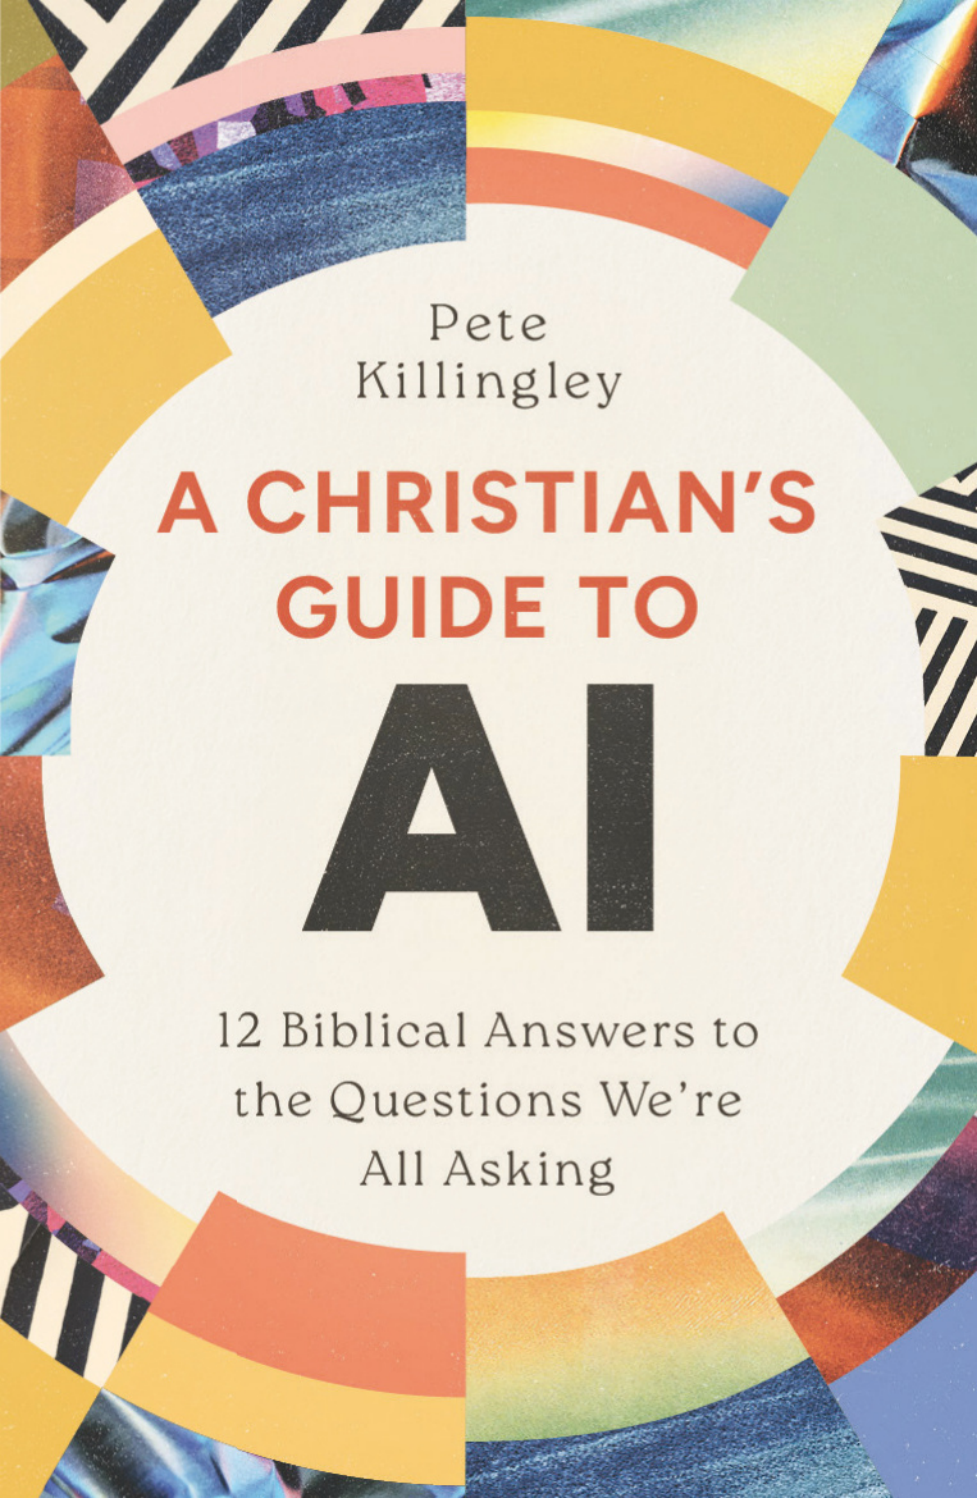
Pete
Killingley
A CHRISTIAN'S
GUIDE TO
AI
12 Biblical Answers to
the Questions We're
All Asking

