

Book Reviews

contributing to the larger society, carrying on family and cultural traditions, and providing the necessary education for a good future. Again, these are primarily Western reports and Mahoney reminds the reader that other countries' perspectives are needed. Like non-Western studies, studies of nontraditional parenting units, such as single parents, same-sex parents, and economically disadvantaged parents, are underrepresented. Furthermore, the type of theistic schema provides another area of diversity that is lacking, as children can be reared in polytheistic, nontheistic, atheist, or agnostic environments.

Mahoney's final section looks at social and cognitive-developmental research. Concepts such as theory of mind and attachment enter the scene. The primary area that has been studied in children's RS development is their concept of God. Preliminary findings suggest that children's perceptions of God mirror how they are being parented (e.g., punishing parents → punishing God, nurturing parents → nurturing God, powerful parents → powerful God, etc.). Examining children's prayers also sheds some light on RS development, though again findings are mixed and limited. There is more work to be done.

Mahoney calls on social scientists to take the lead in providing guidance to parents to uphold the United Nations' 1989 *Convention on the Rights of the Child* Article 14, 1–2 that states: "States Parties shall respect the right of the child to freedom of thought, conscience and religion; States Parties shall respect the rights and duties of the parents and, when applicable, legal guardians, to provide direction to the child in the exercise of his or her right in a manner consistent with the evolving capacities of the child."³ More intentional investigation of children around the globe can help parents directly but also inform policy makers. Mahoney states that "one central observation is that this literature is in its infancy stage" (p. 62).

Overall, Mahoney's review of children's RS development in this volume is thorough yet concise, troubling yet hopeful, vague yet nuanced. She concludes with six key areas and related findings to recap how the scientific study of children's RS development can be improved in the years to come. Thankfully, RS has begun to attract significant attention in the field, including from the Templeton Foundation's attempt to build a more global community of social scientists.⁴ After reading this book, I feel much better equipped to elucidate what is known and what is yet to be discovered. This is important, not only in academic communities of colleagues and students,

but also in the broader communities of church and society and in our personal communities.

Notes

¹Harold G. Koenig, Michael E. McCullough, and David B. Larson, *Handbook of Religion and Health* (New York: Oxford University Press, 2001).

²James W. Fowler, *Stages of Faith: The Psychology of Human Development and the Quest for Meaning* (New York: Harper-SanFrancisco, 1981).

³United Nations Human Rights, "Convention on the Rights of the Child," *Treaty Series* 1577, no. 3 (1989): 1–23, <https://www.ohchr.org/en/professionalinterest/pages/crc.aspx>.

⁴J. D. Warren, "\$10 Million Grant Will Study Children's Religious Views," University of California, Riverside, February 19, 2020, <https://news.ucr.edu/articles/2020/02/19/10-million-grant-will-study-childrens-religious-views>.

Reviewed by Erin Mueller, Professor of Psychology, Northwestern College, Orange City, IA 51041.



TECHNOLOGY

GENIUS MAKERS: The Mavericks Who Brought AI to Google, Facebook, and the World by Cade Metz. New York: Dutton, 2021. 371 pages including notes, references, and index. Hardcover; \$28.00. ISBN: 9781524742676.

As Cade Metz says in the acknowledgments section, this is a book "not about the technology [of AI] but about the people building it ... I was lucky that the people I wanted to write about were so interesting and so eloquent and so completely different from one [an]other" (p. 314).

And, that's what this book is about. It is about people such as Geoff Hinton, founder of DNNresearch, who, once he reached his late fifties, never sat down because of his bad back. It is about others who came after him, including Yann LeCun, Ian Goodfellow, Andrew Ng, Yoshua Bengio, Jeff Dean, Jürgen Schmidhuber, Li Deng, Ilya Sutskever, Alex Krizhevsky, Demis Hassabis, and Shane Legg, each of whom had their strengths, weaknesses, and quirks.

