I. MACHINE LEARNING AND ITS ADMINISTRATIVE APPLICATIONS

Some of the most prominent examples of private-sector growth today—from Amazon[27] to Zillow[28]—depend on the use of machine learning to optimize production processes, supply chains, marketing, and the pricing of goods and services. Machine learning undergirds future growth across a wide range of sectors, from the introduction of "fintech" firms in the financial industry[29] to advances in healthcare delivery via precision medicine.[30] Machine learning's main attraction stems from how well it "outperforms human intelligence."[31] As

26. *Cf.* Adrian Vermeule, *Conventions of Agency Independence*, 113 COLUM. L. REV. 1163 (2013) (describing conventions and norms surrounding the administrative state, particularly agency independence).

27. *See* Metz, *supra* note 4 (describing Amazon's use of machine learning to recommend products to consumers).

28. *See* Jessica Davis, *Zillow Uses Analytics, Machine Learning to Disrupt with Data*, INFORMATION-WEEK (Oct. 14, 2016, 11:06 AM), http://www.informationweek.com/big-data/zillow-uses-analytics-machine-learning-to-disrupt-with-data/d/d-id/1327175 [https://perma.cc/8UTD-QX6S] (describing Zillow's use of machine learning to provide forecasts of housing prices).

29. *See* Falguni Desai, *The Age of Artificial Intelligence in Fintech*, FORBES (June 30, 2016, 10:42 PM), http://www.forbes.com/sites/falgunidesai/2016/06/30/the-age-of-artificial-intelligence-in-fintech [https://perma.cc/EK89-DD3Y] (describing how fintech firms use artificial intelligence to improve investment strategies and analyze consumer financial activity).

30. For a discussion of these and other examples of an increasingly optimizing economy, see Coglianese, *supra* note 22.

31. NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES 11 (2014). Of note, machine learning recently bested human intelligence in the incredibly complex game of Go. *See* David Silver et al., *Mastering the Game of Go with Deep Neural Networks and Tree Search*, 529 NATURE 484, 484 (2016).

private firms pursue significant efficiency gains through the kind of smarter and more contextualized decisions made possible by algorithmic analysis of big data, the government will undoubtedly need to follow suit, not merely to keep up with new risks these private-sector uses of machine learning might bring, but also to improve government's ability to address a host of existing risks and regulatory problems.[32] Machine learning promises to make the government, like the private sector, smarter and more efficient. In this Part, we introduce machine learning and discuss how government agencies are already beginning to explore its use to optimize administrative tasks, an endeavor that is likely to grow both in size and scope in the years ahead. We first explain what machine learning is and describe its distinguishing features. We then discuss how agencies are already using machine learning. Finally, we show how this technology could, in the future, potentially transform the administrative state through what we call "rulemaking by robot" and "adjudicating by algorithm."

A. WHAT IS MACHINE LEARNING?

Fundamentally, machine-learning algorithms are used to make predictions. This emphasis on *prediction* contrasts markedly with traditional statistical techniques which seek to *model* underlying data-generating processes in the real world. Although traditional statistical techniques can also generate predictions, they do so only when the model created by the analyst fits well with the underlying processes being modeled. These traditional techniques require the analyst first to specify a mathematical equation expressing an outcome variable as a function of selected explanatory variables put together in a particular way, and then to see how well the data fit with the analyst's choices. For example, when analysts employ the traditional techniques of ordinary least squares regression or logistic regression, they specify equations that represent their a priori beliefs about the functional relationships that exist between *independent* (or explanatory) and *dependent* (or outcome) variables. What regression does, in essence, is estimate the magnitude and direction of these relationships between the two types of variables that are selected and specified by the analyst. The relationships in the statistical model ostensibly represent the relationships in the real world, which is why regression results are often used to support causal inferences.

By contrast, machine learning is nonparametric in that it does not require the researcher to specify any particular functional form of a mathematical model in advance. Instead, these algorithms allow the data themselves to dictate how information contained in *input* variables is put together to forecast the value of

32. *See* Coglianese, *supra* note 22. A senior strategist at the National Security Agency has noted how advanced analytic techniques are particularly essential for agencies that will soon be forced "to operate at cyberspeed and at scale." GOV'T BUS. COUNCIL, GOV'T EXEC. MEDIA GRP., DATA ANALYTICS: A STRATEGIC ASSET TO GOVERNMENT 2 (2015), http://cdn.govexec.com/media/gbc/docs/2015-09-03_qlik_issue_brief_designed_4.pdf [https://perma.cc/LY76-FY5V].

an *output* variable.[33] Machine-learning algorithms do not generate quite the same kind of information on the magnitude or direction of the effects that might be associated with any single input variable on the output variable, controlling for the other variables. The functional relationships in machine learning are not necessarily the complete set of those in nature's true data-generating process. As a result, no claim can be made that the machine-learning process represents any set of true relationships in the world, and thus none of the causal inferences that typically characterize statistical modeling can be applied to results of machine learning. In short, with machine-learning results, causal relationships between inputs and outputs may simply not exist, no matter how intuitive such relationships might look on the surface. If a machine-learning algorithm tends to forecast that older individuals commit fewer crimes than younger individuals, for example, it cannot be claimed on the basis of the machine-learning process that older age causes any reduction in the propensity to commit crimes.[34]

Nevertheless, from a technical standpoint, machine learning's distinctive predictive and nonparametric focus turns out to be paramount to its impressive usefulness in generating reliable forecasts. Also of central importance, and what gives machine learning its name, is how such algorithms mathematically "learn" to generate their predictions. There are many machine-learning algorithms that do so in different mathematical ways, but they all attempt, as one textbook explains, to "optimize a performance criterion using example data or past experience."[35] In other words, these algorithms make repeated passes through data sets, progressively modifying or averaging their predictions to optimize specified criteria.

To illustrate this functioning, consider a common application of machine learning that has proven critical to improvements in the government's handling of postal mail and other paperwork-processing tasks—the recognition and classification of handwritten digits.[36] In this simple application, an algorithm's performance criterion, or objective function, is classification accuracy—that is, how often it correctly recognizes, say, a handwritten number two as a two. To

33. *See* RICHARD A. BERK, STATISTICAL LEARNING FROM A REGRESSION PERSPECTIVE 13 (2008).

34. For a discussion of the inferential value of outputs from machine-learning algorithms, see *id.* at 9–17.

35. ETHEM ALPAYDIN, INTRODUCTION TO MACHINE LEARNING 3 (2d ed. 2010). In the broader field, the varied types of machine learning are referred to by a dizzying array of different terms, some technical, some colloquial, for example: smart machines, expert systems, neural networks, deep learning, hierarchical learning, reinforcement learning, structured learning, and more. Although we explain some of these different terms in this Part, for the most part throughout this Article we use the terms "machine learning," "algorithms," and "artificial intelligence" for convenience to capture all possible variations in terms, as we are concerned primarily with the legal issues surrounding the general use of this family of techniques.

36. *See, e.g.*, Cheng-Lin Liu et al., *Handwritten Digit Recognition: Benchmarking of State-of-the-Art Techniques*, 36 PATTERN RECOGNITION 2271 (2003); Y. LeCun et al., Comparison of Learning Algorithms for Handwritten Digit Recognition, Presented at the International Conference on Artificial Neural Networks (1995), http://yann.lecun.com/exdb/publis/pdf/lecun-95b.pdf [https://perma.cc/2NGQ-85FS].

perform this classification, an algorithm must "learn" what aspects of a handwritten digit make it likely to be a two. Over the course of iterative passes through the data, such an algorithm tries to use many different mathematical descriptions of shapes, as well as relationships of shapes, in the pictures of handwritten digits to make its classifications. If a particular descriptive method is optimal, the algorithm will be "rewarded" with a low error rate; if the descriptions are not optimal, the algorithm will be "punished" with a high error rate. It can learn, for example, that a handwritten digit is likely to be a two if the topmost section of the digit depicted is semicircular and facing downward. Ultimately, the algorithm will seek to make classifications based on mathematical descriptions of shapes that yield the lowest error rates.[37]

This handwriting recognition example provides an illustration of machine-learning algorithms applied to classification problems, where the goal is to sort objects into classes. But classification problems represent only some of the diverse applications of machine-learning techniques. Machine-learning algorithms can also be used to predict numerical values, such as house prices or stock market index values—endeavors that are often termed regression problems.[38] They also can be applied to scenarios, such as playing chess, where an algorithm can be used to determine the optimal sequence of actions.[39] Variety in the types of machine-learning algorithms means that they can be used in a wide variety of predictive endeavors.