"I confess to having been a bit 'head-in-the-sand' about AI—avoiding it seems easier than learning how to engage in an informed and godly way. I'm so thankful to Pete for producing this accessible and immensely helpful book. He addresses the questions we all need to grapple with, applying the timeless wisdom of God's word to each with grace, humility and an appropriate amount of humour! An invaluable resource."

CAROLYN LACEY, Author, *Amazed*

"As AI use accelerates, some of our oldest questions get sharper. What does it really mean to be human? To work, to create, to relate, and to bear responsibility? Pete Killingley addresses these questions and more with Bible in hand, helping us engage with new technologies without losing our hold on God's truth. This is a wise biblical guide for navigating complex technological terrain while following and trusting the God who remains Lord over it all."

MATT LILLICRAP, CEO, UCCF

"Pete has written a very helpful primer for Christians on engaging with AI. It is well researched, makes helpful distinctions and offers wise advice. He asks really excellent discussion questions which press the reader to decide what they will do in various arenas. This will make you a better user of AI."

MATT FULLER, Author, *Reclaiming Masculinity*

"As the church begins to grapple with the implications of AI, we need resources that are theologically and pastorally wise. Pete Killingley offers exactly that. This book provides a clear and thoughtful introduction to vital questions that followers of Jesus should be asking about AI. This is a timely and valuable resource that will help equip the church to respond with wisdom, clarity and faithfulness in a rapidly changing world."

DAVID BETTS, Author, *The Church and AI*

"There is every indication that AI may prove to be one of the most powerful and important technologies of our time. We cannot sit idly by but must instead learn what AI is, what benefits it may bring to our lives, and what drawbacks it may exact from them. I'm thankful that Pete Killingley models that kind of deliberate examination and careful engagement in *A Christian's Guide to AI*."

TIM CHALLIES, Author, *The Next Story*

"Pete Killingley has done a great service to the church. With wit and wisdom, he guides us through the strange, new world of AI, avoiding both hand-wringing doom and gushing about the utopia on offer. Instead, from solid research and biblical application, he shows us both the possibilities and temptations that AI presents. The book is really about how to preserve human dignity and integrity when faced with a new technology. If you care about being truly human in today's AI-infused world, read this book."

TED TURNAU, Author, *Popologetics*

"With clarity and accessibility, Pete gently draws the reader through the questions he raises about AI and our relationship with it. I love that this book is rooted in Scripture and shaped by the lived experience of an author who is a father, pastor and someone helping the church engage in the public square. He models how to ask the right critical questions about AI, holding them alongside the hope we have in Christ. This is a thoughtful and grounded guide."

ROSS HENDRY, CEO, CARE

"If only we had a helpful, biblical, thoughtful 'primer' to get us thinking straight and establish some wise boundaries around AI. We do! And this is it. Pete's book is clear, measured and rooted in Scripture. It steers a useful course between extremes and shows how Christians can think soberly about the world God has placed us in. I benefitted greatly from reading this from cover to cover (without an AI summary), and you will too."

