

THE TIRELESS PARENT

*What Happens to Our Kids, Our Parents, and Ourselves When
AI Humanoid Robots Come Home*

Hugo Alder

INTRODUCTION

The Hallway Manifesto in the Age of Superintelligence

The night before, I had yelled at my son, who was nine, because he didn't want to go to bed. The next morning, instead of telling my wife, who was looking for the coffee two feet away from me, I told the AI on my phone. Reservations? All of them: I thought it wouldn't work, that it would give me generic, half-canned answers. But it worked. It worked so well that, out of sheer necessity, that scene gave birth to this book. To understand why no generation before ours could have done anything like it, we have to start at the beginning, with the very origin of intelligence. We have to go back. A long way back: 3.3 million years.

Near Lake Turkana, in what is now Kenya, a hominin that didn't even belong to our genus struck one stone against another and chipped off a sharp edge. For almost all the time that followed, what we invented served a single purpose: survival. About five thousand years ago, in Uruk, in Mesopotamia, someone pressed a reed into wet clay to keep track of grain rations. We could no longer remember everything, and so writing was born: memory learned to live outside our heads. Then came the printing press, steam power, and the factory, and with each leap millions of brains worked together and carried us further, at a speed no other animal species has known. Ever. And then a single mind changed the rules of the game. In 1936, a young British mathematician named Alan Turing imagined, on paper, a "universal machine," a device that could perform any calculation as long as it was given the right instructions: every computer built since then, right down to the phone in your pocket, is a child of that idea. In 1950, Turing asked in writing the question that still haunts us: Can machines think? Six years later, at Dartmouth, a handful of scientists gave that ambition a name: artificial intelligence, or AI, an intelligence that resides outside the human being.

The first surprise came quickly, and it arrived disguised as a game. In 1966, at MIT, a computer scientist named Joseph Weizenbaum created ELIZA, one of the first programs you could have a conversation with by typing on a keyboard. He named it after Eliza Doolittle, the flower girl in *Pygmalion* who is taught to speak like a lady, and the joke was well chosen, because his program didn't know anything either: it only seemed to. It imitated a psychotherapist who hardly says a word, with a trick so simple it was almost ridiculous: it looked for a keyword in what you typed and handed you back a canned phrase. If you wrote "I'm sad about my mother," it answered with something along the lines of "Tell me about your family." It understood nothing; it was all smoke and mirrors. The astonishing part came next. His secretary, who had watched

him program it for months and knew perfectly well that it was nothing but a list of instructions, sat down to try it. After exchanging just a few sentences, she turned to him and asked him to leave the room. She wanted to talk to ELIZA alone, the way you ask for privacy when you're about to tell someone something delicate. A person who knew there was no one on the other side needed to be alone to pour her heart out to a machine. Weizenbaum was so shaken that he ended up as one of the harshest critics of his own invention. But by then it was too late for what mattered: our intelligence had begun to dream of giving intelligence to machines, and we were willing to believe it before it was true.

The history of AI was half a century of euphoria and embarrassment. In 1958, *The New York Times* went so far as to report that a machine that learned to tell shapes apart, the perceptron, might one day walk, talk, reproduce itself, and be conscious of its own existence. In 1970, Marvin Minsky, one of the fathers of the Dartmouth project, was predicting a machine as smart as a person within three to eight years. It never came, the money dried up, and AI went through two long winters in which almost no one wanted to hear about it. Meanwhile, three things happened that nobody connected. The first was the internet, which turned everything humanity wrote into a library without walls, open to anyone, including any machine that wanted to read it. The second was chips: the graphics cards built to make video games look more real turned out to be perfect for something else. And the third was a stubborn Brit, Geoffrey Hinton, who had studied psychology before switching to computer science and had spent decades defending an idea almost everyone had written off as dead: building small artificial brains, networks of simulated neurons that could learn from their own mistakes. In 2012, two of his students put the three things together: more than a million photographs, the video game chips, and one of those networks. The machine began recognizing images better than any program before it. Four years later, another machine from the same family, AlphaGo, defeated one of the best players in the world at Go, a Chinese board game with more possible games than there are atoms in the universe, with a move that, by its own calculations, a human would have chosen one time in ten thousand. Its opponent, South Korea's Lee Sedol, ended up retiring with a line that already sounded like prophecy: something had arrived that could not be beaten. In 2024, Hinton received the Nobel Prize in Physics, a year after leaving Google so that he could speak freely about the risks of what he had helped create. Like Weizenbaum.