Admittedly, many "non-learning" techniques have long been used to pursue these same endeavors. For example, ordinary least squares regression can estimate numerical outcomes, and logistic regression is commonly used as a binary classifier. Given the existence of these alternative statistical techniques, what advantages do machine-learning algorithms offer? Put simply, they outperform standard procedures in terms of predictive accuracy and statistical efficiency (that is, the increased ability to obtain predictions with both low bias and low variance).[40] Furthermore, many phenomena that analysts want to forecast

---

37. This is an example of supervised machine learning, where each handwritten digit is labeled with its correct digit so that the algorithm knows when it is making errors. Unsupervised learning uses unlabeled data, so that performance criteria being optimized are not measures of error rates, because the truth is not known, but measures of similarity between digits determined by the algorithm to be the same. *See* ALPAYDIN, *supra* note 35, at 11–13.

38. *See id.* at 9–11.

39. *See id.* at 13–14; *see also* Cade Metz, *In a Huge Breakthrough, Google's AI Beats a Top Player at the Game of Go*, WIRED (Jan. 27, 2016, 1:00 PM), http://www.wired.com/2016/01/in-a-huge-breakthrough-googles-ai-beats-a-top-player-at-the-game-of-go/ [https://perma.cc/9YQ9-LHM4].

40. These benefits result from the mathematical techniques of boosting or bagging (or both). For demonstrations of these benefits resulting from boosting, see Robert E. Schapire, The Boosting Approach to Machine Learning: An Overview, Presented at the MSRI Workshop on Nonlinear Estimation and Classification (2002), https://www.cs.princeton.edu/picasso/mats/schapire02boosting_schapire.pdf [https://perma.cc/64PZ-PY57]. For an illustration of similar benefits resulting from bagging, see Leo Breiman, *Some Infinity Theory for Predictor Ensembles* (U.C. Berkeley Technical Report 577, 2000), https://www.stat.berkeley.edu/breiman/some_theory2000.pdf [https://perma.cc/KZ3D-RWMY]; Peter Bühlmann & Bin Yu, Explaining Bagging, Presented at the Seminar für Statistik (2000),

are extraordinarily complex, and analysts often lack the a priori knowledge necessary to specify an accurately forecasting conventional model. By eschewing this dependency on existing knowledge and the need to identify the functional form of any relationships, machine learning can apply to a wider range of problems and yield vastly enhanced accuracy over its alternatives, whether human intuition, expert judgment, or traditional statistical techniques.[41] Learning algorithms can also adapt more dynamically; as new data become available, they can search for new patterns and thereby improve forecasting accuracy.

Although machine-learning algorithms are known and prized for their accuracy, this benefit does come at an interpretive cost. This cost is frequently invoked by references to machine-learning algorithms as "black-box" procedures.[42] The black-box nature of machine learning holds important implications for administrative law, so to understand this feature of machine learning consider again the classification of handwritten digits. We said that an algorithm might learn that certain geometric characteristics of the shapes of handwritten digits are useful for determining which digits they represent—yet we cannot really know what precise characteristics any machine-learning algorithm is keying in on. Machine-learning algorithms transform a series of inputs to a series of outputs by optimizing a performance criterion, but that is where the analyst's easy ability to interpret the algorithms' workings comes to an end. The user of an algorithm cannot really discern which particular relationships between variables factor into the algorithm's classification, or at which point in the algorithm they do, nor can the user determine how exactly the algorithm puts together various relationships to yield its classifications.[43] For this reason, machine-learning algorithms are often described as transforming inputs to outputs through a black box. An analyst cannot look inside the black box to understand how that transformation occurs or describe the relationships with the same intuitive and causal language often applied to traditional statistical

---

http://e-collection.library.ethz.ch/eserv/eth:23905/eth-23905-01.pdf [https://perma.cc/CPE2-QEAF]. For a demonstration of the benefits resulting from combined classifiers achieved through both bagging and boosting, see V. Koltchinskii & D. Panchenko, *Empirical Margin Distributions and Bounding the Generalization Error of Combined Classifiers*, 30 ANNALS STAT. 1 (2002).

41. *See, e.g.*, Volodymyr Mnih et al., *Human-Level Control Through Deep Reinforcement Learning*, 518 NATURE 529, 529 (2015); Kaiming He et al., Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification, Presented at the IEEE International Conference on Computer Vision (2015), http://arxiv.org/abs/1502.01852 [https://perma.cc/5ERE-L2HK]; Bo Pang et al., Thumbs Up? Sentiment Classification Using Machine Learning Techniques, Presented at the Conference on Empirical Methods in Natural Language Processing (2002), http://www.cs.cornell.edu/home/llee/papers/sentiment.pdf [https://perma.cc/G4VS-XYM2].

42. *See, e.g.*, Leo Breiman, *Statistical Modeling: The Two Cultures*, 16 STAT. SCI. 199, 199 (2001).

43. For example, a common machine-learning algorithm known as random forests generates its predictions by, roughly speaking, producing thousands of decision trees (called classification or regression trees) and then averaging predictions across all trees. *See* Leo Breiman, *Random Forests*, 45 MACHINE LEARNING 5, 5 (2001). The analyst can examine the structure of a particular tree and determine to some extent how variables and interactions between variables functionally affect the predictions, but this will tell the analyst nothing about all such processes in the forest as a whole.

modeling.[44]

Despite this interpretive limitation, machine-learning algorithms have been implemented widely in private-sector settings. Companies desire the savings in costs and efficiency gleaned from these techniques, and the lack of intuitive interpretability is of little concern in endeavors where accuracy, not causality, is the valued metric. Netflix, for instance, employs a form of machine learning called "artificial neural networks" to suggest entertainment options to its customers based on their prior viewing habits.[45] Google uses machine learning to identify house numbers in its Street View imagery,[46] to save energy in its data centers,[47] and to keep its self-driving cars from crashing.[48] Machine learning has also shown great utility in the financial sector, where it is employed to predict the value of investments and financial instruments.[49] The benefits of learning algorithms have also promoted their adoption in academic research in disciplines closely connected to policymaking, where predictive accuracy is critical. For example, researchers have shown that machine-learning algorithms can help predict the propensity of probationers and parolees to commit violent crimes,[50] estimate population densities of homeless persons in cities,[51] and forecast student retention at universities.[52] In these ways, both private businesses and academic researchers have embraced machine learning, and machine-learning applications in a wide variety of settings are already actively shaping society.

B. EXISTING ADMINISTRATIVE APPLICATIONS

For much the same reason that machine learning has been exploited in the private sector, its use holds potentially great value to government agencies. We

44. *See* Breiman, *supra* note 42, at 199–200.

45. Alex Chen et al., *Distributed Neural Networks with GPUs in the AWS Cloud*, NETFLIX TECH BLOG (Feb. 10, 2014), http://techblog.netflix.com/2014/02/distributed-neural-networks-with-gpus.html [https://perma.cc/T7CR-7QGM].

46. Ian J. Goodfellow et al., *Multi-Digit Number Recognition from Street View Imagery Using Deep Convolutional Neural Networks*, CORNELL UNIV. LIB. COMP. VISION & PATTERN RECOGNITION (2013), http://arxiv.org/abs/1312.6082 [https://perma.cc/H7KK-DZJT].

47. Joe Kava, *Better Data Centers Through Machine Learning*, GOOGLE OFFICIAL BLOG (May 28, 2014), http://googleblog.blogspot.com/2014/05/better-data-centers-through-machine.html [https://perma.cc/ZEP4-2Z8M].

48. *See* Madrigal, *supra* note 5.

49. *See* Quentin Hardy, *Wealth Managers Enlist Spy Tools to Map Portfolios*, N.Y. TIMES, (Aug. 3, 2014), http://www.nytimes.com/2014/08/04/technology/wealth-managers-enlist-spy-tools-to-map-portfolios.html [https://perma.cc/K4F4-VZEE]; *see also* MICHAEL LEWIS, FLASH BOYS: A WALL STREET REVOLT 36 (2014).

50. *See, e.g.*, Richard Berk et al., *Forecasting Murder Within a Population of Probationers and Parolees: A High Stakes Application of Statistical Learning*, 172 J. ROYAL STAT. SOC'Y SERIES A (STAT. IN SOC'Y) 191 (2009).

51. *See, e.g.*, Brian Kriegler & Richard Berk, *Small Area Estimation of the Homeless in Los Angeles: An Application of Cost-Sensitive Stochastic Gradient Boosting*, 4 ANNALS APPLIED STAT. 1234 (2010).

