living there. Indeed, the pursuit of science is a fascinating human endeavor, but the point of science is to objectively determine facts. Science does not necessarily provide subjectivity by itself which allows it to be influenced by meaning, moral values, and responsibility.[2] In the moral arena, people with religious beliefs, including Christians, are required to influence the idea of technologic determinism in a positive direction. I highly recommend this book not only to learn about an interesting part of world history but also to appreciate the tragedy of the human condition in the setting of war.

Notes

[1]Michael North, translator, "Greek Medicine," History of Medicine Division, National Library of Medicine, National Institutes of Health, last updated February 7, 2012, https://www.nlm.nih.gov/hmd/greek/greek_oath.html.

[2]Mehdi Golshani, "Science Needs a Comprehensive Worldview," *Theology and Science* 18, no. 3 (2020): 438–47.

Reviewed by John F. Pohl, MD, Professor of Pediatrics, Department of Pediatrics, University of Utah, Salt Lake City, UT 84113.

THE ETHICAL ALGORITHM: The Science of Socially Aware Algorithm Design by Michael Kearns and Aaron Roth. New York: Oxford University Press, 2019. 232 pages. Hardcover; $24.95. ISBN: 9780190948207.

Can an algorithm be ethical? That question appears to be similar to asking if a hammer can be ethical. Isn't the ethics solely related to how the hammer is used? Using it to build a house seems ethical; using it to harm another person would be immoral.

That line of thinking would be appropriate if the algorithm were something as simple as a sorting routine. If we sort the list of names in a wedding guest book so that the thank-you cards can be sent more systematically, its use would be acceptable; sorting a list of email addresses by education level in order to target people with a scam would be immoral.

The algorithms under consideration in *The Ethical Algorithm* are of a different nature, and the ethical issues are more complex. These algorithms are of fairly recent origin. They arise as we try to make use of vast collections of data to make more-accurate decisions: for example, using income, credit history, current debt level, and education level to approve or disapprove a loan application. A second example would be the use of high school GPA, ACT or SAT scores, and extra-curricular activities to determine college admissions.

The algorithms under consideration use machine-learning techniques (a branch of artificial intelligence) to look at the success rates of past student admissions and instruct the machine-learning algorithm to determine a set of criteria that successfully distinguish (with minimal errors) between those past students who graduated and those who didn't. That set of criteria (called a "model") can then be used to predict the success of future applicants.

The ethical component is important because such machine-learning algorithms optimize with particular goals as targets. And there tend to be unintended consequences—such as higher rates of rejection of applicants of color who would actually have succeeded. The solution to this problem requires more than just adding social equity goals as part of what is to be optimized—although that is an important step.

The authors advocate the development of precise definitions of the social goals we seek, and then the development of algorithmic techniques that help produce those goals. One important example is the social goal of privacy. What follows leaves out many important ideas found in the book, but illustrates the key points. Kearns and Roth cite the release in the mid-1990s of a dataset containing medical records for all state employees of Massachusetts. The dataset was intended for the use of medical researchers. The governor assured the employees that identifying information had been removed—names, social security numbers, and addresses. Two weeks later, Latanya Sweeney, a PhD student at MIT, sent the governor his medical records from that dataset. It cost her $20 to legally purchase the voter rolls for the city of Cambridge, MA. She then correlated that with other publicly available information to eliminate every other person from the medical dataset other than the governor himself.

Achieving data privacy is not as simple as was originally thought. To make progress, a good definition of privacy is needed. One useful definition is the notion of *differential privacy*: "nothing about an individual should be learnable from a dataset that cannot be learned from the same dataset but with the individual's data removed" (p. 36). This needs to also prevent identification by merging multiple datasets (for example, the medical records from several hospitals from which we might be able to identify an individual by looking for intersections on a few key attributes such as age, gender, and illness). One way to achieve this goal is to add randomness to the data. This can be done in a manner in which the probability of determining an individual changes very little by adding or removing that person's data to/from the dataset.

A very clever technique for adding this random noise can be found in a *randomized response*, an idea introduced in the 1960s to get accurate information in polls about sensitive topics (such as, "have you cheated on your taxes?"). The respondent is told to flip a coin. If it is a head, answer truthfully. If it is a tail, flip a second time and answer "yes" if it is a head and "no" if it is a tail. Suppose the true proportion of people who cheat on their taxes is p. Some pretty simple math shows that with a sufficiently large sample size (larger than needed for surveys that are less sensitive), the measured proportion, m, of "yes" responses will be close to $m = \frac{1}{4} + \frac{1}{2}p$. We can then approximate p as $2m - \frac{1}{2}$, and still give individuals reasonable deniability. If I answer "yes"

and a hacker finds my record, there is still a 25% chance that my true answer is "no." My privacy has been effectively protected. So we can achieve reasonable privacy at the cost of needing a larger dataset.