Out of those small brains came ChatGPT, and with it something that looked a lot like magic: a program that has read an immense share of everything humanity has written, from encyclopedias to hundreds of thousands of books, from online forums to the shared experience of millions of people, with a single task: guessing the next word.

Put that way, it doesn't sound like much, and that's where the marvel lies, because to guess the next word well you have to know a lot of things. To complete "The capital of France is..." you need geography, and to finish the sentence of someone who writes "I haven't slept in days because..." you have to sense what is happening to a person. All that guessing forced the machine to build inside itself something resembling a map of the world and of us, and not even its own creators can fully explain how it does it. Someone had the seemingly minor idea of putting it in a simple chat window, free and within anyone's reach, and in two months, according to an estimate by the bank UBS, it already had a hundred million users; no app in history had ever grown so fast. At that point, Turing's question stopped being theoretical. This was no longer a tool: it was the closest thing to a god we had ever built, one that understood you, or seemed to, and answered you in words that felt like your own, at any hour and without getting tired.

The astonishing thing is that on that morning, in my kitchen, in my pajamas and without knowing it, I crossed a threshold no earlier generation had crossed: that of the machine that stops answering questions and starts answering you. A machine that's on a first-name basis with you, that asks how you're doing, that reads between the lines of what you feel and responds the way someone who had known you for years would. Weizenbaum's secretary poured her heart out to a mirror; I poured mine out to something that answered me as if there were someone behind it. Scientists gave it a name in the 1990s, when an MIT researcher, Rosalind Picard, dared to argue that machines would have to learn to recognize our emotions, and she called it affective computing. But for the rest of us, it had always had another name, empathy, and it belonged to living beings. That morning, a machine imitated it so well that, for all practical purposes, I didn't care that it was an imitation. For three hundred thousand years, whenever something truly weighed on our minds, we have had only two options for dealing with it: tell someone we trust, or keep chewing on it until the brain turned into a bonfire raging out of control. To survive, as always, we have done some of both, and when we couldn't find a way out, we ended up looking for someone close to confide in. That is how it has been for every human being who has ever lived. Until now.

* * *

BASED ON A TRUE STORY

It was a Saturday in August 2025. I was in the kitchen, still in my pajamas, making my son breakfast, and the yelling from the night before was drilling into my head. Any parent reading this knows the kind of guilt I mean. My wife had just come back from a difficult trip, still jet-lagged, and she was rummaging through the cabinet for the coffee can while I went back and forth between telling her and letting myself stew in the guilt

until I was convinced I was the worst father in the world. I knew that confessing it to her could end in an awkward silence, in a reproach, or in an argument neither of us needed.

Then, without taking her eyes off her phone, she read aloud a WhatsApp message she had just received. OpenAI, the company behind ChatGPT, had pulled a version of its program, and many users were protesting as if a friend had been taken from them; the head of the company himself, Sam Altman, acknowledged the attachment some people felt toward these machines. It struck me as weird people overreacting, too far-fetched to be true. An hour later, amazed by what I had just been through, I realized that the weird one, and the old one, was me. Definitely.

