. Those of us who cook use algorithms often. Aunt Betty’s strawberry cake recipe is an algorithm with a famously delicious result. Patterned jury instructions are a series of algorithms that guide juries through the logic and nuance²⁶⁴ embedded in the law to arrive at a computational result, say, a verdict of guilty or not guilty on a racketeering

262. Christopher McFadden, *15 of the Most Important Algorithms that Helped Define Mathematics, Computing, and Physics*, INTERESTING ENG’G (Aug. 5, 2018), <https://interestingengineering.com/15-of-the-most-important-algorithms-that-helped-define-mathematics-computing-and-physics> (discussing Donald E. Knuth, *Ancient Babylonian Algorithms*, 15 COMM’NS ACM 671, 671–72 (1972), https://dl.acm.org/doi/pdf/10.1145/361454.361514?casa_token=db-UNK0py0kAAAAA:AbLZqGnlStmsBFA0rAC6KYogNE1obvL4rfqPbmZa1xAc0vnRVZIYpyexA9s4oVHj7QsFf02yn-IP).

263. See RUSSELL & NORVIG, *supra* note 41, at 8; *Al-Khwārizmī: Muslim Mathematician*, BRITANNICA, <https://www.britannica.com/biography/al-Khwarizmi> (last visited Dec. 8, 2021).

264. As to nuance and other topics, the field of computational intelligence merits mention as a future law-relevant trend related to artificial intelligence. Computational intelligence examines “how to model, govern, and engage true human behavior within” machine learning and other artificial intelligence systems. Leslie Prives, *Computationally Intelligent*, 13 IEEE WOMEN ENG’G MAG., Dec. 2019, at 6.

charge.²⁶⁵

Stated simply, an algorithm is a series of steps for accomplishing a task.²⁶⁶ Detailed instructions and rules figure into an algorithm's characteristics. The authors of the world's leading textbook on AI explain that an algorithm is a set of detailed step-by-step instructions by which to computationally analyze and solve a problem.²⁶⁷ Defined another way, an algorithm is a "finite ordered set of well-defined rules for the solution of a problem."²⁶⁸

Algorithms may be relatively simple or immensely complex with a nod of mystical appreciation to the people who design them. An algorithm may be as straightforward

265. See, e.g., MANUAL OF MODEL CRIM. JURY INSTRUCTIONS FOR THE DIST. CTS. OF THE NINTH CIR. 8.144, 8.151–.161 (NINTH CIR. JURY INSTRUCTION COMM. 2010), https://www.ce9.uscourts.gov/jury-instructions/sites/default/files/WPD/Criminal_Instructions_6_2021.pdf (associated commentary).

266. Accord DOMINGOS, *supra* note 43, at 1; see also Maura R. Grossman & Gordon V. Cormack, *The Grossman-Cormack Glossary of Technology-Assisted Review*, 7 FED. CTS. L. REV. 1, 8 (2013) (defining algorithm as "[a] formally specified series of computations that, when executed, accomplishes a particular goal.").

267. See RUSSELL & NORVIG, *supra* note 41, at 8.

268. AM. NAT'L STANDARDS INST., AMERICAN NATIONAL STANDARD DICTIONARY OF INFORMATION TECHNOLOGY: ANSDIT (2002) ("algorithm" entry at a2.htm). By "ordered," the definition means that the items within these rule sets are organized within a specified arrangement, e.g., hierarchical, such as with a data tree structure, linear, such as with a sequence. See *id.* ("order" entry at o2.htm#order); see also, e.g., Lucent Technologies, Inc. v. Microsoft Corp., 2007 WL 5734821, at *7 (S.D. Cal. Nov. 13, 2007) (*Markman* order construing algorithm as "sequence of well[-]defined mathematical operations"). But see also, e.g., N.Y.C. AUTOMATED DECISION SYS. TASK FORCE, CHECKLIST FOR DETERMINING WHETHER A TOOL OR SYSTEM IS AN ADS/AGENCY ADS AS DEFINED BY LOCAL LAW 49 (2018) 1, <https://www1.nyc.gov/assets/adstaskforce/downloads/pdf/ADS-TF-Checklist-for-Determining-ADS-Agency.pdf> (statutorily-constituted task force advising city agencies that an algorithm is "[a] set of formal or *informal* rules, processes, or instructions for carrying out a specified operation or solving a problem" (emphasis supplied)). The inclusion of "informal rules, processes, or instructions" in the task force's guidance deviates from the formality and well-defined nature of rules as expressed elsewhere and raises cautionary flags as to the transparency, discipline, accountability, and trustworthiness associated with agencies' AI-mediated decision-making.

