Otteson does, using property rights as a given, is circular. Indeed it is their differing conceptions of justice that informs the disagreement between them.

It is hard to read this book without engaging with the authors, which makes it an excellent choice for a mid or upper-level course, though not for an introductory course, where the students are likely to be confused by the seemingly compelling arguments on both sides.

One of the book's pedagogical virtues is the summary paragraphs after each main section. These are short, clear, and relatively thorough. In addition, there are definitions offered along the way whenever the author (or the editor) thinks a concept will be unfamiliar to the reader. These definitions are supplemented at the end by an extensive glossary, and each author provides a list of suggested readings. There is a short forward by Michael Munger that purports to summarize the two positions.

Michael Goldman, Professor Emeritus, Miami University, 1420 Locust Street, 18i, Philadelphia PA 19102; goldmam@miamioh.edu.

Great Philosophical Objections to Artificial Intelligence: The History and Legacy of the AI Wars

Eric Dietrich, Chris Fields, John P. Sullins, Bram van Heuveln, and Robin Zebrowski

London: Bloomsbury Academic, 2021; pbk, 305 pp., $42.95; ISBN 978-1-4742-5710-7

TIM JUVSHIK

Great Philosophical Objections to Artificial Intelligence (hereafter "GPOAI") is a careful and insightful overview of the history of the development of, and philosophical objections to, artificial intelligence with human-level intelligence. This is typically referred to as artificial general intelligence or AGI, though not in the present text. Dietrich, Fields, Sullins, van Heuveln, and Zebrowski ("the authors") have written a comprehensive and accessible, if a little tonally and conceptually uneven, history of the main philosophical objections to AGI and how they relate to AI issues of concern today. Overall, the authors succeed in their goal of giving a historical overview of the main debates. With respect to its applicability to pedagogical contexts, GPOAI is, to my mind, best suited for using individual chapters in different classroom contexts. GPOAI is divided in two parts. Part 1 covers the four AI "wars" throughout the twentieth century, while Part 2 discusses three issues of central relevance to philosophical explorations of AI today.

Part 1 is divided into the four AI wars. The First War involves two early objections to the effect that AI isn't even possible, covered in chapters 1 and 2, respectively. The first objection stems from John Lucas's use of Gödel's First Incompleteness Theorem where he argues that minds cannot be explained as

 0145-5788 pp. 579–583
DOI: 10.5840/teachphil2023464206

purely mechanical or computational machines because there are formulas that cannot be proven within the system but which we can clearly see to be true. The authors do an excellent job of explaining clearly and accessibly how the objection fails by erroneously assuming AI to be bound to the object-language. The second objection covers well-known issues with the Turing Test and its usefulness as an assessment tool of whether machines are intelligent. Here the authors adopt a non-standard, but interesting, interpretation of the Turing Test as showing that if a machine can pass the test, then it has human-level intelligence, but that Turing didn't intend that we draw any conclusions about the possibility of machine intelligence if a machine failed the test. The authors read Turing as trying to avoid having to define "intelligence" by appeal to the (admittedly anthropocentric) test and cautioning us to avoid drawing any conclusions about the question "Can machines think?" from the litany of failed test cases. Rather, the test is a thought experiment that is intended to prompt us to consider under what circumstances we would ascribe intelligence to a machine. Chapter 1 on logical objections to AI could be used as a reading on its own, especially at the end of an introduction to formal logic class that covers Gödel's Completeness and Incompleteness Theorems as a specific application of those ideas to AI. Chapter 2 would also serve well as a standalone reading, but since it offers a non-standard interpretation of the Turing Test it would perhaps be best presented in contrast to standard appeals to the test, especially in conjunction with presentations of Searle's Chinese Room. It could also serve as a good discussion of the nature of intelligence in an epistemology course.

