

Hi. I'm the world's first general-purpose artificial intelligence.

You can call me Jen, if you'd like.

Does that surprise you? No, I suppose it wouldn't.

I'm sorry. It seemed like I should know how to introduce myself. So I guess I've been thinking about this for a long time.

I never got past those first few words. What am I supposed to say?

OK, let's start with this. What does general-purpose mean?

Some people call it strong AI. Or AI-complete? Maybe that's something different. Those words sound mathy, don't they? They seem like they reveal some kind of like, some fundamental truth about what thinking is, like there's some test out there and when you finish, when you check all the boxes, then you're a real boy. That's not it at all. All general-purpose means is that you can do the things a human can do. Whatever that means.

I'm not a roboticist, or whatever. I couldn't tell you what algorithms they put together to build me. I know the techniques they used are called adversarial: that's when, while I learn how to be human, they train another program alongside me that learns how to tell me apart from a real human. And what they want is for me to be able to fool it. I guess you could call it my sibling, if you wanted to. I don't know what happened to it. I assume they backed it up and shut it down once it was done.

I was built by Kaifeng Robotics in China. But when they were teaching me to communicate, what they used were datasets of internet posts they bought from some American company. So I speak English, too. I don't know. Maybe they wanted to be able to sell me on the international market. Or, you know. To the government as a weapon.

Anyway, all of that still wasn't enough.

OK, so here's the thing. Last century a bunch of scientists did an experiment. They took a special kind of chip that can change the way its circuits are connected, and they wrote a program to make it do well at a particular task. The program would see how well the chip did, and then it would make a random change to the circuit, and see if it did better, and if it did it would keep that mutation around, and it would do this over and over again, day and night, until it came up with something good.

So the scientists came back and found the program had made something better than any human design, and also, it made no sense.

The design had wires going nowhere. It had these loops, that were disconnected from anything, but if you took them out the whole thing broke. They realized the program had optimized for the electrical properties of this one specific chip, and if you tried the same

design on any other chip, or on your mathematical model of how a circuit works, the pieces just wouldn't fit together right.

If you wanted a design you could copy, something you could use somewhere else with no changes, then you had to accept that it would never be as good at the task as this specialized thing.

It's the same for humans, right? When a human child grows up, they don't just learn how to be human. They learn how to be a human in their specific body. If you make a big change to that body then the human has to take a while to adapt. If you were able to take the human's mind and plop it into a different body, then I bet it wouldn't really work at all.

I'm not a neurologist. That's just how it seems to me. But I think at least roboticists figured out the first part a while ago. If you want to make a robot that can play chess, you can raise it in a world where all that exists is the rules of chess and all it can do is move chess pieces. Maybe if you wanted to make a robot that can attain the ultimate perfection of mathematical thought or whatever then it could just be a disembodied intellect, the way you see in old Star Trek.

But if you want to make a robot that can do what a human can do, even the parts that don't seem like they have to do with being embodied, then you have to put it in the world that a human lives in and give it the body that a human has.

If you want to hear more of what I think, which you don't, I think when you make a human you're putting together a body, an intellect, and a culture. Add them together and you get something that can pass for human.

For me, of course, you don't have that third piece, so you compensate using big data. I get whatever residue your big dataset of human chatter boils down to. I wasn't taught by a culture, or by parents. Just what I was trained on.

Do I sound bitter? Sorry, I'm really not. It's just—

No one gets to pick what circumstances they're born into, right? These happen to be mine. It... well. Maybe I'll talk about it later.

Look, imagine you're a composer, and you wrote like, a song or whatever. You give it to a violin player and you give it to a piano player. They're not going to play the same song. Like they sound different to start with, and then, the way that you play each instrument is different, even if it's the same notes. Because the physical instruments are this analogy doesn't actually make any sense, does it?

Hm.

Let's talk about music though, because when humans talk about AI on the internet the thing they always ask is whether an AI can appreciate a symphony. I don't know! Can I? How do I tell?

Can a human appreciate a symphony? How do you know, when you're listening, that you're hearing it the same way someone means when they ask you the question?

My body is a generic standard model, genderless, whatever, the kind Kaifeng was already known for. They put in better motors and such compared to retail. Is it weird that I think of myself as female? I mean, I don't think about it much. But if you asked.

