is the potential for misaligned objectives. Cybernetics founder Norbert Wiener wrote in 1960, “We had better be quite sure that the purpose put into the machine is the purpose which we really desire.” Or, as Tegmark has it, “It’s unclear how to imbue a superintelligent AI with an ultimate goal that neither is undefined nor leads to the elimination of humanity.” In my view, this technological and philosophical problem demands all the intellectual resources we can bring to bear.

Only if we solve it can we reap the benefits. Among these is expansion across the Universe, perhaps powered by such exotic technologies as Dyson spheres (which would capture the energy of a star), accelerators built around black holes or Tegmark’s theorized sphalerizers (like diesel engines, but quark-powered and one billion times more efficient). For sheer science fun, it’s hard to beat the explanations of how much upside the Universe and the laws of physics will allow. We may one day, for example, expand the biosphere “by about 32 orders of magnitude”. It’s seriously disappointing, then, to learn that cosmic expansion may limit us to settling only 10 billion galaxies. And we feel our descendants’ anxiety as “the threat of dark energy tearing cosmic civilizations apart motivates massive cosmic engineering projects”.

The book concludes with the Future of Life Institute’s role in moving these issues into mainstream AI thinking — for which Tegmark deserves huge credit. He is not alone, of course, in raising the alarm. In its sweeping vision, *Life 3.0* has most in common with Nick Bostrom’s 2014 *Superintelligence* (Oxford University Press). Unlike Bostrom, however, Tegmark is not trying to prove that risk is unavoidable; and he eschews dense philosophy in favour of asking the reader which scenarios they think more probable or desirable.

Although I strongly recommend both books, I suspect that Tegmark’s is less likely to provoke in AI researchers a common allergic reaction — a retreat into defensive arguments for paying no attention. Here’s a typical one: we don’t worry about remote but species-ending possibilities such as black holes materializing in near-Earth orbit, so why worry about superintelligent AI? Answer: if physicists were working to make such black holes, wouldn’t we ask them if it was safe?

The Economist has drily characterized the overarching issue thus: “The implications of introducing a second intelligent species onto Earth are far-reaching enough to deserve hard thinking.” *Life 3.0* is far from the last word on AI and the future, but it provides a fascinating glimpse of the hard thinking required. ■

Stuart Russell is professor of computer science at the University of California, Berkeley and co-author of *Artificial Intelligence: A Modern Approach*. e-mail: russell@berkeley.edu

Books in brief

Willy Ley: Prophet of the Space Age

Jared S. Buss UNIVERSITY PRESS OF FLORIDA (2017)

Among the scientists who fled Europe for the United States in the 1930s and 1940s was Willy Ley, first historian of spaceflight. In Jared Buss’s nuanced biography, the German-born rocket expert emerges as a spirited science educator whose promotion of space exploration paved the way for NASA’s triumphs. By 1960, Ley’s romanticized vision of science had been occluded by the harder-edged approach of his fellow émigré and collaborator, aerospace wizard Wernher von Braun — but only partly. As Buss reveals, Ley’s “reenchantment” of big science with a sense of wonder has held.

Uneasy Street: The Anxieties of Affluence

Rachel Sherman PRINCETON UNIVERSITY PRESS (2017)

There have been many cogent analyses of US income inequality. Sociologist Rachel Sherman’s welcome addition probes the psychology and socio-economics of affluence. Sherman interviewed 50 New Yorkers, from academics to financiers, with incomes between US\$250,000 and more than \$10 million. She found complex adaptations to privilege, from ethical acrobatics normalizing extreme consumption to defensive assertions that wealth is “earned”. Most of her interviewees focus on crafting a self-legitimizing moral universe rather than addressing the structural inequalities that trigger their unease.

What It's Like to be a Dog

Gregory Berns BASIC (2017)

The hyper-curiosity of a pet terrier spurred neuroscientist Gregory Berns to probe the dog’s brain — an experiment demanding that she stay awake and unrestrained in a magnetic resonance imaging scanner. Many more dogs and trials later, the project has generated numerous findings on canine brain function, such as how owners’ smells alone activate the reward response. This is just one “adventure in animal neuroscience” that Berns pursues, as he investigates beat synchronization in sea lions, auditory pathways in dolphins and the putative behaviour of the extinct thylacine.

Frankenstein Dreams

Michael Sims BLOOMSBURY (2017)

Victorian upheavals in science and technology jolted the era, and the imaginations of its speculative-fiction writers. This nightshade nosegay of vintage sci-fi tales and excerpts, plucked by writer Michael Sims, richly reflects the shifts. Along with miniature masterworks by the likes of Mary Shelley, Edgar Allan Poe, H. G. Wells and Ambrose Bierce are unexpected gems. William Henry Rhodes’s ‘The Telescopic Eye’, for instance, features chariot-wheel-shaped aliens called Lunarians; and Mary Wilkins Freeman’s ‘The Hall Bedroom’ is a hallucinatory journey into the fifth dimension.

DNA: The Story of the Genetic Revolution

James D. Watson, with Andrew Berry and Kevin Davies KNOPF (2017)

In this update of *DNA: The Secret of Life* (Knopf, 2003), James Watson — the controversially outspoken co-discoverer of DNA’s structure — remaps the genetic landscape. There are new chapters on personal genetics and cancer treatment; key discoveries, such as the CRISPR revolution kick-started by Emmanuelle Charpentier, Jennifer Doudna and Feng Zhang; and findings in epigenetics and agricultural chemistry. Backing nature over nurture, Watson ends by arguing (with caveats) in favour of future germline gene editing. **Barbara Kiser**