DR GPT™

FREE EDITION

Human Judgment in an AI World

Five questions a physician keeps asking about AI and care



DR GPT™ · Harvey Castro, MD, MBA

www.HarveyCastroMD.com · /HarveyCastroMD · #DrGPT

ABOUT THIS BOOK

Human Judgment in an AI World

Five questions a physician keeps asking about AI and care

DR GPT™ · Harvey Castro, MD, MBA

I wrote this book myself, as an independent companion to five talks I gave between 2024 and 2025 at TEDxWarrenton, TEDxPurdueU, TEDxBabsonCollege, TEDxAWC and TEDxSouthlake. It is not an official TED or TEDx publication.

Nothing here is a transcript. I've written the ideas again in new words, and each chapter links to the official recording so you can hear what I said on stage. The book is free. It is educational and does not provide individualized medical advice. If you have a health concern, talk with a clinician who knows you.

Free to read and free to forward. Please share the whole file, unchanged. Not for sale.
www.HarveyCastroMD.com · /HarveyCastroMD · #DrGPT

© 2026 Harvey Castro, MD, MBA. You may share this PDF unchanged and free of charge. All other rights reserved. Web edition, November 2026.

BEFORE YOU READ

Not medical advice. This book is for education. It doesn't diagnose, treat or replace care from a clinician who knows you. If you think you're having a medical emergency, call your local emergency number now.

Technology examples. The tools, devices and studies I mention are examples from my talks. Mentioning one isn't an endorsement, and it doesn't mean the tool is cleared by a regulator, available to you or right for you. Research findings, such as voice biomarkers, aren't diagnostic tests.

Stories and people. I changed patient names on stage. Some stories are composites, scenarios or illustrations, and the text says so. One story can't show that a tool works for everyone.

Numbers. Each figure carries the year and source I could trace. Evidence changes, so check the latest source before you rely on a number. Information is current as of November 2026.

Independence. This isn't an official TED or TEDx publication. TED, the event organizers and the companies named here did not review or endorse it. TED and TEDx are trademarks of TED Conferences, LLC.

Interests. I advise, and have advised, companies that work in healthcare technology. Nothing in this book is a product recommendation.

How this book was made. I wrote it from my own talk transcripts, scripts and slides. AI tools helped me draft, edit and design it. I reviewed the content, and the ideas and stories are mine.

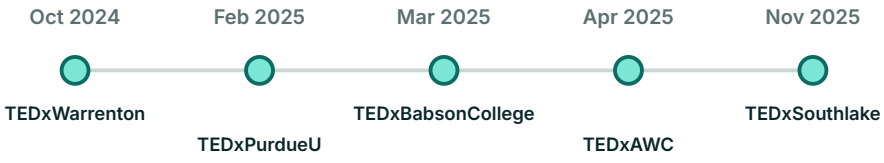
Links. Links lead to outside sites I don't control. The recordings are on the official TEDx Talks YouTube channel.

The exercises use fictional, nonmedical examples on purpose, so you can try them without anyone's health information.

CONTENTS

	Introduction: what this book helps you do
CHAPTER 1	Give the human part more time
CHAPTER 2	Question what you can't see
CHAPTER 3	Prevent harm before it happens
CHAPTER 4	Create time before the emergency
CHAPTER 5	Bring care closer without leaving people out
	Closing: choose one question
	Worksheets
	Selected pearls
	Watch the talks
	Source notes
	About the author

THE TALKS BEHIND THIS BOOK, IN ORDER



Order only. Spacing is not to scale.

THE BOOK AT A GLANCE

Five questions to ask before AI touches care

Tap a card to jump to its chapter.

01 TIME

Give the human part more time

Let AI carry the paperwork, then decide on purpose that the saved minutes go to listening.

02 PERCEPTION

Question what you can't see

Before trusting a signal you can't confirm, ask how it was checked and who it was tested on.

03 PREVENTION

Prevent harm before it happens

Use SAFE: Strategic, Accountable, Fair, Ethical and Early. Name who owns the next step.

04 WARNING

Create time before the emergency

A device can prompt. People and tests decide. Plan the first week after an alert.

05 REACH

Bring care closer without leaving people out

Build data starting with the least represented person, so reach doesn't spread bias.

Introduction: what this book helps you do

I'm an emergency physician. My work happens in short windows, with incomplete information, and with a person in front of me who is frightened. That setting shapes how I think about artificial intelligence. I don't ask whether a tool is impressive. I ask whether it helps me take better care of that person, and whether it leaves me more present or less.

Over about a year I gave five talks that circled the same question from different sides: how can AI improve care without losing the human part? This book gathers what I still believe from those talks and sets it down as five questions you can carry into a clinic, a boardroom, a classroom or a conversation at your kitchen table.

If you're a patient or a family member, this book should help you ask sharper questions when someone tells you a tool is smart. If you're a clinician, it should help you decide where a tool belongs in your day and where it doesn't. If you lead a team or plan an event, it should give you one question worth an hour of honest discussion.

A few ground rules I've held myself to. Stories in these pages are ones I told on stage. Where I described a patient, I've kept the detail thin, and where a story was an illustration or a scenario, I say so. Technologies I mention are examples of what I described at the time. Describing a tool is different from proving it works for everyone, and I've tried to keep that line visible. Numbers appear only where I can point to a dated source, and the source notes say what each source supports.

Each chapter ends with a short exercise. I wrote those exercises for this book. They aren't things I said on stage, and they use ordinary, nonmedical examples on purpose, so you can try them without touching anyone's health information.

Read the chapters in any order. Then pick one question and use it this week.