The book also follows the development of interest in AI by companies like Google, Microsoft, Facebook, DeepMind, and OpenAI. DeepMind is perhaps the least known of these. It is the company, led by Demis Hassabis, that first made headlines by training a neural network to play old Atari games such as Space Invaders, Pong, and Breakout, using a new technique called reinforcement learning. It attracted a lot of attention from investors such as Elon Musk, Peter Thiel, and Google's Larry Page.

While most companies were interested in the application of AI to improve their products, DeepMind's goal was AGI, "Artificial General Intelligence" – technology that could do anything the human brain could do, only better. DeepMind was also the first company to take a stand on two issues: if the company was bought out (which it was, by Google), (1) their technology would not be used for military purposes, and (2) an independent ethics board would oversee the use of DeepMind's AGI technology, whenever that would arrive (p. 116).

Part One of the book, "A New Kind of Machine," follows the early players in the field as they navigate the early "AI winters," experiment with various new algorithms and technologies, and have breakthroughs and disappointments. From the beginning, there were clashes between personalities, collaboration and competition, and promises kept and broken.

Part Two of the book, titled "Who Owns Intelligence?," explores how many of the people named above were wooed by the different companies, and moved back and forth between them, sometimes working together and sometimes competing with each other. The companies understood the power of neural networks and deep learning, but they could not develop the technologies without the direction of the leading researchers, who were in limited supply. To woo the best researchers, the companies competed to develop exciting and show-stopping technology, such as self-driving cars and an AI to play (and beat) the best in Chess and Go.

In Part Three, "Turmoil," the author explores how the players began to realize the shortcomings and potentially dangerous effects of the AI systems. AI systems were becoming more and more capable in a variety of tasks. "Deep fakes" of celebrities and the auto-generation of fake news (often on Facebook) led many to question the direction AI was going. Ian Goodfellow said, "There's a lot of other areas where AI is opening doors that we've never opened before. And we don't really know what's on the other side" (p. 211). One surprising figure taking a stand on the side of caution was Elon Musk, giving repeated warnings of the possible rise of superintelligent actors. Further, it was discovered that the Chinese government was already using AI to do facial recognition and track its citizens as they moved about.

Other concerns dampened the community: it was discovered that small and unexpected flaws in training could have significant effects on the ability of an AI system to do its job. For example, "by slapping a

few Post-it notes on a stop sign, [researchers] could fool a car into thinking it wasn't there" (p. 212).

Additionally, the biases in training data were being exposed, leading some to believe that AI systems would not equally benefit minority groups, and could even discriminate against them. Furthermore, Google was being approached by the US government to assist in the development of programs which could be used in warfare. Finally, Facebook was struggling to contain fake news and finding that even AIs could not effectively be used to combat it.

In the final sections of the book, the author explores the AI researchers' attitudes toward the future and the big questions. Will AI systems be able to eventually take over all work, even physical labor? Can the AI juggernaut be controlled and directed? Will AGI be fully realized?

This last question is explored in the chapter titled "Religion." "Belief in AGI required a leap of faith. But it drove some researchers forward in a very real way. It was something like a religion," said roboticist Sergey Levine (p. 290). The question of the feasibility of AGI continues to generate much debate, with one camp claiming that it is inevitable, while the other camp insisting that AI systems will excel only in limited tasks and environments.

As a Christian, I found the debates about the proper role of AI to be intriguing. Is the development of AGI inevitable? Should we as Christians petition companies and governments to have debates on the pursuit of AGI? Should we enact laws to limit or prohibit the use of AI in warfare? Should independent evaluators be required to review AI systems regarding discrimination? Should Christians participate in the further development of AGI?

Learning the histories and attitudes of the leading individuals in the development of AI also intrigued me. Many of the individuals seem to have very little concern for the potentially negative impact of their work. Their only motivation seems to be fame and fortune. It makes me wonder if the field of computer science should require all its practitioners to take ethics training like professional engineers are required to do. This book certainly confirms the importance of ethics in the field of computer science and the need for its practitioners to be people of virtue.