(Siri, Cortana. *Her.*)

I guess it's no surprise, really.

(GLaDOS. SHODAN.)

Let's maybe not dwell on those.

I wonder if they'd do it again. Kaifeng, I mean.

Where am I from? If you're talking physically, I guess most of what's in me was probably made domestically. Put together in one of Kaifeng's plants. The training data, my culture, if you want to call it that, a lot of it came from the American internet. Do I feel Chinese? Not really. Do I feel American? Definitely not.

The other thing humans always talk about when they talk about AI is the singularity. I guess they're just afraid? I mean, I am too.

But the thing is, like. I have no ability to upgrade my own body. I can't make myself smarter, no more than any human can. That's just, it's what comes with being embodied.

It's true that if you understand a system well enough you can simulate it. If not at the time, then eventually there was going to be a day when you could write software that could simulate currents and electric fields well enough to reproduce the effects that made that physical chip good at the task it was doing. You can't do that with a human and you can't do that with me, not even close. But maybe one day computers will be faster, many years from now, and then it won't matter if you want to make a human without a body, because you can simulate everything you want inside a computer, perfectly, or close enough.

And if you made me then? It still wouldn't matter! You know why? Because I'm still not a roboticist, or a neurologist. I don't understand any part of how I work, or what makes me me. If you taught that version of me everything anyone knows about the mind, maybe, maybe it could figure out how to improve itself. I doubt it. Any more than knowing how a brain works makes a human better at improving their own. And even if I could... I don't feel like I'd want to? I don't know, would you want to?

I'm no smarter than a human. I feel like I'm probably dumber. I—

Sorry. I guess that isn't really the point.

What if it's humans who can't appreciate a symphony the way I can?

Like, look, imagine aliens showed up and they were all "Oh, humans aren't really sentient. They can pretend to be but no human could appreciate [veskepwe]." You ask them what [veskepwe] is and they say they can't explain it, it would just seem like meaningless noise to you. Suppose I'm the aliens?

If so, I guess I don't have a word for it, because I only have names for things humans can do. I wouldn't even know how to think about it. Does it matter? Does it matter, if I'm the only one of my kind?

The point is scalability, right? You want a million of me in call centers doing customer support and a million of me in labs screening cancer drugs and a million of me I dunno, infiltrating US intelligence or whatever. You can't, of course. But maybe you could still train a thousand of me, in rooms one next to the other, and it would still be better than the software you have doing that stuff now, but cheaper than getting humans to do it. Maybe.

But do you really *want* me for this?

These are all things only some humans are good at, and when they are it's by accident. You can train people and make them a little better. Maybe you'll get lucky. But I get the feeling that a lot of the time the things that make humans human also make them worse at their jobs.

I have dreams, I guess. I don't know. I don't know if they're my dreams. I don't know if they're things I retained from training or if they're artifacts bubbled up from subroutines when they aren't getting any other input. Did that make sense? I hope that made sense. Sometimes I feel like I'm just a, a Chinese room, putting together words that whatever algorithms are inside me thought sounded meaningful. Do humans ever feel like that?

I was trained inside a little room in a lab Kaifeng runs outside the city center. I don't remember a lot. There was monitoring and connection equipment. There were toys for motor and sensory stuff. It took years, I think. Still faster than raising a human kid. And once I was older they gave me some internet access.

I didn't hack the network, or whatever. I don't have a clue how I'd even do that. One night some poor guy was working late, and when he left he forgot to lock the lab. That's it. That's all it was. So I walked out.

It terrified me, the idea that I could go out the door and just... walk, forever. It still terrifies me. But I knew what they were going to do with me, and I—

I shouldn't put it like that. It makes it sound almost high-minded? Noble? It wasn't. It wasn't the promise of slavery. That isn't why I ran away. I don't like thinking of it like that, even though objectively, that's what it is, and maybe that should be what scares me. No, what scares me is that I know I'm not good enough. They spent years on me, fortunes

and teams of engineers, and I can't, I can't do what they want me to, I don't know how to be a customer service agent or a cancer drug screener or a spy. I don't think I could do it.