01

Give the human part more time

Time a tool saves helps patients only if someone decides to give it back to them.

Where this chapter comes from: a talk I gave at TEDxWarrenton in 2024, titled How AI and humans can revolutionize medicine together.

Picture a small child in the middle of an asthma attack. If you've been the parent in that room, you know how long a minute can be. The child gets treated, the breathing settles, and then I walk in with discharge paperwork. Pages of it. The paperwork exists for a serious reason, because I don't want that family back in crisis with no idea what to do. But a stack of printed instructions rarely does that job.

So I tried something else. I boiled the instructions down to the few things that mattered: what can set off the asthma, and how the inhaler works. Then I turned those into pictures and sat down with the child to color them, one trigger at a time. I remember thinking this was the best way I knew to deliver care. Then came an uncomfortable thought: most children never get that, and most adults don't either.

That gap is why I started the talk there. Understanding is part of treatment, and a plan the patient can't follow hasn't been delivered yet.

The time problem

Medicine is short on people and short on minutes. The World Health Organization projects a shortfall of about 11 million health workers by 2030, mostly in low and lower middle income countries, and an older Medscape survey, from 2016, found most physicians reporting 13 to 16 minutes with each patient. Two dated readouts, and both describe a system where the conversation gets squeezed first.

In Warrenton I cited a 2011 projection that the doubling time of medical knowledge would fall to about 73 days by 2020. I want to be careful with it here. It was a forecast about volume, published years ago, and I don't have a current measurement to offer you. The practical point survives without the number: no clinician can read everything, and the pile grows faster than anyone's evenings.

When someone asks me what AI is for, I start with paperwork. Let it carry the repetitive part: the summaries and the forms that say the same thing twice. Then spend the recovered minutes on the part only a person can do, which is listening until I understand what you're telling me.

That second half is a choice, and it's the half that worries me. Time saved by a tool doesn't automatically flow back to patients. Someone has to decide it will.

TIME SAVED HAS TO BE SENT SOMEWHERE



Editorial diagram of the argument in this chapter. Not data.

Speaking the patient's language

After the coloring book, I kept going. Older children got comic strips. I showed the audience a short demonstration of AI-generated speech in several languages, including Spanish for the Hispanic families I see, and I described using the same approach to produce sign language for patients who are hard of hearing. Some people in the room found the synthetic voice a little creepy. I understood that, and I still think the idea is worth the discomfort.

These were demonstrations, so I'm describing what I showed, and I'd want any tool like this checked carefully for accuracy before it carried real medical instructions to a family. A fluent mistake is still a mistake.

Why this is personal

When I was 12, I cut my hand on a rusty signpost. My mother took me to the emergency room, where she and I waited eight hours. When the doctor finally stitched me, he used nothing for the pain. The stitches hurt. What stayed with me longer was the feeling of being a number in somebody's queue. I went into medicine, and into emergency medicine, so that fewer people would leave feeling invisible.

Years later my mother called to tell me she'd been diagnosed with cancer. Time stopped. Everything after that one word was noise. She's alive, well and cancer-free, and she was sitting in the audience that night. The lesson I took is practical. In the moment bad news lands, most people stop hearing. A record of the visit, something they can return to at two in the morning when they're ready, lets the conversation happen twice: once in the room, and once when they can take it in.

A recording or transcript of a medical visit needs consent and careful handling, and both belong in the design from the start instead of being added after something goes wrong.

A teammate at three in the morning

I told one more story that night, about a mother who spent three years taking her young son to 17 doctors without an answer. She gathered what those clinicians had produced, the records and the test results, and entered it into an AI tool, which suggested a diagnosis. I said on stage that it did what the doctors couldn't. I also said, and I'd underline it now, that the tool was working from everything humans had already done. That was people and software together.

One family's experience can't tell you how often a tool gets it right, but it does show a role for the tool, which is connecting scattered work that people have already done.

That's the role I want at 3 a.m., when I'm tired and the stakes haven't dropped to match my energy. I think of it as a vigilant teammate, a second set of eyes on what might slip through. A teammate, in my book, still answers to me, and I still answer to the patient.

On stage I described healing healthcare in four parts: humanize it, equip doctors and patients with tools for better-informed decisions, use AI for better diagnosis and patient education, and extend those tools to underserved places. I also suggested something medicine rarely admits. Doctors study disease for years and almost never take a class in how to speak with people. AI could help there too, as a practice partner that reviews how a clinician communicates.

I closed with a line I still use: a human working with AI will do better than the best human alone or the best AI alone. I hold that as a goal to design toward. It only comes true when the human stays engaged.

TRY THIS

An editorial exercise written for this book. It isn't something I said on stage.

Pick a set of instructions you give other people that has nothing to do with health. How to close the shop at night. How to submit an expense report. How to feed the neighbor's cat.

1. Choose one real person who has to follow it.
2. Rewrite it for that person only. Use their words, their reading level, their language.
3. Ask them to explain it back to you without looking.
4. Notice what they left out. That's the part your original version never delivered.

If you use an AI tool for step 2, check its draft line by line before you hand it over.

WATCH THIS MOMENT

The human-touch segment begins at 6:23

Open the official recording at this moment

Worth discussing

- Where could a tool give your team back ten minutes a day, and who decides how those minutes get spent?
- Think of the last time you were handed instructions you didn't follow. What would have made them usable?

02

Question what you can't see

A new way of seeing is only as good as the checking behind it.

Where this chapter comes from: a talk I gave at TEDxPurdueU in 2025, titled Shifting Perspectives: How AI is Revolutionizing Healthcare.