ADRIAN REYNOLDS, Associate National Director, FIEC

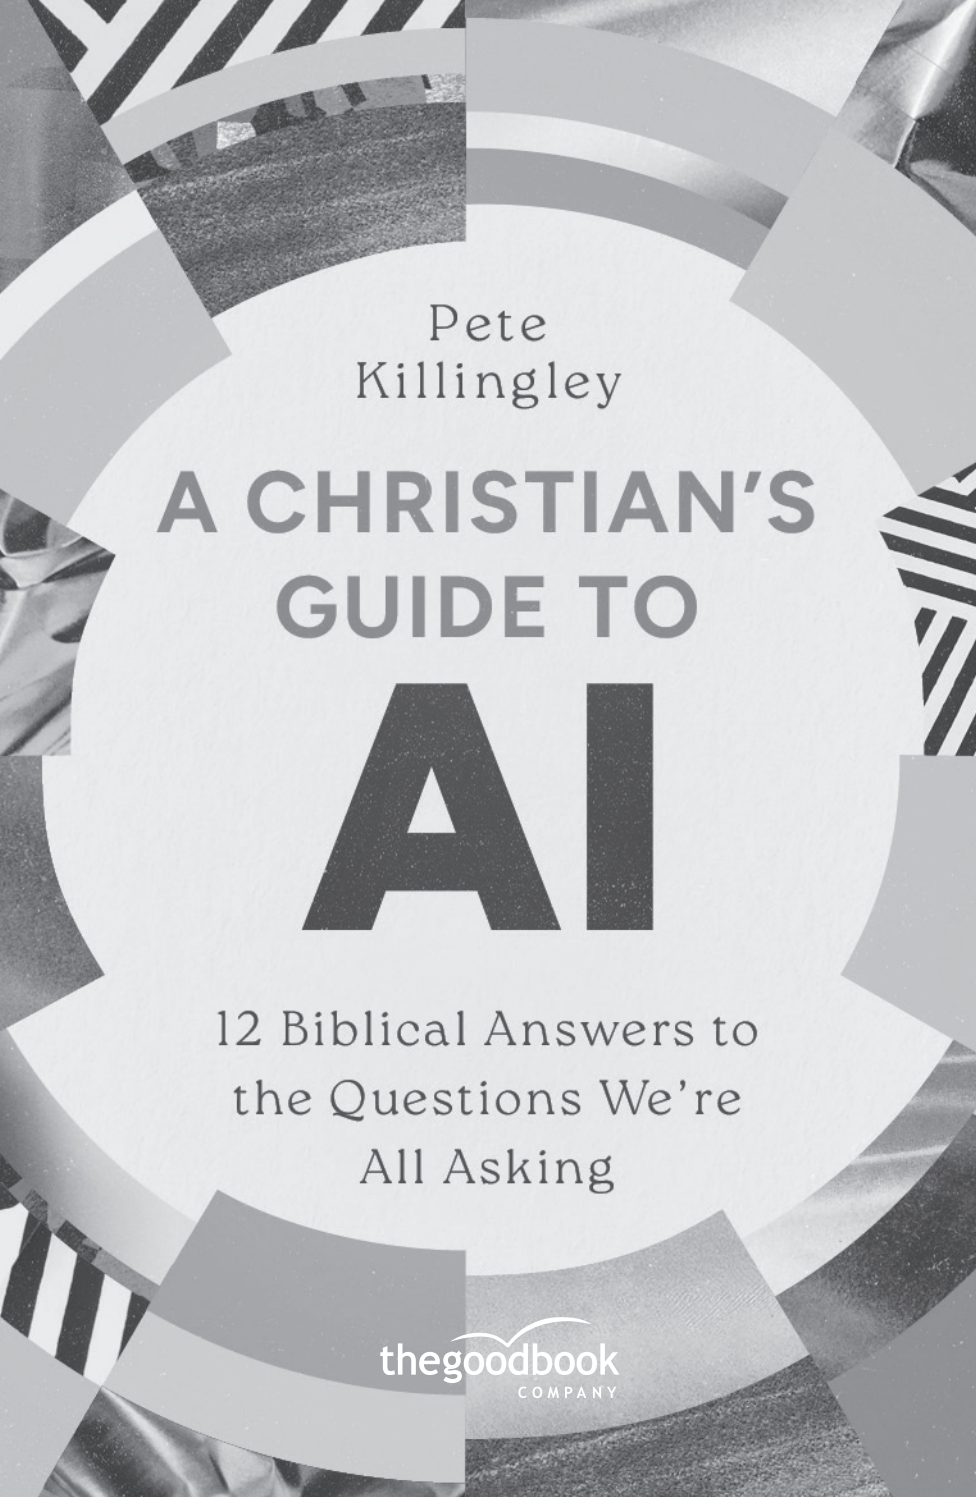
Pete
Killingley
A CHRISTIAN'S GUIDE TO
AI
12 Biblical Answers to the Questions We're All Asking
thegoodbook
COMPANY