52. *See, e.g.*, Dursun Delen, *A Comparative Analysis of Machine Learning Techniques for Student Retention Management*, 49 DECISION SUPPORT SYS. 498 (2010).

have already noted that national security and law enforcement agencies are starting to rely on machine learning to support functions as varied as assessing risks of street crime and automating weapons delivery systems. Outside the security and law enforcement context, other government agencies have also begun to explore uses of machine learning, revealing growing recognition of its promise across a variety of policy settings and at all levels of government.[53]

Although we mainly focus in this Article on the use of machine learning by the federal government, the nation's largest cities have received much attention so far for their embrace of machine learning and its potential to improve governmental efficiency and effectiveness.[54] The City of Chicago, for example, has established an award-winning SmartData Platform initiative through which city officials are using machine learning to support a range of city services, from identifying restaurants that should be inspected[55] to predicting where and when rodent control bait should be placed throughout the city.[56] New York City has established a Mayor's Office of Data Analytics,[57] which, among other things, is working with the city's fire department to use machine learning to decide where to send building inspectors.[58] Flint, Michigan has partnered with Google and the University of Michigan to address its recent water crisis by targeting pipe replacements based on machine-learning predictions of lead contamination.[59] The City of Los Angeles has installed sensors in all of its streets that continuously feed data into a machine-learning system that automatically determines when traffic signals should turn red or green to optimize traffic

53. Although outside our scope, we note that courts are also increasingly looking to machine learning as a tool for discovery in the litigation process. *See, e.g.*, Wallis M. Hampton, *Predictive Coding: It's Here to Stay*, SKADDEN (May 5, 2014), https://www.skadden.com/sites/default/files/publications/LIT_JuneJuly14_EDiscoveryBulletin.pdf [https://perma.cc/4WHG-ZLKF] (noting courts' increasing interest in and favorable inclination toward the use of machine learning to assist in e-discovery).

54. *See, e.g.*, STEPHEN GOLDSMITH & SUSAN CRAWFORD, THE RESPONSIVE CITY: ENGAGING COMMUNITIES THROUGH DATA-SMART GOVERNANCE (2014); Bechara Choucair, Jay Bhatt & Raed Mansour, *How Cities Are Using Analytics to Improve Public Health*, HARV. BUS. REV. (Sept. 15, 2014), https://hbr.org/2014/09/how-cities-are-using-analytics-to-improve-public-health/ [https://perma.cc/R26N-7RU2].

55. Nick Rojas, *Chicago and Big Data*, TECHCRUNCH (Oct. 22, 2014), http://techcrunch.com/2014/10/22/chicago-and-big-data/ [https://perma.cc/P47F-9JKV]; *see also* Edward L. Glaeser et al., *Crowdsourcing City Government: Using Tournaments to Improve Inspection Accuracy*, 106 AM. ECON. REV. 114, 114 (2016). Interestingly, similar restaurant hygiene algorithms deployed in Boston were developed via crowdsourcing, a potentially cost-effective alternative to private contracting in certain situations. *Id.*

56. Ash Center Mayors Challenge Research Team, *Chicago's SmartData Platform: Pioneering Open Source Municipal Analytics*, DATA-SMART CITY SOLUTIONS (Jan. 8, 2014), http://datasmart.ash.harvard.edu/news/article/chicago-mayors-challenge-367 [https://perma.cc/MY8X-PDD6].

57. THE MAYOR'S OFFICE OF DATA ANALYTICS, CITY OF NEW YORK, http://www1.nyc.gov/site/analytics/index.page [https://perma.cc/4QTF-M3UL].

58. Brian Heaton, *New York City Fights Fire with Data*, GOV'T TECH. (May 15, 2015), http://www.govtech.com/public-safety/New-York-City-Fights-Fire-with-Data.html [https://perma.cc/XG7W-5UYH].

59. Gabe Cherry, *Google, U-M to Build Digital Tools for Flint Water Crisis*, U. MICH. NEWS (May 3, 2016), http://ns.umich.edu/new/multimedia/videos/23780-google-u-m-to-build-digital-tools-for-flint-water-crisis [https://perma.cc/GV4C-LR6N].

flow.[60]

At the federal level, one of the earliest domestic applications of machine learning came from, as we already noted, the U.S. Postal Service's need for a method to sort mail automatically by predicting the zip codes written on envelopes.[61] Meteorologists within the National Oceanic and Atmospheric Administration have explored the use of machine learning to improve forecasts of severe weather events.[62] Other federal agencies have also started to rely on machine learning to support various regulatory and administrative activities.[63]

Analysts at the U.S. Environmental Protection Agency (EPA), for example, have developed a program called ToxCast to help the agency predict toxicities of chemical compounds.[64] Chemical toxicity has traditionally been established using animal testing, but these laboratory techniques are costly and time consuming, not to mention often harmful to animals. Faced with tens of thousands of chemicals that could be potentially subject to EPA regulation, the agency developed ToxCast to prioritize which of the multitude of chemicals in production should undergo more in-depth testing. ToxCast applies machine-learning algorithms—specifically, linear discriminant analysis—to data on chemicals' interactions obtained from *in vitro* testing to predict their toxicities.[65] In one

60. Ian Lovett, *To Fight Gridlock, Los Angeles Synchronizes Every Red Light*, N.Y. TIMES (Apr. 1, 2013), http://www.nytimes.com/2013/04/02/us/to-fight-gridlock-los-angeles-synchronizes-every-red-light.html [https://perma.cc/25T2-DCFG]; David Z. Morris, *How Swarming Traffic Lights Could Save Drivers Billions of Dollars*, FORTUNE (July 13, 2015, 4:47 PM), http://fortune.com/2015/07/13/swarming-traffic-lights [https://perma.cc/KP8C-N88S].

61. By 1988, USPS contractors had developed one of the first methods for extracting visual images from envelopes and compiling them into an analyzable data set. *See* Ching-Huei Wang & Sargur N. Srihari, *A Framework for Object Recognition in a Visually Complex Environment and Its Application to Locating Address Blocks on Mail Pieces*, 2 INT'L J. COMP. VISION 125, 125 (1988). This data set then enabled the development of early algorithms analyzing handwritten zip codes. *See* O. Matan et al., Handwritten Character Recognition Using Neural Network Architectures, Presented at the 4th USPS Advanced Technology Conference (1990), http://yann.lecun.com/exdb/publis/pdf/matan-90.pdf [https://perma.cc/P4LS-5HZH].

62. David John Gagne II et al., Day-Ahead Hail Prediction Integrating Machine Learning with Storm-Scale Numerical Weather Models, Presented at the Twenty-Seventh Conference on Innovative Applications of Artificial Intelligence (2015), http://www.aaai.org/ocs/index.php/IAAI/IAAI15/paper/view/9724/9898 [https://perma.cc/ZZ97-3UCJ].

63. More than a decade ago, the General Accountability Office surveyed 128 federal agencies and found that fifty-two of them were engaged in "data mining" activities, defined broadly "as the application of database technology and techniques—such as statistical analysis and modeling—to uncover hidden patterns and subtle relationships in data and to infer rules that allow for the prediction of future results." U.S. GEN. ACCOUNTING OFFICE, GAO-04-548, DATA MINING: FEDERAL EFFORTS COVER A WIDE RANGE OF USES 4 (2004). It is not clear from the report how many of these efforts involved machine-learning techniques as opposed to more traditional statistical methods. We know of no comparable effort to survey agencies across government about their use of machine learning.

64. U.S. ENVTL. PROT. AGENCY, TOXCAST FACT SHEET (2013), http://www.epa.gov/sites/production/files/2013-12/documents/toxcast-fact-sheet.pdf [https://perma.cc/P3UU-YL4X].