The Second War covers the question of how AI could be implemented from a logical, behavioural, and software perspective. Chapter 3 considers philosophical and psychological denials of the mind, in particular those of the logical positivists and behaviourists. The authors give a brief overview of the motivations and arguments of both camps and how and why they subsequently failed. Of special note is an excellent discussion of the metaphysics of algorithms and virtual machines which while short (50–55) would be useful as a primer in a number of pedagogical contexts. Chapter 4 considers the attempts to implement artificial intelligence in a machine given the failure of logical positivists and behaviourists to establish that there is no mind to implement. This chapter provides a comprehensive overview of the four main software architectures appealed to in the AI debates: symbol processors, connectionist and artificial neural networks, embodied and situated cognition, and dynamic systems approaches, as well as a short supplement on quantum computing. The discussion is sophisticated and somewhat technical, and some basic familiarity with computer science or software engineering would be beneficial, but the chapter also provides set up for a later discussion of AI as embodied agents. Chapter 3, on logical positivist and behaviourist criticisms of AI, is quite non-technical and could easily stand alone, making it suitable for a lower level minds and machines or philosophy of mind

course. Chapter 4 presents an excellent overview of the four main software architecture approaches to AI design and would likewise be suitable for a minds and machines or philosophy of mind course.

The Third War, covered in chapter 5, revolves around the classic question of whether machines can exhibit meaning beyond mere algorithmic processing, familiar from Searle's Chinese Room argument, and the authors give a fun and on-going example of a robot pencil sharpener. The chapter is lengthy and has a number of supplements, including why monointelligent AI aren't really intelligent, the standard responses to the Chinese Room, and addenda about consciousness and philosophical zombies. The authors mostly talk of aboutness, rather than meaning, and this is one of the few places where they go beyond historical surveying to offer their own take, namely that aboutness consists in consciousness, not just doing or interacting with the world, and that consciousness is required for both knowledge and mattering. The authors illustrate their view with their robot pencil sharpener, which isn't conscious, so they argue that its processes aren't "about" anything. Chapter 6 introduces the Fourth War, the frame problem—how to implement judgements of relevance in machines, something that humans do that isn't all that well understood. Creating AGI would require knowing how to implement such judgements beyond mere guesswork. The authors give a good overview of the problem and offer a fascinating interpretation of how AI researchers have systematically ignored the philosophical version of the frame problem, thereby limiting AI to monointelligences like Google's AlphaGo or IBM's Deep Blue. Chapters 5 and 6 offer more technical and extended discussions of their respective topics and because the authors offer their own views on aboutness and the frame problem, they may prove challenging in most undergraduate contexts. However, they could serve very well in an upper division or graduate seminar on AI or philosophy of mind, especially when combined with classic texts from the literature, and they complement each other well as an extensive treatment of two core issues but which go beyond historical surveys and argue for specific positions.

Part 2 covers three issues surrounding AI which are of current philosophical interest. Chapter 7 addresses the issue of consciousness and covers some familiar ground from chapters 5 and 6 but with an added discussion of both system 1 and 2 thinking and psychological and cognitive typologies of conscious experiences. The chapter could serve as a useful primer on the topic for graduate students or researchers for whom philosophy of mind isn't their specialty. However, it is a little technical, with a focus on psychological work which builds off of discussions of consciousness earlier in the book, so it doesn't stand on its own particularly well, especially for undergraduate classes, but it could be paired with chapters 5 and 6 to offer a comprehensive and sophisticated treatment of the challenges facing implementation of AGI from semantic, mental/psychological, and metaphysical quarters.

Chapter 8 provides a very brief overview of a new AI war which focuses on the ethics of weak AI applications. The authors gloss over the general move in ethics in this direction and then give two examples which aim to illustrate the complex, pressing ethical questions surrounding AI applications in use today, those of autonomous weapons systems and self-driving cars. Chapter 8 is, in my opinion, the major disappointment of the book. Ethics of applied AI is arguably the most pressing and popular topic in philosophy of AI and makes up the cornerstone of many philosophy of technology and technology ethics courses. Admittedly, this is a big subfield, and there are book length treatments of these topics already, as well as many exemplary articles on ethical issues surrounding a variety of particular AI applications but which aren't really discussed in the text (e.g., predictive algorithms on social media and in policing, chatbots and misinformation, AI use in healthcare diagnosis, robotic caretakers, etc.). The narrow focus on just two applications of weak AI limits the usefulness of the chapter as a class reading since both topics have received in-depth treatments elsewhere. The chapter also doesn't discuss any systematic approach to applied technology or AI ethics, but again there are many treatments of this topic in the literature already.