There's a German word, Umwelt, for the idea that each creature lives inside the world its senses can reach. A bloodhound's world is built from scent. I told the audience that bloodhounds have been known to follow a trail for more than 130 miles. A bat maps a dark room with echoes. You and I walk through the same room and miss most of what those animals perceive.

Clinicians are no different. I trust what I can see, hear and touch. That works until the important thing is something none of my senses can pick up.

The Titanic made the point for me. Its crew steered by stars and lights. About 1,500 of the roughly 2,200 people aboard died in April 1912, on a ship without radar. My argument was that a new instrument changes what a careful person is able to notice.

That talk opened with a dinner story about someone whose eyesight had faded while her sense of taste seemed sharper than mine. The recording tells that story best.

A tour through the senses

The talk walked through sight, hearing, touch, taste and smell, with examples of AI-assisted tools for each. I'll summarize them here, and I'll say this once clearly: these are examples I described on stage. Some are products, some are research, and some are ideas. A description from me isn't evidence that a tool is ready for clinical use, cleared by a regulator, or right for you.

Sight. I described augmented-reality glasses that could give a surgeon a much larger view of a small operating field, and that could put a lab result in front of an emergency physician the moment it posts. I talked about AI as a second reader on mammograms, looking for the tiny calcium deposits radiologists search for, and about research on reading brain scans for early signs of Alzheimer's disease.

The example closest to me is my grandfather, who is legally blind. Losing his sight took much of his independence and some of his spirit. With AI glasses, he can walk into a room and hear a description of what's there: who is on his right, where the table is, where the door is. I also mentioned a cane with a camera that reads the surroundings.

Hearing. Emergency departments are loud. I described headphones that could filter the noise so I can hear the patient, and real-time translation in place of the phone interpreter I used to run and fetch when I worked near an airport. I described AI-assisted stethoscopes that flag irregular heart rhythms, and software that tries to tell one kind of cough from another.

I also said that software can listen to a voice and flag Parkinson's disease, post-traumatic stress disorder or prediabetes. Those claims come from published voice-biomarker research. They're research findings, and none is a cleared diagnostic test. A flag from a voice sample is a reason to examine someone.

Then there was the mother I called Maria, whose baby sounded fine to a human ear while an AI stethoscope picked up early signs of pneumonia. Treat that as an illustration of the idea.

Touch. I described a surgeon operating a robotic arm from far away while a haptic glove returns the feel of tissue, and a rehabilitation glove that gives feedback to a person recovering sensation after a stroke.

I also told the story of a man I called David, who lost his speech to a stroke and, three years later, spoke through a brain-computer interface in a voice rebuilt from old recordings. I told the audience his first words were to his granddaughter, and that they were words of love. I'm repeating it as I told it.

Taste and smell. I talked about a spoon designed to make food taste saltier than it is, which interested me because I ask so many patients with high blood pressure to cut back on salt. I described meal plans tailored to a person's conditions, with groceries ordered before the visit ends. And I described a digital nose that analyzes breath for signs of disease, then imagined the same sensing in a hospital or an airport to catch an outbreak early. Those last ideas were me thinking out loud.

The questions every new sense should face

Underneath the tour sits one argument. A new sense pays off only when I'm disciplined about what it tells me. When a tool reports something I can't confirm with my own eyes or ears, I'm trusting a chain of people I've never met, from the ones who built and tested it to the ones whose data it learned from.

So I ask the same things each time.

BEFORE I TRUST A NEW SENSE

1. **What exactly is it detecting, and how was that checked?**
2. **Who was it tested on? Do they resemble my patient?**
3. **What happens when it is wrong, in each direction?**
4. **What do I do next?**
5. **Whose information does it collect, and where does it go?**

Then decide how much weight the signal deserves.

The questions I ask of any tool that reports something I cannot confirm myself.

I ended that talk with fire. Fire keeps me warm and lights the path. It also burns. Then I asked who controls AI and who benefits: a privileged few, or everyone? I don't think that question has a settled answer. I do think the people asking it early are the ones who'll shape it.

TRY THIS

An editorial exercise written for this book. It isn't something I said on stage.

Run a sense audit on one decision at work that has nothing to do with health. Say you manage a small bakery and you're deciding how much bread to bake for Saturday.

1. List what you observe directly. The line at the door. What sold out last week.
2. List what an instrument tells you. The register report. A weather app.
3. List what nobody measures. Who walked away because the shelf was empty.
4. Now imagine a new tool that claims to predict Saturday's demand. Which list does it draw from? How would you find out when it's wrong?

WATCH THIS MOMENT

The Umwelt segment begins at 2:44

Open the official recording at this moment

Worth discussing

- What might a tool notice that your current senses and instruments miss, and how would you confirm it?
- Who gains access to a new sense first? Who waits?

03

Prevent harm before it happens

Avoiding harm is the floor. Preventing it is the goal, and someone has to own each step.

Where this chapter comes from: a talk I gave at TEDxBabsonCollege in 2025, titled SAFE-AI: Transforming Healthcare with AI's Ripple Effect.

After about 20 years in emergency medicine, there's a sentence I hear in my head. Patients say some version of it after I've told them about the cancer or the heart attack. They say they wish they had known sooner. I've had to deliver that kind of news for two decades, and the wish lands the same way every time. I'd like to be the doctor who says the opposite: you came in time, nothing has been damaged yet, and here's the plan.

