


AI Curbside

*The Practical Guide
for Busy Clinicians*



Kirk D. Wyatt *MD, MAS*

AI Curbside

The Practical Guide for Busy Clinicians

Kirk D. Wyatt MD, MAS

Frontmatter

Copyright © 2026 Kirk D. Wyatt. All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopying, recording, or otherwise—without the prior written permission of the publisher, except for brief quotations in critical reviews and certain other noncommercial uses permitted by copyright law.

Library of Congress Control Number: 2026910640

ISBN 979-8-9958402-4-4 (ebook)

ISBN 979-8-9958402-0-6 (hardcover)

ISBN 979-8-9958402-1-3 (paperback)

ISBN 979-8-9958402-3-7 (audiobook)

Published by The Curbside, LLC, West Fargo, North Dakota

thecurbside.ai · aicurbsidebook.com

For permissions, errata, and inquiries, contact

kirk@aicurbsidebook.com

First edition

Disclaimers

The views expressed are those of the author and do not necessarily reflect those of his employer or research collaborators.

Nothing in this book should be construed as medical advice. This book is intended for educational and informational purposes only and describes the author's personal experiences and perspectives on the use of AI tools in healthcare workflows. Clinical decisions should always be based on a licensed clinician's own professional judgment, institutional guidelines, and the specific circumstances of each patient. The author and publisher disclaim any liability arising from the use or misuse of the contents of this book.

Nothing in this book should be construed as legal advice. AI is an emerging field abounding with uncertainty. Be sure to comply with all applicable laws, policies, and procedures as you utilize AI tools. As I will highlight in this book, AI can be a powerful tool, but you—the user—are ultimately responsible for how it is used and applied in the world we live in.

Many of the AI products described in this book are proprietary, and their inner workings are not always fully transparent—even to advanced users. Tools that appear similar on the surface may behave quite differently under the hood, and individual products evolve frequently. Descriptions in this book reflect my personal understanding of how each tool worked at the time of writing, based on publicly available information and my own use. Nothing in this book is a guarantee or warranty of how any specific tool currently functions. Always consult the vendor's current documentation, their terms of service and privacy policy, and your own institution's policies before using these tools.

The AI landscape is rapidly evolving. Parts of this book are likely to be out-of-date by the time you read it. While I've done my best to ensure everything in this book is accurate as of the time of this writing, I'm sure some lingering errors remain. If you identify any

of these errors, please accept my apology and then email me at kirk@aicurbsidebook.com so that I can address them in a future edition.

AI Assistance

Portions of this book were drafted, edited, and refined with the assistance of a variety of large language models. Cover artwork was developed with AI assistance. The author reviewed all content, made all editorial decisions, and takes full responsibility and accountability for the final text.

Financial Disclosures

Potential financial conflicts of interest related to the content in this book include the fact that I receive royalties for copies sold on all platforms (including ElevenReader and Amazon, which I mention in the book). I have received research funding from the Chan Zuckerberg Initiative through a sub-award. EvolutionaryScale, which is discussed in Chapter 18, was acquired by the Chan Zuckerberg Initiative. I have participated in usability testing and provided feedback to the UpToDate team on their Expert AI feature. No financial or in-kind compensation was exchanged for participation.

Trademarks

Product names, logos, brands, and other trademarks referenced within this book are the property of their respective trademark holders. Use of these names does not imply endorsement.

Table of Contents

Prologue: I Used AI to Write This (and I'm Still Accountable)	8
Part I: Getting Started	10
Chapter 1: AI Is Here	11
Chapter 2: Garbage In, Garbage Out—Prompting as Clinical Communication	17
Chapter 3: Data Privacy and Business Sensitivity	25
Part II: How AI Works	30
Chapter 4: The Clinician's Mental Model of LLMs	31
Chapter 5: Control Knobs and the Reality of Hallucinations	43
Chapter 6: RAG (Retrieval-Augmented Generation)	49
Part III: Your AI Cockpit	54
Chapter 7: My AI Tech Stack (and What Each Tool Does)	55
Chapter 8: The OpenRouter \$5 Lab	65
Chapter 9: Twisting the Dials—Temperature, Nucleus Sampling, and Image Generation	87
Chapter 10: Everyday Accelerators	107
Chapter 11: Advanced Workflows—Where the Real Leverage Hides	119
Part IV: Safety, Security, and Accountability	134
Chapter 12: Practical Security—Shadow IT and Accidental Oversharing	135
Chapter 13: How AI Can Hurt Patients (and You)	141
Chapter 14: The Anatomy of a Perfect Lie (Hallucinated Citations)	145
Chapter 15: The Rules of Disclosure	149
Part V: The Playbook	156
Chapter 16: Healthcare Use Cases Organized by Risk	157

Table of Contents

Chapter 17: Rolling Out AI in Health Systems 161

Part VI: The Future and Staying Current 165

 Chapter 18: What's Next—Proteins, Drugs, and Synthetic Documentation 166

 Chapter 19: Will AI Erode Critical Thinking? Only If You Let It. 170

 Chapter 20: A Closing Challenge 175

 Epilogue: The Final Word 178

Continue the Conversation 182

Glossary 183

Prologue: I Used AI to Write This (and I'm Still Accountable)

Let's get the elephant in the room out of the way.

