Harvard Data Science Review • 1.1

Data Science: What the Educated Citizen Needs to Know

Alan M. Garber

Published on: Jun 23, 2019

Updated on: Oct 04, 2019

DOI: 10.1162/99608f92.88ba42cb

At a recent dinner, the conversation turned to the impressive gains in the power and ubiquity of artificial intelligence. Reflecting both the wonder and anxiety of the times, one of our guests speculated that machines would soon surpass humans at determining guilt and innocence. None claimed to know when that would happen, but nobody doubted that sometime soon – if it has not happened already – algorithms would be able to sift through and interpret the totality of evidence better than a human could. One of the guests, a law professor, questioned whether machines would ever replace human juries, and wondered how influential they would be. Suppose that a machine made a determination of guilt. How would an AI tell the story, explaining its reasoning and persuading the jury? The machine, or rather the argument based upon the machine’s conclusions, would not escape harsh scrutiny by the attorney representing the accused. She would probe every weakness, exposing flaws in the data and algorithm, exploiting the opacity of underlying machine learning techniques, and appealing to emotions in ways that remain difficult for machines to match. Would juries nevertheless believe the AI? Should they? How would they weigh the arguments presented by the defense attorney against the vast knowledge but less than transparent reasoning of the machine? How could anyone be sure that the underlying algorithms were not subject to bias? It seemed unlikely that AI would succeed in court until it could explain better, appeal to emotions, and persuade jurors to shed their skepticism.

And perhaps it shouldn’t – an unjustified faith in the rectitude of a computer-generated determination of guilt would be even more costly than misplaced skepticism. As AI extends its power and reach, all of us will need to learn how to work with it and how to ensure that its influence on human well-being is salutary. Most importantly, we will need to ensure that human judgment not only maintains its primacy but that we, as citizens, are up to the task. That is a responsibility of educators, and one that universities need to take seriously. We increasingly ask how we can best prepare our students for a machine-driven future.

My remarks focus on data science, which shares foundations and in many respects overlaps with AI but is applied to many other areas of endeavor. In fact, the pervasive use of the term “data science” in academic settings reflects both the appeal of the intellectual activities it encompasses and the capaciousness – or vagueness – of its meaning. The collection, organization, analysis, and interpretation of data can all be considered its domain.

As both a descriptor and as an activity, data science has acquired an aura of power and consequence, in no small part due to its association with AI. We now take for granted that our mobile phones identify family members and friends in photos, and that those phones will warn us that we need to leave early for our next appointment because traffic is heavier than usual. Music streaming services point us to unfamiliar tunes that we find at least as appealing as the recommendations of friends. All of these

innovations were enabled by access to more data, along with advances in analytic methods and in the capacity to store information and perform ever more-speedy calculations at a vast scale.

We are told that this is just the beginning: these skills are only previews of far more consequential capabilities. Among them, we can expect, will be recognition of molecular structures that lead to more precisely targeted cancer drugs, improved ability to manage air and automobile traffic, and earlier warnings of impending earthquakes and tsunamis. The economic repercussions will be profound. According to simulations by researchers at the McKinsey Global Institute, “front-runner” companies that develop and adopt AI technologies will increase their cash flow by about 122% between 2017 and 2030. “Follower companies” will experience a 10% growth in cash flow over the same period, while “laggards” that do not adopt AI technologies by 2030 will experience a 23% *decline* in cash flow. The impact worldwide will be enormous: by 2030 the adoption of AI could add on net about 16%, or \$13 trillion, to global output ([Bughin, Seong, Manyika, Chui, & Joshi, 2018, p.13](#)).

If numbers like these attract the attention of businesses, they are not lost on the rest of us. Interest in AI and data science is pervasive on university campuses, where faculty with the right expertise fight off (or not) offers from industry that offer unprecedented research opportunities or compensation, and often both. The demand for classes in computer science, statistics, and related areas of study has soared, along with the number of undergraduate majors in these fields. At Harvard, which is typical in this regard, the number of undergraduate “concentrators” in computer science more than tripled between 2011 and 2017. From 2008 to 2017, the number of statistics concentrators increased tenfold. CS50, introductory computer science, and Stat 110, introductory probability, became among the most popular undergraduate courses. Such behavior hardly conforms to the belief that undergraduates ignore career prospects when choosing their studies, and that their colleges fail to accommodate them. Students know that the market for talent in data science is a seller’s market. Universities are scrambling to meet demand for courses that will prepare their students for a data science-driven future.