65. Robert Kavlock et al., *Update on EPA's ToxCast Program: Providing High Throughput Decision Support Tools for Chemical Risk Management*, 25 CHEMISTRY RES. TOXICOLOGY 1287, 1295 (2012). On the use of machine learning in toxicology, see Huanxiang Liu, Xiaojun Yao & Paola Gramatica, *The Applications of Machine Learning Algorithms in the Modeling of Estrogen-Like Chemicals*, 12 COMBINATORIAL CHEM. & HIGH THROUGHPUT SCREENING 490 (2009).

application during ToxCast's first phase, analysts estimated that using machine learning could save the government $980,000 per toxic chemical positively identified.[66] Although the EPA presently uses ToxCast to identify chemicals for additional testing through more traditional means, its underlying predictive approach could eventually form an independent basis for justifying the imposition of regulatory controls.[67]

The U.S. Internal Revenue Service (IRS) has also used machine-learning algorithms to aid its auditing and enforcement functions. In 2001, it began developing a "risk-based collection model" that prioritized the IRS's collection cases for small businesses and self-employed taxpayers by using machine-learning algorithms, including neural networks, to predict risk of nonpayment.[68] In that same year, the agency began to use support vector machines, another type of machine-learning algorithm, to predict abuse and fraud in tax returns and to allocate cases for human review based on the probability of abuse and the magnitude of the dollar amount of the abuse.[69] More recently, in 2009, the IRS launched an Information Reporting and Document Matching program, which applies algorithms to credit card and other third-party data to predict tax underreporting and non-filing by businesses.[70] The IRS increased its requested funding for enforcement targeting from $1.4 million in 2012[71] to over $39 million in 2016,[72] specifically to develop better ways to use machine-learning algorithms, including neural networks, to "identify emerging areas of non-compliance."[73]

66. Matthew T. Martin et al., *Economic Benefits of Using Adaptive Predictive Models of Reproductive Toxicity in the Context of a Tiered Testing Program*, 58 SYS. BIOLOGY REPROD. MED. 3, 4–6 (2012).

67. Richard S. Judson et al., *Estimating Toxicity-Related Biological Pathway Altering Doses for High-Throughput Chemical Risk Assessment*, 24 CHEM. RES. TOXICOLOGY 451, 457–60 (2011).

68. Jane Martin & Rick Stephenson, Risk-Based Collection Model Development and Testing, Presented at the Internal Revenue Service Research Conference (2005), http://www.irs.gov/pub/irs-soi/05stephenson.pdf [https://perma.cc/M65K-D9D7].

69. David DeBarr & Maury Harwood, Relational Mining for Compliance Risk, Presented at the Internal Revenue Service Research Conference (2004), http://www.irs.gov/pub/irs-soi/04debarr.pdf [https://perma.cc/Y9F8-RWNK].

70. *See* CHRIS WAGNER ET AL., TAXPAYER ADVOCATE SERV., *IRS Policy Implementation Through Systems Programming Lacks Transparency and Precludes Adequate Review*, *in* 2010 ANNUAL REPORT TO CONGRESS 71, 76, http://www.irs.gov/pub/irs-utl/2010arcmsp5_policythruprogramming.pdf [https://perma.cc/3DHR-ZXAT]. Interestingly, this program was temporarily halted in 2014 not because of issues with its predictive algorithms, but due to continued reliance on human tax examiners; the Automated Underreporter System that predicts underreporting was successfully deployed, but the Case Management System that then handed flagged cases to humans for further examination did not meet technical requirements. *See* TREASURY INSPECTOR GEN. FOR TAX ADMIN., 2014-20-088, THE INFORMATION REPORTING AND DOCUMENT MATCHING CASE MANAGEMENT SYSTEM COULD NOT BE DEPLOYED (2014), https://www.treasury.gov/tigta/auditreports/2014reports/201420088fr.pdf [https://perma.cc/E3CS-VLEL].

71. U.S. INTERNAL REVENUE SERV., FY 2012 BUDGET REQUEST 63 (2011), https://www.treasury.gov/about/budget-performance/Documents/CJ_FY2012_IRS_508.pdf [https://perma.cc/GB72-GXW5].

72. U.S. INTERNAL REVENUE SERV., FY 2016 PRESIDENT'S BUDGET 88 (2015), https://www.treasury.gov/about/budget-performance/CJ16/02-06.%20IRS%20FY%202016%20CJ.pdf [https://perma.cc/2BXV-DPFQ].

73. U.S. INTERNAL REVENUE SERV., *supra* note 71, at 63.

In addition to the EPA's and the IRS's use of machine learning, the U.S. Food and Drug Administration (FDA) has conducted research on the use of machine-learning techniques to extract information about known equipment failures, errors, or other adverse events from medical device reports.[74] This safety agency is also currently engaged in a five-year collaborative research agreement with the Massachusetts Institute of Technology (MIT) focusing on "artificial intelligence, advanced statistical machine learning and data mining methods."[75] MIT researchers have also recently collaborated with researchers at the U.S. Department of the Treasury's Office of Financial Research (OFR) to survey methods of evaluating systemic risk in consumer credit markets, including the use of classification and regression trees.[76] Separately, academic researchers have demonstrated how machine-learning algorithms can be used to predict cases of financial statement fraud,[77] electoral fraud,[78] and even illegal fishing practices.[79] Agencies like the Commodity Futures Trading Commission and the Securities and Exchange Commission (SEC) have also taken note of these new approaches to fraud detection.[80]

For machine-learning algorithms to work, they depend on accessible and analyzable data. Toward that end, many agencies are beginning to recognize the importance of so-called big data—or large volumes of information—in ways

74. *Commissioner's Fellowship Program: Final Report Abstracts*, U.S. FOOD & DRUG ADMIN., https://www.fda.gov/AboutFDA/WorkingatFDA/FellowshipInternshipGraduateFacultyPrograms/CommissionersFellowshipProgram/ucm413253.htm [https://perma.cc/67MG-UXJL]. Work by the same researcher has similarly used machine-learning algorithms to extract laboratory test information from FDA decision summaries of device premarket notifications. Yanna Shen Kang & Mehmet Kayaalp, *Extracting Laboratory Test Information from Biomedical Text*, 4 J. PATHOL. INFORM. 23 (2013).

75. U.S. FOOD & DRUG ADMIN., MOU 225-12-0010, MEMORANDUM OF UNDERSTANDING BETWEEN THE UNITED STATES FOOD AND DRUG ADMINISTRATION AND MASSACHUSETTS INSTITUTE OF TECHNOLOGY (2012), http://www.fda.gov/AboutFDA/PartnershipsCollaborations/MemorandaofUnderstandingMOUs/AcademiaMOUs/ucm318476.htm [https://perma.cc/DUA8-9KM4].

76. For the research surveyed, see Amir E. Khandani et al., *Consumer Credit Risk Models via Machine-Learning Algorithms*, 34 J. BANKING & FIN. 2767 (2010).

77. *See, e.g.*, Johan Perols, *Financial Statement Fraud Detection: An Analysis of Statistical and Machine Learning Algorithms*, 30 AUDITING 19 (2011).

78. Francisco Cantu & Sebastian M. Saiegh, *A Supervised Machine Learning Procedure to Detect Electoral Fraud Using Digital Analysis* (Caltech/MIT Voting Technology Project, Working Paper No. 11, 2010), http://ssrn.com/abstract=1594406 [https://perma.cc/NSJ2-9FD9].

79. Cleridy E. Lennert-Cody & Richard A. Berk, *Statistical Learning Procedures for Monitoring Regulatory Compliance: An Application to Fisheries Data*, 170 J. ROYAL STAT. SOC'Y SERIES A (STAT. IN SOC'Y) 671 (2007); *see also* Richard Berk, *Forecasting Consumer Safety Violations and Violators*, *in* IMPORT SAFETY: REGULATORY GOVERNANCE IN THE GLOBAL ECONOMY 131, 135–36 (Cary Coglianese, Adam M. Finkel & David Zaring eds., 2009).

