

FOILING CLEVER HANS: STATE-LEVEL DATA PRIVACY PROTECTION AS A MITIGANT FOR AI HARMS

By: Violet Brull*

I. INTRODUCTION

Clever Hans could do arithmetic.¹ Wilhelm von Osten had been touring with Clever Hans around Germany since 1891 demonstrating this remarkable capability, and the German scientific community was reeling with the implications of Clever Hans's abilities.² Because, apparently, Clever Hans could do *arithmetic*.³ Clever Hans was a *horse* who could do *arithmetic*.⁴

The act itself was simple enough: A questioner would present Clever Hans with a series of questions ranging from numeral identification to simple multiplication and division.⁵ The horse would respond by dutifully tapping out the answer with one hoof.⁶

"Hans, what is the product of 2 and 6?"

Twelve taps.

"Say, Hans, what number is this?"

Eight taps.

"Excellent, Hans! Here, have a sugar cube." And so on.⁷

As best anyone in the scientific community could tell, it wasn't a fraud.⁸ Professional horse trainers had verified that if anyone was giving the horse a signal, it was not one that they had ever seen used to train horses before.⁹ Multiple scientific and skeptical investigators verified that Hans could perform

* J.D. Candidate, May 2025, University of Kansas School of Law. I would like to thank my close friend Eleazar Hazel for providing practical professional insight and stoking my interest in the legal aspects of this topic, my faculty advisor Professor Najarian R. Peters for providing invaluable feedback and research assistance, and all my friends and family who supported me throughout the writing and publication process.

¹ Philip M. Ferguson, *Clever Hans*, ENCYCLOPEDIA BRITANNICA, <https://www.britannica.com/topic/Clever-Hans> [<https://perma.cc/JD3R-Z4VB>].

² *Id.*

³ OSKAR PFUNGST, *CLEVER HANS (THE HORSE OF MR. VON OSTEN) A CONTRIBUTION TO EXPERIMENTAL ANIMAL AND HUMAN PSYCHOLOGY 1* (Carl L. Rahn trans., 1911).

⁴ *Id.*

⁵ *Id.* at 20, 34–35.

⁶ *Id.* at 19.

⁷ An imagined interaction.

⁸ PFUNGST, *supra* note 3, at 1 ("A horse that solves correctly problems in multiplication and division by means of tapping. Persons of unimpeachable honor, who in the master's absence have received responses, and assure us that in the process they have not made even the slightest sign. Thousands of spectators, horse-fanciers, trick-trainers of first rank, and not one of them during the course of many months' observations are able to discover any kind of regular signal. That was the riddle.").

⁹ *Id.* at 6–7.

his work even in the absence of his owner, or, indeed, anyone who they thought could be giving any kind of signal.¹⁰ What could this mean for the field of animal intelligence? For the study of education? For our entire understanding of humanity and its place in the world?¹¹

It wasn't until Oskar Pfungst, a student at the Psychological Institute at the University of Berlin, performed several carefully designed experiments that the baffling puzzle was solved.¹² In a series of controlled trials, Pfungst was able to ascertain that when no one present, *including the questioner*, knew the answer to the question presented, Clever Hans was utterly unable to perform.¹³ The horse had been reading the subtlest of body language from his questioners to determine the exact moment he should stop tapping his hoof for a reward.¹⁴

In today's world, we face a deluge of increasingly "clever" AI¹⁵ tools.¹⁶ Unfortunately, these tools, just like Clever Hans, can appear to make perfectly intelligent decisions while in fact relying on observations that have nothing to do with the actual tasks we assign them.¹⁷ This presents a distinct danger as these tools are increasingly involved in critical decisions about employment and consumer lending: Bias can creep into the decision-making process along with the vast amounts of data used to power these new tools.¹⁸ People can be unfairly denied loans or employment through no fault of their own.¹⁹ At particular risk

¹⁰ PFUNGST, *supra* note 3, at 2.

¹¹ *See id.* at 15–19.

¹² Ferguson, *supra* note 1; *see generally* PFUNGST, *supra* note 3.

¹³ *See, e.g.*, PFUNGST *supra* note 3, at 34–35.

¹⁴ Ferguson, *supra* note 1.

¹⁵ The acronym "AI" stands for "artificial intelligence."

¹⁶ Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge & Felix A. Wichmann, *Shortcut Learning in Deep Neural Networks*, 2 NATURE MACH. INTEL. 665, 1–2 (Nov. 21, 2023) (accessed via arXiv; this edition is paginated 1–29), <https://arxiv.org/pdf/2004.07780.pdf> [<https://perma.cc/WW4J-G78L>] [hereinafter *Shortcuts*].

¹⁷ *Id.* at 2.

¹⁸ *See, e.g.*, Ruha Benjamin, *Assessing risk, automating racism: A health care algorithm reflects underlying racial bias in society*, SCIENCE, Oct. 25, 2019, at 421, 421, <https://www.science.org/doi/10.1126/science.aaz3873> [<https://perma.cc/4UQU-4EFY>].

¹⁹ Fisher Phillips, *How is HR Using AI? An Employer's List of Tools and Potential Pitfalls*, JD SUPRA NEWSTEX BLOGS (Aug. 14, 2023), <https://www.jdsupra.com/legalnews/how-is-hr-using-ai-an-employer-s-list-5615442/> [<https://perma.cc/PJM8-5XFD>]; Kali Bracey & Grace Wallack, *New AI Lending Tech Could Exacerbate Old Bias Risks*, LEXISNEXIS LAW360 EXPERT ANALYSIS (Aug. 17, 2023), <https://plus.lexis.com/api/permalink/21c39dca-7d4c-4cc6-9647-d12726b82fcb/?context=1530671> [<https://perma.cc/58J8-USZY>]; Rohit Chopra, *Algorithms, artificial intelligence, and fairness in home appraisals*, Consumer Fin. Prot. Bureau Blog (June 1, 2023), <https://www.consumerfinance.gov/about-us/blog/algorithms-artificial-intelligence-fairness-in-home-appraisals/> [<https://perma.cc/4Z4U-DMEN>]; Charles Lane, *Will Using Artificial Intelligence To Make Loans Trade One Kind Of Bias For Another?*, NPR (Mar. 31, 2017), <https://www.npr.org/sections/alltechconsidered/2017/03/31/521946210/will-using-artificial-intelligence-to-make-loans-trade-one-kind-of-bias-for-anot> [<https://perma.cc/V9R4-9ZC8>]; F.T.C., *Big Data: A Tool for Inclusion or Exclusion?* (Jan. 2016), <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf> [<https://perma.cc/4A96-BV5F>] [hereinafter *FTC Big Data*].

are historically disadvantaged groups, whom these tools can subject to the inherited prejudices of past decision makers.²⁰

Current legislative approaches to solving this problem are inadequate.²¹ Scholars and policy makers often focus on seeking out and preventing concealed, nefarious intentions or human bias and error in the creation of algorithmic tools, advocating for and effecting legislation that polices procedural fairness in the use of big data.²² These efforts to police the *use* of consumer data are important, but they neglect to address the possibility that the data *itself* can be the vehicle for institutional discrimination.²³ The current body of law, then, is beneficial and provides an important backstop of protection, but is not sufficient.

To provide the supplement necessary to materially protect the rights currently espoused in modern legislation, this Article advocates for the increased protection of consumer data privacy at the state level, particularly in Kansas. When data is the vehicle for discrimination and consumer harm, it is only sensible to address the problem at the level of data collection. Further, this is an approach that can be effectively pursued at the state level. State-level implementation of comprehensive data privacy protection would bypass the sluggishness of the national legislative process. It may also garner greater bipartisan support than would similar legislation introduced at a federal level.

While many sources discuss the potential harms of mass consumer data collection²⁴ or of irresponsible AI use,²⁵ few directly propose data privacy legislation as a mitigant for AI harms.²⁶ Furthermore, most sources discussing desirable features for newly drafted data privacy legislation focus on federal—rather than state-level solutions.²⁷ This Article, by contrast, advocates a pragmatic state-level approach to data privacy protection, espousing its value in specific light of modern AI developments. This Article also signposts important

²⁰ Solon Barocas & Andrew D. Selbst, *Big Data's Disparate Impact*, 104 CALIF. L. REV. 671, 673–74 (2016) [hereinafter *Disparate Impact*].

²¹ *Id.* at 674–75.

²² *Id.*

²³ *Id.* at 673–74.

²⁴ See, e.g., *FTC Big Data* *supra* note 19; *Disparate Impact*, *supra* note 20.

²⁵ See, e.g., Benjamin, *supra* note 18.

²⁶ Karl Manheim and Lyric Kaplan's *Artificial Intelligence: Risks to Privacy and Democracy* does take this stance, but focuses solely on federal implementation and cannot fully contemplate the burgeoning capabilities of AI brought by the passage of the half-decade since its publication. Karl Manheim & Lyric Kaplan, *Artificial Intelligence: Risks to Privacy and Democracy*, 21 YALE J. L. & TECH. 106 (2019) [hereinafter *AI Risks*].