as a flowchart²⁶⁹ or as mind-boggling as those used to calculate municipal bond markups and equity commissions²⁷⁰ or priorities-based allocations of affordable housing.²⁷¹

Some algorithms were formulated decades ago and continue in wide application in non-AI and AI contexts.²⁷² In addition, constant innovations emerge around the optimization of existing algorithms and the creation of new algorithms directed, for example, at increasingly discrete functions within AI systems.²⁷³ In addition, there are five principle schools of thought as to AI algorithm design.²⁷⁴ Algorithms arising from within the intellectual traditions of those schools may reflect alternative approaches to specific AI problems and, in turn, may have implications for legal analysis.

Further, multiple algorithms are linked in play within a

269. See *Hinlicky v. Dreyfuss*, 6 N.Y.3d 636, 639 (2006) (“flow chart, or algorithm” used by anesthesiologist to permit surgery without cardiac evaluation beforehand).

270. See *Grandon v. Merrill Lynch & Co.*, 208 F.R.D. 107, 109 (S.D.N.Y. 2002).

271. See *Allen*, *supra* note 64, at 251 n.169.

272. Khanum et al., *supra* note 40, at 34–35 (K-Means clustering algorithm).

273. See, e.g., Clint P. George, *Convolutional Neural Networks: Alternative Drivers’ Visual Perceptions*, 39 IEEE POTENTIALS, Jan./Feb. 2020, at 19, 19–20 (discussing convolutional neural networks (CNNs), a widely-used pattern recognition algorithm, applied toward creating autonomous vehicle’s machine vision system as substitute of human visual perception). See also generally Bojarski, *supra* note 131 (CNNs in autonomous vehicle project); Khanum et al., *supra* note 40, at 34. CNNs are the most commonly used algorithms for facial recognition. See Musab Coşkun et al., *Face Recognition Based on Convolutional Neural Network* 376 (Nov. 2017), <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=8248937> (2017 International Conference on Modern Electrical and Energy Systems (MEES), Kremenchuk, Ukraine, Nov. 15–17, 2017). See also generally Steve Lawrence et al., *Face Recognition: A Convolutional Neural-Network Approach*, 8 IEEE TRANSACTIONS ON NEURAL NETWORKS, Jan. 1997, at 98, http://www.cs.cmu.edu/afs/cs/user/bhiksha/WWW/courses/deeplearning/Fall.2016/pdfs/Lawrence_et_al.pdf.

274. See Pedro Domingos, *The Master Algorithm*, C-SPAN2 (Sept. 22, 2015), <https://www.c-span.org/video/?328407-1/the-master-algorithm>; DOMINGOS, *supra* note 43, at 51–55.

given AI application²⁷⁵ and within the data structures, systems, and curation processes that form the analytical framework and provide suitable input data.²⁷⁶ For example, a wayfinding AI system disclosed in a series of patents includes a patent for selecting from among several models and, in turn and based upon the selected model, from among several location-determining algorithms, depending upon the numbers of WI-FI access points detected by the larger system.²⁷⁷ The types of algorithms employed depend upon the problem domain to which the system is directed and the complexity, velocity, structure, and other characteristics of the operative data and features.²⁷⁸ Truly elephantine mysteries shroud the study of algorithms.

275. See, e.g., Brown, *supra* note 71, at 264; Tristan Greene, *A Beginner's Guide to AI: Algorithms*, THE NEXT WEB (Aug. 2, 2018, 7:42 PM), <https://thenextweb.com/artificial-intelligence/2018/08/02/a-beginners-guide-to-ai-algorithms/>; Robert Tibshirani, *Regression Shrinkage and Selection via the Lasso*, 58 J. ROYAL STAT. SOC'Y 267: SERIES B (STAT. METHODOLOGY), 273 (1996) (discussing iterative use of ridge regression algorithm for computing parameter for lasso algorithm).