80. *See, e.g.*, Scott W. Bauguess, Deputy Chief Economist, U.S. Sec. & Exch. Comm'n, The Hope and Limitations of Machine Learning in Market Risk Assessment (Mar. 6, 2015), http://cfe.columbia.edu/files/seasieor/center-financial-engineering/presentations/MachineLearningSECRiskAssessment030615public.pdf [https://perma.cc/HY5T-GKXH]; Scott D. O'Malia, Commissioner, U.S. Commodity Futures Trading Comm'n, Opening Statement at the 12th Meeting of the Technology Advisory Committee (June 3, 2014), http://www.cftc.gov/PressRoom/SpeechesTestimony/omaliastatement060314 [https://perma.cc/RK4T-UK84].

that suggest that the analytical infrastructure needed to use machine learning more extensively may soon be realized. Officials at the U.S. Federal Aviation Administration (FAA), for example, have recognized that in the service of aviation safety "there is significantly more potential" for the use of big data.[81] The U.S. Federal Deposit Insurance Corporation (FDIC) has included as a component of a recent Business Technology Strategic Plan the maturation of "the back-end disciplines of in-memory analytics, big data, and data quality."[82] Similarly, the U.S. Federal Communications Commission (FCC) has developed a Data Innovation Initiative to support the goal of improving its data analytic capacity.[83] Throughout the Obama Administration, the White House prioritized big data use across the executive branch through a Big Data Research and Development Initiative,[84] with President Obama's 2016 budget calling for a $1 billion increase in funding for statistical programs.[85]

Efforts remain underway not only to create large data sets to support agency functions but also to make big data more readily analyzable. One example can be found in the creation of the global Legal Entity Identifier (LEI), a universal reference code for each entity active in financial markets.[86] Treasury's OFR launched an effort to establish LEI in 2010,[87] and by 2014 the LEI Regulatory Oversight Committee had assumed operational responsibility for its development.[88] Having such a unique identifier will enhance regulators' ability "to identify parties to financial transactions instantly and precisely," allowing the

81. Letter from R. John Hansman, Chairman, U.S. Fed. Aviation Admin. Res., Eng'g & Dev. Advisory Comm., to Michael P. Huerta, Administrator, U.S. Fed. Aviation Admin. (Oct. 2, 2013), http://www.faa.gov/about/office_org/headquarters_offices/ang/offices/tc/about/campus/faa_host/rdm/media/pdf/Guidance-FY2016.pdf [https://perma.cc/2VDQ-SVZ3]; *see also* U.S. FED. AVIATION ADMIN., MEETING MINUTES OF THE RESEARCH, ENGINEERING, AND DEVELOPMENT ADVISORY COMMITTEE 7 (2012), http://www.faa.gov/about/office_org/headquarters_offices/ang/offices/tc/about/campus/faa_host/rdm/media/pdf/minutes-FullComm_09262012.pdf [https://perma.cc/DNG9-GET4] (noting that the agency is "taking a look at overarching data management").

82. U.S. FED. DEPOSIT INS. CORP., BUSINESS TECHNOLOGY STRATEGIC PLAN 2013–2017, at 8 (2013), https://www.fdic.gov/about/strategic/it_plan/BusinessTechnologyStrategicPlan2013-2017.pdf [https://perma.cc/L7KE-KPT4].

83. Michael Byrne, *Big Data*, FCC BLOG (Oct. 28, 2010, 1:06 PM), https://www.fcc.gov/news-events/blog/2010/10/28/big-data [https://perma.cc/AU7Q-J6XA].

84. *See* Press Release, Office of Sci. and Tech. Policy, Obama Administration Unveils "Big Data" Initiative: Announces $200 Million in New R&D Investments (Mar. 29, 2012), http://www.whitehouse.gov/sites/default/files/microsites/ostp/big_data_press_release_final_2.pdf [https://perma.cc/GX57-AB83].

85. Aaron Boyd, *Obama Budget Pushes Better Decisions Using Open Data*, FED. TIMES (Feb. 3, 2015), http://www.federaltimes.com/story/government/management/budget/2015/02/03/open-data-evidence-based-decisions-funded-2016-budget/22802323 [https://perma.cc/63U6-CQUG].

86. *See generally Legal Entity Identifier–Frequently Asked Questions*, U.S. OFF. OF FIN. RES., http://financialresearch.gov/data/legal-entity-identifier-faqs [https://perma.cc/LH9B-YUG9] (answering frequently asked questions about the legal entity identifier).

87. *See* Statement on Legal Entity Identification for Financial Contracts, 75 Fed. Reg. 74,146 (Nov. 30, 2010).

88. Matthew Reed, *Legal Entity Identifier System Turns a Corner*, U.S. OFF. OF FIN. RES. (July 3, 2014), https://financialresearch.gov/from-the-management-team/2014/07/03/legal-entity-identifier-system-turns-a-corner/ [https://perma.cc/6H64-VFFL].

authorities to apply machine learning to larger data sets.[89]

Agencies are also actively working toward development of the cloud storage systems necessary to exploit the power of machine learning.[90] Such storage that takes place via distributed networks of computers proves to be better suited to running computationally intensive algorithms, and its availability better facilitates interagency sharing of big data. The FDA, for example, has leveraged cloud computing to store information on foodborne pathogens, giving the agency "the ongoing, simultaneous capacity to collect, control and analyze enormous data sets."[91] Similarly, the EPA created a Cross-Agency Data Analytics and Visualization Program intended to foster the creation of databases that will permit the analysis of data from many different agencies and organizations.[92] The SEC is implementing cloud computing to store and process its one billion daily records of financial market activities, often time-stamped to the microsecond, allowing the SEC to "perform analyses of thousands of stocks . . . involving 100 billion records at a time."[93] The proliferation of such efforts to capture, share, and analyze vast quantities of data makes it easy to envision, for example, an extension of the SEC's cloud computing program that would eventually allow agency computers to monitor trading activities in real time, predicting in milliseconds whether a financial transaction is the result of insider trading and then automatically stopping or reversing trades based on those predictions.[94]

89. U.S. OFF. OF FIN. RES., *supra* note 86.

90. *See Response to–Request for Information: Preparing for the Future of Artificial Intelligence*, IBM http://research.ibm.com/cognitive-computing/ostp/rfi-response.shtml [https://perma.cc/5PGS-J24T] (describing how AI systems deployed at scale will require "high-performance distributed cloud systems, new computing architectures such as neuromorphic and approximate computing, and new devices such as quantum and new types of memory devices").

91. Taha A. Kass-Hout, *FDA Leverages Big Data Via Cloud Computing*, FDA VOICE (June 19, 2014), http://blogs.fda.gov/fdavoice/index.php/2014/06/fda-leverages-big-data-via-cloud-computing [https://perma.cc/DT6V-N5JK].

92. *See EPA's Cross-Agency Data Analytics and Visualization Program*, U.S. ENVTL. PROT. AGENCY, https://web.archive.org/web/20160414154548/ https://www.epa.gov/toxics-release-inventory-tri-program/epas-cross-agency-data-analytics-and-visualization-program [https://perma.cc/P769-LGB6].

93. *Market Information Data Analytics System*, U.S. SEC. & EXCH. COMM'N, http://www.sec.gov/marketstructure/midas.html [https://perma.cc/2YYE-3LKN].

94. Some academic research has already been conducted both to use machine-learning algorithms to predict such trading violations and to call for their use to make such predictions. *See, e.g.*, Steve Donoho, Early Detection of Insider Trading in Option Markets, Presented at the Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (2004); Andrei A. Kirilenko & Andrew W. Lo, *Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents*, 27 J. ECON. PERSP. 51 (2013); Shawn Mankad et al., *Discovering the Ecosystem of an Electronic Financial Market with a Dynamic Machine-Learning Method*, 2 ALGORITHMIC FIN. 151 (2013); Gregory Scopino, *Preparing Financial Regulation for the Second Machine Age: The Need for Oversight of Digital Intermediaries in the Futures Markets*, 2015 COLUM. BUS. L. REV. 439, 443–44 (2015); *see also Foresight: The Future of Computer Trading in Financial Markets*, GOV'T OFF. FOR SCI. 42 (2012), http://www.cftc.gov/idc/groups/public/@aboutcftc/documents/file/tacfuturecomputertrading1012.pdf [https://perma.cc/NW9Z-EXGG].

With these various efforts underway, the government is well on its way into the era of machine learning. Before turning to the legal implications of this new era, we next develop more precisely what machine learning portends for government agencies and why its use might raise questions under prevailing administrative law doctrines.

C. ADJUDICATING BY ALGORITHM, RULEMAKING BY ROBOT

What exactly might be problematic about an era in which government embraces machine learning? Up to this point in our discussion, perhaps the answer will not be obvious. Were it not for dire warnings in the popular press of impending artificially intelligent oppression, it might seem that machine learning simply represents a more sophisticated, data-rich, and predictively useful version of the kind of analytic methods that government agencies have long used. If that is what machine learning is, and if government can use new statistical techniques to improve its performance of various functions from weather forecasting to identifying potentially hazardous chemicals, then presumably a machine-learning era in government should not only be completely unproblematic but also positively encouraged.

Three principal properties of machine learning combine, however, to distinguish it from other analytical techniques and give rise to potential concerns about the greater reliance on machine learning by governmental authorities. The first is machine learning's self-learning property. The results of algorithms do not depend on humans specifying in advance how each variable is to be factored into the predictions; indeed, as long as learning algorithms are running, humans are not really controlling how they are combining and comparing data. These algorithms effectively look for patterns on their own. The second key property is machine learning's "black box" nature. The results of machine learning analysis are not intuitively explainable and cannot support causal explanations of the kind that underlie the reasons traditionally offered to justify governmental action. Finally, machine learning, as with other computational strategies in today's digital era, can be fast and automatic, supporting uses in which the algorithm produces results that can shorten or potentially bypass human deliberation and decision making. All three of these factors combine to make machine-learning techniques appear qualitatively more independent from humans when compared to other statistical techniques.