That wish made me think of my grandfather. He's my father figure. He sat me down at five and taught me chess, and along with it how to knot a tie, how to drive, and to carry my wallet in my front pocket, which I still do. Later he developed macular degeneration, which blacks out much of the visual field. I watched him tilt his head and lean toward the board to find the pieces, and sometimes make the wrong move because he couldn't see the right one. It was hard to watch, and then one day he caught me with a move I never saw coming and laughed that he could still beat me half blind.

I carry a gene variant linked to a higher risk of his condition, and I said on stage that others in my family carry it too. That turns prevention from a slogan into a family calendar. It's also why I gave him AI glasses that can describe a scene or tell him where the chess pieces sit.

From do no harm to prevent harm

Every medical student learns the old vow: do no harm. It made sense in an era when treatments themselves often injured people, and I argued on stage that the time has come to add to it. Avoiding harm with my own hands is the floor. The better goal is to keep harm from arriving at all.

Think of a fire department. Nobody sees a firefighter until something is burning, and too often nobody sees a doctor until something hurts.

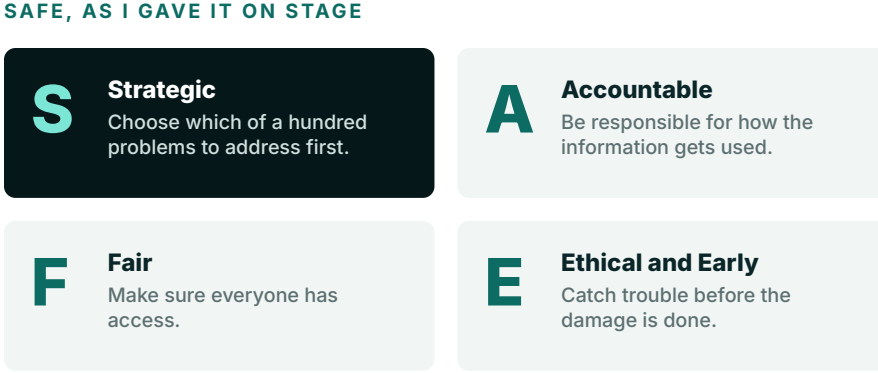
Two numbers from that talk need careful handling. On stage I cited about 41 million deaths a year from noncommunicable diseases such as heart disease, cancer and diabetes, and the World Health Organization's current fact sheet, updated in 2025, puts the toll at no fewer than 43 million deaths in 2021. I also called those deaths preventable. The source describes a category of disease, and I shouldn't imply every one of those deaths could have been avoided. Many could have been delayed or averted, and that's the case for acting earlier. The American Diabetes Association estimated the economic cost of diabetes in the United States at $412.9 billion in 2022, which I rounded to 400 billion.

Prediabetes was my example of a quiet stage where a signal could matter. I described tools that might surface it: voice analysis, and wearables that send vital signs to a care team, including during pregnancy. Those were descriptions of what I'd seen and read. A prediction of risk is a reason to look closer, and it's still a long way from a diagnosis.

Two images carried the rest. A rock dropped in a still pond, with rings spreading outward, which is how I think one good use of AI travels through a health system. And the asteroid that ended the dinosaurs, as a picture of scale. In the talk I turned the image around: this force could arrive to help, and I'd add today that whether it does depends on the people steering it.

SAFE

Plenty of people in that audience were uneasy about AI. They're the ones I had in mind when I offered a short checklist called SAFE, and I've kept the wording from the talk.



Wording kept from the talk.

A checklist earns its place when it changes a decision. I'd use SAFE before a purchase, when the room is excited and nobody has yet asked who owns the mistakes.

Three pillars and three conditions for trust

I gave the audience three working pillars. First, detect: find risk before symptoms, the way a retinal scan could have warned my grandfather. Second, build the detection into the workflow, so it shows up where clinicians already work. Third, personalize. I pointed at a pair of twins in the audience and said that even they shouldn't get identical care by default.

Then I talked about trust, through a patient I called Sarah, who likes AI and worries about it in the same breath, and I think she's right on both counts. If medicine gets this wrong, people will refuse tools that could have helped them.

Three conditions, as I gave them:

- **Privacy.** Data has to be secure. I raised the question of what happens to sensitive data, genetic data above all, when the company holding it fails.
- **Fairness.** The technology has to reach everyone.
- **Regulation.** This is the thorny one, and I admitted I see both sides. I want progress on cures and new drugs. I also want rules strong enough that nobody is harmed and trust survives.

I ended by asking everyone to put a hand over their heart and feel it beat, and to imagine using these tools to give people more beats and more years with the ones they love. My grandfather was always three moves ahead of me. I said the technology is too, and that waiting too long is how a game ends in checkmate.

> **TRY THIS**
>
> *An editorial exercise written for this book. It isn't something I said on stage.*
>
> Run a prevention review on a fictional, nonmedical case. A neighborhood library plans to switch on an automated system that suspends borrowing privileges after three late returns.
>
> 1. Strategic: what problem is the library solving? Is this the first one it should solve?
> 2. Accountable: name the person who reverses a wrong suspension. If you can't name one, stop there.
> 3. Fair: who is most likely to be suspended? Families without a car? People working nights?
> 4. Ethical and Early: what could the library do at the first late return that would make the third one less likely?
>
> Then ask what harm you could catch earlier in your own work, and who would own the next step.

WATCH THIS MOMENT

The SAFE segment begins at 7:41

Open the official recording at this moment

Worth discussing

- Which of the four letters does your organization skip when it's in a hurry?
- If a tool flagged a risk tomorrow, who exactly would act on it?

04

Create time before the emergency

An early warning buys time only when a person acts on it and explains it.