I used artificial intelligence to help me write the book you are reading right now. I'm telling you this up front, not as a gimmick, but in full transparency. I didn't use AI because I'm lazy. I used it because I am a full-time pediatric oncologist and clinical informatician with a wife and two young children. Like you, my time is my most fiercely guarded resource.

AI didn't do the work *for* me. I didn't just tell it "write a book on AI for physicians" and use that for the book. I worked *collaboratively* with AI through a lengthy back-and-forth process over seven weeks to develop the ideas, write the text, go through several iterations of revisions, copyedit the text, and develop the logo on the cover. AI amplified my work. It helped me put out better work, faster.

Here's a concrete example: as I listened to a draft of the audiobook narrated by an AI voice as I did chores around the house, I would quickly jot down rough notes on things I wanted to change—a phrase that sounded awkward, a point that needed clarification, a factual claim I wanted to double-check. I then fed those rough (and messy) notes to an AI model along with the full manuscript and asked it to find each quote, describe the problem, and suggest a revision. I reviewed each suggestion, accepting some, rejecting others, and revising many. The AI did the tedious work of locating passages and proposing fixes so that I could review and edit the book on the fly, and I made the editorial decisions. That workflow

—human judgment steering AI capability to improve the quality and efficiency of work—is precisely what this book teaches.

Another motivator for using AI is that if I had written this book the traditional way—from start to finish, completely from scratch during spare time on evenings and weekends—it would have taken me months to get a draft to a publisher. In the AI space, 12 to 18 months is a lifetime, so by the time you got around to reading it, much of what is in this book would be outdated. It would be a historical artifact.

But here is the most important thing you need to know: *I* am ultimately accountable for every single word, idea, and recommendation in this book. I used AI as a tool to help me write better and as an accelerator to help me write faster, but the ideas are genuinely mine. If there are mistakes in here (as I'm sure there are), they are my mistakes, not the algorithm's. In the end, you can be the judge of the final product—is this book just a bunch of “AI slop,” or is it a polished book that helps you meaningfully improve your life? Did AI degrade the quality of the book, or improve it?

In this book, I'm going to break down how these tools work in terms you can understand. I will show you how I use AI in my day-to-day life as a physician and researcher, and explain how you can practically apply AI as a busy healthcare professional. My promise to you is simple. By the time we finish our brief time together, you won't just understand the hype around AI. You will get hands-on experience with a practical, actionable framework for how to use it safely and effectively in your daily life—starting today.

Part I: Getting Started

Chapter 1: AI Is Here

Your Use Is Optional. Your Exposure Isn't.

Medicine is, by design, a conservative profession. We are taught to be skeptical of the "new and shiny" because, historically, the "new and shiny" hasn't always been proven to be safe and comes with many unknowns. We demand evidence, we demand trials, and those things take time.

This is, in part, why healthcare has lagged in adopting and leveraging exciting new technological advancements. Major technology companies have become impressively adept at using data spread across public databases and left behind as crumbs of our internet browsing history. They're so good at this that they can show us targeted ads that we'll reliably click on (and make them money in the process). If tech companies are so good at this, why can't healthcare be so savvy? Why can't we put these algorithms to good use to improve human health, instead of reducing the role of this technology to tempting us to buy things online that we don't need?

With the emergence of **generative AI** (that is, artificial intelligence that generates novel content) over the past few years, things are changing rapidly. AI is now one of the most powerful general-purpose technologies that have ever existed, and it can be wielded by anyone with plain-language prompts. It's now everywhere: AI is here, whether you like it or not.

My hope is that if you are reading this book, you're curious about AI and looking to seriously use it in your everyday life. If you're

one of those people, congratulations! You are ahead of many of your peers.

If you're still skeptical and hoping AI will just go away as a passing fad, I hope that I can convince you in these pages that you risk being left behind. I'm not saying I think AI is going to replace you in your job; instead, you will be outpaced by your peers who are using it and—crucially—using it *better* than you. I want you to protect your time, your relevance, and your agency. I want *you* to be the one who stays relevant because you've effectively wielded these tools. That's what this book is about. I'm going to teach you everything you need to know to integrate AI into your workflows to help you put out better work, faster, and hopefully reclaim some personal time in the process.

Reps Beat Theory

If you take one thing away from this book, follow this advice: the only way to understand what AI is truly capable of is to use it. Use it early, often, and on a variety of tasks.

You cannot learn to be a doctor just by reading a medical textbook. That's why medical training doesn't end in the classroom—it takes years of clinical rotations, residency, and for many, fellowship to develop the judgment and expertise required to care for patients.