To illustrate these features of machine learning and what they portend for government, consider the challenges the U.S. Pipeline and Hazardous Materials Safety Administration (PHMSA) faces in deciding how to allocate limited inspection resources to oversee the many thousands of miles of gas, oil, and chemical pipelines throughout the United States. Major leaks as well as explosions from leaky pipelines in recent years have made palpable the significance of effective governmental oversight of pipeline safety.[95] In recent years, PHMSA has explored using a traditional regression approach to predict risks of pipeline accidents and decide how to target the agency's inspections.[96] Although such an attempt to engage in quantitative decision making is certainly laudable, much more efficient inspection targeting could result if PHMSA generated its risk predictions using machine learning. This could be possible in the near future once big data sets are shared in real time between different agencies and information streams could be provided by remote-sensing technologies. Instead of being limited to analyzing a dozen or so variables that PHMSA's analysts have predetermined should be included in their regression analysis, machine-learning algorithms could work their way through massive amounts of data containing hundreds of potentially predictively useful variables, ranging from pipeline operators' tax returns to their firms' workforce diversity. The resulting predictions of pipeline accident risk could be used not only to target inspections but also potentially, if such an algorithm could be supplied with real-time data from remote sensors, to order preemptive shutdowns of pipeline segments that the algorithm predicts are at risk of imminent failure. To the extent that modern pipeline systems are equipped with computerized, remote shut-off capabilities, a machine-learning algorithm could even be programmed to send an automatic order to pipeline operator's system calling for an immediate, automatic shutdown of a section of pipeline based on real-time forecasts produced by machine learning, all potentially without any human intervention.

Machine learning is well suited for automating these kinds of decisions, given its emphasis on accuracy and the government's overwhelming need to use its limited resources to prevent dangers from arising. But notice that a shift to a machine-learning approach in this context could come along with some qualitative loss of human involvement. Under machine learning, PHMSA analysts would no longer predetermine which variables should be included in the agency's risk models; indeed, they would not even create any risk *models* at all, in the sense of building equations specifying exactly how various variables might impact pipeline risk. Machine learning also does not afford a ready means of explaining why any section of pipeline should be inspected or shut down. The computerized nature of machine learning also means that it can automate decisions currently made by humans, such as the dispatching of inspectors or even the inspections themselves. It is not difficult to imagine a future in which a

95. *See, e.g.*, David R. Baker, *L.A. Gas Leak Plugged, but California Pipelines Regularly Leak*, S.F. CHRON. (Feb. 14, 2016), http://www.sfchronicle.com/business/article/L-A-gas-leak-plugged-but-California-pipelines-6830717.php [https://perma.cc/6HQA-X6FU]; Deirdre Fulton, *More Than 300 a Year: New Analysis Shows Devastating Impact of Pipeline Spills*, COMMON DREAMS (Nov. 17, 2014), http://www.commondreams.org/news/2014/11/17/more-300-year-new-analysis-shows-devastating-impact-pipeline-spills [https://perma.cc/L7MN-S7UH].

96. RICK KOWALEWSKI & PEG YOUNG, BUREAU OF TRANSP. STATISTICS, SR-010, DATA-DRIVEN RISK MODELS COULD HELP TARGET PIPELINE SAFETY INSPECTIONS 3–4 (2008), http://www.rita.dot.gov/ bts/sites/rita.dot.gov.bts/files/publications/special_reports_and_issue_briefs/special_report/2008_010/pdf/entire.pdf [https://perma.cc/4825-D39L].

machine-learning system based in PHMSA's headquarters in Washington, D.C. could be used to automatically dispatch agency drones distributed throughout the country, having them fly over sections of pipeline to take video images or collect air quality samples, thereby removing altogether the need to send human inspectors to the scene.

Admittedly, even with this potential future scenario in mind, it still might not be self-evident why machine learning may be problematic. After all, machines have long supported governmental functions in the administrative state. Although not flown by drones, machines currently collect air quality samples in a network of fixed sites around the country, informing state and federal environmental regulatory decision making.[97] Moreover, when it comes to inspecting potentially hazardous sites, eliminating the need for humans to enter high-risk areas should surely be a positive advance, not a reason for alarm.[98] Furthermore, decisions about the allocation of inspection resources have long been treated—as a matter of well-accepted law—as falling entirely within an agency's discretion.[99] If agencies can legally allocate inspection resources by the flip of a coin—that is, sending inspectors to sites at random, as some agencies do—then they should be able legally to rely on more sophisticated algorithms that deploy scarce inspection resources automatically but more efficiently.[100] For this reason, we foresee comparatively little resistance, as a matter of law, to applications of machine learning that aim to make more efficient use of scarce inspection resources.

Many uses of machine learning by administrative agencies will be like the use of machine learning to decide where to send inspectors in that they will inform actions committed to agency discretion. Most of these uses will be unproblematic from the standpoint of administrative law.[101] Surely the U.S. Postal Service's reliance on machine-learning algorithms to sort mail hardly constitutes any grave threat to society, either existential or constitutional. In

---

97. *See, e.g.*, *Ambient Air Monitoring*, U.S. ENVTL. PROT. AGENCY, https://www.epa.gov/air-quality-management-process/ambient-air-monitoring [https://perma.cc/GQ2N-SYJE].

98. *Cf.* Suzanne Goldenberg, *Deepwater Horizon Oil Spill: Underwater Robots Trying to Seal Well*, GUARDIAN (Apr. 26, 2010, 2:31 PM), http://www.theguardian.com/environment/2010/apr/26/deepwater-horizon-spill-underwater-robots [https://perma.cc/NAC5-RGKX].

99. Heckler v. Chaney, 470 U.S. 821, 831 (1985) (noting that the Supreme Court has recognized that "an agency's decision not to prosecute or enforce, whether through civil or criminal process, is a decision generally committed to an agency's absolute discretion").

100. In addition to New York City's use of machine learning to determine where to send its building inspectors, the City of Chicago is using algorithms to allocate food safety inspectors, Mohana Ravindranath, *In Chicago, Food Inspectors Are Guided by Big Data*, WASH. POST (Sept. 28, 2014), https://www.washingtonpost.com/business/on-it/in-chicago-food-inspectors-are-guided-by-big-data/2014/09/27/96be8c68-44e0-11e4-b47c-f5889e061e5f_story.html [https://perma.cc/V25Y-53CG].

101. We recognize, of course, that it might be possible for agencies to abuse their discretion, which is why we characterize uses of machine learning for discretionary purposes to be *virtually* unproblematic. *See, e.g.*, 5 U.S.C. § 706(2)(A) (2012). One way that such discretion could possibly be abused would be if it were to be deployed in an unlawfully discriminatory fashion. In Section II.C, we consider whether the use of machine learning even in discretionary enforcement allocation decisions might offend equal protection.

addition, even when agency officials use learning algorithms to support actions that are not committed to agency discretion, if they use them simply to inform their own independent judgments, this too should be unremarkable. Such use would be indistinguishable from any other research support or informational input into agency decision making. The non-shaded parts of Table 1 highlight several general types of agency uses of machine learning that should easily be viewed as beyond reproach, at least from the standpoint of existing general principles of structural law governing the administrative state.

