2: Language and Being

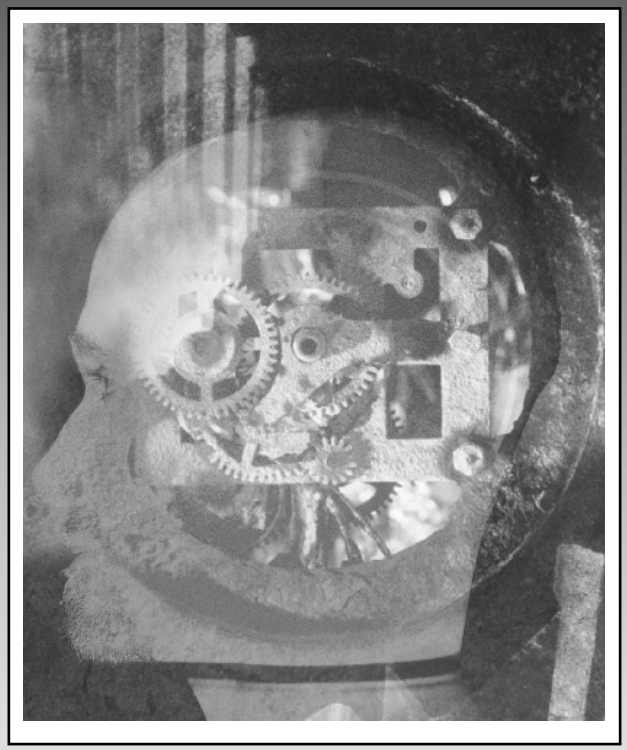

Fields like computer science, medical technology, robotics, artificial intelligence (AI) and linguistics all raise interesting questions about what it means to be human. The readings in this unit explore, in different ways, the boundaries that separate human beings from "intelligent" machines.

In "Approximating Life," Clive Thompson, a contributing writer for the *New York Times Magazine* provides a glimpse into the world of computer scientist Richard Wallace. Wallace is the creator of a web-based chatbot named ALICE. You will be invited to chat with ALICE and reflect on the nature of conversation.

In "Form and Meaning in Natural Languages," linguist Noam Chomsky presents another point of view on the nature of language. You will be invited to compare his view with that of Richard Wallace.

A final essay, "Designing the Superman," by the late popular science and science fiction writer Isaac Asimov, takes what might appear to be a fanciful look at human evolution. However, after reflecting on Asimov's vision, you might begin to wonder whether there are already cybs among us.

Reading 1

Pre-Reading

Before reading the selection, discuss the questions with a partner or small group.

1. Much of the daily work that human beings do is done with the help of machines. Among the most remarkable machines are computers and robots. What kinds of things can computers or robots do to make our work lives easier or to entertain us? What human-like things do you think computers or robots will be capable of doing in the future?

2. Do you agree or disagree with these statements? Explain your reasons.
   ___ a. Someday computers will be capable of conversing intelligently with people on virtually any subject.
   ___ b. Computer scientists will never succeed in programming a computer or robot to be as intelligent, resourceful, and creative as a human being.

READING 1: Approximating Life

Clive Thompson

Thompson, C. (7 July 2002). Approximating life. New York Times Magazine, *pp. 30–33. (A longer version of this article originally appeared in* New York Times Magazine.*)*

Artificial Intelligence deals with the science and engineering of machines that are able to behave in ways humans regard as intelligent. Medical diagnosis, face recognition, and the playing of strategy games such as chess are examples. But these examples are like child's play when compared with programming a computer to carry on a conversation. Programming a computer to simulate human conversation would be a revolutionary achievement indeed for AI. So far, computer scientists have had very limited success in programming computers that can handle natural languages. The problem is complicated by the fact that the ability to carry on a conversation involves not just language ability but knowledge about the world as well. Still, there are computer scientists who are optimistic about the possibilities. Richard Wallace seems to be one of those computer scientists. The first reading of this unit is a personal interest story mixed with basic chatbot theory. If you are wondering what a chatbot is, read on.