276. See JAMES A. STORER, AN INTRODUCTION TO DATA STRUCTURES AND ALGORITHMS 3 (2002). See also generally, e.g., Stonebraker et al., *supra* note 202; MIT, *Auto-Tuning*, *supra* note 159.

277. See Location-based Services that Choose Location Algorithms Based on Number of Detected Access Points Within Range of User Device, U.S. Patent No. 7,305,245, (issued Dec. 4, 2007) [hereinafter '245 patent]. This '245 patent claims methods of selecting from among several location-determination algorithms that embody a simple signal strength weighted average or a nearest neighbor model; a triangulation technique; or an adaptive smoothing technique that accounts for the velocity of the location-detection device, e.g., mobile phone; or as the algorithm selection method is further refined, with routing or other inputs from the user's wayfinding application. See *id.* at claims 6–10 (discussing so-called “Chinese Postman” optimized routing algorithm model as claimed invention's preferred embodiment); Skyhook Wireless, Inc. v. Google, Inc., No. 10-11571-RWZ, 2012 WL 4076180, at *7, *9–10 (D. Mass. Sept. 14, 2012) (discussing “Chinese Postman” model); see also *id.* at *3 (stipulated construction of “simple signal strength weighted average” as an “algorithm that includes taking a simple average of the calculated locations of identified Wi-Fi access points weighted according to a function of their received signal strengths,” and construction of “triangulation technique” as an “algorithm that includes (1) estimating the distances from the user device to at least two identified Wi-Fi access points using their received signal strengths and (2) determining a location based on the estimated distances”).

278. See STORER, *supra* note 276, at 161.

2. Common Classes of Artificial Intelligence Algorithms

Remember, statistical models only approximate relationships between features within the input data to produce AI outputs. These models are expressed or embodied within corresponding algorithms. Thousands of types of algorithms are used in and are being newly directed toward AI tasks.²⁷⁹ This Section discusses three common classes of AI algorithms used in machine learning,²⁸⁰ classification, regression, and clustering algorithms.²⁸¹ It also provide examples of the application of each class. To introduce the terminology, it also names some subtypes within each class of algorithm.²⁸²

a. Classification Algorithms

Classification algorithms are used in supervised machine learning scenarios.²⁸³ Under those models, a classification algorithm-equipped AI system is “taught” how to sort input variables into outcome classes. The output target of classification algorithms AI is qualitative, such as a class or a label, or tag, associated with a particular class.

279. See, e.g., sources cited *supra* note 273.

280. Output generation is not the sole role of algorithms in AI. For example, algorithms may be applied to data curation tasks prior to the output-focused computation. See, e.g., MOHAMMED ET AL., *supra* note 21, at 138 (*k*-means cluster algorithm used to provide initial estimates for subsequent iterative computation using expectation maximization algorithm); HASTIE, ET AL., *supra* note 94, at 43 (input transformations by linear regression methods, considerably expanding their scope); Stonebraker, et al., *supra* note 202, at 1 (machine learning for data curation tasks).

281. Some algorithms may fall within more than one of these classes, depending upon the desired AI output, see, for example, MOHAMMED ET AL., *supra* note 21, at 83–84 (*k*-nearest neighbors algorithm used in classification and regression model-driven systems), and there are many more types of algorithms than discussed herein, see, for example, Tibshirani, *supra* note 275, at 274–86.

282. Each algorithmic type may be iteratively broken down into hierarchies. For example, a random forest classifier algorithm is a type of bootstrap aggregator algorithm, which is a type of decision tree algorithm, which, in turn, is a type of classification algorithm.

283. See *AI Algorithms Guide*, *supra* note 149.

Classification is a process in which a model, also called a “function,” is found that will distinguish between data classes, also called “concepts.”²⁸⁴ The purpose, therefore, of classification algorithms is to sort data into binary²⁸⁵ or multiple classes of output.²⁸⁶ To illustrate, consider a child welfare ADSS directed toward distinguishing a data class, or concept, e.g., “high child risk,” from other classes.²⁸⁷ The system uses data about the child and his or her family, medical, and educational circumstances. An AI system could use a classification algorithm to sort those collections of data and circumstances into two categories: one in which the subject child is predicted to be exposed to an unacceptable level of risk of injury or other maltreatment and should be considered for removal from the home and placement into foster care; and one in which he or she is not predicted to have that level of risk exposure and should not be considered for removal.