Even if the predictions of the coming impact of AI are overblown, all the evidence suggests that data science will only grow in importance in the coming years. What are the implications for university education? Is our main task to ensure that we can train enough specialists in this field, or do we need to view data science as a field of study that everyone needs to learn something about, much as universities require instruction in mathematics, language, and writing skills? And if data science is becoming, in essence, a liberal art, what do non-specialists need to know?

In answering these questions, it is helpful to distinguish among three groups of students. First are those who intend to become full-fledged experts by specializing in data science, with ambitions to contribute to advances in the field through their teaching or research, whether in universities or

within companies, government, or nonprofit organizations. They may pursue doctoral degrees in data science or longer established disciplines like computer science and statistics.

A second group has the potential to be larger: students who are intrigued by the power of data science as a tool that can lead to advances in their field of primary interest, which might be physics, chemistry, medicine, public health, climate science, political science, sociology, economics, or many others, including the humanities. Theirs is often a bridging role. Building on the well-recognized advantages of multidisciplinary teams, and sometimes as an alternative to such teams, students trained in both data science (often at the master's degree level) and another discipline will be among the best positioned to lead their fields of application. Many universities are devising curricula and degree programs to prepare them to become applied data scientists.

Differences between the first and second groups are not sharp; they are on a spectrum ranging from the more general-purpose and methodological aspects of data science to those that are more focused on the area of application. A third category is broader and less focused on research. It is comprised of students who do not aim to acquire deep expertise but need to gain a basic understanding of data science, whether they recognize it or not, and for whom traditional introductory statistics courses are an incomplete solution. Existing data science curricula have focused heavily on the first two groups, but nearly everyone else falls into this one. Today's students live in a world shaped by the application of data analytics. Because the influence of data science will be even greater in the future, there will be ever more compelling reasons to learn about it. Students may be motivated to do so out of simple curiosity about this important part of their world, a desire to interpret information better, or as career preparation. While experts debate whether automation will lead to displacement of high-skill jobs in such fields as radiology ([Obermeyer & Emanuel, 2016](#); [Thrall et al., 2018](#)) and long-term increases in unemployment rates ([Autor, 2015](#); [Autor & Salomons, 2018](#); [Pew Research Center, 2014](#)), students are well aware that successful careers will require an evolving set of high-level skills and knowledge ([Aoun, 2017](#); [National Academies of Sciences, Engineering, and Medicine, 2017](#)). Narrowly directed training will not provide a solid foundation for a successful life of work. Insofar as data science becomes an element of 21st century liberal arts, what do students in the third group – that is, every educated citizen – need to know about it?

The need

When crafting a curriculum, the first step is to define the field. That basic exercise is seldom simple or harmonious, since it has implications for both what is taught and who teaches it. It is more complicated when the field is claimed by people trained in different disciplines. For the purposes of these remarks, my approach to defining data science is practical and cursory. I consider the field to encompass an

array of activities related to the collection, linkage, and storage of data, but focus here on analytics that build upon modern computer science and statistics, along with related fields such as operations research, decision sciences, and mathematics, to classify and draw predictions and inferences from large, often high-dimensional data sets. Data science is characterized more by modern technical resources and opportunities than by a set of methods or even specific applications. However its boundaries are drawn, it is clear that data science powers some of the most impressive contemporary accomplishments of artificial intelligence. For the most part, these are feats of classification and prediction, enabled by increasingly massive data sets (Goodfellow, Bengio, & Courville, 2016).

