humanitarian law may 'lose legitimacy', this does not necessarily affect their right to consent to foreign interventions under the *jus ad bellum* (though a supporting State might be in violation of IHRL and international humanitarian law).

The important role of the elusive concept of legitimacy in the recognition of governments that Redaelli discusses has most recently been illustrated in the context of Afghanistan, where the Taliban violently seized control on 15 August 2021. Three months later, no State has formally recognised the Taliban government and, at the time of writing, the Taliban's request to represent Afghanistan in the United Nations (UN) General Assembly is being considered by the UN Credentials Committee.

Redaelli does an excellent job establishing positions and trends in relation to legitimacy of origin, effectiveness, and legitimacy of exercise in the context of civil wars. However, it would have been helpful to explore further how economic and social rights are addressed. Discussion of 'gross and systemic' human rights violations tend to centre on civil and political rights, but it is important to remember that oppression can take many forms. For example, it would be interesting to consider whether and how the legitimacy of a government would be affected where it prioritises foreign investments over its people's right to water. Or, indeed, where a government that has seemingly endless resources available for its military capabilities still allows large numbers of its people to go hungry. These are certainly questions to grapple with for the future.

Chiara Redaelli should be congratulated for this impressive and valuable contribution to international legal scholarship, which will be returned to for years to come.

Marie Aronsson-Storrier*

The Reasonable Robot: Artificial Intelligence and the Law by Ryan Abbott [Cambridge University Press, Cambridge, 2020, viii + 156pp, ISBN: 978-1-108-47212-8, £85 (h/bk), £23 (p/bk)]

It is a curious feature of the history of artificial intelligence (AI) that its successes have often been measured in games. Early programs were taught bounded problems like tic-tac-toe and draughts. These were novelties, but the defeat of chess world champion Gary Kasparov by IBM's Deep Blue in 1997 was presented as a threat to the intellectual dominance of humanity—comparable, perhaps, to the Cold War rivalry that had played out in the match pitting Bobby Fischer of the United States against the Soviet Union's Boris Spassky quarter of a century earlier. Another 25 years on, and the machines have beaten us in even more complex games, such as *Go*, as well as idiosyncratic ones, such as *Jeopardy!*.

Ludology offers a relatable measure of machine achievement. Yet it is curious because such games are, by definition, meant to be *fun*. Deep Blue, AlphaGo, and other AI systems have many qualities, but the ability to have fun is not among them. Another explanation might be that we focus on trivial measures because it makes the advances of our metal and silicon creations seem less threatening.

As Ryan Abbott's *The Reasonable Robot* makes clear, those advances will affect every aspect of human society and economy. That much we have heard before, from the World Economic Forum's breathless talk of a Fourth Industrial Revolution to prophetic warnings of the coming singularity. Abbott's contribution is to try to offer clarity in how law should respond.

Many attempts tend to follow Isaac Asimov, articulating rules to shape AI behaviour. The past five years has seen hundreds of lists, most failing to understand that Asimov's literary career was built on the fact that his 'Three Laws of Robotics' might have been wonderful in theory but did not work in practice. Abbott predicates his own approach not on *what* the rules should be so much as *how* regulators should develop them. The new guiding tenet, he argues, should be 'AI legal neutrality'.

*Lecturer, School of Law, University of Reading, e.m.l.aronssonstorrier@reading.ac.uk.

doi:10.1017/S0020589321000415

That tenet, appropriately enough, works in both directions: the law should not discriminate against AI nor against humans. At present it does both. AI systems, for example, will (eventually) be safer drivers than humans, but are prohibited from plying the roads. Humans, by contrast, may be better at customer-facing jobs, but are being replaced by machines to save on taxes. Neutral legal treatment, he argues, 'would ultimately benefit human well-being by helping the law better achieve its underlying policy goals'.

Abbott is not claiming that AI systems should be treated as if they *are* persons, with rights or legal personality. His more subtle argument is that a presumption of neutrality helps clarify areas of the law in which AI should be treated more like humans, and where humans may sometimes need to be treated more like AI. This, he argues, will promote competition, improve safety, incentivise innovation, and reduce antisocial behaviour.

In addition to being a law professor, Abbot is a physician and a patent attorney. Both qualifications are on display in his research. In medicine, diagnostics requires the analysis of signs and symptoms to identify a disease, condition, or injury. A human doctor relies on his or her training and experience, which might include years of study and decades of seeing hundreds of patients or more. An AI system can be programmed with every textbook ever printed and millions of patient records. Misdiagnoses and accidents will occur in either situation, but how should the law respond?

Abbott warns that imposing strict liability on AI will discourage innovation. In areas where AI may ultimately be safer than humans—medical care, transportation—a negligence standard that treats AI 'like a person' would more appropriately weigh the costs and benefits of automation. A decade ago, Ryan Calo went further and argued that robot manufacturers needed immunity to remove the uncertainty of potential lawsuits. Neither approach has been embraced by any major jurisdiction, and yet there does not appear to have been any appreciable slowdown in research and development. (China may offer a counter-example, where its dominance in AI is often attributed to minimal restrictions on data collection and tort law is comparatively underdeveloped. Even there, however, data protection laws and limits on technology companies are being strengthened, rather than weakened.)