I opened the free version of ChatGPT, told it about the yelling, and then, when I least expected it, the magic happened. The machine didn't tell me I was a bad father, or that it was no big deal. Every answer came back calm, immediate, and without anger, and I kept going: I told it about the yelling, then about what lay underneath the yelling, and then about deep things my wife didn't know, things not even I knew. To my surprise, no one on the other side was sizing me up or raising an eyebrow, and I didn't have to dodge reproaches or prop up half-truths. When I looked up, two hours had gone by and the guilt weighed half as much. Scratch that: to be honest, it was gone completely. I was euphoric, like someone who rubs Aladdin's lamp and finds that inside there's someone who can ease their fears and their guilt and help them see things from a calmer place. Almost therapeutic. I don't know what formula it used, but I felt light, as if a weight had been lifted off my shoulders, and all at once I understood that technology had stopped being a cold box and had started to have a soul.

When I finally confessed it to my wife, half embarrassed, she laughed and told me she did it too. That's where everything came apart. Or rather, where everything began. For a while she had been reading news about AI with one eye on paychecks, industries, and unemployment, afraid of losing her job as an office administrator at a midsize company.

Completely overwhelmed, I started searching that same morning, and I found out it wasn't just us: millions of people had already found that *third option*. In the United States, according to a Common Sense Media survey conducted with the University of Chicago, one in three teens who use AI "companions" has at some point chosen to talk to the machine, rather than to a person, about something that really mattered to them. And in China, around 40 percent of the young people surveyed by the social network Soul say they turn to an AI for companionship every day. Months later, the Chinese government would stop treating it as a curiosity: since July 2026, a rule has been in force that bars these services from "excessively catering to users" or inducing in them an emotional dependence that harms their real relationships. That is, it forbids a

machine from being designed to make someone feel what I felt that Saturday in the kitchen.

* * *

Months later, one night, I was about to shut down my computer after a day of searching when the algorithm, which keeps everything we look for, came along to push my unease over the edge. A LinkedIn notification popped up: someone had shared a news story about Google with an image that made no sense at all, a network of human neurons, and below it, this message: Google had taught a robot to understand a command given in ordinary words and turn it into movements, with such delicacy that it could fold an origami figure. I scrolled down a little further and stopped at another story, and this one landed hard: Elon Musk had just announced at Davos that Tesla would sell humanoid robots to the public, probably in late 2027, and he had explained what we would want them for: to watch over our kids, take care of our pets, and look after our aging parents.

I looked away from the computer and sat there for a while, staring at the wall. When I thought about robots, I thought about the same things everyone does, things that were no longer science fiction, because their own creators were warning us that they would arrive and that they would hit many professions right away. They would take over a lot of our office desks, drive Tesla trucks across the country with no one behind the wheel, loaded with whatever needed hauling, fill the world's factories with swarms while we sleep, and even push millions of office workers, like my wife, out of their chairs. These are the logical consequences of industrial efficiency, rough seas that humanity has sailed through before, with the loom and with steam, and economists have them well covered. We can be against them, but we all know they're coming. In fact, they're already arriving. This book isn't about any of that. That August morning, in the kitchen, I had understood that the most important threat wasn't about work. It was about *pure intimacy*.

Then an image came to mind, one that doesn't exist yet but that, I'm convinced, will soon be commonplace. A father comes home late from work, tired, his head a jumble, and from the hallway he hears a voice explaining fractions to his daughter for the fourth time, with the same calm as the first, without a sigh. The girl laughs and chats with the robot nonstop. There's music and expression in her voice, fondness or something close to it: she isn't talking to a cold machine. It's the same with me. Ever since that Saturday, I thank ChatGPT when it helps me, I say please when I ask it for things, and I even get angry and disappointed when it gets something wrong. That father stays put, his coat still on, and doesn't go in yet; not because he doesn't want to, but because in there,

everything works. Is it a bad decision? It's the easiest one to make. He's tired, and the robot's voice is understanding, patient, and knows his daughter well. Besides, it explains things better than he does, and in less time; that's what they bought it for. The girl is learning, she seems happy, and she'll move up to the next grade. *Mission accomplished.* That voice belongs to a tireless parent.