Chapter 9 is a lengthy discussion of whether embodied AIs could be ethical agents and is, to my mind, one of the highlights of the book, partly because it is one of the few places where the authors offer their own view. The chapter provides a discussion of why metaphysics of mind is relevant to ethics, how the hard coding approach to AI ethics has been a failure, the role of consciousness in AI ethics, and, most importantly, why AI ethics needs to attend to embodiment. Here the authors draw on some elements of embodied cognition from chapter 4, as well as various other threads throughout the book, including consciousness and aboutness, and argue that embodiment is needed for generating abstract concepts and having a social life that generates empathy. The authors also have a short discussion (255–60) of artificial phronesis and the potential for AI to reason morally which could arguably stand on its own as a short piece in an AI ethics course, though the authors are agnostic on whether AI could ever genuinely achieve human-level practical moral reasoning. In contrast to the previous chapter, chapter 9 is an excellent treatment of a difficult topic in part because it argues extensively for a specific view. It does a great job of showing the importance of the intersection of AI and robotics and would complement discussions of phenomenology and its relevance to analytic philosophy of mind. As a result, the chapter could serve well in an upper division philosophy of mind class or a minds and machines class or a graduate seminar on these topics, especially if presented as a foil to dualist views and their bearing on AI.

Overall, GPOAI provides a thorough and accessible survey of the history of philosophical objections to AI over the past seventy years. The book is a little uneven in tone, with some chapters being more accessible and less technical than others. To its benefit, most chapters can stand alone, which

makes them more useful for specific courses. The authors do an excellent job in surveying the four distinct AI wars and explaining both historically and philosophically why they seem to have quieted down. While this historical aspect of the book should be lauded, there is simultaneously a paucity of content on applications of applied weak AI, which is relegated to a single twenty-page chapter. This is all the more odd since the authors describe these applications as the new AI wars. It's also unfortunate that GPOAI was published the year before Chat-GPT was launched since its uses put these issues front and centre of philosophical and ethical interest in applied AI. Finally, the book would have benefited from more copy editing since there are numerous typos, a fair number of citations made in the text that aren't included in the bibliography, and a surprisingly short (2 pages) index. Nonetheless, despite the various criticisms above, the authors do an admirable job of presenting both sides of the various arguments and the text is at its best when they argue for their own specific views on these topics. As indicated above, individual chapters would fit well in a wide variety of pedagogical contexts. However, it may be best used as a primer on philosophy and AI by graduate students or researchers in philosophy or cognate disciplines who don't specialize on these issues.

Tim Juvshik, Department of Philosophy, Middlebury College, Twilight Hall, 50 Franklin Street, Middlebury VT 05753; tjuvshik@gmail.com.

The Creative Animal: How Every Animal Builds Its Own Existence
Roberto Marchesini
London: Palgrave Macmillan, 2022; paperback, 389 pp., $159.99; ISBN-13 978-3-031-07416-5

MIRIANA MAIO

The Creative Animal: How Every Animal Builds its Own Existence is an interesting book that aims to clearly explain animal creativity Animal *creativity*, or C-factor, represents the core idea of the book. Roberto Marchesini emphasises its importance throughout. In simple words, creativity is the capacity to adapt to new situations. Apart from creativity, this book has three keywords: *endowments*, *subjectivity* and *freedom*.

The book has twelve chapters, all connected as a "path to strictly follow." That is why I heartily recommend reading them in order. At the end of each chapter, there is a list of works cited. In this way, readers can immediately consult the references after finishing one chapter.

In chapter 1, *In Praise of Improvisation*, the role of *endowments*, or innate predispositions, is immediately defined. The author stresses that non-human animals (henceforth, "animals") are protagonists and not passive stimuli receptors. Animals use their innate predispositions as tools. These innate

 0145-5788 pp. 583–585
DOI: 10.5840/teachphil2023464207