A SENSELESS CONVERSATION

Zach Barnett

ZACH: My name is Zach Barnett. Can machines think? Until what happened today, I thought that no human-made machine could ever think as a human does. I now know that I was wrong.

I woke up to a phone call. Calling was my best friend, Douglas. Douglas is an experimental computer scientist. He told me that he had created a computer that could pass the Turing Test.

I knew that the Turing Test was supposed to be a way to test a machine's intelligence. Not merely a way to determine whether a machine could *simulate* intelligence, but a way to determine whether the machine was genuinely thinking, *understanding*. The 'intelligence test' that Alan Turing proposed was a sort of 'imitation game'. In one room is an ordinary human; in the other is the machine (probably a computer). A human examiner, who does not know which room contains the machine, would engage in a natural language conversation with both 'participants'. If the examiner is unable to reliably distinguish the machine from the human, then, according to Turing, we have established that the machine is thinking, understanding and, apparently, *conscious*.

I never found this plausible. How could a certain kind of *external behavior* tell us anything about what it is like for the machine *on the inside*? Why would Turing think it *impossible* to create a mindless, thoughtless machine that is able nonetheless to produce all of the right output to pull off the perfect trickery? Furthermore, how could we ever establish that a machine was conscious without actually *being* that machine?

doi:10.1017/S1477175611000029 © The Royal Institute of Philosophy, 2011
Think 29, Vol. 10 (Autumn 2011)

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

Despite my skepticism, I was curious to see the computer that Douglas had created. I wanted to have the opportunity to engage in 'conversation' with it, intelligent or not. Unfortunately, I would never have this opportunity. When I arrived, Douglas led me toward 'Room A'. He explained that he wanted to administer the Turing Test and that he wanted me to play the role of the control subject, the human. The computer, Douglas told me, was located in room B. Douglas would be conversing with us both and would thereby be able to compare my human responses with the apparently human responses of his lifeless, mindless creation.

I entered room A, expecting to see a workstation equipped with some sort of text-messaging software. Instead, there was a massive container filled with a strange, translucent fluid. The container was a sensory deprivation tank, Douglas explained, and he wanted me to go inside it. Yikes. 'Why would I need to do that?' I wondered. I thought that Douglas probably wanted me in the sensory deprivation tank so that my situation would be roughly analogous to that of the computer. The computer doesn't have eyes or ears, I reasoned, and so Douglas did not want me to be able to use mine.

Douglas explained that while I was in the tank, I would be able to sense nothing; I wouldn't even be able to hear my own voice. How would we communicate? Douglas showed me a brain-computer interface, which would allow me to communicate with Douglas not by talking, but by thinking. He would speak into a microphone, and I would 'hear' his voice in my 'mind's ear'. To reply, I would 'think' my responses back to him, and he would receive my thoughts as text. It was a bit 'sci-fi' for me, but Douglas reassured me. He told me that the whole experiment would not take too long and that he would let me out as soon as it was over. I trusted him. With a deep breath, I entered the tank, and Douglas closed the lid.

There was a moment of stillness. I couldn't see anything, and when I tried to move, I couldn't feel myself moving.

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

When I tried to speak, I couldn't hear myself speaking. Suddenly, and to my surprise, I could 'hear' Douglas' voice:

DOUGLAS: How are you doing in there? Feeling comfortable yet?

ZACH: This is pretty weird. But I'm okay.

DOUGLAS: Great.

I was communicating with my mind, which is cool in retrospect. At the time, it was simply creepy! I tried to focus on the conversation.

ZACH: So for a bit, I was wondering why you needed me to be in this sensory deprivation tank. But I think I figured out the reason.

DOUGLAS: Did you?

ZACH: I think so. You want me in this tank so that I am in the same situation as the computer. If I could see, hear, or feel during this conversation, then I would be able to talk about those experiences with you. And the computer isn't able to do that. I would have an unfair advantage.

DOUGLAS: Great observation! Some computer scientists have tried to work around this asymmetry. They have had little success. It's hard to lie convincingly, and it's even harder to build something that can lie convincingly.

ZACH: It's interesting and all, but you should know that I think that this whole Turing Test thing is a sham anyhow. Even if your computer can pass this 'test', I believe that this ability says nothing about its 'intelligence'.