You cannot learn to use a **large language model (LLM)** by just reading articles about it. You have to get behind the wheel. When you use it daily for low-stakes tasks, you start to develop an intuitive sense for its strengths and its limitations. This will help you feel more comfortable using it for more advanced and complex tasks. If you make these tools part of your everyday workflows,

through that process you'll understand the good, the bad, and the ugly.

As you build confidence, you'll discover that the real power isn't in asking quick questions—it's in having AI do complex and time-consuming tasks for you. AI can generate custom briefings, help you analyze datasets, build slide presentations, and even write code to create tools you didn't know you could build. It's those things that will be the real game-changers in how you use AI.

Don't Drive Last Year's Model

Before we go any further, we need to discuss a small investment I'm going to ask you to make. There are free versions of AI tools, and there are paid versions. If you're one of the growing majority of Americans who have used an LLM, you're probably saying, "I already have enough subscriptions I don't use. Why would I pay for this if I can get it for free, and it already does what I need?"

If you are using the free version of these services, you need to understand that you are driving last year's model. Actually, in AI-years, you might be driving a car from a decade ago. Or, if the vendor lets free subscribers access the current model, it provides only a small amount of usage—just enough for a test drive.

The companies that built these models reserve their smartest and most capable models for their paying subscribers. The models provided in the free versions are relatively inexpensive to operate, but they hallucinate more often, and they don't reason as well. If you're using these models, you are not experiencing what today's AI can do.

You might be hesitant to spend your hard-earned money on an AI subscription. After all, you're probably still trying to pay off your student loans. But as a busy clinician, if an advanced model saves you just thirty minutes of administrative grunt work a week, the money you invested in that subscription will be repaid many times over. You don't realize what you're missing out on until you use the most advanced **frontier models**. Consider the paid subscription a professional investment.

In a short while, we'll be rolling our sleeves up and testing out some different models. This AI speed-dating exercise is going to cost \$5 (I'm keeping the stakes low for the skeptics). After that, you'll have tried several models developed by different companies and can decide which one (or ones) to subscribe to. More on that in a bit.

The Mindset Shift

As we walk through the coming chapters, I want you to keep one rule in mind: focus on your workflow, not the tool.

The tools will change. New companies will emerge. The model that mesmerized you last year will quickly become "last year's news" and a new model, developed by a different company, will become the new crown jewel. But the workflow—the process of identifying a bottleneck in your day, passing that task along with relevant context and constraints to an AI model in the form of a carefully structured prompt, and verifying the output—will remain the same. It doesn't matter which model you are using.

There is one more mindset shift I want to plant early because it becomes increasingly important the more you use these tools: *stay mindful* as you use these tools.

When you send a prompt to an AI model, the response streams back to you in real time, word by word at a pace that feels breathlessly fast. I've found that this speed creates a subtle psychological pressure—a feeling that I need to keep up with the machine, fire off my next question before I've even finished reading the last answer, and keep the momentum going. And when you have a paid subscription with a usage allowance that resets every week, you find yourself racing to use as much of that allowance as possible by the end of the week so that none of it is “wasted.”

Resist these urges. You are not in a conversation with a colleague who will get impatient if you pause to think, and the model will still be there tomorrow if you don't use it today. The model is obedient. It will wait for you. It has no feelings to hurt and no schedule to keep. Take the time to carefully read what it produced, evaluate whether it's accurate and useful, and *then* decide what to ask next, and don't feel like you have to use AI for *everything*. In our technology-driven lives, we are constantly bombarded with emails, text alerts, and notification popups that fragment our attention and pull us toward mindless reactivity. One of the major criticisms of social media is that while technology initially promised to connect us, some argue that it has degraded our focus and shifted our focus toward performing shallow, distracted, and mindless interactions.

You can fall into this same trap with AI tools if you let them. You ask the model to perform a task, and while it works in the background, you try to multitask—checking your inbox, glancing at

your phone, and starting another chat. But multitasking is largely an illusion; what we're really doing is rapidly bouncing between tasks, doing each one worse than if we'd given it our full attention.

As I've used AI more over the past year, I've made a deliberate effort to treat my interactions with these tools as opportunities for *focused* work, not *fragmented* work. When I'm working with an AI model, I try to be fully present in that task—reading its output carefully, thinking critically about what it said, and crafting my next prompt with intention. This is where the real value lies. The clinicians who get the most out of AI won't be the ones who use it the fastest and outsource everything to AI; they'll be the ones who use it the most thoughtfully and collaboratively.

Key Takeaways

- AI is now a general-purpose technology that every clinician will encounter, whether they choose to use it or not. Early adopters will have an advantage over their peers.
- The only way to understand what AI can do is to use it—early, often, and on a variety of tasks.
- Invest in a paid subscription to access the most capable frontier models; free tiers use older, less capable models that do not represent the state of the art.
- Focus on the *workflow*, not the tool. The tools will change, but the process of identifying a bottleneck, crafting a prompt, and verifying the output will remain the same.
- Stay mindful when using AI. Resist the urge to multitask. Read outputs carefully, think critically, and craft your next prompt with intention.