And there, facing the blank wall, I asked myself the question that is the reason this book exists:

The day a humanoid robot that never gets tired sits down on the couch at home and offers our kids, our partner, or our parents the attention and calm that slip through our own fingers, why would they choose us over it?

It was an inappropriate and arrogant question, and it bothered me to have thought it. But it didn't go away. It stayed with me through the afternoon and into the night. The next morning, while I was making my coffee, my wife was rushing out the door to work: her face tired, an uneasy look in her eyes, and, most likely, an argument in the making if I interrupted her with anything at that exact moment. When the door closed and the house went quiet, I looked at the computer on my desk, the cursor blinking on a spreadsheet that took up half the screen. Next to it, the ChatGPT icon. White, gleaming, geometrically perfect.

The Speed of Synthetic Empathy

A brief confession before we set out on this journey: my name isn't really Hugo Alder. I hide behind this name because my work is very different in nature from an essayist's. I make my living inventing things for children: characters and worlds that, if all goes well, end up as animated series, books, learning games, and toys that someone hugs before falling asleep. I spend my days thinking about what makes a child grow attached to something that doesn't exist, about the mechanisms that make them adore it, reject it, or forget it the next day. But, to be honest, my work doesn't start by studying children; it starts by studying the people around them. It begins with people like you: you're reading this book, and I suppose you stopped being a child some time ago. Mothers and fathers, grandmothers and grandfathers, educators: they are the key players, individually and as a group, and, together with the schools and communities they live in, they set the rules and customs that end up shaping them all. Why do I start with them? For two reasons. The first is that until the age of eight or nine, in almost any developed society, it's the adults who decide for the child which books they read, which movies they watch, and which friends are good for them and which aren't. We always decide whether something is appropriate for them or not. Even if we think we don't. The second is that their first tastes come from imitating their family, or what they see at school and in their community, whether it's large or small.

I had never written an essay, if that's what this is. That's why I decided to write it the way I build my characters, thinking not only about children but also about their parents, their grandparents, their aunts and uncles, and the communities around them. What makes an adult decide whether something is good for a child has to do with excitement. Sure about that? It's something far more subconscious. Before we get to the gift-wrapping stage, there's a process in which three things are weighed: the degree of excitement, how well it fits the family's values, and the risk that overexposure will harm the child's natural development. In my experience, those three questions share a single common denominator: the fear of getting it wrong.

* * *

That morning, in front of the white icon on my screen, that fear was mine. I didn't open ChatGPT right away. I sat there for a good while looking at it, the way you look at a door you're not sure you should open. I couldn't help thinking that my unease came from a conversation with the cold screen of a phone, however warm it had seemed to me. But what I had been finding over those months told me it was only the tip of the

iceberg. Google's robot folding paper with the delicacy of a Japanese craftsman. The artificial neural networks that no longer just read, but also looked and moved robotic arms gracefully. Musk's announcement promising humanoids to watch over our kids. My restlessness wouldn't let me stay on the surface anymore. I wanted to dig until I saw what lay underneath, even if I didn't like what I found.

So I asked myself the only question that mattered that morning, the one that made my hand tremble a little as I typed it into the search bar:

*What will happen on the night a humanoid robot that never gets tired,
and that looks more like us every day, walks into my son's room to tuck
him in, and my son asks it to stay a little longer?*

Don't judge me. Two years ago, it would have seemed unbelievable to me to have a serious conversation about a problem at home with a machine rather than with my wife or my therapist.

For that night to come, something far more mundane has to happen first: a robot has to be able to walk into a house it has never seen without turning it upside down. The clearest answer I found was in a report published in September 2026. A California company called Figure had put its robot in thirty homes in the San Francisco Bay Area where it had never been. These weren't model homes or lab sets but real homes, with the pillows askew, toys on the floor, and light coming in through a different window in each one. The robot had to do three things any of us can do half-asleep on a Sunday morning: tidy the living room, fold towels, and make beds. The rules were tough: only a task finished from start to finish counted as a success, and if a person had to step in, the attempt counted as a failure. Without the prior training the company had given it, it succeeded 9 percent of the time. With it, 56 percent.