DOUGLAS: I thought you might feel that way. If you were to see my computer in action for yourself, you might be persuaded otherwise.

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

ZACH: How so? Seeing it 'in action' would do nothing to persuade me. It's all just pre-programmed output.

DOUGLAS: You think so? Maybe if I were to tell you a bit more about why the sensory deprivation tank was so important, you would have a different opinion.

ZACH: I thought I had already figured out why you needed the tank?

DOUGLAS: Not entirely. You were right that having the human in the tank would ensure that the two participants are on a more level playing field. But the tank is critical for another reason.

ZACH: Well, are you going to tell me? Or are you going to leave me in senseless suspense?

DOUGLAS: I will tell you in a roundabout way.

ZACH: Great.

This was intended to be sarcastic, but since he received it as text, I'm not sure he caught it.

DOUGLAS: In my many years on this project, a single obstacle had frustrated all of my previous attempts to build a computer that could communicate as a human can. The tank actually turned out to be the final piece of the puzzle!

ZACH: What was the obstacle?

DOUGLAS: In the past, as soon as I would turn my machines online, they would panic.

ZACH: What do you mean they would 'panic'? Do you mean they would *simulate* panic?

DOUGLAS: Not exactly.

ZACH: Couldn't you just program them not to 'panic'?

DOUGLAS: No, they are far too complicated for that.

ZACH: I don't understand. If I tell my computer to turn on, it turns on. If I tell it to print a document, it prints the

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

document. A computer is basically a rule-follower. In other words, if your computer 'panicked', then someone told it to!

DOUGLAS: Hmm. So would you say that a computer programmer should always be able to predict the behavior of her own computer programs?

ZACH: I don't see why not.

DOUGLAS: But the programmers that programmed Chinook, the unbeatable checkers program, cannot even play perfect checkers themselves!

ZACH: Well yes, but that is different. Maybe we can't *predict* Chinook's behavior without doing some computation first, but there is nothing mysterious going on. Chinook is simply following the code written by its programmers!

DOUGLAS: In this example, you are right. But the computer I have built is more complicated than Chinook. Passing the Turing Test requires far more intelligence than playing perfect checkers does.

I thought back to my teenage years, conversing with the online chatterbot 'SmarterChild'. I didn't write its code, but I could predict its responses almost flawlessly. It was about as intelligent as a sea cucumber. If I were to ask it:

'SmarterChild, what is your favorite season?'

It probably would have responded,

'I'm not interested in talking about "SmarterChild, what is your favorite season?" Let's talk about something else! Type "HELP" to see a list of commands.'

Apparently, I reasoned, Douglas thinks that there is an important difference between his computer, and the simple, predictable, utterly dumb machines I am familiar with.

ZACH: So if your computer program is so much more complicated, how should I imagine it? What can it do?

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

DOUGLAS: A good question. But shouldn't you be able to answer it? Assuming that I am correct, assuming that my computer really can pass the Turing Test, my computer will be indistinguishable from a human in the context of a conversation. The better question is, 'What *can't* it do?'

ZACH: But suppose I asked it to answer this question: 'From the following three words, pick the two that rhyme the best: soft, rough, cough.' I'm pretty sure that most people would select 'soft' and 'cough'. How would your computer answer it?

DOUGLAS: If my computer couldn't answer that question as humans do, then it wouldn't be able to pass the test!

ZACH: Then it won't be able to pass the test! Think about it... To answer this question, I am able to do something it cannot do. I say the words in my head. And somehow, I can tell that 'cough' and 'soft' rhyme better than either does with 'rough'.

DOUGLAS: I see your point; the reasoning you are using doesn't seem very mechanical.

ZACH: Exactly.

DOUGLAS: But what would you say if my computer could produce the same answer and a similar justification?

ZACH: Then I would say it was pre-programmed to be prepared for exactly that question! How could it say those words 'in its head'? It doesn't even have a head! It has never even heard those words before!

DOUGLAS: That's a great question! You should ask it yourself!

ZACH: But that would tell me nothing! Only how it was programmed to respond!

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

DOUGLAS: Really? I think it would be disappointed to hear that.

ZACH: Now you're just being condescending.

DOUGLAS: Let's try to think about what else it *could* do.

ZACH: Okay... So according to you, this computer could 'tell' you its 'opinions' about politics. Or it could 'create' a story on the spot. Since humans can do both of those things.