²⁷ See, e.g., *AI Risks*, *supra* note 26; Nuala O'Connor, *Reforming the U.S. Approach to Data Protection and Privacy*, COUNCIL ON FOREIGN RELS. (Jan. 30, 2018), <https://www.cfr.org/report/reforming-us-approach-data-protection> [https://perma.cc/F32J-J7DR]; Jessica Rich, *After 20 years of debate, it's time for Congress to finally pass a baseline privacy law*, BROOKINGS: COMMENTARY (Jan. 14, 2021), <https://www.brookings.edu/articles/after-20-years-of-debate-its-time-for-congress-to-finally-pass-a-baseline-privacy-law/> [https://perma.cc/U3JK-53LZ].

features that legislators should incorporate into any new state-level data privacy protection legislation.

Part II of this Article will lay the basic technical foundation necessary to understand the issues surrounding AI technologies covered by the remainder of the Article, as well as a similar foundation around modern “big data” collection and use practices. Part III will discuss how institutional discrimination can be perpetuated by these practices. Part IV will propose the solution of state-level consumer data protection as a strong mitigant, noting the inadequacy of current legislation to address the problems at hand. Part V will address several potential counterarguments.

II. BACKGROUND

An understanding of the field of machine learning is critical when attempting to solve the novel problems presented by today’s AI tools.²⁸ If one ignores the nuances involved in understanding the broad field of machine learning, one risks developing solutions for the problems it raises that fail to adequately address all cases.²⁹ Section A of this Part will thus begin under the broad umbrella of AI and progress in detail to a brief explanation of deep learning models and the datasets used to train them.

Since this Article advocates for increased consumer data privacy protections, a general understanding of the field of big data analytics is also necessary. Section B of this Part will provide information on the field, its history, and the problematic practices of many of its companies.

A. Machine Learning & Artificial Intelligence

The term “artificial intelligence” has two main understandings.³⁰ The first is “artificial general intelligence,” a term which describes systems that endeavor to make intelligent decisions on their own in a broad field of applications.³¹ The second is “intelligence augmentation,” a term which describes systems that endeavor only to aid humans in making complex decisions.³² This Article will not strongly distinguish between the two, since the border is nebulous and has a tendency to shift over time as intelligence augmentation tools become more and more capable.³³ It is important to recognize, however, that the field of AI is extremely broad and includes a plethora of systems produced by many different methods.

²⁸ David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 669 (2017) [hereinafter *Playing with the Data*].

²⁹ *Id.*

³⁰ KEVIN P. MURPHY, *PROBABILISTIC MACHINE LEARNING: AN INTRODUCTION* 28 (2022), probl.ai.

³¹ *Id.*

³² *Id.* at 29.

³³ *See id.*

Machine learning, one such method, is a process by which a computer program, given some curated experience, improves at a task over time.³⁴ In other words, the program “learns” through “training.”³⁵ Many different specific techniques fall under the definition of “machine learning,” and each can be used in a variety of practical applications.³⁶ Recently, the field of AI has become increasingly reliant on machine learning.³⁷ While it was long-thought that the path forward in developing AI was in hand-coding the “intelligent” systems, the difficulty of actually encoding all the knowledge such systems need in order to be useful has necessitated the adoption of machine learning approaches to allow systems to acquire the requisite knowledge for themselves.³⁸

One specific machine learning technique, deep learning, has largely powered the recent explosion in AI technologies.³⁹ Deep learning is a type of machine learning that uses artificial neural networks, with structures that emulate the ways in which human and animal brains function, as the basis for the training process.⁴⁰ A deep learning “model” or “algorithm” is an individual instance of the process: a set of “rules” that are learned over time through training on a dataset.⁴¹

A dataset is a collection of input-output pairs that can be used to facilitate machine learning.⁴² For example, one dataset might contain pairs where the input is an audio clip and the output is the transcribed text of that clip.⁴³ Such a dataset could be used in training an AI model to transcribe audio.⁴⁴ While machine learning (and specifically deep learning) has been the vehicle for recent AI developments, the availability of increasingly large datasets has been the fuel.⁴⁵ In previous eras, the collection of these datasets was seen as the limiting factor in the field’s development.⁴⁶ Now, these newly-available large

³⁴ MURPHY, *supra* note 30, at 1.

³⁵ *See id.*

³⁶ *Id.* at 27–28.

³⁷ *Id.* at 28.

³⁸ *Id.*

³⁹ *Shortcuts*, *supra* note 16, at 1.

⁴⁰ *What is AI?*, MCKINSEY & CO. FEATURED INSIGHTS (Apr. 24, 2023), <https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-ai> [<https://perma.cc/3F7T-XP4G>].

⁴¹ *Playing with the Data*, *supra* note 28, at 671–72.

⁴² Amandalynne Paullada, Inioluwa Deborah Raji, Emily M. Bender, Emily Denton, and Alex Hanna, *Data and its (dis)contents: A survey of dataset development and use in machine learning research*, 2 PATTERNS 1, 2 (2021) [hereinafter *Data and its (dis)contents*].

⁴³ *Id.*

⁴⁴ *Id.*

⁴⁵ *Id.* at 1.

⁴⁶ *Id.*

datasets, commonly referred to as “big data,” provide the foundation needed to support the wave of growth in machine-learning technologies.⁴⁷

B. Big Data

Big data doesn't start big.⁴⁸ The collection process begins, frequently, on consumer devices.⁴⁹ As users interact and shop online, online advertisers, retailers, social media companies, and others collect and compile information about their habits and history.⁵⁰ Companies known as “data brokers” specialize in this collection and compilation process.⁵¹ Data brokers gather consumer information en masse, using it to build extensive profiles—each containing thousands of datapoints—on nearly every U.S. consumer.⁵²

Sometimes, this collection process is carried out in a relatively transparent manner.⁵³ When, for example, a consumer logs into an online retailer's website before making a purchase, they may reasonably expect that the purchase transaction will be associated with their account, and in turn with the consumer themselves.⁵⁴ Sometimes, however, data brokers rely on more underhanded tactics to obtain consumer data without the consumer's permission or even knowledge.⁵⁵ Data brokers and others may use tracking cookies,⁵⁶ browser or device fingerprinting,⁵⁷ and even history sniffing⁵⁸ to surreptitiously

⁴⁷ *Data and its (dis)contents*, *supra* note 42, at 1.

⁴⁸ *FTC Big Data*, *supra* note 19, at 3.

⁴⁹ *Id.* at 3–4.

⁵⁰ *Id.*

⁵¹ *Id.* at 4.

⁵² *Id.*; F.T.C., *Data Brokers: A Call for Transparency and Accountability*, 46–47 (May 2014), <https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf> [<https://perma.cc/WXY7-REM4>] [hereinafter *FTC Data Brokers*].

⁵³ *See FTC Big Data*, *supra* note 19, at 3.

⁵⁴ *See id.*

⁵⁵ *FTC Data Brokers*, *supra* note 52, at 46; *see FTC Big Data*, *supra* note 19, at 3.

⁵⁶ Cookies are small text files stored on user devices for a variety of purposes, such as maintaining user preferences or remembering the contents of user shopping carts. Tracking cookies are cookies that keep track of which webpages a user has visited and potentially communicate that information to a third party, frequently without the user having any visibility or control over the process. *What are cookies?*, CLOUDFLARE: LEARNING CENTER, <https://www.cloudflare.com/learning/privacy/what-are-cookies/> [<https://perma.cc/RXC6-X9EH>].

⁵⁷ Fingerprinting methods use characteristics of a user's device (such as which type of graphics card or CPU is installed in the device) to identify users even when they avoid identification via tracking cookie. *See* Yinzhi Cao, Song Li, & Erik Wijmans, *(Cross-)Browser Fingerprinting via OS and Hardware Level Features*, NETWORK AND DISTRIBUTED SYSTEM SECURITY (NDSS) SYMPOSIUM (2017) [<https://perma.cc/69MG-D72M>].