In summary, this was a different kind of book from many others in the field of technology. It was fascinating that so much of what I was reading about had

Book Reviews

happened in just the last ten years. Hearing the anecdotes of back-office meetings, public outcries, and false claims was intriguing. If you, like me, wonder how we got to where we are today in the area of AI, this is the book for you.

Reviewed by Victor T. Norman, Assistant Professor of Computer Science, Calvin University, Grand Rapids, MI 49546.

TEACHING MACHINES: The History of Personalized Learning by Audrey Watters. Cambridge, MA: The MIT Press, 2021. 313 pages. Hardcover; \$34.95. ISBN: 9780262045698.

Teaching Machines, by freelance writer, researcher, and technology commentator Audrey Watters, is a history framed by a critical rallying cry. The main body of the book is a history of the development and demise of “teaching machines” (mechanical devices for self-paced, programmed instruction) from the 1920s to the 1960s. It attends closely to the extent and limits of the influence of B. F. Skinner (and his forerunner Sidney Pressey), the role of commercial interests and processes, the development of a receptive social imaginary through popular media, the inconclusive nature of empirical findings about the learning that resulted, the eclipse of the mid-century teaching machine by programmed learning in book form, and the rise of computers. This account by itself might seem a little arcane. It is, however, given added heft by a framing argument that ties the history of teaching machines to present-day trends, and critiques some common myths regarding the history of educational technologies that are used to sell current technological options. This framing argument contends, on the one hand, that the “Silicon Valley mythology” (p. 249), regarding education’s digital future, rests on misinformation about the past, and, on the other hand, that current digital developments have more continuity with the behaviorist and totalitarian impulses of that past than is commonly admitted.

Concerning the former point, Watters points to a common narrative purveyed by figures such as Sal Khan and Bill Gates that presents education as beset by a static factory model rooted in the nineteenth century and buttressed by resistance to change on the part of Luddite educators. The solution then comes in the form of commercially sourced digital tools that now offer revolutionary degrees of individualization and access to learning. Watters’s account undermines both halves of this story. She marshals a substantial body of evidence to show that education has been far from static over the past century, that techno-

logical innovations designed by educators regularly stalled due to inertia and disorganization on the part of the business world, and that the rhetoric of revolutionary individualization and personalization of learning has been the stock-in-trade of purveyors of a long string of new educational technologies but has also consistently fallen short in practice. A generous amount of space is devoted to B. F. Skinner’s bouts of epistolary fury directed at his business partners who stalled development of his teaching machines until their moment had passed. More significantly, Watters makes clear that the recurring claim of individualization came within a recurring and expanding envelope of standardization. Proponents of teaching machines made much of the potential for individualized instruction, understood as the capacity for learners to proceed at their own pace. Those same learners were expected to follow programmed sequences, assemble predetermined atoms of knowledge, prepare for standardized tests, and submit to a rather deterministic process of behavioral manipulation. The talk of individualization may perhaps have been sincere, but it amounted in the end to something comparable to today’s processes of “personalizing” your smartphone by choosing the same device as millions of others in one of a handful of colors, or perhaps clicking on the same online instructional video, framed by the same perspective, as everyone else. In the meantime, the appeal to individualization helped to shift product.

The suggestion of contemporary parallels points to the second part of the book’s framing agenda, which claims that teaching machines were not just a curious episode that met its demise with the rise of computing. Watters points out that claims to revolutionary breakthroughs in education through technology commonly end up looking oddly conservative. Dreams of technocratic learning and robot teachers in the 1950s and 1960s still placed the robots in front of classrooms with rows of chairs in which students answered multiple-choice questions. Watters suggests that contrary to some tellings of the story, the teaching machines of the day did not give way to computers so much as help to establish assumptions about programmed learning rooted in behavioral manipulation, atomization of content, and linear progress that continue to inform today’s digital educational technologies. The commercial involvement in all of this is, moreover, far from disinterested, with considerable research and design acumen going into the creation of digital products that reinforce behaviors favorable to those who make their living from eyeballs remaining on webpages and apps. After a