An obvious first goal for a basic education in any field would be to learn about the methods, or at least it would be if the field were united by common methods. Defining a canon or list of required courses is contentious, for the same reason defining the field is. To appreciate the added challenges of doing so, just compare data science textbooks written by statisticians (Efron & Hastie, 2016; Hastie, Tibshirani, & Friedman, 2009) to those written by computer scientists (Goodfellow et al., 2016). The pedagogical approaches and the assumed backgrounds of the students vary, for good reasons, and the language used to describe the same or closely related techniques differs by discipline. The difficulties in putting together a curriculum informed by the contributing disciplines are real, but the rewards are greater still. Convening faculty with diverse backgrounds, assumptions, methods, and areas of focus will enhance dialogue across boundaries and may lead to new collaborations and ultimately to research advances. It also offers the best hope of creating a curriculum of enduring value, serving students for years beyond graduation.

We can and should have spirited debates about which methods need to be taught – deep learning, linear and nonlinear regression techniques, resampling methods, and so on – but in my view the greater challenge, and opportunity, is to pursue a more ambitious agenda: to change how our students think about the world they inhabit, by helping them gain probabilistic and statistical literacy.

The need for this agenda has been clear for decades. Probability and statistics offer well-developed tools to inform decision making under uncertainty, but they have been available for fewer than three centuries (Stigler, 1986). For the rest of human history, people have somehow managed to make consequential decisions without the benefit of formal methods, and they continue to do so every day. Those decisions are or should be informed by evidence, but most of us are not particularly skilled at it, nor do we always get better with experience.

The Nobel-prize winning psychologist Daniel Kahneman, building on decades of research on cognitive errors, has framed the problem as the human predilection to rely heavily on the mind’s “System 1,” which “operates automatically and quickly, with little or no effort and no sense of voluntary control.” He contrasts it with the slower but more reliable “System 2,” which “allocates attention to the effortful mental activities that demand it, including complex computations.” The principal weakness of System

1, which presumably conferred survival advantages throughout most of our evolutionary history, is its tendency to bias and error. As he wryly notes, System 1 “sometimes answers easier questions than the one it was asked, and it has little understanding of logic and statistics ... [and] it cannot be turned off” (Kahneman, 2011, p.25).

The list of faults of System 1 is long. In 1974, decades before he wrote about the two systems, Kahneman and his long-term collaborator Amos Tversky published a paper in *Science* that changed views about human rationality and our ability to make complex decisions ([Tversky & Kahneman, 1974](#)). The first type of error they identified was what they called the “representativeness” heuristic. They recounted the results of a study in which subjects were asked to assess the probability that an individual was engaged in a particular occupation. The study described Steve, who was “...very shy and withdrawn, invariably helpful, but with little interest in people, or in the world of reality. A meek and tidy soul, he has a need for order and structure, and a passion for detail.” Asked whether Steve is more likely to be a librarian than a farmer, most study subjects will say yes, even though many more men are farmers than librarians. Kahneman and Tversky refer to experiments showing that people ignore the base-rate frequency or prior probability when answering such questions, assigning probabilities based on representativeness – fit with a stereotype – instead. In essence, they violate the norms of probability by failing to apply Bayes’ Rule correctly.

Tversky and Kahneman’s paper also described the heuristics of availability (assessing probability based on “the ease with which instances or occurrences can be brought to mind”) and adjustment and anchoring (a tendency to adjust probabilities from an initial anchoring value, which implies that the final probability estimate can be manipulated by changing the phrasing of a question). This paper and the large literature that it spawned have yielded copious insights into how people deal with complex information and into the nature of errors in judgment and in interpreting probabilities, especially in areas such as legal decisions ([Kaye & Koehler, 1991](#); [Kelman, Rottenstreich, & Tversky, 1996](#); [Sunstein, 1997](#)) and medical judgments ([Casscells, Schoenberger, & Graboys, 1978](#); [Elstein, 1999](#)). Above all else, this literature casts doubt on the ability of even experts to avoid cognitive errors.

Intelligent systems, built upon deep learning and other data science techniques, may not perform as well as the human mind at some pattern recognition tasks or at recognizing the emotional states of other people. Some skills that humans acquire without apparent effort remain beyond the capabilities of machines. But intelligent systems are making rapid progress toward gaining such skills and are already superior at many other tasks. Sometimes their access to massive data sets enables them to draw predictions and inferences that apply far more broadly than those that could be drawn from an individual human’s experience. And they have another important advantage: they do not fall prey to the cognitive limitations of System 1. An intelligent system would not be vulnerable to endowment effects, availability, or adjustment and anchoring, nor would it have trouble applying Bayes’ rule,

unless it were designed to do so in order to mimic flawed human decision making. In other words, it would come closer than humans to “ideal” rational judgment ([Parkes & Wellman, 2015](#)).