Numerical scores produced through the algorithm’s computation are compared, perhaps against a threshold value for the class associated with unacceptably high risk or to the class associated with acceptable risk.²⁸⁸ To be considered together with this concept of thresholds as driving class assignment are the types of errors that can occur in those classifications: false positives, or Type I, errors in which low risk data concepts are improperly assigned to the high risk class; and false negatives, or Type II, errors in which high risk data concepts are improperly assigned to the

284. HAN ET AL., *supra* note 43, § 1.4.3, at 18.

285. See Amir E. Khandani et al., *Consumer Credit-risk Models via Machine-learning Algorithms*, 34 J. BANKING & FIN. 2767, 2781 (2010).

286. See Brown, *supra* note 71, at 237–38 (multiple classes); HAN ET AL., *supra* note 43, § 1.4.3, at 18.

287. See generally, e.g., Kyle Jennison, Allegheny Cty. Dep’t of Human Svc’s Off. Analytics, Tech. & Planning, Guest Lecture in Artificial Intelligence & Social Justice Class (Apr. 12, 2021) (video on file with author) (discussing Family Screening tool); Glaberson, *supra* note 64 (discussing Eckerd Rapid Safety Feedback tool).

288. See Khandani et al., *supra* note 285, at 2781.

low risk class.²⁸⁹ Central to the legal implications that follow, a threshold for the acceptable rate of each type of error is established with the aim of balancing the system's performance as to these errors.²⁹⁰ The precision of that balance is intentional around these two types of errors and must be especially so when, for example, the individual and societal risks involved in improperly attributing a crime to an innocent person, that is, a Type II error. To evaluate the reliability of evidence, for example, from classification AI systems, lawyers need to be able understand and examine the systems' rates of both types of errors.

Some subtypes of classification algorithms that are frequently used in AI²⁹¹ include naïve Bayes,²⁹² hierarchical algorithms known as decision trees, or DTs,²⁹³ random forest,²⁹⁴ support vector machines, or SVMs,²⁹⁵ and *k*-nearest neighbours algorithms.²⁹⁶

289. *See id.*

290. *See id.*

291. *See AI Algorithms Guide, supra* note 149.

292. *See* DOMINGOS, *supra* note 43, at 151–53; MOHAMMED ET AL., *supra* note 21, at 73–82; *see also* Tony Yiu, *Understanding Bayes' Theorem: Understanding the Rationale Behind the Famous Theorem*, TOWARDS DATA SCI. (Oct. 19, 2019), <https://towardsdatascience.com/understanding-bayes-theorem-7e31b8434d4b>.

293. *See* STORER, *supra* note 275, at 127; MOHAMMED ET AL., *supra* note 21, at 37–48; Nagesh Singh, *Decision Tree Algorithm, Explained: All You Need to Know about Decision Trees and How to Build and Optimize Decision Tree Classifier* (Dec. 24, 2019), <https://www.kdnuggets.com/2020/01/decision-tree-algorithm-explained.html> *see also, e.g.*, STORER, *supra* note 275, at 132–50, 267 (discussing binary search trees); *id.* at 237–47 (2–3 trees); *id.* at 248–51 (red-black trees); *id.* at 254–57, 267 (Adelson-Velskii and Landis, or AVL, trees).

294. *See* DOMINGOS, *supra* note 43, at 238; Tony Yiu, *Understanding Random Forest: How the Algorithm Works and Why It Is So Effective*, TOWARDS DATA SCI. (June 12, 2019), <https://towardsdatascience.com/understanding-random-forest-58381e0602d2>.

295. *See* DOMINGOS, *supra* note 43, at 190–96; MOHAMMED ET AL., *supra* note 21, at 115–28.

