Irish Math. Soc. Bulletin
Number 96, Winter 2025, 93–95
ISSN 0791-5578

Richard Susskind: How To Think About AI, Oxford University Press, 2025.

ISBN: 978-019-8941-92-7, GBP 9.75, 123+xiv pp.

REVIEWED BY ZHENWEI LYU

In a time when artificial intelligence is reshaping industries, education, and everyday life, this book stands out as a clear and thoughtful guide to what AI truly is, what it might become, and how we should respond. It is not just about how AI works today, but about how it could evolve over the coming years, and what that means for all of us – individuals, professionals, governments, and society as a whole. Drawing from every part of the book's argument, especially as seen in the conclusion, the author urges us to stop focusing only on tools like ChatGPT or image generators and instead think about where these technologies are leading us. This book is not meant just for AI experts. It is written in a way that anyone interested in the future – teachers, students, business leaders, or policymakers – can understand and learn from. It offers a roadmap for navigating the coming changes with clarity, responsibility, and deep reflection.

One of the book's most useful ideas is the distinction between "process-thinking" and "outcome-thinking". These two ways of thinking help make sense of many of the disagreements and debates people have about AI. Process-thinkers are those who focus on how AI systems work internally – data, algorithms, training models – while outcome-thinkers care more about what AI systems actually do in the real world – the results, the impact, the usefulness. The author shows how many arguments between experts happen because they are thinking in these different ways without realising it. For example, some thinkers worry that AI isn't "really" intelligent because it doesn't think like a human. Others respond that it doesn't matter how it works, as long as it gets the job done. The book doesn't say one view is right and the other is wrong. Instead, it helps us see that both are useful in different ways and that combining them gives us a clearer picture. This framework helps explain why big-name thinkers such as Noam Chomsky and Henry Kissinger often seem to disagree – they are coming from different "AI cultures".

In the middle chapters, especially Chapters 3 through 5, the author clears up many of the most common misunderstandings about AI. One of the strongest points made is about the "AI Fallacy" – the wrong belief that, for AI to be good at something, it must work like a human brain. This mistake leads people to underestimate machines that perform well but don't "think" like we do. The author argues that we should stop comparing AI to the human mind and instead judge it based on performance and outcomes. If a machine can solve complex problems, write useful reports, or help diagnose diseases – even if it does so in a completely different way from us – then it is still powerful and important. Another key idea is what the author calls "not-us thinking". This is when professionals assume that AI might replace or disrupt other people's jobs, but not their own. For example, some doctors and lawyers believe their work is too human, too personal, or too complex to be touched by AI. But the book explains that most clients don't care about the process – they care about the result.

Received on 30-11-2025; revised 10-12-2025.
DOI: 10.33232/BIMS.0096.93.95.

If an AI system can deliver the same or better results, people will use it, regardless of whether it works like a human or not.

In Chapter 5, the author looks at some of the abilities we often think only humans can have, such as judgement, empathy, and creativity. Many believe that machines will never be able to do these things. But instead of simply agreeing or disagreeing, the author introduces new terms: quasi-judgment, quasi-empathy, and quasi-creativity. These help us understand that AI doesn't need to feel emotions or have imagination in the human sense to perform similar roles. For example, a chatbot can give comforting replies to someone in distress. It may not "feel" the emotion, but the effect on the person can still be real. Similarly, AI can generate music, art, or stories that are creative in output, even if the process is different from how a human creates. The book argues that we need new words and new ways of thinking to describe what AI does. Just like the Industrial Revolution brought new vocabulary – such as "factories" and "middle class" – the AI era needs its own language. Without it, we are stuck using misleading terms such as "hallucination" to describe machine errors, which only adds confusion.

Chapters 6 and 7 shift the focus to how AI will change work. The author identifies three major effects: automation (doing current tasks faster or cheaper), innovation (creating new ways to do things), and elimination (removing certain tasks entirely). These changes will affect not just blue-collar jobs but also many white-collar professions. A powerful example in the book is about neurosurgeons. The author explains that patients don't really want a neurosurgeon – they want to be healthy. If AI can help deliver health more effectively, people will choose that, even if it means human surgeons become less necessary. This is a key point: people care about outcomes, not job titles. This part of the book is especially important for professionals who assume their work is safe from automation. It challenges all of us to rethink our value not in terms of what we do, but in terms of the results we deliver.

One of the most important topics in the book is the possibility of AGI—Artificial General Intelligence. This means AI that can do any intellectual task a human can do. The author doesn't say AGI is definitely coming soon, but suggests it could arrive between 2030 and 2035. This is not a prediction but a warning: we should prepare for this possibility, just in case. The author introduces a helpful way to think about this called "What-if-AGI?" thinking. This is a thought experiment where we imagine what the world would look like if AGI existed. What rules would we need? How would we share the wealth it creates? What jobs would be left for humans? How would we protect people's rights and dignity? Even if AGI never fully arrives, thinking this way helps us prepare for many changes that are already starting. It encourages long-term thinking in a world that usually only looks at the next big app or gadget.

Chapter 8 is one of the most practical and useful parts of the book. It lays out seven types of risks that AI brings – some short-term, some long-term, and all serious. These include political risks (such as misinformation), economic risks (such as job losses), psychological risks (such as identity confusion), and even existential risks (if AI becomes uncontrollable). The author doesn't try to scare the reader but helps us think clearly. This chapter gives readers a framework for understanding risk that is easy to apply and very needed in today's noisy debates. Chapter 9, titled "Harnessing AI," is a strong message that we are late to this conversation. Many of the issues we're facing now were already visible decades ago, but we failed to act. Now, AI systems are becoming more powerful, and we need smart rules, ethical planning, and collaboration across all parts of society.

Later chapters explore some of the biggest and hardest questions: Can machines ever be conscious? What does AI mean for the future of life itself? The author doesn't

pretend to have all the answers, but offers deep, balanced thinking. He quotes philosophers including A.C. Grayling and Bryan Magee to help us understand that even human consciousness is still a mystery. So it's okay to be unsure about machine consciousness. One especially powerful idea is in Chapter 12, which discusses the "Great Schism" – a possible future where biological and artificial intelligences split into separate paths. This is not just science fiction. The book treats it as a real possibility and asks: What if our role as humans is to begin the next era of intelligence, even if it goes beyond us?

The conclusion of the book is one of its strongest parts. The author tells us not to get stuck thinking only about today's tools, such as GPT-4. We must use our imagination to think about GPT-7 and beyond. The systems of the future will likely be far more powerful than we expect. If we don't prepare now, we will be caught off guard. The book calls for a new kind of thinking – one that is brave, ethical, and forward-looking. AI is not just another invention. It is a change in how we think, work, live, and understand ourselves. We must be ready.

This book is clear, original, and deeply important. It avoids hype but doesn't downplay the risks. It is not too technical, but it is serious and smart. Most of all, it gives us the tools we need to think clearly about the future of AI. Whether you are a policymaker, teacher, doctor, engineer, student, or simply curious, this book will help you understand where we are headed and how to prepare. It offers a new way of thinking that we all need right now.

Zhenwei Lyu is an assistant professor of Physics at Dalian University of Technology. He earned his PhD in gravitational physics from the University of Guelph. Since joining Dalian University of Technology, his research interests have expanded to include numerical relativity.

Leicester International Institute, Dalian University of Technology
E-mail address: zwlyu@dlut.edu.cn