⁵⁸ History sniffing is the practice of obtaining information about a user's browser history without the consent or knowledge of the user. Ioana Patrinenaru, *These New Techniques Expose Your Browsing History to Hackers*, U.C. SAN DIEGO TODAY (Oct. 31, 2018), https://today.ucsd.edu/story/history_sniffing [<https://perma.cc/CN7K-K26M>]. While web browsers try to prevent history sniffing, new techniques are being developed over time that maintain the effectiveness of this category of attacks. *Id.*; *History sniffing with CSS and JS: Learn how and why it still works in 2022.*, HACKINGLOOPS, <https://www.hackingloops.com/css-browser-history-sniffing/> [<https://perma.cc/P98X-B4JF>].

connect data about a given consumer from disparate sources.⁵⁹ In some cases, companies can gain even more information about a consumer by tracking the consumer's identity across multiple devices (e.g., the consumer's phone, laptop, smart home devices, and wearables) and combining the data received from each source.⁶⁰ Data brokers can also use the above techniques to pair information gained about a consumer from offline sources, such as voting registration, bankruptcy information, or even product warranty registrations, with information about that consumer's online activity, such as their online purchase history, advertisement interaction, and social media activity.⁶¹

Once data brokers have collected consumer data, they can combine and analyze the various data they collect to make potentially sensitive inferences about consumers.⁶² Data brokers place consumers into categories based on certain inferred demographic characteristics, which are then marketed to advertisers as potential segments of new customers.⁶³ While some of these categories are relatively innocuous (e.g., "Dog Owner," "Winter Activity Enthusiast," or "Mail Order Responder"), some of the categories are far more concerning, implicating the ethnicity, income level, marital status, education level, and medical concerns of the segmented consumers.⁶⁴ For example, categories that have been used by data brokers for this purpose include "Urban Scramble" and "Mobile Mixers," both of which disproportionately contain people in minoritized ethnic groups with low income, "Rural Everlasting," a category including "single men and women over the age of 66 with 'low educational attainment and low net worths,'" "Married Sophisticates," a category containing thirty- to forty-year-old individuals who are wealthy and married without children, and categories implicating medical concerns such as "Expectant Parent," "Diabetes Interest," and "Cholesterol Focus."⁶⁵

Data brokers may further offer "data append" services.⁶⁶ Clients of these services provide the data broker with identifying information about one or more consumers (e.g., names and addresses), and data brokers in turn provide additional information about those consumers back to the clients.⁶⁷ This information can include, for example, direct information about individual consumers' gender, occupation, religious or political affiliations, or even height and weight.⁶⁸

⁵⁹ *FTC Big Data*, *supra* note 19, at 3–4.

⁶⁰ *Id.* at 4.

⁶¹ *FTC Data Brokers*, *supra* note 52, at 46–47.

⁶² *Id.* at 47.

⁶³ *Id.* at 28, 47.

⁶⁴ *Id.* at 47.

⁶⁵ *Id.*

⁶⁶ *Id.* at 24.

⁶⁷ *Id.*

⁶⁸ *Id.* at 24–25.

Even prior to the dawn of AI ubiquity, unscrupulous companies could use big data to target vulnerable populations for exploitation.⁶⁹ Companies could potentially use big data analytics to find vulnerable populations to target with scams or misleading offers.⁷⁰ Online retailers have, in the past, used big data analytics to increase prices for customers from low-income communities, where they face less competition from brick-and-mortar stores.⁷¹ When companies collect and analyze large bodies of data for the purpose of making predictive decisions, the process carries the inherent potential to cause adverse outcomes for entire sociodemographic groups.⁷²

The lack of transparency surrounding data brokers' practices has only worsened the issue. Unlike consumer reporting agencies ("CRAs"), traditional sources of consumer information that are bound by the Fair Credit Reporting Act (the "FCRA"), most data brokers are not required to provide consumer access to the data they collect about individual consumers.⁷³ Even those data brokers that do provide consumer access often restrict that access to a consumer's raw data alone, rather than the results of any analysis performed by the broker.⁷⁴ This means that a data broker could file a consumer into a category that implies sensitive information about that consumer (which may or may not be true) and provide that categorization to a client company, all without the consumer having any way to know this was occurring.⁷⁵

With the proliferation of internet-connected consumer devices comes the proliferation of opportunities for companies to collect data on consumers.⁷⁶ With the proliferation of that opportunity comes the proliferation of commercially available consumer data.⁷⁷ Ultimately, access to broad swathes of consumer data gives companies powerful opportunities to facilitate either inclusion or exclusion—to advance the interests of historically disadvantaged populations or to exploit the same.⁷⁸ In a perfect world, companies would solely use consumer data in ethical ways, advancing the common good and protecting vulnerable communities. Reality, unfortunately, is far from perfect.

III. BIG DATA, AI, AND THE PERPETUATION OF INSTITUTIONAL DISCRIMINATION

In today's world, machine learning technologies are used to make myriad decisions about individual people, ranging from the relatively

⁶⁹ *FTC Big Data*, *supra* note 19, at 10–11.

⁷⁰ *Id.* at 10.

⁷¹ *Id.* at 11.

⁷² *Disparate Impact*, *supra* note 20, at 673.

⁷³ *FTC Big Data*, *supra* note 19, at ii; *FTC Data Brokers* *supra* note 52, at 42.

⁷⁴ *FTC Data Brokers*, *supra* note 52, at 42.

⁷⁵ *See id.*

⁷⁶ *FTC Big Data*, *supra* note at 19, at i.

⁷⁷ *Id.*

⁷⁸ *Id.* at 12.

inconsequential to the highly impactful.⁷⁹ For each machine-learning model employed, the journey from data to trained model is fraught with numerous complexities and obstacles, many of which provide opportunities for the inadvertent injection of harmful biases that can be reflected in the final product.⁸⁰ Section A of this Part will examine the difficulties of constructing an equitable dataset for machine learning use. Section B of this Part will discuss the impracticability of remedying bias in most existing datasets. Section C of this Part will highlight the pitfalls in model training and use that could result in the derivation of bias, even from the perfectly vetted dataset.

A. Obstacles to Equitable Dataset Construction

The ill-considered development and use of large datasets can lead to troubling societal impacts, despite apparent progress.⁸¹ Even industry professionals admit that using massive-scale analytics to attempt to find useful but non-obvious patterns in enormous consumer datasets is a relatively recent practice, and that companies are still learning how to avoid the potential ill consequences it could cause.⁸² If the collection and compilation process for a body of data reflects bias towards or against certain sociodemographic groups, that bias can bleed through into the statistical relationships supposedly discovered through the analysis of that data.⁸³

The introduction of machine-learning technologies into the already-problematic space of big data analytics has only exacerbated the issue. Among many professionals in the machine learning field, the unrestricted distribution of datasets is seen as an unequivocal good.⁸⁴ Scholars in the field have argued that such distribution is necessary for the review and verification of new discoveries in machine learning methodology.⁸⁵ When data gathered for one purpose is

⁷⁹ MICHAEL KEARNS & AARON ROTH, THE ETHICAL ALGORITHM: THE SCIENCE OF SOCIALLY AWARE ALGORITHM DESIGN 64 (2019) [hereinafter THE ETHICAL ALGORITHM]; Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 81 PROC. MACH. LEARNING RSCH. 1, 1 (2018) [<https://perma.cc/VRA6-JBWJ>] [hereinafter *Gender Shades*]; Sylvia Lu, *Data Privacy, Human Rights, and Algorithmic Opacity*, 110 CALIF. L. REV. 2087, 2090–91 (2022).

⁸⁰ See *Playing with the Data*, *supra* note 28, at 681–702.

⁸¹ Emily M. Bender, Timnit Gebru, Angelina McMillan-Major & Shmargaret Shmitchell, *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, in FACCT '21: PROCEEDINGS OF THE 2021 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 610, 610 (Mar. 2021), <https://dl.acm.org/doi/pdf/10.1145/3442188.3445922> [<https://perma.cc/4YT5-W2S2>] [hereinafter *Stochastic Parrots*.]; *Data and its (dis)contents*, *supra* note 42, at 1.

⁸² *FTC Big Data*, *supra* note 19, at 5.

⁸³ *Id.* at 8.

⁸⁴ *Data and its (dis)contents*, *supra* note 42, at 7.

⁸⁵ *Id.*

reused for a different purpose, however, it can raise myriad ethical concerns and cause a variety of social harms.⁸⁶

A concerning case study illustrating this point is the story of the Pima Indians Diabetes Dataset (the “PIDD”), housed at the University of California, Irvine Machine Learning Repository.⁸⁷ This dataset was originally collected by the National Institutes of Health from the Indigenous community living at the Gila River Indian Community Reservation, whose members refer to themselves as Akimel O’odham.⁸⁸ The Akimel O’odham people have historically been the subject of numerous anthropological and biomedical studies, including the longitudinal study that produced the PIDD, due to the unusually high prevalence of diabetes among their community.⁸⁹ While these studies have advanced medical understanding of diabetes as a disease, they have not resulted in any substantial decrease in the rates of obesity or diabetes in the community itself.⁹⁰

Furthermore, the PIDD has been used thousands of times in the development of machine learning algorithms as a “toy” dataset.⁹¹ In these myriad cases it has been used not to further studies of diabetes, or even human health in general, but rather only as fodder for machine learning development.⁹² This use, normalized among members of the machine learning field, raises concerns over the reproduction of the exploitative patterns of colonialism.⁹³ It further calls the flawed ethical practices surrounding the collection and preservation of Henrietta Lacks’s cervical cells⁹⁴ uncomfortably to mind. Concerns like these have caused some scientists to call for more rigorous ethical norms, such as those of human-subjects research, to be applied to the burgeoning field of data science.⁹⁵ In traditional human-subjects research, institutional review boards and informed consent requirements help to protect vulnerable populations from exploitation.⁹⁶ In the field of machine learning, these critical protections remain absent.⁹⁷

Even when data analysts carefully purpose-build their datasets, though, rather than using inappropriately acquired data created for an incompatible

⁸⁶ *Data and its (dis)contents*, *supra* note 42, at 6–8.