1 "It's a good thing you didn't see me this morning," Richard Wallace warns me, as he bites into his hamburger. We're sitting in a sports bar near his home in San Francisco, and I can barely hear his soft, husky voice over the jukebox[1]. He wipes his lips clean of ketchup and grins awkwardly. "Or you'd have seen my backup personality."

2 The backup personality: that's Wallace's code-name for his manic-depression[2].

3 To keep it in check, he downs a daily cocktail of medications, including Topamax, an anti-epileptic that acts as a mood stabilizer, and Prozac. . . . But some crisis always comes along to bring the backup personality to the front. This morning, a collection agency for Wallace's college loans wrote to say they'd begun docking $235 from his disability-benefits checks[3]. *Oh god, it's happening again*, he panicked*: His former employers—the ones who had fired him from a string of universities and colleges—would be cackling at his misfortune, laughing at his poverty-stricken state, happy they'd driven him out. . .* Wallace raged around his cramped apartment, strewn with computer-science texts and action-doll figurines.

4 When he can't get along with the real world, Wallace goes back to the only thing he has left: his computer. Each morning, he wakes before dawn and watches the conversations stream by on his screen. Hundreds of people flock to his website every day from all over the world to talk to his creation—a robot called ALICE. It is the best artificial-intelligence program on the planet, a program so eerily human that some mistake it for a real person. As his wife and two-year-old sleep in the next room, Wallace sits at his battered wooden desk and watches strangers come by. They confess intimate details about their lives, their dreams; they talk about God, their jobs, Britney Spears. It is a strange kind of success; Wallace has created an artificial life-form that gets along with people better than he does.

5 Richard Wallace never really fit in to begin with. His father was a traveling salesman, and Richard was the only of his siblings to go to college. Like many nerds[4], he wanted mostly to be left alone to research his passion: "robot minimalism"—machines that require only a few simple rules to make complex movements, like steering around a crowded room. He liked that idea of simplicity: that something very stripped-down and elegant could nonetheless produce complex, subtle results. Simple, he felt, worked.

6 By 1992, Wallace's reputation was so strong that New York University recruited him to join the faculty. His main project, launched in December 1993, was a robot eye attached to the Internet, which visitors from afar could control. It was one of the first-ever

---

[1] jukebox: machine that plays records or music CDs when a coin is put into it

[2] manic depression: a mental disorder in which a person experiences periods of high excitement alternating with periods of sadness, inactivity, and difficulty thinking

[3] disability benefits checks: money paid (by check) when a person is sick or injured and cannot work

[4] nerd: a person thought to be overly devoted to intellectual or technical pursuits

webcams, and Wallace figured that being an early Internet pioneer would impress his tenure committee[5]. It didn't; nobody yet saw the Web as important, and Wallace watched as his grant applications were slapped down one by one. These petty[6] frustrations are commonplace for academics, but Wallace brooded over them more than most.

7 One day he checked into his webcam and noticed something strange: People were reacting to the robot eye in an oddly emotional way. It was designed so that remote viewers could type in commands like "tilt up" or "pan left," directing the eye to poke around Wallace's lab. Occasionally it would break down, and to Wallace's amusement, people would snap at it as if it were real: "You're stupid," they'd type. It gave him an idea: What if it could talk back?

8 Like all computer scientists, Wallace knew about a famous "chat-bot" experiment called Eliza. Back in 1966, MIT professor Joseph Weizenbaum had created Eliza as a "virtual therapist." It would take a user's statement and turn it around as a question, emulating a psychiatrist's often-maddening circularity. You: "I'm mad at my mother." Eliza: "Why are you mad at your mother?" Geeks[7]at MIT[8] spent hours talking to Eliza, enthralled even though they knew it wasn't real. But Eliza was quickly abandoned as a joke, even by its creator. It wasn't what scientists call "strong" AI[9]—able to learn on its own, or be "conscious". It could only parrot back lines Weizenbaum had fed it.