296. *See* HASTIE ET AL., *supra* note 94, § 2.3, at 30. *But cf.* Brown, *supra* note 71, at 280 (characterizing *K*-nearest neighbor algorithm as of clustering, rather than classification, type).

b. Regression Algorithms

Regression analysis, and the algorithms that carry out these computations, is a tool by which the relationship between two or more variables, meaning input data, may be isolated and identified while “controlling for,” that is, holding constant, the effects of other variables that impact on the subject variables.²⁹⁷ The purpose of employing regression analysis is to hypothesize and test the proposed hypothesis that there is a strong, that is, statistically significant, relationship between the subject variables that may not be attributable to relationships with other variables.²⁹⁸ The mere existence of a statistically significant relationship between subject variables does not illuminate the underlying reason for that relationship, however.²⁹⁹ Importantly, the existence of such a relationship does not establish that the presence or value of one variable causes the presence or value of the other.³⁰⁰

For example, a regression analysis in a housing lending application may show a statistically significant correlation between mortgage applicants’ zip code and late mortgage payments or defaults. Such a correlation may be the basis for the decision, autonomously made by an algorithmic system or an AI-mediated decision made by humans, to deny mortgage lending to applicants living within particular zip code areas or to grant that lending but at substantially higher rates. The correlation determined by regression analysis does not signify a causal relationship between zip code and mortgage risk, however.³⁰¹ Rather, the data as to

297. CHARLES WHEELAN, *NAKED STATISTICS: STRIPPING THE DREAD FROM THE DATA* 11 (2013).

298. *Accord id.* at 12.

299. *Accord id.*

300. *Accord id.*

301. *E.g.*, *EEOC v. Sears, Roebuck & Co.*, 839 F.2d 302, 360 (7th Cir. 1988) (Cudahy, J., concurring in part, and dissenting in part) (“Regression statistics by themselves only demonstrate correlations between variables; to move from correlation to causation, there must be some independent theory about the causal

which the correlation is derived may be profoundly tainted by historical racial discrimination by government and other actors and nevertheless wrongly used to continue the deeply discriminatory “redlining”³⁰² and its enduring and propagating legacy.³⁰³

Regression algorithms are also commonly used in supervised machine learning models.³⁰⁴ Here, however, the aim is to predict a numeric computational result,³⁰⁵ rather than a qualitative classification result.

The point of regression models is to fit a curve to a set of input data points, including a curve, for example, in the form of a straight line for linear regression.³⁰⁶ That curve is fitted, including overfitted or underfitted, based upon relationship criteria in the regression model.

There are many regression models. Among the regression algorithms often used in AI³⁰⁷ are linear regression, least absolute shrinkage and selection operator, also referred to as “lasso,” regression,³⁰⁸ logistic regression, and multiple or multivariate regression algorithms.³⁰⁹ From a viewpoint 30,000 feet above the elephant, the differences between these types of regression analyses and the algorithms that express them include: (1) the geometry of the curve that explains the function, that is, the mathematical expression, of the relationship predicted to exist between the

relationships of the variables.”).

302. See, e.g., Allen, *supra* note 64, at 235–53.

303. See PASQUALE, *supra* note 111, at 23; Havard, *supra* note 241, at 247.

304. See *AI Algorithms Guide*, *supra* note 149.

305. See HAN ET AL., *supra* note 43, § 13.2.1, at 599.

306. As to linear regression, see *id.* § 1.4.3, at 19; *id.* § 13.2.2, at 90; *id.* § 3.4.5, at 105–06; HASTIE ET AL., *supra* note 94, at 43; NEILL A. WEISS, *INTRODUCTORY STATISTICS 745* (1997) (Def. No. 13.2, “regression line”).

307. See *AI Algorithms Guide*, *supra* note 149.

308. Tibshirani, *supra* note 275, at 267.

309. See, e.g., WHEELAN, *supra* note 297, at 198–204; HASTIE ET AL., *supra* note 94, at 106 (discussing multiple linear regression as yielding single output modeled as linear function of two or more inputs); WEISS, *supra* note 306, at 778.

independent and dependent variables within the subject dataset;³¹⁰ and (2) the numbers of variables that the analysis is attempting to correlate.

c. Clustering Algorithms

Cluster analysis, also known as “data segmentation,”³¹¹ is a statistical method of analysis that employs algorithms and other tools³¹² by which to organize similar objects or data items into groups, or “clusters.”³¹³ Clustering algorithms are designed to maximize the similarities between objects belonging to the same group and minimize the similarities between those objects and objects in other groups.³¹⁴ The similarity indicators are referred to as “proximate measures.”³¹⁵ Clustering may be especially useful when analyzing a large data set by grouping together into smaller meaningful groups similar, those being more relationally proximate items, as distinguished from less similar or dissimilar, those being less proximate items.³¹⁶