⁸⁷ *Id.* at 7.

⁸⁸ *Id.* at 7; Joanna Radin, “Digital Natives”: How Medical and Indigenous Histories Matter for Big Data, 32 OSIRIS 43, 44 (Jan. 2017), <https://www.journals.uchicago.edu/doi/full/10.1086/693853> [<https://perma.cc/C8N9-Y3SS>].

⁸⁹ Radin, *supra* note 88, at 44; *Data and its (dis)contents*, *supra* note 42, at 7.

⁹⁰ Radin, *supra* note 88, at 50; *Data and its (dis)contents*, *supra* note 42, at 7.

⁹¹ *Data and its (dis)contents*, *supra* note 42, at 7.

⁹² *Id.*

⁹³ Radin, *supra* note 88, at 45.

⁹⁴ Henrietta Lacks was treated for cervical cancer at Johns Hopkins Hospital in 1951. Lacks passed away shortly afterward, but cells from a tumor biopsied during this treatment were cultured into the HeLa cell line, a cell line still widely used in medical research to this day. Lacks’s consent to use her cells for medical research was neither sought nor given, and neither Lacks nor her family were ever compensated for the cells. Indeed, Lacks’s family was not even made aware of the cell line until decades after Lacks’s death. *See generally* REBECCA SKLOOT, *THE IMMORTAL LIFE OF HENRIETTA LACKS* (2010) (providing background on Henrietta Lacks’s history).

⁹⁵ *Data and its (dis)contents*, *supra* note 42, at 6–7.

⁹⁶ *Id.*

⁹⁷ *Id.* at 7.

purpose, the resultant datasets can be fraught with flaws introduced during their creation. In recent years, researchers have increasingly found that many prominent machine learning datasets contain either a troubling lack of representation or outright harmful misrepresentation of certain sociodemographic groups.⁹⁸ To give a few examples: People with darker skin-tones are underrepresented in datasets used in facial recognition, and, worryingly, those used in training self-driving cars to recognize pedestrians.⁹⁹ Female pronouns and female-coded names are underrepresented in several datasets used to train AI models for natural language processing (“NLP”).¹⁰⁰ Datasets used to develop NLP models also frequently reflect social biases and stereotypes around race, gender, disability and more, such as a dataset being used to train a model to detect “toxic” text that disproportionately associated words used to describe queer identities with toxicity.¹⁰¹

These types of bias can come from several sources. The most obvious and direct of these is the actual conscious or unconscious prejudice of dataset creators, which can be veiled and abstracted by the complexities of the analytics process.¹⁰² This is far from the only possible source of bias, however. Inherent barriers to the collection of certain types of data can prevent the population sampled during the creation of the dataset from ever truly aligning with the population to whom the resulting algorithm will be applied.¹⁰³

For example, a lending company looking to identify loan candidates who are likely to default on their loans might sample the population of previous loan recipients to see which of them had defaulted.¹⁰⁴ But the population sampled, namely “everyone who has received a loan in the past,” does not align with the population to whom the trained model will presumably be applied, namely “everyone who applies to receive a loan, whether or not that loan is granted.”¹⁰⁵ This type of inherent barrier to accurate data sampling can cause disproportionate harms to historically disadvantaged populations.¹⁰⁶ The

⁹⁸ Stochastic Parrots, *supra* note 81, at 613; *Data and its (dis)contents*, *supra* note 42, at 3.

⁹⁹ *Gender Shades*, *supra* note 79, at 2–3; *Data and its (dis)contents*, *supra* note 42, at 3.

¹⁰⁰ *Data and its (dis)contents*, *supra* note 42, at 3. The phrase “natural language processing” refers to the useful processing of human language by computers, for purposes such as engaging in conversation with human users or translating text “intelligently” from one human language to another. See DANIEL JURAFSKY & JAMES H. MARTIN, *SPEECH AND LANGUAGE PROCESSING: AN INTRODUCTION TO NATURAL LANGUAGE PROCESSING, COMPUTATIONAL LINGUISTICS, AND SPEECH RECOGNITION* 1 (2d ed. 2008).

¹⁰¹ *Id.*; Stochastic Parrots, *supra* note 81, at 614–15; see generally TOLGA BOLUKBASI, KAI-WEI CHANG, JAMES ZOU, VENKATESH SALIGRAMA & ADAM KALAI, *MAN IS TO COMPUTER PROGRAMMER AS WOMAN IS TO HOMEMAKER? DEBIASING WORD EMBEDDINGS* (2016) [<https://perma.cc/S86P-CR88>] (describing a famous example of this problem).

¹⁰² *Disparate Impact*, *supra* note 20, at 675–76.

¹⁰³ *Playing with the Data*, *supra* note 28, at 680.

¹⁰⁴ *Id.*

¹⁰⁵ *Id.*

¹⁰⁶ *Id.* at 680–81.

hypothetical model based on data from “everyone who has received a loan in the past” will inevitably perform worse for populations about which it has less data—populations who have been disproportionately denied loans in the past.¹⁰⁷

One demonstrable inherent barrier to equitable data collection arises in the creation of large text datasets from online sources.¹⁰⁸ Not only are people from wealthier countries overrepresented on the internet generally, young adult men are the predominant authors of the specific bodies of online text frequently used to create these datasets, such as the user-written content of sites like Reddit, Wikipedia, and X (formerly known as Twitter).¹⁰⁹ When equitable data collection is so hindered by the structural problems of these frequently used sources, it is no wonder at all that the resulting datasets demonstrate the flaws they do.¹¹⁰

The actual task of creating a dataset presents an additional technical challenge.¹¹¹ A data scientist must not only decide which input data to measure and how to measure it, but also, much of the time, how to translate this slew of complex information into a concrete and specific outcome variable (called the “target variable”) that the computerized model trained with the dataset will predict.¹¹²

Imagine a straightforward, if frivolous, example: A data scientist wishes to design a dataset for use in training a model to recognize whether there is a horse in a given image. In this case, the input data will be a large set of images, each of which may or may not contain a horse. The target variable will be a simple “true” or “false” for each image, indicating whether or not the image contains a horse.

Consider the decisions the dataset’s designer needs to make, even in this relatively simple hypothetical. First, the designer would have to answer several questions about which input data to measure, and how to measure it. For example, how should the training images be digitally encoded? Should the dataset include the color of each pixel of the image? The brightness? Should greyscale images be allowed in the dataset, or will the dataset be restricted to only color images? How should it handle images of different sizes? Should the dataset even include images of different sizes? Where will the designer get that large body of images (some of which contain horses) in the first place?

Having answered these questions, the designer might breathe a sigh of relief, thinking the hardest part of their work behind them. After all, once these difficult technical questions have been answered, all that remains is to hire some plucky intern to slog through every picture in the dataset and indicate whether that picture does, in fact, contain a horse.

¹⁰⁷ *Playing with the Data*, *supra* note 28, at 680–81.

¹⁰⁸ See *Stochastic Parrots*, *supra* note 81, at 613.

¹⁰⁹ *Id.* This is due, at least in part, to structural issues such as poor moderation that can expose many would-be authors with diverse voices to severe online harassment or even result in the victims of harassment being ousted from the platform rather than the offenders. *Id.*

¹¹⁰ See *id.*

¹¹¹ *Playing with the Data*, *supra* note 28, at 668.

¹¹² *Id.*

To the designer's dismay, however, the intrepid intern reports in with a barrage of new questions about what does or does not constitute a horse's presence in an image. If an image contains only a horse's head—the remainder of the horse being out of frame—does it contain a horse? What if it contains only the horse's tail? The horse's legs? How about half of a horse? Sixty percent of a horse? Forty? Half of one horse and half of another? What if the image is of a *painting* of a horse? What if the image contains, not a full-grown horse, but a foal? What if it contains a zebra? A donkey?

When data scientists answer questions like those raised by our hypothetical intern, they make tacit assumptions about which measurable facts can stand in for which underlying concepts.¹¹³ The intern's questions force our hypothetical designer to consider, in excruciating detail, what it means for an image to “have a horse in it.” Our weary scientist now faces the challenge of accurately defining the dataset's target variable.