Thus we might ask, if we are hard-wired to rely on a faulty system, why not replace it with one that is not only free of those faults but can learn rapidly from its mistakes and constantly incorporate new information?

Intelligent systems with data analytic foundations do just that. And in some settings, those systems will lessen or eliminate the need for routine human intervention. But much of the time they will aid, not replace, human judgment. As sophisticated as image recognition has become, in many radiology and pathology applications AI is used to identify areas of images for more detailed human scrutiny, rather than completely automating image interpretation. Machines do not vote or serve on juries, and as the law professor observed, the suggestion that they should do so is odious to many, even if they are convinced that a machine can or soon will do better than humans at determining guilt or culpability. Perhaps we will cede more authority to machines even in these domains when they are able to tell a story, explain their reasoning, and convey empathy and compassion along with their abilities to predict and classify. Until then, their role will not be to replace humans in making sensitive decisions, despite our cognitive limitations, but to aid them.

Most of our students – that is, those who will use the products of data science but neither advance data science methodologies nor apply those methodologies to research in other fields or to develop new products and services – will need to learn how to be smarter consumers of data and data analytics, and better decision makers. For the most part, that will mean empowering System 2 by gaining a better understanding of probability and statistics.

What we can teach

What aspects of probability and statistics are most important for a well-educated citizen to know?

Many readers will have strong views on this subject. The brief and incomplete list I propose here is meant to start conversation, not to lay out an educational program. Because it is intentionally parsimonious, it might justly be criticized as insufficiently ambitious. Similar proposals have included more topics (e.g., [Utts, 2003](#)) and a strong case can be made for the value of instruction in computational thinking ([Wing, 2006](#)). I do not presume that every student will take two semesters or three quarters of probability and statistics, computer science, or data science. Any call for new requirements must recognize the competing demands of different areas of the curriculum. Those demands are not the same in all institutions; the set of courses required in a small liberal arts college

might not be those that would be required in a science-and-technology focused research university. And the more time-intensive the requirement, the less likely it is to be adopted.

First is recognition of the pervasiveness of uncertainty along with mastery of probabilistic thinking as a framework for understanding uncertainty and its implications. Every well-educated person should understand the concepts undergirding conditional and joint probabilities, along with Bayes' Rule, how probabilities are combined, and the reasons why intuition can be so misleading. For example, the tendency to ignore small probabilities or to fail to draw distinctions between, say, a probability of 0.01 and of 0.001 leads both to logical inconsistencies and to incorrect inferences. There should be two goals – to understand at an operational level the simple principles of probability, and to be aware of the errors that we are prone to make when guided by our intuitions, and their consequences.

Second is the relationship between the sample and the relevant population. At a basic level, the concept of the sample is easy to understand. You don't need to take a statistics course to recognize that a study to characterize the market for lawn mowers is more likely to be informative if its sample is drawn from suburban Connecticut than from Manhattan's Upper East Side. But people often draw broad conclusions from limited data, whether the limitation is one of sample size or representativeness. Our current approaches to teaching concepts such as external validity and selection bias may not be enough, unless we believe that our students who can define the concepts and answer related exam questions also apply their knowledge in everyday life.

The implications of false positives and false negatives, whether our students remember them as Type I and Type II errors or not, constitute a third area of basic knowledge. Every student who takes a statistics course is exposed to this concept early – it is, after all, fundamental to statistical inference. But they may not have a deep understanding of the consequences of each kind of error, nor may they be able to conceptualize the value of added information. When assessing whether a new wing design for an airplane will be subject to greater fatigue and the possibility of a flight disaster, the failure to discover an increase in accident risk is far more costly than a mistaken inference that the design is hazardous when in fact it is not. No great technical sophistication is needed to understand why we would accept different probabilities of Type I and Type II errors in different circumstances, that we can decrease the rate of one type of error at the cost of an increase in the rate of the other, and that we can usually reduce the rate of both by gathering more data. We all apply these concepts, even when we aren't fully aware that we're doing so. They have everything to do with the reason why American courts apply a standard of "beyond a reasonable doubt" in criminal trials and "preponderance of evidence" in civil cases. Formal education about such errors can help us think more clearly about related issues, such as when it is crucial to gather different information or increase a sample size.