Where this chapter comes from: a talk I gave at TEDxAWC in 2025, titled Thriving Before the Emergency: The Future of Medicine with AI.

I opened with a question. Suppose I could tell you that you'd have a heart attack a year from now. Would you do anything differently?

Two people are behind that question. My father died of a heart attack when I was 20. He had no warning signs. One day he was here and the next he wasn't. My mother's story ran the other way. She raised me with very little money, and unless I was seriously ill, a doctor's visit was out of reach. She trained as an ultrasound technician while working three jobs, and one day she scanned her own neck and found a small nodule in her thyroid. It was cancer, and her doctors then found a second cancer in her throat. She found it early, and that made the difference.

One parent had no time. The other had some. In medicine, time is the most valuable thing there is.

A ship and a lighthouse

I described the body as a ship taking on water while healthcare patches holes after the damage is done. What a ship needs is a lighthouse, a warning before the rocks.

The figure I gave came from a 2014 study in BMJ Quality and Safety, which estimated about 12 million outpatient diagnostic errors a year among adults in the United States. On stage I said misdiagnosed. The study's term is diagnostic error, and it describes outpatient care in that period.

Three stories about time

I told the audience I change patients' names, and I've kept details thin here for the same reason.

A man I called Mark came in with nausea and fatigue. By the time I made the diagnosis, his diabetes was advanced and his organs had paid for it. I remember wishing I could rewind to the stage before, when his blood sugar was high and nothing was yet broken.

The second story was mine, and I'd never told it in public. I had just finished training. On my first shift as an emergency physician I started having migraines. I had imaging, a spinal tap and no answer, and I went home with a pill to take at the next headache. I took it. Half an hour later I told my wife to call an ambulance. My blood sugar had crashed. I asked the paramedic to secure my airway because I thought I was about to die. In the emergency department, colleagues who knew me asked my date of birth, and all I could get out was the same single word, over and over. My left side wouldn't move. I spent seven days in intensive care, able to hear everything and say almost nothing. It turned out to be an adverse drug reaction that caused a spasm in the blood vessels of my brain and gave me the symptoms of a stroke.

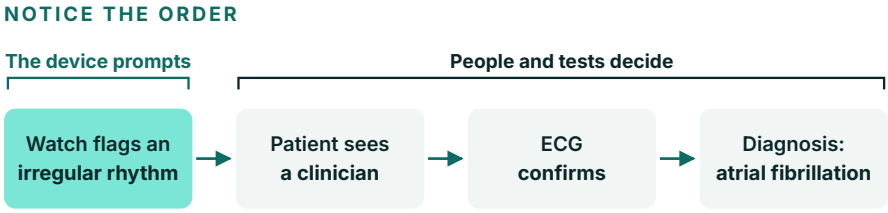
I told that story because I wished, lying there, for something that could have raised a flag earlier, though I can't say any device would have caught it. I can say I know what it's like to be the patient nobody has an answer for.

The third was about a woman I called Maria (I reuse a few names across talks, and each story here is separate), a teacher whose screening mammogram was read as normal. Her hospital used AI as a second reader for mammograms, and when it marked something on hers, she was called back, and a biopsy found an early cancer. She's doing well, I said. One good outcome proves nothing about every patient, though it shows what a second look can be worth.

Everyday devices

On stage I held up my smartwatch and an old glass thermometer. The watch had told me that morning that I'd slept well and that my stress was high, which I blamed on the audience. Wearables now estimate things like temperature trends and heart rhythm, and I described three ways that could buy time.

A runner I called James felt palpitations and blamed coffee. His watch suggested an irregular rhythm, his doctor ran an electrocardiogram, and it confirmed atrial fibrillation. Notice the order. The watch prompted the visit. The clinician and the test made the diagnosis.



From the runner story in this chapter, as I told it on stage.

Jenna, as I called her, is a single mother who kept landing in the emergency department with high blood sugar because she looked after everyone except herself. A continuous glucose monitor paired with AI coaching gave her feedback on food, water and walking, and she came in less often.

And I imagined a hike in the Grand Canyon with no signal, where glasses or a phone running AI on the device could coach a bystander through CPR. That was a what-if, and I said so with the glasses in my hand.

I offered a line I liked enough to put on a slide: what happens when AI listens to your symptoms before your symptoms speak? It's a good question because it has two honest answers. Sometimes you gain time. Sometimes you get an alarm about nothing, and worry you didn't need. An alert is a reason to check, and the checking is still a person's job.

Small steps, and who gets left behind

My friend Juan arrived at business school without English, studied with a dictionary in one hand and a textbook in the other, and finished first in his class. He now holds a senior role at a large company. I told it because prevention scares people when it looks like a mountain. It works as small steps repeated.

Then I turned to the risk. I said plainly that AI makes me nervous too. Done badly, it widens the gap between people who have access and people who don't. I cited a study on how little AI training data comes from low-income populations. The principle doesn't depend on it: a tool trained on some people holds up poorly for everyone else.

I also said that patients are more willing to trust AI when their doctor explains it. That puts the work on clinicians. If I recommend a tool, I owe you an explanation of what it does and what it can't do, in words that make sense to you.

CARE

The argument closed with a second acronym, CARE.

CARE, AS I GAVE IT ON STAGE

C — Compassion
Use the technology with compassion.

A — Access
Everyone should have it, here and around the world.

R — Respect privacy
My data should not be visible to the neighbors.

E — Execute it correctly
On stage I put it as: implement it correctly.

Wording kept from the talk, with the E word as it appears on my website.