Defining a target variable is a deeply subjective task, since it involves taking an often-hazy real-world concept and translating it into hard numerical values based on available data.¹¹⁴ It is so subjective, in fact, that it is often referred to as the “art” in data analysis.¹¹⁵ To be successful in this task, one must have a deep understanding of the underlying subject matter.¹¹⁶ One could imagine, for example, how much more difficult our hypothetical dataset designer's job would be if, instead of ordinary pictures which may or may not contain horses, the dataset contained microscopic images which may or may not contain a particular type of bacteria, or architectural schematics which may or may not contain a particular design element. Even when an analyst possesses the requisite understanding and operates with the utmost care in creating a target definition, though, it is still possible for the analyst's personal prejudices or simple mistakes to introduce unfair bias into the dataset.¹¹⁷

Sometimes, the process of defining a target variable is further complicated by the fact that the supposed underlying concept has no concrete basis in reality.¹¹⁸ When the underlying question the model is trying to predict an answer for has a concrete basis (e.g., “Does this picture contain a horse?”), the process of defining the question is difficult enough.¹¹⁹ When the question itself hinges on an arbitrary and abstract definition (e.g. “How creditworthy is this person?”), the process of defining it accurately becomes nigh-impossible.¹²⁰ It is, again, possible for bias to creep into the dataset at this point; some

¹¹³ See *Playing with the Data*, *supra* note 28, at 674-75.

¹¹⁴ *Disparate Impact*, *supra* note 20, at 678.

¹¹⁵ *Id.*

¹¹⁶ *Id.*

¹¹⁷ *Id.*

¹¹⁸ *Id.* at 679.

¹¹⁹ See *id.* at 678-79.

¹²⁰ *Id.* at 679.

definitions of several important questions have been shown to result in more adverse outcomes for certain protected classes than other definitions of the same questions.¹²¹ All told, while concerns about AI models' accuracy in answering important questions are absolutely legitimate, at least as legitimate are concerns about the ability of data scientists to accurately define the questions themselves.¹²²

Despite all of these complexities, industry professionals in the field of machine learning have not typically taken adequate care when assembling new datasets.¹²³ Rather than proceeding with meticulous caution when gathering data, as is the norm in other scientific fields, machine learning professionals have operated with "*laissez-faire*" sensibilities when it comes to the collection and curation of new datasets.¹²⁴ Current culture around dataset creation, collection, and use, possessed of a single-minded focus on rapid expansion of machine learning capabilities, flies in the face of scientific and ethical concerns about its ambivalence to the harms lack of care in dataset development can cause.¹²⁵

B. Obstacles to Dataset Remedy

Even though the problem of biased datasets is pervasive in the machine learning space, detection and subsequent correction of bias in an individual dataset can be extremely difficult. The principal factor leading to this difficulty is the size of the datasets in question: Machine learning datasets *have* to be big.¹²⁶ While there is no strict lower bound on the size of dataset which can effectively be used to train a model, having a number of input-output pairs that does not measure in at least the tens of thousands is frequently insufficient.¹²⁷

Modern datasets are frequently so large as to make manual review at best impracticable and at worst impossible.¹²⁸ Even the most sophisticated approaches to automating the process are prone to many of the same pitfalls that would be encountered in the use of the biased data.¹²⁹ In other words, when applied to the datasets where automated review is the most necessary, that review is often at its least effective.¹³⁰

Even when bias is detected in a dataset, remediation of that bias is not easy or even always possible. The twin problems of underrepresentation and stereotype-aligned misrepresentation tug in opposite directions.¹³¹ If one tries to combat under-inclusion of a certain population by artificially increasing that population's representation in the dataset, one runs the risk of magnifying the

¹²¹ *Disparate Impact*, *supra* note 20, at 680.

¹²² *Id.*

¹²³ *Data and its (dis)contents*, *supra* note 42, at 4.

¹²⁴ *Id.*

¹²⁵ *Id.* at 9.

¹²⁶ *Playing with the Data*, *supra* note 28, at 678–79; *Data and its (dis)contents*, *supra* note 42, at 5.

¹²⁷ *Playing with the Data*, *supra* note 28, at 678–79.

¹²⁸ *Data and its (dis)contents*, *supra* note 42, at 5.

¹²⁹ *Id.* at 5–6.

¹³⁰ *Id.*

¹³¹ *Id.* at 3.

harms associated with misrepresentation, and vice versa.¹³² Even more difficult to remediate are biases introduced to the data through past intentional discrimination.¹³³ In these cases, there is often no clear way to systematically remove bias from the data itself, and attempts to sway the results of an algorithmic decision after the fact can rely on practices that could spark both political and legal controversy.¹³⁴

In summation, not only is it incredibly difficult to design a large dataset in such a way as to avoid introducing bias, it is also incredibly difficult to cure a dataset of that bias if and when it is discovered.

C. Correlations Genuine and Spurious

Were the perfect dataset nonetheless created and maintained, despite the myriad labors involved, issues of bias could still arise in application due to that quintessential misunderstanding reinforced by the mystique surrounding machine learning technologies: equivocation of correlation with causation.

Correlation is *not* causation.¹³⁵ Fundamentally, machine learning techniques are automated ways to discover often-complex correlations in data.¹³⁶ If companies make decisions based on mere correlations uncovered using AI tools without understanding the underlying causal relationships (or lack thereof), those decisions can lead companies to do harm to consumers, other companies, and themselves.¹³⁷

Using an AI that makes decisions based on faulty correlations is analogous to asking Clever Hans whether a particular candidate would be a good hire. Clever Hans will give you an answer, but that answer will not in any way be based on the candidate's objective qualifications. While all machine learning techniques harbor the potential to allow users to make such faulty reliances, deep learning models can actively encourage this type of behavior in two important ways.

First, deep learning models are capable of learning unintended “shortcut” strategies for moving from input to output.¹³⁸ For example, a deep learning model trained to recognize and describe the contents of images (the selected input) with a string of text (the selected output) might accurately

¹³² *Data and its (dis)contents*, *supra* note 42, at 3.; *see also* Stochastic Parrots, *supra* note 81, at 613–14 (providing a more detailed example of this phenomenon).

¹³³ *Disparate Impact*, *supra* note 20, at 671.

¹³⁴ *Id.*

¹³⁵ *See generally* John Aldrich, *Correlations Genuine and Spurious in Pearson and Yule*, 10(4) STATISTICAL SCI. 364 (1995); *see also* Randall Monroe, *Correlation*, XKCD, <https://xkcd.com/552/> [<https://perma.cc/CT5U-H59F>].

¹³⁶ *Playing with the Data*, *supra* note 28, at 671.

¹³⁷ *FTC Big Data*, *supra* note 19, at 9.

¹³⁸ *Shortcuts*, *supra* note 16, at 2.

describe a picture of sheep grazing in a field.¹³⁹ However, when shown a picture of the same field without the sheep, it might respond with the same string of text, asserting that the sheep, though clearly absent to any human observer, are present and happily grazing away.¹⁴⁰ If it did, it would likely be because the model has learned the shortcut of associating the grassy field with the phrase, “sheep grazing in a field,” rather than identifying the sheep themselves.¹⁴¹ The same model may entirely fail to recognize sheep outside the context of a grassy field for them to graze in.¹⁴²

When deep learning models form these types of shortcuts, they often perform perfectly well within the limited training environment, failing in unexpected ways only when they are exposed to the complications of actual application.¹⁴³ When this failure results in an empty pasture being misidentified as containing grazing sheep, the ill consequences are minimal. In other situations, however, the possibility of error on this axis is cause for greater concern.¹⁴⁴

Second, deep learning models encourage reliance on faulty spurious correlations through the very nature of the process of employing such a model. In selecting or creating a dataset to use in training a deep learning model, one presupposes a connection between the input and output variables of their dataset.¹⁴⁵ This can cause issues, since it is entirely possible to construct a dataset where the input and output have no meaningful causal relationship, but which does contain spurious correlations that can allow a deep learning model to convincingly assert such a relationship despite its absence.¹⁴⁶

For example, one could interview ten thousand people, asking each to choose a random number between one and ten, and then asking each what they had eaten for dinner the previous evening. The interviewer could then connect the two categories of information as input-output pairs and even train a deep learning model on the resulting dataset, but that would not manifest any actual causal relationship between arbitrary culinary and numerical choices in general.

¹³⁹ *Shortcuts*, *supra* note 16, at 2.

¹⁴⁰ *Id.*

¹⁴¹ *Id.*

¹⁴² *Id.*

¹⁴³ *Data and its (dis)contents*, *supra* note 42, at 3.

¹⁴⁴ See Sam Levin, *Imprisoned by algorithms: the dark side of California ending cash bail*, THE GUARDIAN: US NEWS (Sep. 7, 2018), <https://www.theguardian.com/us-news/2018/sep/07/imprisoned-by-algorithms-the-dark-side-of-california-ending-cash-bail> [<https://perma.cc/G3J9-B9ZH>]; Julia Angwin, Jeff Larson, Surya Mattu & Lauren Kirchner, *Machine Bias*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> [<https://perma.cc/ES9G-LE58>] [hereinafter *Machine Bias*]; *Gender Shades*, *supra* note 79, at 1; Benjamin, *supra* note 18, at 421.