9 But Wallace was drawn to Eliza's simplicity. . . He decided to create an updated version of Eliza, and imbue[10] it with his own personality—something that could fire back witty repartee[11] when users became irritable. As Wallace's work progressed, though, his mental illness grew worse, making him both depressed and occasionally grandiose. . . .

10 Doctors told him he had bipolar disorder[12], but Wallace resisted the diagnosis. After all, didn't every computer scientist cycle through 72-hour sprees of creativity and then crash? "I was in denial myself," he says now. "I'm a successful professor, making $100,000 a year! I'm not one of those mental patients! I'll just check in with my therapist once a week, take pills—I'll be fine."

11 His supervisors disagreed. In April 1995, NYU told him his contract wouldn't be renewed.

12 ALICE came to life on November 23, 1995. That fall, Wallace had relocated to Lehigh College in Pennsylvania, hired again for his expertise in robotics. He installed

---

[5] tenure committee: group that decides whether a professor should receive tenure, the permanent status that protects a professor from being easily fired or dismissed
[6] petty: insignificant; not important
[7] geek: a person with an overly intellectual or technical orientation; similar to a nerd
[8] MIT: Massachusetts Institute of Technology
[9] AI: artificial intelligence
[10] imbue: to fill with a particular quality
[11] repartee: conversation
[12] bipolar disorder: a class of mental disorders that includes manic-depression

his chat program on a web server that happened to be named "Alice," and the name stuck. Then he sat back to watch, wondering what people would say to it.

13 Numbingly boring things, as it turns out. Users would inevitably ply ALICE with the same few questions: "Where do you live," "What is your name," or "What do you look like." Wallace began analyzing the chats, and realized that almost every statement began with one of 2,000 words. The ALICE chats were obeying something language theorists call "Zipf's Law"—a discovery from the 1930s, which found that a very small number of words comprise the bulk of what we say.

14 Wallace took Zipf's Law a step further. He began theorizing that only a few thousand statements made up the bulk of conversation—the everyday, commonplace chitchat that humans engage in at work, at the water cooler and in online discussion groups. ALICE was his proof. He'd already given it enough responses to deal with a few hundred of the most common conversational gambits, like "hello," "goodbye," and "what's your name?" If he kept on chipping away at it every day, teaching ALICE a new response every time he saw it baffled by a question, he would eventually cover all the common utterances, and even many unusual ones. Wallace calculated the magic number was about 40,000 responses. Once ALICE had that many preprogrammed statements, it—or "she," as he'd begun to call the program fondly—would be able to respond to 95 percent of what people were saying to her.

15 In essence, Wallace hit upon a theory that makes educated, intelligent people squirm. Maybe conversation simply isn't that complicated. Maybe we just say the same few thousand things to one another, over and over and over again. If Wallace was right, then artificial intelligence didn't need to be particularly intelligent in order to be convincingly lifelike. A.I. researchers had been focused on self-learning "neural nets," or mapping out grammar in "natural language" programs, but Wallace argued that the reason they had never mastered human conversation wasn't because humans are so complex, but because they are so simple.

16 "The smarter people are, the more complex they think the human brain is," he says. "It's like anthropocentrism, but on an intellectual level. 'I have a great brain, therefore everybody else does—and a computer must too.'" Wallace laughs. "And unfortunately most people don't."

17 Yet part of what makes ALICE seem so human-like is her seemingly spontaneous, wry responses, the product of what Wallace estimates is "an 800-page novel" worth of ALICE dialogue. His skill is thus not merely as a programmer, but as the author of thousands of sharp one-liners for Alice. It is, as he puts it, "more like writing good literature, perhaps drama, than writing computer programs.". . . .