My grandfather closed that talk too. He was about to turn 93 and still playing chess with me after some 45 years. Because the gene runs in my family, I have AI send me a weekly summary of new research and clinical trials on macular degeneration. That's my small step. It gives me hope, and it gives me something to do with the hope.

I had three asks. Patients, expect more from your care. Doctors, start learning these tools, even if you'd rather not. Policymakers, get the rules right, because the cost of getting them wrong is high.

TRY THIS

An editorial exercise written for this book. It isn't something I said on stage.

Think about an early warning in ordinary life. Your car shows a low tire pressure light on a cold morning.

1. What do you do in the next ten minutes? The next week?
2. Who do you tell, and who can check it properly?
3. What does a false alarm cost you? What does ignoring a real one cost?
4. After the third false alarm, do you still look?

Now pick one signal in your own work that arrives early and gets ignored. Write down the smallest step that would make someone act on it.

WATCH THIS MOMENT

The access, privacy and trust segment begins at 13:30

Open the official recording at this moment

Worth discussing

- If a device could warn you a year ahead, what would you want to happen in the first week after the warning?
- Who explains the tool to the patient, and how would you know the explanation worked?

05

Bring care closer without leaving people out

Connected care can reach farther, and so can any bias built into it.

Where this chapter comes from: a talk I gave at TEDxSouthlake in 2025, titled How AI, Satellites and Edge Tech Keep your Doctor Close.

I began with a scenario. It's 3 a.m. in rural Texas. A woman I called Maria wakes with chest pain. She's alone, there's no cell coverage, and the nearest hospital is 70 miles off. In the scenario she's wearing a small patch on her chest that sends her heart rhythm by satellite to a cardiac team hundreds of miles away, and her own doctor calls to say help is coming.

Let me be exact about what that was. It was a picture of how connected care could work, built from technologies that exist in pieces. It wasn't a report on a service you can sign up for, and Maria isn't a patient of record.

I built it because of something that bothers me as a physician. Where you live can decide whether you live.

The patients who never reach the door

The emergency department is organized chaos. At Southlake I told a version of the coloring-book story from Chapter 1, this time about a seven-year-old I called Jimmy, who came in fighting for air with asthma. I

drew the dust mites as monsters and his inhaler with its spacer as a blaster to fight them, and that was the first time I saw him smile.

It felt good, and then it bothered me. I can only help the people who make it through the doors. A rancher came to mind, a man who was about five hours away when his stroke began, and since stroke care runs on a clock, he spent that clock on a dusty road.

The question for the talk became whether care could reach people before they reach me.

A nervous system for care

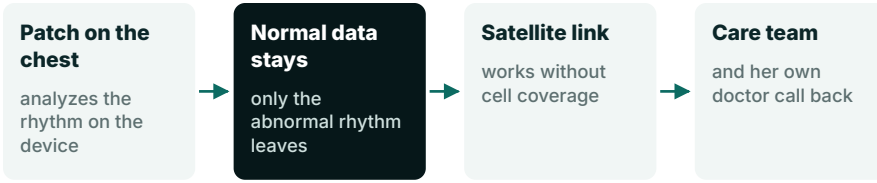
I suggested thinking of healthcare as a nervous system that extends to wherever a person is. I named three parts.

Satellites. Most of the audience had found the venue with a map app that leans on satellites overhead, so I asked them to imagine a link to the sky that also carried a parent's heart monitor. Ground networks fail in hurricanes, earthquakes and wildfires, and some places never had coverage. A link through the sky is a second route when the first one is down.

Edge computing. A body produces an enormous stream of data. Sending all of it somewhere far away and waiting for a reply is slow, and in dead zones it's impossible. Edge computing means the device does the first round of analysis itself and transmits only what looks abnormal. I compared it to a Polaroid camera that develops the picture in your hand. In the scenario, the patch sent the worrying rhythm and kept the rest.

AI. People ask whether it will replace me. I told the audience I'm not worried about that. I picture it tapping my shoulder to ask whether I saw the small finding on a scan, because no human is perfect, and in my work an error can cost a life. I want the tap.

A NERVOUS SYSTEM FOR CARE



The scenario from the talk. A picture of how the parts could connect, not a deployed service.

Scaling care, or scaling bias

This is the part of the talk I'd keep if I had to cut everything else. A good AI lets my care reach more people. A biased AI reaches more people too, and it spreads the bias along with the care.

I pointed to one side of the room and asked the audience to imagine a system trained only on them. It would perform well for that side and dangerously for everyone else. Nobody intended harm. The data was incomplete, and that's all it takes. My answer on stage was to build data that includes everyone, and to start with the least represented person.

Privacy sits in the same design. In the scenario, the patch sends only the abnormal rhythm, and the rest of the information about her stays on the device, which makes the message faster and exposes less of her. Sending less is the point.

I finished the scenario this way: she reaches the emergency department, goes home with a plan, and keeps the patch. The system stayed close, and her doctor stayed closer. I care most about that last clause. The technology exists to shorten the distance between a patient and a person who is responsible for her.

Three symbols closed the talk: the satellite as a nervous system, the edge device processing close to the patient, and AI there to extend my reach and never to shrink my role. Then I asked for something that has little to do with hardware: more unity, more listening, and less distance between people.

What I'd ask before believing a connected-care pitch

These questions are mine, written for this book.

- Does it work where there's no coverage, or only where coverage is already good?
- What does the device send, what does it keep, and who can see it?
- Who was it tested on? Were the people it's meant to serve part of the testing?
- When it raises an alarm at 3 a.m., which human being responds?
- What happens to the patient when the subscription, the battery or the company runs out?