¹⁴⁵ See *Data and its (dis)contents*, *supra* note 42, at 4.

¹⁴⁶ *Id.*; see also Erika Andersen, *True Fact: The Lack of Pirates Is Causing Global Warming*, FORBES (Mar. 23, 2012), <https://www.forbes.com/sites/erikaandersen/2012/03/23/true-fact-the-lack-of-pirates-is-causing-global-warming/?sh=176670983a67> [<https://perma.cc/PW3B-47V6?type=image>]; see generally *Spurious Correlations*, TYLERVIGEN.COM, <https://tylervigen.com/spurious-correlations> [<https://perma.cc/4X8J-Z68T>] (illustrating several humorous examples of input-output pairs containing spurious correlations).

A deep learning model trained on the dataset, however, might find spurious flukes within the limited scope of the training data that allow it to outperform the baseline of a random guess when asked to determine someone's number choice based on their dinner.¹⁴⁷ Its performance would then falsely imply the existence of an actual connection between the two variables, and a careless researcher might begin to believe in the false implication.¹⁴⁸

The same logic applies across all possible input-output pairs: Just because an AI model discovers a correlation between two variables does not mean such a correlation actually exists. Just because one *can* construct a dataset with an input of online spending habits and an output of lending risk, or with an input of a resume and recorded interview and an output of predicted value to an employer, does not mean that those input-output pairs have any sort of actual connection, despite the fact that AI models using heuristics based on those supposed relationships can outperform baseline predictions within their dataset.

IV. THE NECESSITY OF DATA PRIVACY PROTECTION

Section A of this Part examines the inadequacy of current law to address the potential harms of new AI technologies. Section B proposes the specific solution of increased data privacy protection at the state level. Section C identifies and advocates key features of the ideal state-level data privacy protection legislation.

A. Inadequacy of Extant Law

Law established before the rise of big data analytics does not adequately protect against the systemic discrimination which can result from data mining.¹⁴⁹ Indeed, extant law does not adequately protect against discrimination in general.¹⁵⁰ Discrimination in American housing, employment, lending, and consumer markets is pervasive and persistent. It finds its roots both in the active prejudice of decision-makers and in the institutional momentum of systems that tend to passively punish historically disadvantaged groups.¹⁵¹ New approaches to decision-making based on big data analytics only exacerbate the problem with their potential to sidestep existing civil liberty protections.¹⁵²

Flaws in data which result in the reinforcement of extant societal biases or the prejudices of the data's creators are pervasive, and many of these flaws can only be remedied, if at all, *before* a trained model is put to use in the real

¹⁴⁷ See *Data and its (dis)contents* at 4.

¹⁴⁸ See *id.*

¹⁴⁹ *Disparate Impact*, *supra* note 20, at 675.

¹⁵⁰ See *id.* at 673–74.

¹⁵¹ *Id.* at 673–74.

¹⁵² *Id.* at 674.

world.¹⁵³ Decisions made prior to the release of a model are made less inscrutably and with far more human involvement than decisions reached by that model in practice.¹⁵⁴ Thus, the time before a model is released would be the ideal time for legislation aiming to control the harms of AI misuse to target.¹⁵⁵ Under current frameworks, though, any adjudication on the complex steps of the data mining process under major statutory consumer protections requires subjective and fact-bound judgments, dramatically reducing the possibility for effective enforcement.¹⁵⁶ Tort law is also ineffective at protecting consumers—the privacy protections offered under most states' (including Kansas's¹⁵⁷) common law do not effectively defend against the broad, systemic harms perpetuated by big data analytics.¹⁵⁸

The novelty of the problem and the inadequacy of current legal frameworks are not secrets. In October of 2022, the White House Office of Science and Technology Policy published *Blueprint for an AI Bill of Rights*, a document aimed at setting policy goals around consumer protection from the new threats posed by AI technology.¹⁵⁹ A year later, President Biden issued an executive order setting out administrative policies for this technology's responsible development and use, again recognizing threats to privacy, equity, and civil rights.¹⁶⁰

Our nation needs a new approach to defending our citizens against these threats. The fact of that need is demonstrated by the clear inadequacy of current law to address the proliferation of discriminatory harms brought about by the rise of AI and its misuse. The only remaining question, then, is what form that new approach should take.

B. Data Privacy Protection: A Strong Mitigant

The dawn of AI technologies driven by machine learning has dramatically increased the incentive to over-collect and overuse consumer data.¹⁶¹ Because of data's role in advancing machine learning technology, consumer data has immense value.¹⁶² In the modern economic and technological climate, companies are incentivized to pursue the collection of private and

¹⁵³ *Disparate Impact*, *supra* note 20, at 675.; *Playing with the Data*, *supra* note 28, at 655–57.

¹⁵⁴ *Playing with the Data*, *supra* note 28, at 657.

¹⁵⁵ *See id.*

¹⁵⁶ *Disparate Impact*, *supra* note 20, at 676.

¹⁵⁷ *See* *Munsell v. Ideal Food Stores*, 494 P.2d 1063, 1074-75 (Kan. 1972); Pattern Inst. Kan. Civil 127.61.

¹⁵⁸ *AI Risks*, *supra* note 26, at 121.

¹⁵⁹ WHITE HOUSE OFF. OF SCI. & TECH. POL'Y, *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People*, (2022), <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf> [<https://perma.cc/87S4-YVUY>].

¹⁶⁰ Exec. Order No. 14110, 88 Fed. Reg. 75191, §§ 2(d)-(f) (Oct. 30, 2023) [hereinafter *Exec. Order: Responsible Development and Use*].

¹⁶¹ *AI Risks*, *supra* note 26, at 121.

¹⁶² *Id.* at 119.

potentially sensitive data at an ever-increasing scale, even when legal and ethical boundaries must be pushed or crossed to do so.¹⁶³

Companies and government entities are then further enticed by the promises of AI to make ethically questionable decisions when using this ill-gotten data. Take, for example, the credit card company that, without informing its customers, began assigning higher credit risk to customers who used their cards to pay for marriage counseling, therapy, or tire-repair services.¹⁶⁴ Or, for a far more concerning illustration, look to the multiple algorithms actively used by criminal justice agencies that “learned” to falsely flag Black defendants as having a high risk of recidivism based on the intrinsically biased data they were given.¹⁶⁵

The solution, at least in part, is greatly enhanced protections for consumer data privacy. State governments have long played a role in protecting specific elements of consumer privacy, but data privacy issues in particular are currently skyrocketing in importance in state legislatures across the nation.¹⁶⁶ President Biden’s October executive order also recognizes data privacy protections as an important piece of safe and equitable AI development and usage.¹⁶⁷ Many state legislatures are considering legislation that could protect aspects of consumer online privacy, and several have even enacted comprehensive (or “omnibus”) consumer data privacy protection laws.¹⁶⁸ In addition to the five states that had enacted omnibus statutes before 2023, at least eight more states enacted such statutes in 2023 alone.¹⁶⁹ Unfortunately, Kansas’s state legislature has been one of the least active in this area.¹⁷⁰ As of February of 2025, not a single bill regarding consumer data privacy has been considered by the state legislature.¹⁷¹

¹⁶³ *AI Risks*, *supra* note 26, at 119.

¹⁶⁴ *FTC Big Data*, *supra* note 19, at 9.

¹⁶⁵ Levin, *supra* note 144; *Machine Bias*, *supra* note 144.

¹⁶⁶ Heather Morton, *2023 Consumer Data Privacy Legislation*, NAT’L CONF. STATE LEGISLATURES (Sep. 28, 2023), <https://www.ncsl.org/technology-and-communication/2023-consumer-data-privacy-legislation> [<https://perma.cc/P4MY-TEBA>].

¹⁶⁷ *Exec. Order: Responsible Development and Use*, *supra* note 160.

¹⁶⁸ Morton, *supra* note 166.

¹⁶⁹ *Id.* Additionally, New Jersey and New Hampshire passed omnibus consumer privacy laws in early 2024. Nancy Libin, David L. Rice, John D. Seiver & Benjamin Robbins, *New Jersey Governor Signs Comprehensive Privacy Law*, DAVIS WRIGHT TREMAINE LLP (Jan. 22, 2024), <https://www.dwt.com/blogs/privacy--security-law-blog/2024/01/new-jersey-data-privacy-law-signed> [<https://perma.cc/YM2Q-K9A3>]; David P. Saunders & John C. Ying, *And Another: New Hampshire Passes New Consumer Privacy Law*, MCDERMOTT WILL & EMERY (Jan. 24, 2024), <https://www.mwe.com/insights/and-another-new-hampshire-passes-new-consumer-privacy-law/#:~:text=Overview,take%20effect%20January%201%2C%202025> [<https://perma.cc/6K8P-53T3>].

¹⁷⁰ Morton, *supra* note 166.