This short book discusses privacy, fairness, multiplayer games (such as using apps to direct your morning commute), pitfalls in scientific research, accountability, the singularity (a future time where machines might become "smarter" than humans), and more. Sufficient detail is given so that the reader can understand the ideas and the fundamental aspects of the algorithms without requiring a degree in mathematics or computer science.

One of the fundamental issues driving the need for ethical algorithms is the unintended consequences that result from well-intended choices. This is not a new phenomenon—Lot made a choice based on the data he had available: "Lot looked about him, and saw that the plain of the Jordan was well watered everywhere like the garden of the LORD, like the land of Egypt …" Genesis 13:10 (NRSV). But by choosing that apparently desirable location, Lot brought harm to his family.

I have often pondered the command of Jesus in Matthew 10:16 where he instructs us to "be wise as serpents and innocent as doves." Perhaps one way to apply this command is to be wise as we are devising algorithms to make sure that they do no harm. We should be willing to give up some efficiency in order to achieve more equitable results.

Reviewed by Eric Gossett, Department of Mathematics and Computer Science, Bethel University, St. Paul, MN 55112.

A WORLD WITHOUT WORK by Daniel Susskind. New York: Metropolitan Books, 2020. 305 pages. Hardcover; $28.00. ISBN: 9781250173522.

Will AI systems inevitably displace humans from employment? While computer and AI technology continue to advance at astronomical rates, the popular concern is often of an apocalyptic future where highly intelligent robots have taken over (e.g., *Terminator, Matrix*, etc.). In his book, *A World without Work*, Daniel Susskind predicts the current capabilities of technology will lead to a future in which powerful AI systems can do many of the jobs held by humans. Susskind therefore believes that the proliferation of AI systems will lead to a future "world without enough work for everyone to do" (p. 5). With his expertise in economics, Susskind explores how the continued advanced of technology will have profound effects on future employment, growing inequality, and the methods whereby humans find meaning and purpose.

The book is divided into three sections. In the first section, Susskind sets out the historical context of technological advancements and their effects on employment and economics. He highlights how the early advancements of computer technology were often met with disappointment as creators found it exceedingly difficult to create a machine that could replicate human intelligence. However, this early disappointment led humans to underestimate the efficiency of AI systems in performing tasks that are easy to automate (or what Susskind refers to as "routines").

In the second section, the discussion shifts to exploring how the increased power and affordability of machines enable them to perform more human roles. The fear of increasing unemployment due to technological advancement is a real fear. Susskind differentiates between two types of technological unemployment: frictional and structural. While frictional unemployment (humans not having the skills to perform a job) is certainly an issue, structural unemployment (there actually being too few jobs for everyone) is the more pressing problem. The threat of rising unemployment leads Susskind to predict that economic inequality will grow since only certain people will be able to acquire well-paying jobs.

In the third and final section, Susskind tries to provide a solution to the growing unemployment problem. He claims the attempted solution of technology education fails as a long-term response since not all people have the disposition to learn about technology, nor will there be enough jobs. A potential solution is to provide a UBI (universal basic income) for all people so that the economic inequality will not be so severe. However, Susskind rejects the UBI solution in favor of his proposed CBI (conditional basic income) which still provides income but with requirements that must be met. Susskind believes his proposed CBI solution has the added benefits of solving the inequality problem and providing meaning and purpose that a job once held.

Computer and AI technology are certainly advancing at a rapid rate. Susskind is not alone in his warnings regarding the potential dangers of technological advancements. However, Susskind helpfully points out that the danger does not come from machines gaining sentience and oppressing humans but, rather, the danger is one of machines gradually replacing us in our employment due to their overwhelming speed and efficiency. While there is relief that such an apocalyptic future is unlikely, the prediction of a future without enough work to go around ought to be a significant concern.

While Susskind's prediction of a future with significantly reduced employment is well founded, his potential solution of implementing a CBI to provide the meaning and purpose lost from unemployment seems incomplete. With jobs no longer providing the sense of meaning and purpose, it is difficult for Susskind to find a solution to fulfilling these existential longings can be fulfilled. Unfortunately, he is unwilling to seriously consider a religious answer to these existential