Fourth would be an understanding of basic inference and the distinction between association and causation. This is a particularly daunting challenge, and not because people have trouble

distinguishing between cause-and-effect and mere correlation. It is daunting because even people who should know better frequently draw causal inferences from correlations. The problem is exacerbated by incentives to emphasize causal interpretations. We are far more interested in causation than in association alone. How many times a week does the press report a medical or epidemiological study of a disease association? Often the story will include a caveat that the association may not be causal. But the attention the media give to the study often casts doubt on the sincerity of the disclaimer. Why would the finding be interesting if the relationship were *not* causal? In epidemiology, as in many other areas of health and medicine, associations can point to mechanisms and promising interventions, but for the most part, public and scientific interest is overwhelmingly due to the possibility that the association is causal. We are interested in the association between consumption of organic foods and cancer because we want to know whether organic foods are safer. For most non-academics, a report that people who have other health-enhancing habits also tend to eat organic food makes for a less compelling headline.

Good empirical studies always offer a balanced discussion of the degree to which the association is likely to be causal, but press reports are less circumspect. It falls to the reader or viewer to bring informed skepticism when assessing the meaning of such reports. A deeper understanding of causality will not simplify the world, but it can help make more sense of it, and avoid costly mistakes.

Prospects for success

These observations about the gaps in knowledge of probability and statistics, and the reasons why people need to know more, are not new; they could have been written decades ago. Why are they more salient today? And doesn't the proliferation of machine learning and its integration into daily life, by providing a superior alternative to flawed human decision making, obviate the need to understand?

Deep learning and other machine learning techniques have inherent advantages, but providing a simple explanation is not among them. They are powerful tools, especially for prediction and pattern recognition. Their primary purpose is not to elucidate mechanisms or to explain why they give the results they do. By avoiding the functional form restrictions and relatively simple model specifications of traditional applied statistics, they optimize for the tasks they are known for. But the features that enhance their abilities to predict and classify also obscure the reasons for the prediction, and often preclude simple explanation. Ironically, they are most useful when their results differ greatly from those of simpler statistical models; if the difference in results is due to, say, high order interactions and nonlinear effects, it will be challenging to tell a comprehensible story about the reasons for the results. That can have damaging repercussions for the adoption of AI; concerns that algorithms are

biased and can lead to discrimination (e.g., less accurate recognition of faces of racial minorities) are only heightened when they are complex and opaque.

The machine learning community is taking steps to address this problem. They know that even stunningly accurate prediction will not carry the day in settings in which process and explanation matter, most visibly in courtrooms and legislatures. An emerging body of research addresses the fairness, accessibility, and transparency of machine learning ([Chouldechova, 2017](#); [Doshi-Velez & Kim, 2017](#); [Dwork, Hardt, Pitassi, Reingold, & Zemel, 2012](#); [Kleinberg, Mullainathan, & Raghavan, 2016](#)). A better understanding of probability and statistics will not be enough to make an opaque technique transparent to educated citizens. But it will help consumers of the results of machine learning techniques – and all statistical techniques – formulate basic questions that will enable them to put the results in context. That will be key to acceptability.

Thus far I have emphasized that we can draw on interest in data science and AI to give our students the tools they have long needed to think probabilistically and to weigh evidence. But today there is an even more pressing policy challenge: How can we prepare workers to adapt to a world in which automation takes over jobs in many different segments of the workforce, including white collar positions? And what are the implications for data science education?