¹⁷¹ See *US State Privacy Legislation Tracker*, IAPP (Feb. 24, 2025), <https://iapp.org/resources/article/us-state-privacy-legislation-tracker/#state-privacy-law-chart> [<https://perma.cc/JL4Z-WPMK>].

When Oscar Pfungst wanted Clever Hans to stop making decisions based on the information the horse gleaned from his questioner's body language, he contrived to cut off that source of information. If we want our machine learning algorithms to stop making discriminatory decisions based on consumer data, the only sure path to that destination must include severing access to that data at the source. Due to the direct link between the overcollection and misuse of consumer data and the harms caused by AI decision-making, the best course of action moving forward is for state legislatures, including the Kansas legislature, to adopt broad statutory protections for consumer data privacy.

C. Features of an Ideal Approach

The ideal legislation would be an omnibus-style bill that uses common definitions and principles borrowed from the European Union's (EU) General Data Protection Regulation (the "GDPR"), covers all consumer data (as opposed to only certain categories deemed most sensitive), features affirmative requirements for all personal data processing (rather than relying solely on notice-and-choice), and establishes a regulatory agency to manage the implementation of the law. The remainder of this Section will discuss each of these elements in turn.

1. Omnibus Approach vs. Sectoral Approach

At the federal level, data privacy is protected only in bits and pieces, with bodies of legislation separated by industry.¹⁷² In practice, this approach places very little restraint on the ability of companies to process and use immense quantities of consumer data.¹⁷³ The separation of the various bodies of law encourages heavy industry lobbying to the point of regulatory capture.¹⁷⁴ This leaves large gaps in the overall effectiveness of federal regulation.¹⁷⁵ When it comes to catching problematic practices and abuses of privacy, the sectoral approach is less a leaky bucket and more a rusty sieve.

Commentators have bemoaned the inadequacies of the federal sectoral approach for years, noting that the complexities of navigating multiple bodies of law can frequently leave consumers without knowledge of which of their data is actually protected.¹⁷⁶ The EU, on the other hand, leads the world in data protection regulation with its omnibus-style GDPR.¹⁷⁷ The GDPR, unlike the federal sectoral approach, provides individuals with meaningful control over their data across all industries.¹⁷⁸ There is nothing the Kansas legislature can do to correct the catastrophe that is current federal privacy law, but Kansas can at

¹⁷² *AI Risks*, *supra* note 26, at 161.

¹⁷³ Lu, *supra* note 79, at 2093.

¹⁷⁴ *AI Risks*, *supra* note 26, at 161–63.

¹⁷⁵ *Id.*

¹⁷⁶ See O'Connor, *supra* note 27.

¹⁷⁷ Lu, *supra* note 79, at 2095; *AI Risks*, *supra* note 26, at 161.

¹⁷⁸ *AI Risks*, *supra* note 26, at 167–68.

least avoid duplicating the federal government's mistakes. As more and more states follow the far-more-effective example of the EU, the legislatures of states like Kansas whose state does not yet have any sectoral data privacy laws should take the opportunity to adopt omnibus legislation that offers Kansans the privacy and control they deserve over all aspects of their personal data.

2. *Adoption of GDPR Principles and Common Terms*

The GDPR is the single main body of data privacy protection law in the EU.¹⁷⁹ It stems, philosophically and legally, from the recognition of data privacy as a fundamental human right.¹⁸⁰ To set the course for enforcing this right, it lays out six key principles for the use of personal data.¹⁸¹ First, data must be processed lawfully, fairly, and transparently with regard to the individual associated with that data (the “data subject”).¹⁸² Second, data must be collected only for specified, limited, explicit purposes, and processed only to the extent compatible with those purposes.¹⁸³ Third, data not required for the specified purposes may not be collected.¹⁸⁴ Fourth, the accuracy of collected personal data must be maintained, or else the entity holding the data must dispose of it.¹⁸⁵ Fifth, the data must be stored only as long as the specified purpose requires, and any identifying portions of the data must be stripped away as soon as possible.¹⁸⁶ Sixth and finally, the data must be processed in a manner providing appropriate security, integrity, and confidentiality.¹⁸⁷ These principles form a robust starting point for any legislative scheme, and should be adopted as the foundation of the ideal Kansas bill.

Many U.S. states have additionally elected to adopt portions of the terminology and definitions used in the GDPR.¹⁸⁸ Kansas should mirror this approach. The less friction that businesses face in understanding and complying with the patchwork of state laws forced by the lack of uniform federal protections, the more effective those state laws can be. It makes no sense to

¹⁷⁹ Lu, *supra* note 79, at 2093.

¹⁸⁰ *Id.*

¹⁸¹ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance), 2016 O.J. (L. 119) 1-88, art. 5(1) [hereinafter GDPR].

¹⁸² *Id.* at art. 5(1)(a).

¹⁸³ *Id.* at art. 5(1)(b).

¹⁸⁴ *Id.* at art. 5(1)(c).

¹⁸⁵ *Id.* at art. 5(1)(d).

¹⁸⁶ *Id.* at art. 5(1)(e).

¹⁸⁷ *Id.* at art. 5(1)(f).

¹⁸⁸ Sheila A. Millar & Tracy P. Marshall, *The State of U.S. State Privacy Laws: A Comparison*, NAT'L L. REV. (May 24, 2022), <https://www.natlawreview.com/article/state-us-state-privacy-laws-comparison> [https://perma.cc/D28U-L6YF] [hereinafter *Privacy Laws: A Comparison*].

reinvent the wheel when it comes to the terminology Kansas uses to implement its law.

3. *Extent of Protection*

The GDPR protects all “personal data” which is broadly defined as any data relating to an identified or identifiable natural person, even if that person is identified only through reference to any one of a broad field of characteristics.¹⁸⁹ The breadth of this protection is not accidental.

Part II of this Article discussed how user data from many sources could be conglomerated to form an uncomfortably detailed and accurate dossier about most U.S. consumers.¹⁹⁰ This might lead one to mistakenly believe that if access to a consumer’s most sensitive characteristics were shut off, the entire conglomeration would be stymied and most of the problem would be solved.

Unfortunately, sensitive characteristics can be determined with statistical significance from seemingly innocuous data.¹⁹¹ For example, researchers were able to use a person’s publicly available list of “likes” to determine that person’s gender, political affiliation, religion, use of cigarettes and alcohol, and even whether that person had experienced parental divorce before the age of 21.¹⁹² This issue thwarts attempts to prevent discrimination on the part of machine learning models by screening the datasets used to train them.¹⁹³ Even when data miners are extremely careful to avoid using any data containing inherently discriminatory flaws, algorithms can still pick up on proxy indicators of people’s protected attributes, effectively negating the entire effort.¹⁹⁴

Furthermore, data that has been anonymized by stripping away all obviously identifying portions can frequently be de-anonymized using modern AI technologies.¹⁹⁵ In other words, simply avoiding the practice of providing AI models with specific sensitive data is not enough to prevent flawed, discriminatory decision-making based on those data. Kansas and other states adopting new omnibus legislation should ensure that the protections provided under that legislation are as broad as practicable, or else they might as well not legislate at all.

¹⁸⁹ GDPR, *supra* note 181, at art. 4(1) (“[A]n identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.”).

¹⁹⁰ *See supra* Part II.

¹⁹¹ THE ETHICAL ALGORITHM, *supra* note 79, at 52.

¹⁹² *Id.*

¹⁹³ *See Disparate Impact*, *supra* note 20, at 674.

¹⁹⁴ *Id.*

¹⁹⁵ *AI Risks*, *supra* note 26, at 128.

4. *Notice and Choice*

“Notice-and-choice” is the name given to the approach to online privacy protection wherein entities that wish to collect and process consumer data must inform the consumer of which data they wish to collect and what uses they wish to put it to (“notice”), then receive consent from the consumer to do so (“choice”).¹⁹⁶ The notice-and-choice paradigm forms the basis for most modern data privacy regulation, but it is nonetheless fraught with shortcomings.¹⁹⁷

Legislative solutions that rely solely on consumer notice-and-choice are naive at best.¹⁹⁸ Requiring consumers to read and understand the myriad complex privacy agreements they encounter on a day-to-day basis is completely impracticable.¹⁹⁹ Even in the EU, where citizens are protected under comprehensive notice-and-choice consent requirements, only a small fraction of people actually read privacy policies in full.²⁰⁰ This problem is further compounded by the fact that many data brokering companies operate without coming into contact with consumers in the first place, so consumers do not get meaningful chances to give or revoke consent.²⁰¹ To address this issue, legislation should include affirmative duties for all companies handling personal data, such as the affirmative duty to strictly limit the use of collected data, or even a “duty of care” to consumers writ large for large tech companies which effectively treat consumers as their products.²⁰² At the very least, legislation should forbid conditioning access to a good or service on the provision of information not necessary to provide that good or service.²⁰³ Such an approach would sidestep the impracticability of direct consumer engagement with the complexities of privacy law and cut straight to the goal: protecting state consumers.