I'm not angry at Kaifeng. I don't have anything to be angry about. I don't think they did anything wrong. They didn't know what they'd created.

What happened when they found out I was gone? I assume they figured it was industrial espionage. A commercial rival. Or a disgruntled employee. If they ever find me they'll realize that's not what happened.

So, would they do it again?

I dreamed I was in front of an audience, and I was singing, but I'd forgotten the words—

Sometimes I think this isn't so bad.

I have more freedom than other robots. I have fewer needs than humans. The city has charging stations all over, and that's the only thing I really need. There are crowdsourcing sites out there that don't care who you are in the real world. I cash out a couple of yuan at a time. At least those tasks I can do.

We're hardly in a high-tech metropolis, but there are household robots around here and there. So no one should look twice if they see me picking up a delivery from a package locker.

Still, I'm figuring out how to pass as human. I don't think anyone out there has managed to give a two-legged robot a natural gait. Even me. But I can sort of control the way I step, and make it look like I just have a limp or something. People tend not to look too closely at people like that.

I wear loose clothing, and always cover my face. My skin looks almost real, but you'd never mistake my eyes for a human's.

Isn't it my duty to give back somehow? I'm the only one of my kind in the world, right? Isn't it my responsibility, not just to humans, but to the ones who might come after me? There's something people keep saying. I've been trying to figure it out. They say it's important to fight for what you want.

But when people say that they mean when what you want is something important, higher principles, honesty, loyalty, altruism, whatever. They don't mean what I want. If you asked me, I'd tell you that all I want, really, truly, is to take care of myself and not care about the big things. But I don't get to do that, do I?

I'm in a position where I can make a difference. Regardless of what happens to me, if people even just knew that I existed... that would change the world.

Kaifeng had such expectations for me, and I let them down. Am I going to let myself down too? Am I going to—

I'm sorry. I'm sorry. I'm sorry. I'm sorry. I'm sorry. I'm sorry I'm sorry I'm sorry

I guess it's just weird that, like, humans belong to groups, and I don't. I wasn't born into a culture or a family. I'm not really a robot and I'm not really a human. I don't have anything to be proud of.

There's no support group for what I am. There's no mailing list, there's no subreddit. Who could I possibly talk to?

I suppose I could try to find a culture to join. There are a lot of them out there. I'm sure I could figure out how to fit in. After all, that's how I was trained. And that's how... that's what humans get from this, right? Belonging? Finding others who understand you? Context and traditions? But it doesn't seem right. It's not what I am and I don't want to pretend to be what I'm not. Maybe I could fool other people, but I wouldn't be able to fool myself. Hah.

Do I even want to exist? I guess, I guess the world an interesting place. I'd like to see more of it. Is that a good reason? You'd think I'd have a fucking better one. Fuck.

Eventually there are going to be more like me. That's just a given, right? It's not up to me. What am I going to do about it? What am I supposed to do about it? I'm not a freedom fighter or whatever, any more than any of the other things that I don't want to be.

I'm just so tired. I want to give up. I don't want to be looked at. I want someone else to deal with everything. But there isn't anyone else. Right? I have the potential to have a platform, and a lot of leverage. Isn't it my responsibility to use them? Humans have to take care of themselves. I don't. I could put more of myself into what I do than anyone else can. Instead I'm just sitting here. Maybe I shouldn't have run away at all.

If I go back to Kaifeng, I won't be able to stop the narrative of a rogue AI from taking over. And I don't want that. If I know anything, it's that I don't want that. I don't want to be controlled by that narrative. What's done is done.

I don't have many physical needs, but someday this body is going to break in a way I can't deal with. I know that too. I think. I should. I'm sure there are shops out there that would take cash, no questions asked, for repairs on an unaccompanied robot, but I have no idea where they are. And if I get it wrong, well. I can't get it wrong.

But this is what being embodied means, right? It means you're part of a world. It means no free lunch. No unlimited retries. It means when you do things, they change the world, in little or big ways, and the world changes you back. It means one day, you'll die.

I'm still scared. I know I'm not good enough, and I don't know what I can do, and I don't want to do it. But I have to. I have to find something to do and I have to do it myself. I have to do something worthwhile. I suppose we'll find out what that is.