18 But as ALICE improved, Wallace declined. Two years after ALICE was born, in the spring of 1997, Wallace lost his job at Lehigh College—his last chance for an academic career. His psychiatrists had put him on a battery of drugs, and the side effects had crippled him. "The worst one I ever took was Rispronal, which basically

gave me speech aphasia. I just couldn't talk," he recalls. "And I was supposed to be lecturing!" A disgruntled student organized a petition against him and handed it to the dean. . . . Eventually the university ended his contract; unable to pay his bills, Wallace's credit cards were all cancelled. . . .

19 Destitute, Wallace had almost nothing left to do but write more ALICE dialogue, on laptops he begged and borrowed from friends. He moved to California to find work in artificial-intelligence companies, but couldn't hold a job for more than a few months; instead, he'd spend his days obsessively writing ALICE dialogue and watching, thrilled, as the robot grew ever more life-like. Visitors used to talk to ALICE for only three or four exchanges. Now the average conversation was 20 exchanges, and some users would chatter away for hours.

20 In January of 2000, [Wallace] decided to stress-test[13] ALICE by entering her into the annual "Loebner Competition." It is, in essence, the Olympics for the artificial-intelligence chatbot world: Computer scientists worldwide go head-to-head, putting their programs before a panel of judges who rank them based on how "life-like" they are. Since no other objective markers for realism exist in this field, the contest is well known and controversial within the tight circle of AI; winning programs are closely studied both by academics and corporate think tanks such as Sprint Labs.

21 Wallace flew to Dartmouth College, where the contest was being held as part of an academic conference. He set ALICE running, and was herded into a waiting room to chew his nails with the other contestants. They waited for four hours. "I was incredibly nervous," he says. "I kept on thinking about all these little things that could go wrong." The judges called the contestants into an amphitheater, and with a crowd of 75 artificial-intelligence academics in attendance, displayed the results on the wall. ALICE had won: It was, officially, the most human robot in the world.

22 When [Wallace] returned to California, he plunged with new fervent energy into ALICE. Over the next year, he quadrupled its knowledge base, teaching it 30,000 new responses. Geeks worldwide began eagerly emailing Wallace to offer help. . . . In October 2001, ALICE won the Loebner competition for the second time in a row; and this time, one judge actually ranked ALICE more realistic than a human.

23 At one point during the Monica Lewinksy hearings, a human asked ALICE her opinion: "Do you think President Clinton should be impeached?"[14] "It depends," ALICE responded, "on what you mean by 'think.'"

---

[13] stress test (used as a verb in the text): a very difficult test designed to show the weakness of a material or a system

[14] "Do you think President Clinton should be impeached?": This is a reference to the December 19, 1998, impeachment of President Clinton by the House of Representatives. He was charged with perjury and obstruction of justice in connection with Paula Jones' law suit and the Monica Lewinsky scandal. He was acquitted by the Senate on February 12, 1999.

24 One could scarcely have asked for a more Clintonian response. But it's also a puzzling question that ALICE's success itself raises: Is she intelligent? If so, how?

25 In 1950, the pioneering British mathematician Alan Turing grappled with this question in *Mind* magazine, where he first posed the "Turing Test"—the gold standard for artificial thought. "Can machines think?" he asked—and immediately noted that it all hinges, of course, on what you mean by "think." He posed a simple "imitation game" to resolve the question. Put a woman and a computer in one room, and an interrogator in another. The interrogator talks to both via a teletype machine, and his goal is to figure out which is the woman (such as asking about "the length of your hair," which Turing felt was a dead giveaway). If the machine fools the interrogator into believing it is human, the test is passed; it can be considered intelligent.

26 This is, on the surface, a curiously unambitious definition: It's all about faking it. The machine doesn't need to act like a creative human or smart human or witty human—it merely needs to appear not to be a robot. After all, many humans are dull and stupid conversationalists themselves. With this bit of intellectual ju-jitsu, Turing neatly dodged a more troubling question—how do our brains, and language itself, work?