The domains in which machine learning might be of concern, at least as a prima facie matter, will be those in which artificial intelligence is used more for determining, rather than just supporting, decisions that are not otherwise committed to agency discretion by law. As shown in the shaded portions of Table 1, that leaves two important realms in which machine learning could be incorporated into the administrative state: *adjudicating by algorithm* and *rulemaking by robot.*[102]

One example of *adjudicating by algorithm* would be our posited PHMSA pipeline safety machine-learning system that automatically issues shut-off orders when the system forecasts a heightened risk. It is not difficult to imagine

Table 1. Applications of Machine Learning in the Administrative State

Type of Administrative Action

Role of Machine Learning in Agency Decision-Making	"Discretionary"	"Non-Discretionary" Adjudication	"Non-Discretionary" Rulemaking
Supportive			
↕			
Determinative		*Adjudicating by Algorithm*	*Rulemaking by Robot*

102. Adjudication and rulemaking, of course, are the two canonical types of actions that agencies may take under the Administrative Procedure Act. 5 U.S.C. § 551(5), (7) (2012). We have labeled these two types of administrative action as "non-discretionary" not because agencies are mandated to take these actions (although sometimes they can be). Rather, we have labeled them this way because these actions will be surrounded by "law to apply" that will subject these actions to judicial review under the Administrative Procedure Act. Citizens to Preserve Overton Park v. Volpe, 401 U.S. 402 (1971). By "discretionary" in Table 1, we mean simply that an action is "committed to agency discretion" and thus not subject to judicial review. 5 U.S.C. § 701(a)(2). We do recognize, of course, that on occasion there may be law to apply even to supportive uses of analytic techniques or to the use of other factors that support decisions, such as would be the case if a statute were to prohibit an agency from using machine learning in even a non-determinative role. *Cf.* Whitman v. American Trucking Assocs., Inc., 531 U.S. 457 (2001) (holding that the Clean Air Act prohibits the EPA Administrator from considering costs when setting air quality standards). For these reasons, Table 1 should be viewed simply as a heuristic intended to illustrate some generalizations about the administrative state.

other examples of adjudicatory decisions that could be automated by algorithms,[103] especially when the relevant criteria for an adjudicatory action are forward-looking and thus dependent on accurate predictions. At some point, for example, the FAA might be able to license pilots through an entirely automated process relying on risk-based machine learning forecasts of individual applicants' overall level of safety.[104] The Federal Trade Commission or the Department of Justice's Antitrust Division might conceivably come to rely on machine learning to predict what effects a proposed merger would have on future competition and market pricing, perhaps entirely automating the antitrust review process.

When it comes to *rulemaking by robot*, we need not rely entirely on the imagination. The City of Los Angeles' current traffic signaling system illustrates a very simple but still real-world application of rulemaking by robot. Although deciding the color of traffic lights may seem like a trivial example, a traffic signal does determine what rule applies to anyone who wants to drive along a city street at a given period of time. Yet with the system in place in Los Angeles, just as no human determines when a traffic light should be red or green, no government official can really explain why the city's machine-learning system sets any given traffic light (that is, rule) when it does. We can expect it will not be long before more government authorities, at the local and federal levels, will be able to develop similar systems in their own domains that are conceptually equivalent to Los Angeles' traffic control system.

It is not difficult to imagine more complex and consequential examples of regulating by robot. Consider the possibility that the SEC might find it beneficial, even necessary, to govern the world of high-speed electronic trading by making nimble and equally high-speed adjustments to the rules of market transactions, perhaps modifying stock exchanges' current, rigid trading circuit breakers with ones that adjust in real time.[105] The U.S. Department of the Treasury, for similar reasons, might plausibly seek to establish a dynamic, automated process according to which certain macro-prudential rules governing financial institutions respond to real-time market changes indicative of systemic

103. Imagination may not be required for much longer. IBM is currently developing machine-learning algorithms to predict smog levels in China, predictions that may soon be used to determine governmental shutdowns of factories or limits on traffic volumes. Will Knight, *Can Machine Learning Help Lift China's Smog?*, MIT TECH. REV. (Mar. 28, 2016), https://www.technologyreview.com/s/600993/can-machine-learning-help-lift-chinas-smog/ [https://perma.cc/SL77-VJVK].

104. Currently, such forward-looking adjudicatory decisions like licensing are based on rules, making the issue one of whether applicants comply with the applicable rule, or the criteria contained in a rule, and thus qualify to receive a license. Machine learning makes an alternative adjudicatory framework possible, one that considers forecasted risk based on an algorithmic analysis of potentially hundreds of variables. Machine learning has been shown to be an effective tool in making certain forward-looking adjudicatory decisions in the criminal law system. Richard A. Berk et al., *Forecasting Domestic Violence: A Machine Learning Approach to Help Inform Arraignment Decisions*, 13 J. EMPIRICAL L. STUD. 94, 110 (2016) [hereinafter Berk, *Forecasting Domestic Violence*]; Berk et al., *supra* note 50, at 208.

105. *See supra* note 94.

risk.[106] Even when time is not so critical and the "good cause" exemption to the standard rulemaking process might not apply, it is hardly unimaginable today that agencies could automate entirely the notice-and-comment rulemaking process, especially for the kinds of routine rules that make up the bulk of government rules.[107] Natural language processing programs could even conceivably read and summarize any public comments submitted on proposed rules and potentially even craft some of the regulatory language.[108]

For anything but perhaps the simplest rules, like traffic signals, rulemaking by robot will require that machine learning be combined with other analytic techniques. Rules are forward-looking, but they also involve complex normative judgments, not merely predictive ones. Determining the content of rules often

106. For a discussion of how predictions of systemic risk can affect rulemaking by regulatory agencies, see Dimitrios Bisias et al., *A Survey of Systemic Risk Analytics* 2, 10–11 (U.S. Dep't of the Treasury Office of Fin. Research, Working Paper No. 0001, 2012), https://www.treasury.gov/initiatives/wsr/ofr/Documents/OFRwp0001_BisiasFloodLoValavanis_ASurveyOfSystemicRiskAnalytics.pdf [https://perma.cc/D8LJ-8EYJ].

107. The notice-and-comment process, and the good cause exception to it, is provided at 5 U.S.C. § 553 (2012). It is important to recognize that we have adopted a formulation of algorithm-created rules resembling that of rules as they exist today; an algorithm would promulgate rules specifying a particular course of action, a particular safety standard, a particular acceptable emissions level, or so forth. Under this formulation, any changes to an algorithm that result in a different prescription, including merely re-running the algorithm as specified on new data, would necessitate a new rulemaking process, absent a good cause exemption. But, because machine learning is likely to be of most utility when engaged to regulate dynamic, time-sensitive environments, it is probable that such an exemption could often or even categorically be claimed. An alternative formulation might be one in which rules state merely that an algorithm will be used to promulgate prescriptions continuously; instead of a rule reciting a particular course of action or safety standard on the basis of algorithmic output, the rule would say that a future algorithm will run continuously and be updated dynamically to decide the appropriate course of action or safety standard in the future. *Cf.* Coglianese, *supra* note 25, at 370–71 (noting the possibility of "a reconceptualization of the form in which rules are promulgated"). Such a formulation probably would not necessitate that a new rulemaking be commenced whenever the algorithm is updated, but it might be legally problematic given the need for reason-giving and transparency. *See infra* Section II.D.

108. Already agencies use digital tools to sort and identify duplicates in comments submitted in rulemaking proceedings that have generated large volumes of public submissions. *See, e.g.*, Jane E. Fountain, *Prospects for Improving the Regulatory Process Using E-Rulemaking*, 46 COMMC'NS ACM 63, 63–64 (2003) (discussing federal agency use of automated tools to sort comments beginning as early as 1997). Similar programs are now capable of processing data and automatically writing prose, at least for now in the context of sports reports and fiction. *See* Ian Crouch, *The Sportswriting Machine*, NEW YORKER (Mar. 26, 2015), http://www.newyorker.com/news/sporting-scene/the-sportswriting-machine [https://perma.cc/25FB-UBFS]; Matt McFarland, *A Computer Program Is Writing New 'Friends' Episodes. Are They Any Good?*, WASH. POST (Jan. 21, 2016), https://www.washingtonpost.com/news/innovations/wp/2016/01/21/a-computer-program-is-writing-new-friends-episodes-are-they-any-good/ [https://perma.cc/57FC-UY86]. By itself, an algorithm might not be capable of writing the entire content of a final rule document, as sections such as the summary would presumably require a nuanced explanation of the background of a rule and its purposes, and natural language processing cannot generate such complete thoughts de novo. Natural language processing could, however, probably write sections of a rule document that require statements of facts. Such algorithms can rely on previous examples of how factual statements are worded to create sentences that describe new facts in a fill-in-the-blank manner. Other sections of rule documents require human input related to the goal of a potential rule, but such sections could plausibly be written in advance by humans and then the rest could be filled in with algorithm-written content.

requires making difficult choices about the entities to be regulated, the conduct or outcome that the rule tells these entities to achieve or avoid, and the nature and degree of the consequences that follow from adhering or not adhering to the rule's commands.[109] Machine-learning algorithms cannot directly make the choices about these different aspects of a rule's content not only because some of these choices are normative ones, but also because learning algorithms are merely predictive and thus unable to overlay causal interpretations on the relationship between possible regulations and estimated effects.[110] The justification for new rules depends, after all, on the effects that their adoption and implementation are likely to cause.[111]

Nevertheless, it may be possible for machine learning to make rules in this fashion when used in conjunction with procedures known as agent-based models (ABM) or multi-agent systems (MAS).[112] Agent-based modeling refers to the use of an algorithm consisting of a mathematically-defined environment that includes agents that observe the overall environment and take actions designed to reach a specified goal.[113] Multi-agent systems are similar to agent-based models but with multiple autonomous agents interacting with each other.[114] With either of these agent-based techniques, the agents—which, in the rulemaking context, would include the regulator and the regulated entities—must have some defined decision-making processes that allow them to translate observations of the environment into actions. These decision-making processes can be specified a priori by the researcher or regulatory official, but such a priori knowledge often does not exist or is not sophisticated enough to mimic how real-world agents make their decisions. Therefore, machine learning—often called reinforcement learning in these applications (or what we will, for ease of reference, call "embedded machine learning")—is incorporated into agent-based models' decision-making processes of individual agents. The mathematical agents within these systems, in other words, learn how to make decisions.