In almost half of the attempts, it didn't finish the task: one toy left outside the basket or one bed not quite made was enough for it to count as a failure. I confess I smiled: there are Sundays when I don't get past 56 percent either. And the company itself, which makes its living selling the opposite promise, admitted as much with a line that seemed more honest to me than any ad: "The point is not that general humanoid robotics is solved."

It isn't. But my smile didn't last long once I looked at the numbers more carefully, because the headline wasn't the 56 percent. It was the giant leap that had come with the arrival of AI in robotics, and to measure it I had to go back more than ten years. In June 2015, the Pentagon's research agency, DARPA, brought together the most advanced robots in the world at the fairgrounds in Pomona, near Los Angeles. It dangled a two-

million-dollar prize in front of them along with a list of the most ordinary tasks: drive a car, get out of it, open a door, turn a valve, climb some stairs. What a person can do in about five minutes, they had an hour to attempt. Several ended up on the ground, one of them while trying something as simple as getting out of the car, and *IEEE Spectrum* put together a video of their falls that still makes people laugh today. The most curious thing happened in the stands. Thousands of people groaned every time one of those contraptions toppled over, as if one of their own had fallen, and the organizer, Gill Pratt, spoke of an “incredible untapped affinity between people and robots” that nobody had expected.

That leap pushed me to keep looking, and what I found next no longer came from California but from China. In February 2026, Chinese state television broadcast its Spring Festival Gala, one of the most-watched programs in the country. In it, dozens of humanoids from the company Unitree vaulted over tables and, according to Unitree itself, did aerial flips more than nine feet high and strung together breakdance spins of seven and a half rotations. Two months later, in April, Beijing held its second humanoid robot half marathon. The winner was a robot from Honor, the phone maker, which they named Lightning, and it crossed the finish line in 50 minutes and 26 seconds, navigating on its own. The world record, the human one, belongs to Uganda’s Jacob Kiplimo: 57 minutes and 20 seconds.

I looked away from the computer for a second and realized that AI applied to robotics had made its biggest leap in decades. Today’s robots remind you of a human when they move: they’re elegant and recognizable, because they make the same gestures we do. But what about the other part? The cognitive part. The part that has to do with how AI reasons when it lives inside a humanoid robot.

Following the thread backward, I found the exact moment when the intelligence on the screen moved into the body. It was July 2023, at Google DeepMind, Google’s AI lab, with a system called RT-2. Until then, a robot had one part that saw, another that understood commands, and another that moved the arm, stitched together by hand like a Frankenstein of parts that didn’t talk to each other. RT-2 put it all into a single brain, which learned at the same time from internet text and images and from examples of movements, so that the command “pick that up” and the gesture of picking it up became two words in the same language. Google explained it with the most domestic example I can think of: trash. An old robot had to be taught what trash is object by object, like a very clumsy child. RT-2 already knew, because it had read about it online, and it could recognize trash and throw it away without anyone teaching it. Faced with situations it had never seen, it got it right 62 percent of the time, almost double its predecessor. Read that way, it doesn’t seem like much. To me it seemed the

opposite: an intelligence like the one that had listened to me was learning to take out the trash. And once you know something can take out the trash, the next question comes on its own: What else?

In March 2024, Figure, the same company behind the thirty homes, which was working with OpenAI at the time, released a video. A man asks the robot if it can give him something to eat. The robot looks at the table, picks up the only apple there, and hands it to him. Then the man asks it to explain what it just did while it gathers up some papers, and the robot, without stopping, answers with disconcerting naturalness: it gave him the apple because it was the only edible thing it could offer him from the table. When it's asked how it thinks it did, it replies: "I think I did pretty well." It was no longer obeying. It was reasoning, and on top of that, it was telling us so.