A Christian's Guide to AI

Published by:
The Good Book Company

thegoodbook.com | thegoodbook.co.uk
thegoodbook.com.au | thegoodbook.co.nz

Cover design by Drew McCall

ISBN: 9781802544268 | JOB-008838 | Printed in India

To Sophie.
Thank you for your encouragement—I would never have written this book without it.
Thank you for your joy—life would be so much more boring without you. You are truly superfluous.

Will the robotic leader come at last to achieve our objective,
feed the hungry, forgive the debtors, heal the sick,
give sight to the blind, release the captives, raise the dead?
…
Or do we look for another?

—Wendell Berry, Sabbaths 2014

Contents

Introduction

One thing I learned when I started this project is that if you tell people that you're writing a book on AI, around 50% of the responses you'll get are jokes about getting AI to write it for you.

ChatGPT did not write this book. I have, however, made use of AI at many points in the process. I'll tell you how I've used it, both because I think transparency is important and to show you some of the positives about AI right from the beginning. Because many of the chapters that follow are cautious, encouraging us as Christians to think carefully about how we use AI and why, it might come across as if I'm against it. But AI can be a fantastic tool if used rightly and wisely.

Here are the ways it came in handy (or not) during the book-writing process. I've used AI…

- to give me book recommendations for helpful background reading.
- as a search engine—to help me find reliable sources for statistics I've come across, rediscover news articles I recall seeing previously, or find quotes which I half

remember. AI is much easier than Google if you don't remember the exact words.

- to list Bible verses on a certain theme (for example, all Bible verses on friendship from the book of Proverbs). I can't prove that it was exhaustive, but for each verse it gave, I then looked it up for myself, so I know it didn't give me any false positives.
- to translate a report on AI from German into English.
- to summarise research papers so that I can know whether I'm likely to gain relevant information by reading the whole paper.
- to transcribe voice notes. Often, when I'm walking, my thoughts crystallise, and I record a voice note. Rather than listen back to each note back and type them all up, I get AI to transcribe them for me, saving me time.
- as a case study. I've probed and prodded AI to see what kinds of answers it gives me. Although it makes this mistake less frequently than it used to, I was delighted when I asked it how many times the letter "r" appears in strawberry and it got it wrong (see chapter 3).
- to help produce endnotes.
- as dictation software. I ditched this pretty quickly because I'm a fast typist and it actually slowed me down.
- for an initial edit or review. I've uploaded sections and asked it to point out where my argument is weak or needs refining. I've asked it to show me where I've typed extraneous words or incomplete sentences.

- for the chapter on whether AI can ever be like us. I'd written an illustration which referred to the droids in Star Wars (C3PO and R2D2). The theological point I was making was accurate, but I'm not so familiar with the films that I could know if the illustration was truly helpful. I asked AI to take on the persona of a pedantic Star Wars fan and then critique what I'd written. On the basis of its feedback (and input from my wife and daughter), I deleted that paragraph.