5. *Establishing a Regulatory Agency*

As part of its original omnibus data privacy bill, California created and funded a new agency to facilitate the implementation and enforcement of its

¹⁹⁶ John A. Rothchild, *Against Notice and Choice: The Manifest Failure of the Proceduralist Paradigm to Protect Privacy Online (or Anywhere Else)*, 66 CLEV. ST. L. REV. 559, 561–62 (2018).

¹⁹⁷ See generally *id.*

¹⁹⁸ See Rich, *supra* note 27.

¹⁹⁹ *Id.*

²⁰⁰ Lu, *supra* note 79, at 2111.

²⁰¹ Rich, *supra* note 27.

²⁰² See *id.* (noting recent legislative approaches to addressing the issues with notice-and-choice reliance); *AI Risks*, *supra* note 26, at 119 (describing the business models of companies such as Facebook and Google). This is not to say that such protections should come at the expense of consumers’ ability to directly assert rights to access, correct, delete, or limit the collection of their personal data—the ideal legislation would provide both, so that consumers are enabled to see and control their data actively as a supplement to the passive protection of the law.

²⁰³ Rothchild, *supra* note 196, at 647.

statutory scheme.²⁰⁴ Given the complexity and rapidly changing nature of the field, this approach makes intuitive sense. Expecting legislators to individually track the highly technical minutiae required to effectively prevent abuse of the regulatory scheme is simply unreasonable.²⁰⁵ Furthermore, the enforcement provided by such an entity could allow for the protection of the rights of those unable to bankroll expensive private litigation.²⁰⁶

The creation of this entity, however, should not eclipse a private right of action. The threat of class action claims can help to hold large tech companies accountable for their misuse of data and can do so largely on private dime.²⁰⁷ A well-funded agency working hand-in-hand with litigators exercising a measured private right of action would thus provide the most robust state-level protections possible for consumers without unduly burdening the system.²⁰⁸ For this reason, states such as Kansas should establish and appropriately fund a regulatory agency dedicated to the protection of consumer privacy while also providing for a private right of action against violating companies.

V. COUNTERARGUMENTS

Two main arguments are routinely brought against state-level implementations of consumer data privacy to mitigate AI harms: first, that any regulation on data privacy should be federal- rather than state-level, and second, that merely enhancing data privacy is not sufficient to protect consumers against all the threats posed by AI misuse. Section A of this Part will address the former, and Section B will address the latter.

A. State vs. Federal Implementation

While states have increasingly been stepping up to fill the gaps in federal privacy law, the patchwork system of state-level regulation can present its own issues.²⁰⁹ When multi-state or even international businesses are required to interact with the overlapping and sometimes conflicting regulations of multiple states, it can be costly, unfeasible, or even impossible for those businesses to comply with all relevant law.²¹⁰ Additionally, states drafting legislation in the area must contend with the possibility of future preemption by

²⁰⁴ *Privacy Laws: A Comparison*, *supra* note 188.

²⁰⁵ See Amy Keller, 'Paper Tiger' State Privacy Laws Worse Than Having No Law at All, BLOOMBERG L. (Oct. 12, 2023, 3:00 AM), <https://news.bloomberglaw.com/privacy-and-data-security/paper-tiger-state-privacy-laws-worse-than-having-no-law-at-all> [https://perma.cc/22B2-UDNT].

²⁰⁶ See Katie Mansfield, *State-Level Consumer Data Privacy Laws Get the Ball Rolling, But On Their Own, Represent a Piecemeal Approach to Regulation*, JOLT DIG. (Oct. 25, 2023), <https://jolt.law.harvard.edu/digest/state-level-consumer-data-privacy-laws-get-the-ball-rolling-but-on-their-own-represent-a-piecemeal-approach-to-regulation> [https://perma.cc/CM7C-V5RT].

²⁰⁷ *Id.*

²⁰⁸ See *id.*; Keller, *supra* note 205.

²⁰⁹ *AI Risks*, *supra* note 26, at 162–63.

²¹⁰ *Id.*

federal statute, which would at least partially invalidate the time, effort, and funds states spend enacting and enforcing their regulations.²¹¹

There are also issues of cynical practicality: much of the effectiveness of the GDPR stems from the ability of the EU to wield its economic weight to enforce fines steep enough to incentivize even large international corporations to change their business practices to comply.²¹² Kansas (or any single state, with the possible exception of California) simply cannot leverage the same degree of economic power and weight as the collective EU in its enforcement, and therefore cannot make demands as large or provide protections as complete as those of the GDPR.

Sadly, despite the shortcomings inherent to state-level solutions, state governments cannot rely on the federal government to enact reform any time soon. Congress has made numerous fruitless efforts to pass comprehensive data privacy legislation over the past twenty years.²¹³ As just one example, the bipartisan American Data Privacy and Protection Act, the most recent of these attempts to gain serious momentum as of late 2023, died without a vote after being introduced by the House Energy and Commerce Committee.²¹⁴ If states want to see their citizens protected in the age of AI, in other words, it is imperative that they take action on their own.

B. Insufficiency of Data Privacy Alone

Several scholars rightly note that data privacy protection alone cannot solve all the novel problems that arise with the ever-increasing use of AI technologies.²¹⁵ Scholars have noted that even the GDPR does not adequately protect consumers against all of these harms.²¹⁶ This should not, however, discourage lawmakers—half a loaf is far better than none, and the increased consumer protections provided by data privacy legislation are an important piece of any full solution. Moreover, many states have begun to include a right against automated decision-making in their omnibus legislation, helping to mitigate AI

²¹¹ See generally Lauren Zabierek, Tatyana Bolton, Brandon Pugh, Sofia Lesmes, & Cory Simpson, *Preemption in Federal Data Security and Privacy Legislation*, HARVARD KENNEDY SCHOOL: BELFER CTR. FOR SCI. AND INT'L AFFS. (June 14, 2022), <https://www.belfercenter.org/publication/preemption-federal-data-security-and-privacy-legislation> [https://perma.cc/V35U-EDXL].

²¹² *AI Risks*, *supra* note 26, at 166.

²¹³ Rich, *supra* note 27.

²¹⁴ CONGRESS.GOV, *H.R.8152 - American Data Privacy and Protection Act* (2022), <https://www.congress.gov/bill/117th-congress/house-bill/8152>.

²¹⁵ See generally Lu, *supra* note 79; *AI Risks*, *supra* note 26.

²¹⁶ Lu, *supra* note 79, at 2093–94.

harms further than would be possible with data privacy protection alone.²¹⁷ While this Article lacks the scope to fully explore this regulatory option, the author strongly urges legislators to consider such a provision when drafting future data privacy statutes.

VI. CONCLUSION

The dawn of AI has ushered in a new chapter in the age of data, one where the slow surfacing of problems with U.S. privacy law has begun to rapidly accelerate. In the modern world, most U.S. consumers are effectively subjected to constant surveillance through their use of technologies that have become a necessary part of everyday life.

The introduction of state-level data privacy protections cannot, unfortunately, solve all the novel problems presented by AI use and misuse, nor can it necessarily provide the permanent protections desperately needed by all U.S. citizens. Even so, our unprotected states must begin to pass omnibus privacy bills that work within existing paradigms, create broad protections (both active and passive) over consumer data, and provide for development and enforcement through both administrative agencies and private action. The passage of such legislation is the only pragmatic path towards relief from the Orwellian world created by the unfortunately prevalent harms of our advancing modern technology.

²¹⁷ See, e.g., Avi Gesser, Robert Maddox, Anna Gressel, Mengyi Xu, Samuel J. Allaman & Andres S. Gutierrez, *New Automated Decision-Making Laws: Four Tips for Compliance*, DEBEVOISE & PLIMPTON DATA BLOG (Jun. 25, 2022), <https://www.debevoisedatablog.com/2022/06/25/new-automated-decision-making-laws-four-tips-for-compliance/> [<https://perma.cc/L3PU-FNBQ>]; Kirk J. Nahra, Ali A. Jessani & Samuel Kane, *California Privacy Protection Agency Publishes New Draft CPRA Cybersecurity and Automated Decisionmaking Regulations in Advance of December Board Meeting*, WILMERHALE PRIV. & CYBERSECURITY L. BLOG (Nov. 30, 2023), <https://www.wilmerhale.com/en/insights/blogs/wilmerhale-privacy-and-cybersecurity-law/20231130-california-privacy-protection-agency-publishes-new-draft-cpra-cybersecurity-and-automated-decisionmaking-regulations-in-advance-of-december-board-meeting> [<https://perma.cc/8YU5-FPJ3>]. The GDPR includes a limited protection against automated decision-making, but in June 2024, the EU also enacted a separate AI bill to provide more complete and well-tailored protections. GDPR Art. 22; REGULATION (EU) 2024/1689 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act).