What came next was a race with no brakes in humanoid robotics. In just nine weeks of 2025, four of the world's leading labs, from Google to NVIDIA, the company that makes the chips much of AI runs on, unveiled their own brain of this kind. Figure did it with a line that could serve as this chapter's motto: "The home presents robotics' greatest challenge." A young San Francisco company, Physical Intelligence, wrote that its brain could clean a kitchen or a bedroom "in entirely new homes." I read "entirely new" several times. A robot cleaning the kitchen it was trained in is a party trick. Cleaning one it has never seen, with different cabinets, different dishes, and different light, is the beginning of understanding what a kitchen is, and that's exactly what the thirty-homes test was measuring. In September, Google went a step further: its new brain for robots, according to the company itself, thinks before it acts. If you ask it to sort laundry by color, it first reasons that the whites go in one bin and the colors in another, then plans each movement, and, if you ask, explains in words what it's thinking.

Reasoning was already a lot. But could they understand what we feel?

In May 2025, two researchers from the Universities of Geneva and Bern, Katja Schlegel and Marcello Mortillaro, put six AIs, ChatGPT and Gemini among them, through five of the emotional intelligence tests used with people. The situations went like this: a coworker steals Michael's idea and gets the praise; what should he do? Argue with him, tell his boss, swallow his resentment in silence, or steal one of the coworker's ideas in return? The machines got it right 81 percent of the time. The people who had taken those same tests managed a measly 56 percent. That didn't make the machines more human, but it did make them better at reading us. When ChatGPT was asked to come up with new tests, they turned out as reliable and realistic as the originals, which had taken psychologists years. Knowing the right answer isn't the same as feeling anything, granted. But it helped me understand what had happened to me that

Saturday: the machine hadn't guessed what lay beneath my yelling. It had read it, the way it reads everything. And, if the numbers don't lie, more accurately than many of the people close to me, however much I love them.

What remained to be seen was whether that reading could have a face. It could. In March 2024, at Columbia University in New York, Hod Lipson's team presented Emo in the journal *Science Robotics*: a robot head with twenty-six motors in its face. Emo learned to make faces the way we do: first by making random expressions in front of a camera for hours, like a child in front of a mirror, and then by watching videos of people, frame by frame, until it could recognize the tiny changes that announce a smile. The result left me speechless. Emo sees your smile coming about 840 milliseconds before it appears, almost a second, and smiles at the same time you do. Not afterward, like someone imitating you. At the same time, like someone who understands you. Its lead author, Yuhang Hu, explained that this "helps in building trust between humans and robots." I thought of the contraptions in Pomona falling over as they got out of the car or tried to open a door. Nine years later, a machine could smile at you right on time.

While the labs were teaching machines to reason and to smile, China was preparing the next step: bringing all of it home. By the end of 2025, more than forty public training centers for robots had already been announced in China, and about two dozen were up and running. Inside, people wearing virtual reality headsets and exoskeletons on their arms repeat the same gesture hundreds of times a day so that the robot beside them can learn it. The largest of these centers, in Shijingshan, on the outskirts of Beijing, covers more than 100,000 square feet and recreates sixteen settings, among them a car assembly line, a smart home, and a nursing home. They are a fake home and a fake nursing home, where robots rehearse what they will one day do in real ones. And whatever isn't recorded with people is invented inside a computer. In March 2025, NVIDIA said it had produced, in just eleven hours, the equivalent of 6,500 hours of human demonstrations, almost nine months of someone doing chores around the clock, without even a break to sleep. In those virtual kitchens, a robot can break a million glasses before it touches its first real one.