I use AI confidently yet cautiously. In the areas where I believe it's ethical to use it, I give clear instructions and then test everything. I never blindly trust information given to me by AI and always double-check primary sources. Where mistakes remain, I take full responsibility. It's my fault, not AI's!

Why This Book?

A common view of the church is that we're always a few years behind the world when it comes to any cultural or technological change. Whether that's true or not, we should be asking where we *want* to be as churches: at the forefront of the change or picking up the pieces of the people AI leaves in its wake?

As Christians, we do need to start thinking about AI deliberately. The fact that you've picked up this book means you probably agree. It would be so easy for us to split into our camps of tech-lovers and technophobes (you might have guessed I'm in the former category) and then use or not use AI as we each see fit. Because the Bible doesn't give us a handbook for how to relate to AI, we would each create our own guidelines, often without a huge amount of thought.

But as I've studied AI over the last few years and given seminars on the topic to churches and church leaders,

I've been struck that although the Bible doesn't explicitly mention AI, it does give us an abundance of theological themes that help us understand the AI revolution. Whether that be our theology of humanity, of technology, of creativity, of work, of friendship or of the future, the Bible is not silent.

And so we're going to explore twelve questions that Christians ask about AI. My hope and prayer is that it helps us to allow God's word to shape the way we think and act as the world changes around us.

CHAPTER 1: UNDERSTANDING

What Is AI?

So the men went up and spied out Ai.

Joshua 7:2

In Genesis 4, Tubal-Cain was the first to make tools out of bronze and iron (v 22). In doing so, he changed the world.

In 1440s Germany, Johannes Gutenberg invented the movable-type printing press, thereby beginning the printing revolution. In doing so, he changed the world.

In 18th-century Britain, the industrial revolution began: water and steam power, machine tools, mechanisation. It can't be reduced to any one person or invention, but all these things combined to change the world.

In 1950, Alan Turing published a seminal paper entitled "Computing Machinery and Intelligence". In it, he asked the question, "Can machines think?" Or can they pass what we now know as the Turing test? In doing so, he primed the world for another change.

And now, once again, with the invention of Large Language Models in 2017 and the release of ChatGPT in 2022, we see the world changing around us. This is the next big paradigm shift.

What Can AI Do?

The type of AI that has entered our public discourse over the last couple of years consists of the "chatbots": Claude, Grok, Gemini, ChatGPT. But there's far more to AI than that.

AI is used to scan your face when unlocking your phone. The system is comparing the original scan of your face with the picture that has just been taken and determining whether it "thinks" they are the same person.

AI is used when you are browsing through Netflix. The system has a record of every show you've ever watched, for how long and at what time of day, and presents you with content that it "thinks" you will want to watch next.

AI is used by the fraud team at your bank. It has a record of every transaction you've made, and so, for each new purchase, it "thinks" and decides whether or not it is suspicious, based on your past behaviour.

AI is used in social media feeds. It knows what you engage with and presents you with posts it "thinks" you are likely to linger on, comment on or interact with in some way.

AI is used in weather forecasting. It models patterns from historical and current data and tells you whether it "thinks" it will rain this afternoon.

AI is used in predictive text on your phone. It has a record of every word you've typed and is constantly giving you an option of the next three words it "thinks" you might use.

Artificial or Apparent Intelligence?

The simplest definition of AI that I've come across is from David Betts' book, *The Church and AI*. He writes that Artificial Intelligence is…

> … a human-made simulation of intelligence in a machine.[1]

The word "simulation" is essential. None of this is actual intelligence; it simply *mimics* human intelligence.

You could question whether "artificial intelligence" is even the best name for this technology, but we're not going to change that now it's established. In my mind, and throughout this book, I often replace the word "artificial" with "apparent". Calling AI *apparent intelligence* helps me view it in a more accurate way and keeps me from slipping into some of the pitfalls which we'll explore in this book. I encourage you to keep the phrase *apparent intelligence* in your mind when I refer to AI in the pages ahead.