Will automation and the widespread adoption of AI lead to rising unemployment? There has been no shortage of speculation about this topic, but it is too early to offer a definitive answer ([Acemoglu & Restrepo, 2018](#); [Aoun, 2017](#); [Autor, 2015](#); [Autor & Salomons, 2018](#); Brynjolfsson & McAfee, 2014; Lee, 2018; [Pew Research Center, 2014](#)). Historically, technological change has led to the creation of new jobs as existing jobs were rendered obsolete. But if future technological innovations have more dramatic and far-reaching effects, jobs may be lost to automation more quickly than attractive new opportunities become available. Our ability to predict the specific features of technological change beyond a decade in the future is limited, so the employment consequences remain a matter of speculation. However, there is evidence that some kinds of technological change will increase the demand for skilled workers ([Acemoglu & Restrepo, 2018](#); [Autor, 2015](#); [Autor & Salomons, 2018](#); [National Academies of Sciences, Engineering, and Medicine, 2017](#); [Pew Research Center, 2014](#)). At least in the past, a robust educational system seemed to ensure that technological change would not threaten employment, since new jobs arose and the need for human skills did not diminish (Goldin & Katz, 2008).

If predictions that AI will disrupt work more completely and rapidly than previous technological advances are correct, understanding its strengths and limitations will be a strategic advantage. The well-informed and well-educated will be among the best positioned to take advantage of AI and to apply it to achieve their goals.

As data science increases the reliability and usefulness of automated tools, there is a risk that we will assume they don't make mistakes. The correct drug will be chosen for the cancer, only a fraudulent financial transaction will be blocked, the smart bomb will reach its target rather than firing on an ally, the new building will withstand an earthquake that measures 6.7 on the Richter Scale. When they are based on strong analytics applied to extensive, appropriate data, confidence in the predictions may be justified, but that does not mean they're never wrong. We will need to recognize errors and correct them quickly, and to remain vigilant even as they become less frequent. Indeed, the very rarity of major mistakes will make it harder to stay alert to them and minimize damage when they do occur.

The benefits of data science are real and have never been more salient or important. Increasingly accurate predictions will make the products of data science more valuable than ever, and will increase interest in the field. The advances can also breed complacency and blind us to flaws. Workers of the future need to recognize not only what data science does to assist them in their work, but also where and when it falls short. Their education should prepare them for such a future. Many of the skills needed to work with data-science based tools will be specific to the occupation – what an accountant needs to know will be different from what an air traffic controller needs to know. But a deeper understanding of probabilistic reasoning and the evaluation of evidence is a general skill that will serve all of them well.

Although they may never replace humans on juries, machines will become skilled at more than prediction. They will continually improve their ability to think, judge, and explain. We need to ensure that the same will be true of humans.

Acknowledgements

Many thanks to Xiao-Li Meng, Joe Blitzstein, David Parkes, and Francesca Dominici for their very helpful comments. Any errors, and all opinions, are my own.

References

Acemoglu, D., & Restrepo, P. (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, 108(6), 1488-1542.
<https://doi.org/10.1257/aer.20160696>

Aoun, J. E. (2017). *Robot-Proof: Higher Education in the Age of Artificial Intelligence*. Cambridge, MA: MIT Press.

Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30. <https://doi.org/10.1257/jep.29.3.3>

Autor, D., & Salomons, A. (2018). Is Automation Labor Share-Displacing? Productivity Growth, Employment, and the Labor Share. *Brookings Papers on Economic Activity*, 2018(1), 1-87. <https://doi.org/10.1353/eca.2018.0000>

Brynjolfsson, E., & McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: WW. Norton & Company.

Bughin, J., Seong, J., Manyika, J., Chui, M., & Joshi, R. (2018). *Notes from the AI Frontier: Modeling the Impact of AI on the World Economy*. Retrieved January 6, 2019, from <https://www.mckinsey.com/>.

Casscells, W., Schoenberger, A., & Graboys, T. B. (1978). Interpretation by physicians of clinical laboratory results. *New England Journal of Medicine*, 299(18), 999-1001. <https://doi.org/10.1056/nejm197811022991808>

Chouldechova, A. (2017). Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data*, 5(2), 153-163. <https://doi.org/10.1089/big.2016.0047>

Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. Retrieved from <https://arxiv.org/abs/1702.08608>.