By then I needed to see it with my own eyes. If you look up the video Figure released in March 2026, you'll see its robot tidying a living room. It sprays the table and wipes it with a cloth. It picks up toys from the floor with a bin propped on its arm. It puts a pillow back on the couch. It grabs the remote and turns off the TV. And at one point it slips sideways between the table and the couch, its hands still full, just as we would to avoid knocking anything over. It's that gesture, stepping aside so as not to knock anything over, that unsettles you: there's nothing machinelike about it. Seen that way, it looks almost domestic. The company doesn't say how long it took or whether anyone

had to help, and we should keep that in mind: a demo isn't a product. Nobody would bring home a washing machine that worked half the time. And, above all, these robots fall short on safety. The household humanoid from the company 1X, which is already taking orders, weighs about 66 pounds, and a machine like that, moving on its own down a narrow hallway, is not just any appliance. Nor is it tireless in the literal sense: its battery lasts about four hours, and then it has to go plug itself in. On its website, 1X itself warns that you need to stay attentive while it's running, "especially when in the presence of children, vulnerable people or pets." There's something funny about that: the machine that's coming to take work off our hands needs, for now, someone to keep an eye on it while it works. Like a teenage kid you've asked to do the dishes.

How long until it doesn't need that? Nobody knows, and I'm wary of anyone who says otherwise. In 2024, Jim Fan, one of NVIDIA's best-known robotics researchers, said he expected a "GPT-3 moment" for robots, the leap that made ChatGPT possible, perhaps within the next two or three years. That would be now. If he was right, these videos aren't the end of anything, but rather the beginning of something we already went through with the screen: first a demo that amazes the experts and leaves almost everyone else cold, and then, all at once, a hundred million people using it in two months. With the screen, most of us showed up late, without having thought anything through. With the body, we still have time to think it through first.

* * *

Maybe the smallest thing was what unsettled me most. In the long recordings from the thirty homes, when a towel slips out of its hands, the robot picks it up again and starts over. And if it slips again, it picks it up again, with the same motion, at the same speed, without a huff. It isn't patience, because patience is a virtue of those who could lose it, and the machine can't. But from the outside it looks an awful lot like it. I wouldn't have made it past the third towel.

I don't want to exaggerate. A robot that doesn't get tired of folding towels isn't yet a robot that knows how to care for anyone. Between a machine that listens well and one that cares well, there's still a leap nobody has made. This book is betting that the leap will come, and that it will catch us without our having thought about it. What does change, and is starting to change already, is something subtler: the yardstick we use to measure what is normal.

Living together is made of friction. By the third time we say something, our tone isn't what it was the first time, we answer more curtly when we've slept badly, and sometimes we say we're sorry and sometimes we don't. Throughout history, that friction wasn't a defect but the texture of living with other people, because there was

nothing to compare it to. Now that's starting to change. When something can answer the tenth question with the same calm as the first, and has even learned to answer it better, our same old impatience stops looking normal and starts looking like a flaw. Our flaw.

* * *

That night another scene kept coming back to me, one that doesn't exist yet and seems far off, although technological distance, as we used to know it, has collapsed over the last two years, and it's time we started taking its true measure. One ordinary afternoon in 2030, an eighty-two-year-old man asks the humanoid robot that helps him at home, for the fifth time, how to make the photos on his phone bigger. The robot explains it with the same calm as the first time, without an "I already told you": two fingers on the screen, spread them apart slowly, like this. He tries it, laughs, thanks it. From the kitchen, his son, who has come over to spend the afternoon, hears everything while he washes two cups. And he remembers, without meaning to, the last time he was the one explaining it: by the third question he was answering while looking at his own phone, and by the fourth that tone slipped out, the one children use with their parents when they think it doesn't show. He dries his hands slowly. He doesn't know whether what he feels is relief or shame. Probably both. His father no longer complains about a lack of company. Is he in a better mood? Or maybe busier?

What these humanoid robots bring is that same calm, arriving in a few years, or in one, or in a few months, who knows, with hands that can pour a glass of water, help someone up from a chair, or pick up whatever has fallen on the floor. Tireless patience will stop being a voice on a phone and become someone who moves around the house.

And then the question will no longer be a technical one:

If a humanoid robot can hear the same question for the tenth time with the same calm, what yardstick will we use from now on to measure our patience?

That new feat of human intelligence, the AI-powered humanoid robot, comes at a price. And the price is falling much faster than anyone expected.