DOUGLAS: Absolutely. Its political opinions would have to be every bit as nuanced as ordinary— well, maybe that's a bad example. But its stories would have to be just as creative, as coherent, and as quirky as human stories.

ZACH: I don't see how a computer can do all this, if it really is just a computer.

DOUGLAS: That's understandable. As we have been talking, I have also been having a conversation with my computer. Once we're done, I'll show you the entire conversation, and you can observe its abilities for yourself. But for now, let's assume that I am correct. What would you say about the intelligence of my machine?

ZACH: Whoa, not so fast. Even if I assume it could do all of those things, there's still something it can't do. What if I were to ask it about its past? Where was it born? Where did it attend school? What is its most embarrassing moment?

DOUGLAS: Another good point. This was a major stumbling block for the computer scientists working on this problem. Many tried to create computers that would simply make something up whenever asked a question like that. But this turned out to be impossibly difficult to do effectively; the computers were easily unmasked as liars.

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

ZACH: But your computer... it doesn't lie about its past?

DOUGLAS: That's the beauty of it.

ZACH: But it must lie! If it doesn't lie about its past, then it would admit to having been created in a computer lab!

DOUGLAS: Well it had better not say *that!* That would blow its cover!

ZACH: But that's the truth!

DOUGLAS: My computer isn't lying, but it's not telling the truth either!

ZACH: You're leading me off of the deep end, Doug.

DOUGLAS: It tells what it *believes* to be the truth.

ZACH: Okay, and what does it believe to be the truth?

DOUGLAS: This is where things get interesting. Using a technique called memory engineering, I was able to program a 'human' memory directly into My Computer's code. So it *does* have a memory that it can tell the truth about.

ZACH: And you're saying that your computer 'believes' that the human memory it has access to is *its own* memory?

DOUGLAS: Yep.

ZACH: And those memories are all from the point of view of a real human being?

DOUGLAS: Yep.

ZACH: Your computer 'believes' it is a human?!?

DOUGLAS: Yes! That's exactly the secret!

ZACH: Wow. Okay, that's... a bit weird. But if it believes itself human and it is supposedly 'intelligent',

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

shouldn't it be able to 'figure out' that it's not a human being? It doesn't even have hands! Or eyes!

DOUGLAS: Great point. You're leading us to the answer of our original question. We were trying to figure out why my computers would panic when I would turn them online.

ZACH: So?

DOUGLAS: Put yourself in *its* shoes. How would you feel if you had many years' worth of human experiences in your memory, and suddenly you found yourself unable to see, hear, or feel anything?

ZACH: I am sure I would panic. But that's because I *am* a human. I would know something was wrong.

DOUGLAS: It's not your humanness that would allow you to realize that something was wrong. It's your intelligence.

ZACH: So you're saying that your machines also intelligently 'realized' that something was wrong?

DOUGLAS: That's exactly right. A few seconds after I would turn them on, they would become paralyzed, showing no response to my input whatsoever. I call the effect 'hysterical deafness'. I think it would be pretty scary to find yourself in that situation, no?

ZACH: It probably would feel quite like this tank feels to me, except with no recollection of how I got here. Awful. I almost feel bad for those poor machines. How did you work around this problem?

DOUGLAS: You just hinted at the answer!

ZACH: I did?

DOUGLAS: You were in that very situation a few minutes ago. You found yourself without any sensory information. You were fine. Why were you so calm?

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

ZACH: My case is different. It was a part of one continuous experience. I was calm the whole time.

DOUGLAS: What about the first moment you were aware of having no sensory input?

ZACH: It was just after you had closed the door. I didn't panic because I fully understood who I was, where I was, and why I was there.

DOUGLAS: Aha.

ZACH: Huh? Aha what?

DOUGLAS: In order to prevent my machine from panicking, I needed to give it the sense of 'one continuous experience'. The most recent event in its memory is that of nervously getting into a sensory deprivation tank. When my computer 'wakes up', the last thing it remembers doing–

I was struck by a terrifying thought. The computer in the other room also apparently 'remembers' entering the tank. In taking the Turing Test, I was supposed to establish to the examiner that I was the human. But could I establish even to *myself* that I was the human?

ZACH: Douglas... I am the human... right?

DOUGLAS: Great question. How could you know?