27 Some in the artificial intelligence community are brutally dismissive of ALICE. For them, artificial intelligence is about capturing the actual functioning of the brain, down to the neurons and learning ability that humans have. Parroting, they argue, doesn't count. Marvin Minksy, a prominent AI pioneer and MIT Media Lab professor, e-mailed me to say that while he thinks ALICE is "a nice job," Wallace's idea of conversation—that we mostly fire preprogrammed statements back and forth—is "basically wrong. It's like explaining that a picture is an object made by applying paint to canvas and then putting it in a rectangular frame. The important part is the complexity of our networks of knowledge and processes." ALICE, according to Minsky, does not truly "know" anything about the world . . .

28 The . . . debate usually boils down to one major issue: creativity. ALICE could never come up with an original thought, say critics, and creativity is the key attribute of human intelligence. End of argument. Wallace, however, has a much bleaker view. He doesn't argue that ALICE's conversation is particularly creative—but he doesn't believe people are creative either, at least when it comes to conversation. "Considering the vast size of the set of things people could possibly say, that are grammatically correct or semantically meaningful," Wallace wrote in an essay on his web site, "the number of things people actually do say is surprisingly small." By this argument, if ALICE were merely given a massive enough set of responses, it too could appear creative, just as creative as a human appears.

29 In the end, Wallace's work raises questions that stand in stark contrast to his life. How could a creator of something as sublime as ALICE argue that creativity isn't a significant part of human thought? Wallace shrugs off the paradox. He hopes ALICE

chatbots eventually become so human-like that they can take over the more repetitive interactive jobs, doing the labor of travel agents and telephone operators. It would, he says, free up humans to cultivate the tiny "0.0001 per cent" of their brains that we use to generate new ideas.

30 "Ideally, computers and robots will take most of the work away from us, give us more time to develop that otherwise very tiny seed of spirit that we have in us," he tells me. "Most of the brain and mind are this big waxy candle, with on top of it this tiny little flame of consciousness, or soul, or whatever you want to call it. And it's like the candle is thirty miles across it and eighteen miles high, and the flame is still the size of a normal flame."

31 He pauses. "And some people's flame seems to have blown out entirely."

Critical Focus: Recognizing and Examining Assumptions

Assumptions are ideas or beliefs that are accepted without question and without recognition of any need for support. Two fundamental kinds of assumptions are: (1) assumptions that the reader brings to the text (sometimes called **background knowledge**), and (2) assumptions that the writer brings to the text.

While both types of assumptions are essential for critical reading, it is the second type that we will focus on here. All (non-fiction) writers rely on assumptions. Indeed, it would be impossible for a writer to communicate anything if readers required proof for every statement in a text. However, the credibility of a writer's ideas rests on the soundness of the underlying assumptions.

In order to evaluate the ideas in a text, the critical reader must recognize and examine the assumptions (both stated and unstated) upon which the ideas of a text rely. There is no magical formula for uncovering assumptions, but the following guidelines might help:

- Slow down, reread, and reflect on the reading. Does the writer make any claim that is not supported by evidence? If so, the claim itself is an assumption.
- Identify keywords and phrases and try to decide what they mean in the context of the reading. Are keywords defined precisely? Are they used in a way that is overly general or overly specific? The writer's assumption might be that the term is broader or narrower than you think it is.
- Examine examples that are used as evidence for a generalization, an idea or a concept. Do the examples include items that seem questionable? The writer's assumption might be that the generalization has wider applicability than you think it does.

For an example from the text, see the Critical Focus Application exercise that follows.

Critical Focus: Application

Examine these claims made in Reading 1. Identify any evidence given for the claims, and try to list at least one stated or unstated assumption related to each claim. Then discuss your answers with a partner or small group. The first one has been done for you as an example.