It is in the field of intellectual property that Abbott's work spills off the page and into the real world. Even as he wrote the book, his argument that AI already generates intellectual property was being presented not just to Cambridge University Press but in courts around the world. Again, the nuance of his position is important: he is not proposing that AI systems should 'own' their creations but highlighting a gap in the law that AI-created patentable inventions are not recognised for want of a natural person who qualifies as the 'inventor'.

Working with computer scientist Stephen Thaler, Abbott named the AI system DABUS as the inventor of a functional container design and a type of emergency signal and applied for patents in various jurisdictions. Initial responses were not promising. The British Intellectual Property Office, the European Patent Office, and the US Patent and Trademark Office all rejected the application on formal grounds: the relevant legislation presumes a human inventor. Only the first of these cases had been decided when his book went to press. The second and third might have been disheartening.

Then in July 2021, a year after publishing *The Reasonable Robot*, Abbott's team prevailed in both South Africa and Australia within the space of a week. If he is correct and AI inventions become more common, these cases will be seen as landmark decisions.

While I agree with much of what Abbott argues, his guiding tenet of neutrality—as he himself acknowledges—cannot answer many of the regulatory problems thrown up by technology. The final section of the book includes a case-by-case examination of different areas of practice to determine whether specific sectors warrant change or not. But in areas such as tax, the lens he offers does help clarify policy questions that governments may be loath to confront. One might have a debate about whether AI should replace workers in a particular area of the economy, for example, but the more immediate argument should be whether governments ought to be encouraging that replacement through taxes that are higher on labour than on capital.

AI systems are destined to surpass humans in many areas, perhaps most areas. Abbott is sanguine about this, comparing it to long-distance running. Once of practical importance for transmitting messages or packages, the advent of telecommunications and transportation made distance running redundant. Yet people still compete in marathons. Similarly, people still play games like chess and *Go*—even if we know that a computer somewhere could thrash us at both.

The shifting reasons *why* we do such things are questions for philosophy and psychology. Here Abbott stays in his lane, noting in a deadpan aside that 'it is beyond the scope of this book to establish the meaning of life'. Fair enough, but his emphasis on a purely utilitarian approach to AI regulation is itself a moral position.

Though well-written and admirably concise, the prose sometimes slides into techno-utopianism: 'Once superintelligent inventive AI is run-of-the-mill, the financial costs of innovating will be trivial, the push to incentivize will be unnecessary, and future innovation will be self-sustaining.' It is atypical for a lawyer (or a doctor) to be so optimistic, but perhaps he is looking forward to a future in which academics continue to write in the same way we continue to play chess: for the fun of it, rather than because there is anything truly new to say that the machines have not thought of already.

Until that time, *The Reasonable Robot* is an accessible and illuminating account of the problems AI poses for law—and those that law, for the time being, might pose for AI.

Simon Chesterman*

The United Nations Convention against Torture and its Optional Protocol: A Commentary by Manfred Nowak, Moritz Birk and Giuliana Monina [2nd edn, Oxford University Press, Oxford, 2019, 1306pp, ISBN: 978-0-19-884617-8, £262.50 (h/bk)]

The first edition of this Commentary was published to much acclaim in 2008. This second edition appeared in 2019 and now takes account of the work of both the Committee against Torture (CAT) and the Sub-Committee for the Prevention of Torture (SPT) to March 2017, with some later material also drawn upon. To describe this as an 'update' would fail to do justice to its significance. When the first edition appeared, the SPT had only just come into being and there was little real practice for it to take account of at all. Part II of this Commentary which addresses the Optional Protocol now takes full account of the nearly ten years of work by the SPT and represents the first systematic account of this. Similarly, the work of the CAT has continued to expand and evolve since the publication of the first edition and this too has been carefully worked into the already expansive and detailed text.

First and foremost, this is a commentary on the articles of the Convention and the Optional Protocol. Each article is presented in turn and is considered according to a common format. Following a short contextual introduction, the travaux préparatoires are presented and analysed. This is then followed by a section exploring 'issues of interpretation' arising from the text. Some of these relate to the interpretation of the terms found in the article in question, others relate to more holistic issues concerning each article as a whole. So, for example, the section addressing Article 1 of the Convention, its definition, looks at both the meaning of each of the principal definitional elements, but also whether the article actually establishes a discrete obligation not to torture—something which is usually assumed but, as the Commentary rightly points out, is not expressly provided for.

The first edition was structured slightly differently, with issues of 'interpretation' presented following the travaux préparatoires, as now, but then followed by a section looking at the work of the committee in relation to the article in question separately. In retrospect this was somewhat unfortunate, since the issues identified were often looked at twice: from a textual perspective and then from the committee's perspective. Not only did this lead to some unnecessary repetition in

*Dean and Professor, National University of Singapore Faculty of Law; Senior Director of AI Governance, AI Singapore, chesterman@nus.edu.sg.

doi:10.1017/S0020589321000452