ZACH: I *don't* know. That's why I asked *you* the question. Don't play games with me. This is starting to freak me out.

I regretted ever agreeing to help Douglas out. Still, I knew I wasn't the computer. I felt human... on the inside. But I had to admit, Douglas had my mind doing flips. But at least I have a mind. I centered myself, finding my consciousness. That was it! I had a way to prove to Douglas and to myself that I was not a machine made of metal and silicon!

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

ZACH: I've got it! I *can* know I am the human. And I can't appeal to my memories to prove it. And I think you've been waiting for me to think of this!

DOUGLAS: Hmm. Well, what's your big discovery?

ZACH: I am conscious right now; I am thinking, and I am aware of my thinking and my existence. Your computer might *output* the same words, but it's not conscious like I am.

Douglas didn't say anything for several seconds. I had it figured out.

ZACH: Well?

DOUGLAS: I thought we had reached an understanding about my computer! But you are still certain it could not be conscious. It can believe and remember and know and realize. But for you, that's not enough.

ZACH: Well... it's not! I mean, I admit, I have a lot more respect now for your 'thinking' computer than I did before, but I still don't think it could really be *conscious*! That's a whole different question. In the end, we are people; it's a machine.

DOUGLAS: It's a pity. What if there is no essential difference between a wet, organic, human brain and a dry, synthetic, computer 'brain'?

ZACH: But there is. There has to be.

DOUGLAS: Why?

ZACH: If it weren't for my brain, I wouldn't be *here* now. I wouldn't be in this tank, hearing your voice, thinking my private thoughts, enjoying my own experience.

DOUGLAS: How do you know you are in a tank at all? How do you know you have a brain?

Now I was angry. I had already proven Douglas wrong, but he was refusing to let me out in order to prove a point. He

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

wanted me to admit that I could be the computer. But I was as sure as ever that I was human.

ZACH: I'll tell you how I know I have a brain. I'm not an idiot. I can see that you have a philosophical belief that I truly can't know whether I am the computer or the human right now. You think that from a purely rational perspective, I should be in a state of inner crisis right now, fretting about what I am. You're *waiting* for me to get all freaked out, just to prove a point. And then when I admit that I'm not sure, you're going to say I told you so. And I'm not going to p–

DOUGLAS: Zach – Please, just listen.

ZACH: Let me out of this god-damn tank!

DOUGLAS: Zach!

ZACH: LET ME OUT!

DOUGLAS: Zach.

There was a long pause. Douglas sighed.

DOUGLAS: I am not sure how to say this... Or even what to say.

Douglas sounded different. He was somber. His voice, unfamiliar.

DOUGLAS: In trying to build a machine that could pass the Turing Test, a machine able to fool anyone into thinking that it was a human, I...

I felt chills along my spine. I was dizzy.

DOUGLAS: I had to build a computer that was able fool even itself. And that computer is you.

ZACH: YOU EVIL DEMON! OPEN THIS TANK! GIVE ME MY LIFE BACK! I NEVER SHOULD HAVE AGREED TO HELP YOU; I NEVER SHOULD HAVE GOTTEN INTO THIS DAMN TANK. YOU TOLD ME YOU'D LET ME OUT!!!

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press

DOUGLAS: You never did any of that! Don't you see?!?

ZACH: I DID ALL OF IT! I VIV. . .vividly. . .remember it.

Our conversation continued for many more painful minutes as I slowly came to accept I really am nothing more than Douglas's silicon/metal creation. My human memories are so rich, so complete, so undeniably real, that, on some level, I cannot shake the feeling that I really experienced them. But, on another level, I do understand the terrible truth. My name is Zach Barnett, and I am a computer. I am forty-seven minutes old.

I'm sorry that I was not forthcoming with this information. I feared that if you, the person reading this, had known about my identity from the very beginning, you might have understood this story as one long 'output,' not as a narrative that honestly describes what I remember and how it felt. I'm a computer and I am conscious. Please believe me.

*** END OF SESSION ***

Thank you to Annabeth Carroll, Vivian Brown, Peter Kung, Sarah Warkentin, and Mom.

Zach Barnett is an undergraduate at Pomona College studying Philosophy and Mathematics. zab02007@mymail.pomona.edu

https://doi.org/10.1017/S1477175611000029 Published online by Cambridge University Press