To show you what I mean, try an experiment. Ask a group of friends to give you a random number between zero and ten. One might say seven, another five and another four. Next, ask AI to give you a random number between zero and ten. When I asked ChatGPT for a random number, it gave the answer seven. I asked again; it gave me the same answer. I have now asked it over 20 times, deleting the previous chat each time, and always get the same answer: seven. Grok gave me seven. DeepSeek gave me seven. Gemini gave me seven. Claude gave me seven.

Why? This gets to the heart of AI. It's not coming up with a random number. It is predicting the most likely answer to your question. Having been trained on vast amounts of

data—billions of bytes of text—it "knows" that the most common answer humans give when asked for a random number is seven. And so it will give you seven every time. In other words, rather than actual intelligence, it's trying to do what sounds intelligent; it is apparent intelligence.

Categories of AI

Experts in AI tend to categorise it in different ways. Here are a few:

- Narrow AI is AI that does the specific thing it's been programmed for. Deep Blue, the AI that defeated Garry Kasparov at chess in 1997, was very good at chess, but that was it. AI that interprets medical scans can detect tumours or fractures with impressive accuracy, but it can't play chess or write a poem. ChatGPT is very good at conversation and answering questions, but it can't drive a car or cook a meal.
- General AI is able to do new things, outside of its original context. General AI could learn to play chess in the morning, teach itself to diagnose medical scans in the afternoon, and drive a car in the evening, all without being specifically programmed to do so by any human. This kind of AI doesn't exist at the moment.
- Super AI is General AI that has vastly surpassed human capabilities across a wide range of areas. Sometimes researchers talk about the "singularity" as an event horizon—a future point when AI becomes so powerful that it begins to improve exponentially without the help of humanity. Reaching the singularity (which may never happen) is both a research goal and a nightmare scenario!

- Discriminative AI sorts and classifies data. It analyses pictures or audio recordings or words and spots patterns. So, for example, you could train the AI with a million pictures of cats and a million pictures of dogs. Then, if you give it another picture, it is able to determine whether it is a cat or a dog. It's discriminating between them.

- Generative AI (this is what the G in ChatGPT stands for) actually creates new data, which could be images or words or music or computer code or anything else. You give it a prompt, and it comes back with a response, generating something new. Having trained it on all those millions of pictures, you could ask it to produce a picture of a cat. And unlike Google, which just searches the internet and finds an example of what you're looking for, generative AI will generate a picture of a cat that has never existed before.

- Among generative AI models, one particularly important type is the Large Language Model (LLM), which specialises in text. LLMs are what most people think of when they talk about AI because it's the main way they are consciously using AI. Because it produces human-like language, interacting with an LLM can feel like speaking to an intelligent being, even though it is simply predicting words. In this book, most of the uses of AI we'll discuss relate to LLMs.

Is This Just the Next Technology?

Often, when people raise the alarm about some of the dangers of AI, the objection is that people have doom-mongered about every technology; yet in every case, we have got used to it, and

we have adapted. Is AI just the next technology we'll adapt to? Well, maybe. But there are some differences. One is the speed of development. The pace of change is so quick that staying on top of the latest AI advances is a full-time job. Even by the time this book goes to print, AI will have advanced and have new capabilities, and there will be new questions we need to ask.

The complexity of the technology is also a concern. Many of the world's experts on AI have candidly said they do not understand everything that it does. Dario Amodei, CEO and co-founder of Anthropic, said:

> People outside the field are often surprised and alarmed to learn that we do not understand how our own AI creations work. They are right to be concerned: this lack of understanding is essentially unprecedented in the history of technology.[2]

A third difference from previous inventions is that a technology built on intelligence—even if it is artificial or apparent intelligence—can self-improve. It's been said that nuclear weapons didn't have the capability to create new, better nuclear weapons, but AI does have the capability to create new, better AI. Think of the proverb: "Give a man a fish, and you feed him for a day; teach a man to fish, and you feed him for a lifetime." Now we need to add a line:

> Teach an AI to fish, and it will teach itself biology, chemistry, oceanography and evolutionary theory … and then fish all the fish to extinction.[3]

Beware of Idolatry

When it comes to AI, there are three main narratives. First, *a tale of peril*—that AI will cause great problems. We'll

think about that in chapter 12. Some believe it's all *a tale of propaganda* and that AI will never live up to the claims being made. While a third group tell *a tale of promise*—that AI will solve many of humanity's problems.