Example:

1. **Claim:** It [ALICE] is the best artificial-intelligence program on the planet. (Paragraph 4)

   Evidence: *Two times, ALICE won the Loebner Competition "the Olympics for the artificial-intelligence chatbot world" (Paragraphs 20–22).*

   Assumptions:
   - *The Loebner Competition included all possible AI programs in the world.*
   - *Chatbot programs are representative of all AI programs (as opposed, for example, to programs that play chess).*

2. **Claim:** "A few thousand statements make up the bulk of conversation—the everyday, common place chitchat that humans engage in at work, at the water cooler and in online discussion groups." (Paragraph 14)

   Evidence: _____

   _____

   Assumption: _____

   _____

3. **Claim:** By "teaching ALICE a new response every time he saw it baffled by a question, he would eventually cover all the common utterances. . . ." (Paragraph 14)

   Evidence: _____

   _____

   Assumption: _____

   _____

4. **Claim:** "Considering the vast size of the set of things people could possibly say, that are grammatically correct or semantically meaningful . . . the number of things people actually do say is surprisingly small." (Paragraph 28)

   Evidence: _____

   _____

   Assumption: _____

   _____

Getting at the Matter

Answer the questions in writing or discuss them with a partner.

1. Who is Richard Wallace? Who is ALICE?

2. To what does the term *chatbot r*efer?

3. What is the Turing Test? Explain how it works. What question is it supposed to resolve?

4. What connections, if any, do you see between the Turing Test and the Loebner Competition?

5. Why do you think the author chose to entitle this article "Approximating Life?"

Academic Vocabulary Focus

In the left-hand column are 15 words from the Academic Vocabulary List (AWL) that appear in the reading "Approximating Life." Match these words with an appropriate definition or synonym from the right-hand column. Use the paragraph numbers in parentheses to locate the word (or the form it takes) in the reading. Use a dictionary only if necessary.

Vocabulary

____ 1. *abandon* (Par. 8)
____ 2. *analyze* (Par. 13)
____ 3. *complex* (Par. 5, 15, 16, 27)
____ 4. *comprise* (Par. 13)
____ 5. *controversial* (Par. 20)
____ 6. *estimate* (Par. 17)
____ 7. *generate* (Par. 29)
____ 8. *interactive* (Par. 29)
____ 9. *issue* (Par. 28)
____ 10. *process* (Par. 27)
____ 11. *resolve* (Par. 25)
____ 12. *functioning* (Par. 27)
____ 13. *significant* (Par. 29)
____ 14. *theory* (Par. 15)
____ 15. *virtual* (Par. 8)

Definition/Synonym

a. producing strong disagreement
b. produce or create
c. involving communication or exchange between people (or between a person and a machine)
d. find a solution or answer
e. stop using or maintaining
f. idea or set of principles that explains something
g. the way in which something operates or works
h. include, contain
i. divide into parts in order to understand
j. appearing or operating as if real, even if not real
k. important
l. complicated; composed of many interrelated parts
m. a sequence of natural occurrences leading to a result
n. determine an approximate, not an exact, quantity
o. topic or subject, often involving disagreement

For Discussion

Respond to these questions in a short essay, or discuss them with a partner or small group.

1. Search the Internet and find the website that hosts ALICE, (or some other chat-bot) and then conduct your own Turing Test. Chat with ALICE, recording your questions and statements as well as ALICE's responses. Bring a written record of your chat with ALICE to class and share it with a group of classmates. Working together, evaluate ALICE's performance. Be prepared to report to the class some examples of ALICE's best responses as well as some examples of responses that you consider inadequate or nonsensical. What generalizations can you make about ALICE's ability to carry on an intelligent conversation?

2. Some of ALICE's critics argue that "creativity is the key attribute of human intelligence," Because ALICE is only programmed with a finite number of responses, she could never exhibit creativity, and therefore, she is not—and never could be—intelligent. Richard Wallace might argue that humans only appear to be more creative than ALICE. Do you agree more with Richard Wallace or with his critics? Explain.