TRY THIS

An editorial exercise written for this book. It isn't something I said on stage.

A fictional city launches an app that predicts where buses are needed. It learned from riders who tap in with a smartphone downtown.

1. Who rides the bus and never shows up in that data?
2. What will the app recommend for their routes?
3. If you could add one group to the data first, who would it be?
4. Who at the transit office hears about it when the app is wrong?

Then ask the same about a tool your own team relies on. Whose data is missing?

WATCH THIS MOMENT

The representation and bias segment begins at 9:35

Scan to open the recording at this moment

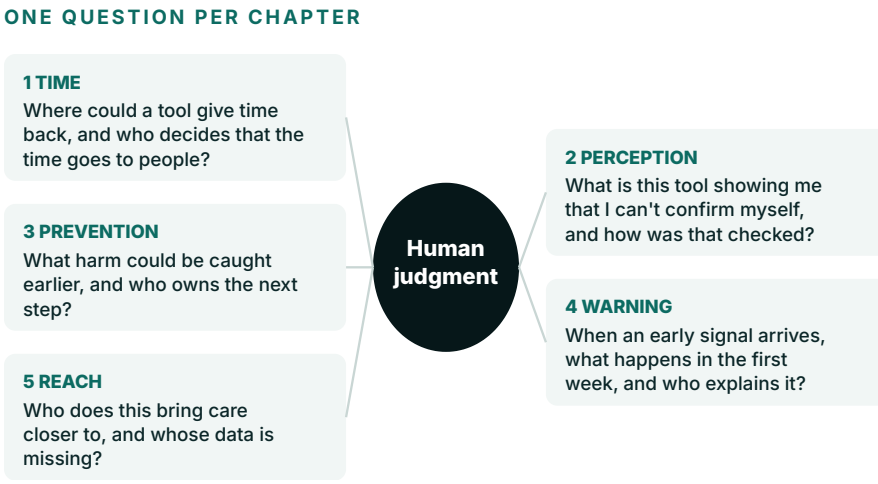
Worth discussing

- How would you judge whether a connected-care proposal improves access for people who are poorly represented in its data?
- In your setting, who is the person the system must stay close to, and who is the person who must stay closer?

Closing: choose one question

Five talks, one worry. Each time I stood up, I was trying to hold two things together. These tools can help people I couldn't help before. None of that help is automatic.

The five questions, in one place:



Pick one and bring it to your next meeting, demo or appointment.

My most recent talk, at TEDxFactoria in 2026, carried the same thread forward. Its recording hasn't been published as I write this, so I'll say only what I've already said publicly about its themes: question a confident AI answer before acting, check what evidence could change the decision, and keep a person accountable for the final call.

If you take one thing from this book, take one question. Bring it to your next staff meeting, your next vendor demonstration, or your next appointment with your own doctor. Ask it out loud. Then listen to what comes back, including the silence.

Human judgment is where responsibility lives. I don't expect technology to engineer it out, and I wouldn't want it to. My hope is that these tools give that judgment more time and better information, and that you and I stay willing to do the judging.

Worksheets

These pages are for printing or filling in on screen. Use them in a team meeting, a vendor demonstration or a conversation with your own doctor. They are thinking tools for education, not clinical protocols, and they don't replace your organization's policies or a clinician's judgment.

WORKSHEET 1

One tool, five questions

Name the tool or decision: _____

- [ ] Time. Where could a tool give time back, and who decides that the time goes to people?

- [ ] Perception. What is this tool showing me that I can't confirm myself, and how was that checked?

- [ ] Prevention. What harm could be caught earlier, and who owns the next step?

- [ ] Warning. When an early signal arrives, what happens in the first week, and who explains it?

- [ ] Reach. Who does this bring care closer to, and whose data is missing?

SAFE and CARE checks

WORKSHEET 2 · SAFE, BEFORE A DECISION TO ADOPT

- ☐ **S** **Strategic.** Which problem is this tool solving first, and why this one out of all the others?
- ☐ **A** **Accountable.** Who is responsible for how the information gets used? Write a name.
- ☐ **F** **Fair.** Can everyone who needs it get it, or only some people?
- ☐ **E** **Ethical and Early.** Is it ethical, and does it act early enough to catch trouble before the damage is done?

Letters and words as I gave them on stage. The questions turn each word into something you can check.

WORKSHEET 3 · CARE, WHEN A TOOL REACHES A PATIENT

- ☐ **C** **Compassion.** Is the technology used with compassion for the person in front of me?
- ☐ **A** **Access.** Can everyone who needs it get it, here and elsewhere?
- ☐ **R** **Respect privacy.** Does the person's data stay private, and who else can see it?
- ☐ **E** **Execute it correctly.** Is it implemented correctly, and has it been explained in words the person understands?

Letters and words as I gave them, with the E word as it appears on my website.

WORKSHEET 4

Sense check for a new tool

Sense	What I can confirm myself	What the tool claims to detect	How was that checked?	Who was it tested on?
Sight				
Hearing				
Touch				
Taste and smell				

Selected pearls

Every line below is a **paraphrase**. None is a verbatim quotation. Each link opens the official recording near the passage so you can hear the original wording.