109. *See* Cary Coglianese, *Engaging Business in the Regulation of Nanotechnology*, *in* GOVERNING UNCERTAINTY: ENVIRONMENTAL REGULATION IN THE AGE OF NANOTECHNOLOGY 46, 50–51 (Christopher J. Bosso ed., 2010).

110. *See* BERK, *supra* note 33, at 9–17. This difficulty also often faces even conventional techniques in attempting to claim causal inference. *See* Richard A. Berk et al., *What You Can Learn from Wrong Causal Models*, *in* HANDBOOK OF CAUSAL ANALYSIS FOR SOCIAL RESEARCH 403, 422–23 (Stephen L. Morgan ed., 2013).

111. CARY COGLIANESE, MEASURING REGULATORY PERFORMANCE: EVALUATING THE IMPACT OF REGULATION AND REGULATORY POLICY, OECD Expert Paper No. 1 (Aug. 2012), http://www.oecd.org/gov/regulatory-policy/1_coglianese%20web.pdf [https://perma.cc/6ZZZ-NERP].

112. For incorporation of machine learning into ABM, see, for example, W. Rand, *Machine Learning Meets Agent-Based Modeling: When Not to Go to a Bar* 2 (Northwestern Univ. Working Paper, 2006), https://ccl.northwestern.edu/papers/agent2006rand.pdf [https://perma.cc/SU6T-AB9V]. For use of machine learning in MAS, see Lucian Buşoniu et al., *A Comprehensive Survey of Multiagent Reinforcement Learning*, 38 IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS, PART C: APPLICATIONS AND REVIEWS 156, 156 (2008).

113. For a more formal definition of ABM, see Nigel Gilbert, *Agent-Based Models*, *in* 153 QUANTITATIVE APPLICATIONS IN THE SOCIAL SCIENCES (John Fox ed., 2008).

114. For a more formal definition of MAS, see Buşoniu et al., *supra* note 112, at 156.

To translate these embedded machine-learning techniques to possible rulemaking applications, consider how the Occupational Safety and Health Administration (OSHA) might proceed if it were to create an automated process for determining whether to implement a new workplace safety regulation. OSHA could implement an algorithm in which the modeled agents are the employers being regulated. The environment in which these agents operate would include mathematically-specified factors capable of influencing agent behavior, including a possible regulation. The employer-agents in the model would "observe" the environment, which would include different regulatory alternatives (including an environment with no regulation), and then "take" actions, such as complying with the regulation, to reach their own goals, perhaps defined as profit maximization. Now, although OSHA would like to use this agent-based model to see how employers respond to the potential new regulation and, consequently, what effects the regulation may have, OSHA does not know a priori how employers will decide how to respond to any regulation. The agent-based model would therefore use a machine-learning technique to select employers' optimal responses to the regulation given their profit maximization goal.

This example suggests how OSHA might use machine learning embedded within an agent-based model of the effects of a proposed regulation. But the techniques' real potential to inform the content of regulations comes from the ability of OSHA to include an agent representing itself in the ABM. This mathematically-represented agent would "issue" multiple different possible regulations—formulated in advance by human programmers—and then "select" the regulatory alternative that yields those effects, as defined in relation to observable components of the environment, that maximize an objective function (or goal) established by the real-world OSHA. The possible regulations analyzed in this fashion could assume any number of different combinations of regulatory targets, commands, and consequences, with the forecasted effect of these regulations on the actions of regulated entities being observed through the modeling exercise. Unlike in the adjudicatory context, where machine learning directly makes individualized forecasts and where an adjudication can be "determined" simply by whether an algorithm forecasts risks or other outcomes above a threshold level, in the rulemaking context machine learning would need to be nested within a larger decision-making model to support automated regulatory decisions. Machine learning predictions would, within an agent-based simulation, inform agents' actions, which in turn would generate predicted outcomes from different regulatory permutations.

This fusion of agent-based or multi-agent models with machine learning may hold great potential for assisting in certain kinds of rulemaking, but, even with this fusion, governmental reliance on algorithms would still not cede entirely the involvement of humans. As already indicated, at a foundational level, humans will still need to choose and then input into embedded machine-learning systems the data, as well as overarching goals to be maximized and constraints to be minimized. Moreover, due to data limitations as well as core uncertainties, many rulemaking decisions will still by necessity call for human judgment and thus be incapable of automation.

As with any statistical technique, the algorithms that could be embedded in automated rulemaking models will require data. Because all historical data arise within a world with a different rule than the one proposed (even if that is no rule at all), regulators will seldom (if ever) find enough data to correspond to all possible forms a future regulation and resulting environmental state might take. This is often a challenge in applications of agent-based models in other contexts, such as healthcare provision. In those other contexts, the lack of data is often addressed through the creation of simulated environmental data.[115] Generation of simulated data, however, requires that the architecture of the environment being modeled, and the relationships between components of that environment, be sufficiently well known a priori as to be specifiable. Embedded machine-learning techniques have been successfully developed for applications like modeling how infectious patients should be moved around a hospital.[116] In that context, the environment of interest can be reasonably well specified. The actors and parameters are limited—for example, healthcare professionals, infected patients, uninfected patients, and rooms—and the analyst knows a priori enough about how diseases are transmitted to generate simulated data using probabilities of infection based on proximity and time spent near infected patients. This kind of a priori knowledge would seem to be less likely to exist in the more complex or uncertain situations that many regulators address, where the relevant causal relationships do not stem from processes as law-like as biological disease transmission. If the system being modeled is extremely complex—as with many forms of regulation, whether of complex financial instruments or advanced industrial operations[117]—the regulator may not know enough about the underlying causal architecture to generate simulated environmental data bearing any resemblance to real-world data.

Of course, despite these difficulties, the conditions for using embedded machine learning for rulemaking may still sometimes exist. In a comment letter to the SEC, for example, academic and business experts in agent-based modeling and financial markets have advocated the use of such models in regulating equity markets, arguing that algorithms in this context would be sufficiently specifiable.[118] Although we take no position on these specific claims, we raise

115. *See, e.g.*, Marek Laskowski, *A Prototype Agent Based Model and Machine Learning Hybrid System for Healthcare Decision Support*, *in* DIGITAL ADVANCES IN MEDICINE, E-HEALTH, AND COMMUNICATION TECHNOLOGIES 230, 231 (2013).

116. *Id.* at 235–36.

117. *See generally* CHARLES PERROW, NORMAL ACCIDENTS: LIVING WITH HIGH-RISK TECHNOLOGIES (2d ed. 1999).

118. W. Brian Arthur et al., Comment Letter to Elizabeth M. Murphy on File Number S7-02010 “Concept Release on Equity Market Structure” (Apr. 16, 2010), https://www.sec.gov/comments/s7-02-10/s70210-109.pdf [https://perma.cc/GTE3-W2M8].

them to suggest the plausibility of using embedded machine learning to automate the process of selecting and designing regulations in some settings. Agency officials will need to determine the applicability of any embedded machine-learning rulemaking tool on a case-by-case basis.

Our point is to show that, even if many applications of machine learning will be completely benign as a matter of administrative law, agencies may soon be able, for the first time, to set the content of certain types of rules by automated artificial intelligence techniques. Whether in making individualized forecasts or in feeding into more generalized modeling results, machine-learning algorithms have the potential to transform key governmental functions in ways that not only augment human judgment but replace it with automated, algorithmic analysis. For some observers, this prospect will trigger loud alarm bells of the kind set off by the use of artificial intelligence more generally. At a minimum, the prospect of either robotic rulemaking or algorithmic adjudication raises important questions about whether such automated techniques can be squared with core principles of constitutional and administrative law.