Dwork, C., Hardt, M., Pitassi, T., Reingold, O., & Zemel, R. (2012). *Fairness through awareness*. Paper presented at the Proceedings of the 3rd Innovations in Theoretical Computer Science Conference. <https://doi.org/10.1145/2090236.2090255>

Efron, B., & Hastie, T. (2016). *Computer Age Statistical Inference: Algorithms, Evidence, and Data Science*. Cambridge: Cambridge University Press.

Elstein, A. S. (1999). Heuristics and biases: selected errors in clinical reasoning. *Academic Medicine* 74(7):791-794. <https://doi.org/10.1097/00001888-199907000-00012>

Goldin, C., & Katz, L. F. (2008). *The Race Between Education and Technology*. Cambridge, MA: Belknap Press.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. Cambridge, MA: MIT Press.

- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction* (2nd ed.). New York: Springer.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus, and Giroux.
- Kaye, D. H., & Koehler, J. J. (1991). Can Jurors Understand Probabilistic Evidence? *Journal of the Royal Statistical Society. Series A (Statistics in Society)*, 154(1), 75-81. [doi:10.2307/2982696](https://doi.org/10.2307/2982696)
- Kelman, M., Rottenstreich, Y., & Tversky, A. (1996). Context-Dependence in Legal Decision Making. *The Journal of Legal Studies*, 25(2), 287-318. [doi:10.1086/467979](https://doi.org/10.1086/467979)
- Kleinberg, J., Mullainathan, S., & Raghavan, M. (2016). Inherent trade-offs in the fair determination of risk scores. Retrieved from <https://arxiv.org/abs/1609.05807>.
- Lee, K.-F. (2018). *AI Superpowers: China, Silicon Valley, and the New World Order*. Boston: Houghton Mifflin Harcourt.
- National Academies of Sciences, Engineering, and Medicine. (2017). *Building America's Skilled Technical Workforce*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/23472>
- Obermeyer, Z., & Emanuel, E. J. (2016). Predicting the future—big data, machine learning, and clinical medicine. *The New England Journal of Medicine*, 375(13), 1216. <https://doi.org/10.1056/nejmp1606181>
- Parkes, D. C., & Wellman, M. P. (2015). Economic reasoning and artificial intelligence. *Science*, 349(6245), 267-272. [doi:10.1126/science.aaa8403](https://doi.org/10.1126/science.aaa8403)
- Pew Research Center. (2014). *AI, Robotics, and the Future of Jobs*. Retrieved from <http://www.pewresearch.org/wp-content/uploads/sites/9/2014/08/Future-of-AI-Robotics-and-Jobs.pdf>
- Stigler, S. M. (1986). *The history of statistics: The measurement of uncertainty before 1900*. Cambridge, MA: Harvard University Press.
- Sunstein, C. R. (1997). Behavioral Analysis of Law. *The University of Chicago Law Review*, 64(4), 1175-1195. [doi:10.2307/1600213](https://doi.org/10.2307/1600213)
- Thrall, J. H., Li, X., Li, Q., Cruz, C., Do, S., Dreyer, K., & Brink, J. (2018). Artificial Intelligence and Machine Learning in Radiology: Opportunities, Challenges, Pitfalls, and Criteria for Success. *Journal of the American College of Radiology*, 15(3, Part B), 504-508. <https://doi.org/10.1016/j.jacr.2017.12.026>
- Tversky, A., & Kahneman, D. (1974). Judgment Under Uncertainty: Heuristics and Biases. *Science*, 185, 1124-1131. <https://doi.org/10.1126/science.185.4157.1124>

Utts, J. (2003). What educated citizens should know about statistics and probability. *The American Statistician*, 57(2), 74-79. <https://doi.org/10.1198/0003130031630>

Wing, J. M. (2006). Computational thinking. *Communications of the ACM*, 49(3), 33-35. <https://doi.org/10.1145/1118178.1118215>

This article is © 2019 by Alan M. Garber. The article is licensed under a Creative Commons Attribution (CC BY 4.0) International license (<https://creativecommons.org/licenses/by/4.0/legalcode>), except where otherwise indicated with respect to particular material included in the article. The article should be attributed to the author identified above.