From the 2024 Warrenton talk

- Paraphrase: Turn discharge instructions into something a child will color, and color it with them. Watch 1:08
- Paraphrase: AI can be a vigilant teammate at 3 a.m., when a tired clinician can't afford a miss. Watch 7:42
- Paraphrase: A record of the visit lets a family return to hard news when they're ready to hear it. Watch 9:29

From the 2025 Purdue talk

- Paraphrase: Every creature lives inside the world its senses can reach. A new instrument enlarges it. Watch 2:41
- Paraphrase: AI is like fire. It warms and lights the way, and it can burn. Watch 14:00
- Paraphrase: Ask who controls the technology and who benefits from it. Watch 14:14

From the 2025 Babson talk

- Paraphrase: Nobody sees a firefighter until there's a fire, and medicine too often works the same way. Watch 8:08
- Paraphrase: The old vow is to do no harm. The next step is to prevent harm. Watch 9:46
- Paraphrase: Trust depends on private data, fair access and sensible rules. Watch 10:54

From the 2025 AWC talk

- Paraphrase: In medicine, time is the most valuable thing there is. Watch 1:51

- Paraphrase: Healthcare patches the ship after it's taking on water. It needs a lighthouse first. Watch 2:23
- Paraphrase: Big changes in health come from small steps repeated. Watch 12:22

From the 2025 Southlake talk

- Paraphrase: Where a person lives shouldn't decide whether they live. Watch 1:12
- Paraphrase: A biased AI spreads its bias to every patient it reaches. Watch 9:02
- Paraphrase: Build the data starting with the least represented person. Watch 9:35

Watch the talks

Every published recording is on the official TEDx Talks YouTube channel. Official titles are unchanged. Scan a code or tap a link.

TEDXWARRENTON 2024

How AI and humans can revolutionize medicine together

youtube.com/watch?v=SW_2NjIvuPU

TEDXPURDUEU 2025

Shifting Perspectives: How AI is Revolutionizing Healthcare

youtube.com/watch?v=_I0Knrj_k5E

TEDXBABSONCOLLEGE 2025

SAFE-AI: Transforming Healthcare with AI's Ripple Effect

youtube.com/watch?v=Ts3Chh09070

TEDXAWC 2025

Thriving Before the Emergency: The Future of Medicine with AI

youtube.com/watch?v=6pPQmW_4Ctw

TEDXSOUTHLAKE 2025

How AI, Satellites and Edge Tech Keep your Doctor Close

youtube.com/watch?v=Y-dnIPhFkLA

Source notes

How this book was written. I worked from transcripts of my own talks and rewrote the ideas in new language. The book contains no verbatim quotations from the recordings. Where wording matters to you, use the watch links.

For Babson, AWC and Southlake I worked from caption-derived transcripts, which were not checked against the audio line by line. Watch links use timestamps taken from the public recordings.

Numbers used in this book, and what each source supports.

What I said on stage	What the source supports
A global shortage of 10 million health workers	World Health Organization, Health workforce: a projected shortfall of 11 million health workers by 2030, mostly in low and lower middle income countries.
About 13 minutes with the doctor	Medscape Physician Compensation Report 2016: most physicians reported 13 to 16 minutes per patient.
Medical knowledge doubles every 72 days	Densen, Transactions of the American Clinical and Climatological Association, 2011: a projection of doubling about every 73 days by 2020. A historical forecast, not a current measurement.
41 million preventable deaths a year	World Health Organization, Noncommunicable diseases fact sheet (updated 2025): at least 43 million deaths in 2021. Older editions gave about 41 million a year. The source doesn't say all were preventable.
$400 billion a year on diabetes in the United States	American Diabetes Association, Economic Costs of Diabetes in the U.S. in 2022: $412.9 billion.
12 million Americans misdiagnosed each year	Singh and colleagues, BMJ Quality and Safety, 2014: about 12 million outpatient diagnostic errors a year among US adults.

What I said on stage	What the source supports
Voice can flag Parkinson's, PTSD and prediabetes	Published voice-biomarker research. Research findings, not cleared diagnostic tests.
About 1,500 people died on the Titanic	Encyclopaedia Britannica and Royal Museums Greenwich: about 1,500 of roughly 2,200 aboard, April 1912.
A bloodhound can follow a scent 130 miles	PBS Nature: bloodhounds have been known to follow a trail for more than 130 miles.

Other figures I mentioned on stage are left out of this book because I couldn't tie them to a dated source.

About the people in these stories. Patient names used on stage were changed, and identifying detail has been kept out. I reused a few names, such as Maria, across talks, and each story is separate. Some stories were illustrations or scenarios, and the text says so where that applies. A single story doesn't establish that a tool works for everyone.

Medical boundary. This book is educational. It doesn't diagnose, treat or advise on any individual's care.

About the author



DR GPT™ · Harvey Castro, MD, MBA

I'm a board-certified emergency physician who still works clinical shifts, a healthcare AI keynote speaker and the author of over 45 books, including ChatGPT and Healthcare (2023) and AI in Emergency Medicine (Wiley). Between 2024 and 2026 I spoke at six TEDx events. I tailor keynotes to the program and audience, from clinicians and health system leaders to patients and students.

WEBSITE

www.HarveyCastroMD.com

Speaking and contact:
harveycastromd.com/speaker
DRGPT@harveycastromd.com

FIND ME EVERYWHERE AS /HARVEYCASTROMD

LinkedIn · /in/harveycastromd
YouTube · @harveycastromd
X · @HarveycastroMD
Instagram · @harveycastromd
Facebook · /HarveyCastroMD
Substack · harveycastromd

#DRGPT

Was this forwarded to you?
This book is free. Read it, pass it along, and find the talks and more at www.HarveyCastroMD.com.

© 2026 Harvey Castro, MD, MBA. This independent companion is not an official TED or TEDx publication. TED and TEDx are trademarks of TED Conferences, LLC. Educational content only, not individualized medical advice.