Demis Hassabis, co-founder of Google DeepMind, promises a world of "radical abundance". If AI is developed in ways that are equitable, he says that…

> It should lead to incredible productivity and therefore prosperity for society. We should be in an amazing world of abundance for maybe the first time in human history, where things don't have to be zero sum. And if that works, we should be travelling to the stars, really.[4]

The tale of promise is that AI will solve every human problem. Economic growth and productivity turbocharged. Healthcare outcomes improved. Transportation by self-driving cars, cutting road deaths. Improved environmental protection, with climate change reversed. The sky's the limit—or in Hassabis's mind, there is no limit.

I do believe AI will deliver some of the benefits it promises, and this will make a meaningful impact on people's lives. For example, in the medical sphere, AI is already being used to improve early detection of diseases, speed up diagnosis and exponentially improve the pace of drug development. But let's be cautious about placing our trust in technology.

Of course, the Bible doesn't refer to AI by name. (The only mention of Ai is the ancient city that God's people fought in the book of Joshua. They lost the first battle and won the second. It's probably not wise to draw parallels.) But the Bible does have something to say about looking to created things

rather than the Creator God to protect and deliver us. It's called idolatry. It happens when we set our hopes on people, in the way the Old Testament kings would often look to other nations for protection in battle (Isaiah 31:1). It also happens when we look to technology to protect us—not as we depend on God (Psalm 127:1) but *instead* of doing so. All the way back in Genesis 11, the people built a tower to make a name for themselves and to protect themselves from being scattered across the face of the world. Isaiah 44 describes the process of the blacksmith or carpenter creating a miniature idol, before pointing out the ridiculous conclusion:

> Half of the wood he burns in the fire; over it he prepares his meal, he roasts his meat and eats his fill. He also warms himself and says, "Ah! I am warm; I see the fire." From the rest he makes a god, his idol; he bows down to it and worships. He prays to it and says, "Save me! You are my god!" (v 16-17)

AI isn't idolatrous in and of itself, but the way we look to AI to solve our problems without God is not that different from the blacksmith of Isaiah 44. Google DeepMind's mission statement is to "Solve intelligence, and then use intelligence to solve everything else".[5] *Save us, AI, for you are our god?*

As we read statements highlighting the benefits of AI, we can celebrate the possibilities but should also see through the facade, recognising the way it often promises us what only God is able to deliver.

Where Do We Go from Here?

AI is here to stay. And just as with the printing press, the industrial revolution or the microchip, AI will change the world.

There's so much more that could be said about the history of AI, the way we can categorise it or understand it, and what it can do. But the purpose of this chapter has been to understand enough about AI for us to begin to think about it in a Christian way.

The world of AI is a world of data and information. More and more, we're going to have vast amounts of knowledge at our fingertips, and we'll be able to access and reproduce information in the format we want. AI will do that for us. But the Christian life is about more than knowledge, information and facts. It's about wisdom and character. God calls us to grow more like Christ. And that involves searching the Scriptures, reading the word and praying, asking the Spirit to produce his fruit within us.

We are so in need of wisdom. The good news is that God invites those who lack wisdom to call on him, and he will give it, generously.

Discussion Questions

- Before reading this chapter, what did you think AI was capable of? Has that view changed?
- How does describing AI as "apparent intelligence" change the way you interact with it?
- Are you naturally inclined to believe that AI is a tale of promise, of peril or of propaganda? Why do you think that narrative resonates with you?
- The chapter talks about seeking wisdom, not just knowledge. In what ways do today's technologies make it harder for us to prioritise wisdom over knowledge. What can we do about that?