Humans have long used clustering as an analytical method in science and in everyday life.³¹⁷ To illustrate, the children in my family practiced clustering analysis when we played “buttons,” a game that involved emptying our *abuela’s* button jar and then sorting and grouping the buttons by

310. See, e.g., Tibshirani, *supra* note 275, at 270–72, Fig’s 1–4.

311. HASTIE ET AL., *supra* note 94, at 501.

312. Jorge Bacallao Gallestey, *Cluster Analysis*, BRITANNICA., <https://www.britannica.com/topic/cluster-analysis> (last visited Dec. 8, 2021).

313. Shimon Ullman et al., Ctr. for Brains, Minds, and Machines, MIT, *Unsupervised Learning: Clustering 2* (2014), <http://www.mit.edu/~9.54/fall14/slides/Class13.pdf>.

314. See Gallestey, *supra* note 312; HASTIE, ET AL., *supra* note 94, at 501–03 (cluster analysis).

315. Ullman et al., *supra* note 313, at 6.

316. See *id.* at 3; Gallestey, *supra* note 312.

317. See, e.g., JOHN SNOW, ON THE MODE OF COMMUNICATION OF CHOLERA 12–26 (1849), <https://collections.nlm.nih.gov/ext/cholera/PDF/0050707.pdf> (famous historical clustering analyses linking cholera cases to contaminated wells).

color, size, and other characteristics. Certain characteristics of the input data may be weighted differently, rather than all uniformly, to produce clusters that may be more suitable for various reasons.³¹⁸ In the buttons game, fabric-covered buttons may have been considered more desirable than bone, metal, or plastic buttons and, therefore, weighted more heavily in terms of their desirability in the game.

Unsupervised machine learning uses cluster analysis and various types of clustering algorithms.³¹⁹ Fraud detection systems designed to protect consumers' debit and credit cards are an example of a clustering algorithm-based AI systems. A subject consumer's payment card transactions may be grouped together based upon the general geographical location in which her transactions are usually made, for example, in the Washington, D.C. area. A clustering algorithm-based AI system may flag transactions coming from, say, Barbados, as potentially fraudulent and the transaction may be declined, pending confirmation by the consumer.

Among clustering algorithms commonly in AI use³²⁰ are *k*-means clustering,³²¹ fuzzy C-means,³²² expectation maximization, or EM,³²³ hierarchical clustering,³²⁴ and hidden Markov model algorithms.³²⁵

318. See HASTIE ET AL., *supra* note 94, at 504.

319. See Harry Surden, *Machine Learning and Law*, 89 WASH. L. REV. 87, 113 (2014); Khanum et al., *supra* note 40; *AI Algorithms Guide*, *supra* note 149.

320. See *AI Algorithms Guide*, *supra* note 149; MOHAMMED ET AL., *supra* note 21, at 17, 145–48 (hidden Markov model).

321. See MOHAMMED ET AL., *supra* note 21, at 31–36.

322. See HAN ET AL., *supra* note 43, § 11.1.1, at 499–501.

323. See MOHAMMED ET AL., *supra* note 21, at 138.

324. See HAN ET AL., *supra* note 43, § 10.1.3, at 449; *id.* § 10.3, at 457–59.

325. See MOHAMMED ET AL., *supra* note 21, at 145–48.

V. CONCLUDING FORWARD

This Article has lent its efforts to providing readers with a single concise, accessible, but comprehensive source for better understanding the complex, fascinating, and potentially terrifying elephant of artificial intelligence. A majestic mystery, AI has entered crashing into legal domains and, if not insightfully governed, threatens people, civil society, and humanity in an unbridled, market-driven frenzy.

By pondering these contributions and continuing their respective journeys to learn more, lawyers throughout the profession will be more ethically competent, intellectually rigorous, and authoritative in their profoundly critical and urgent work regarding artificial intelligence. These gains will work to improve the governance, security, justness, and well-being of the increasingly algorithmic world.