Chapter 2
The *Eliza* Effect

Meeting *Eliza*

When I was a teenager—in the 1980s—my mother bought a personal computer. It was an impressive machine for the day, decked out with two floppy drives, a dot matrix printer, a Hayes modem, and a monochrome amber display.

At first I only used the machine for some minor programming experiments (in Basic and later Pascal), writing for school (in WordStar), and a few games. But that mysterious modem sat there. Probably intended to let my mother exchange data with the big Digital Equipment Corporation machines she had in her university lab, I knew modems could also be used for other things.

This was about a decade before the Internet began to make its way into homes like ours, and I had no interest in the manicured gardens of services like the Source or CompuServe. Rather than any long-distance journey, I wanted to use the modem to explore the local wilderness, to visit the unruly bulletin board system (BBS) scene sprouting in the dens and basements of my neighbors' homes.

While largely forgotten today, a BBS was the online destination of choice for 1980s teenagers.[1] Most were run by individuals out of their homes: computer enthusiasts with machines much more powerful than ours, hooked to one or more dedicated phone lines. A user like me could call into a BBS, read messages, leave messages, download and upload files, play text-based games, and (if

1. For more on BBS culture, I highly recommend Jason Scott's *BBS: The Documentary* (2005).

2. *Eliza* was not the first system to give audiences the impression of meaningful exchange with a computer. Matthew Kirschenbaum's *Mechanisms* (2008) offers the intriguing example of "Professor RAMAC"—a …

Notes continued at end of this chapter.

the owner of the BBS was at the computer, or if someone called in to one of the other phone lines) have real-time conversations, with total strangers, in text. In other words, the BBS wasn't just a file repository. It was a window into what has now become obvious: the incredible social potential of combining computers and networks, which has given us email, instant messaging, wikis, blogs, social networking web sites, and much more.

Given the glimpse of this potential, a BBS with multiple lines could feel a little lonely when no one else was on. But then one day I was over at the house of a childhood friend (we no longer went to the same school), and he showed me that on his computer, conversation was always waiting. He showed me a program he'd downloaded from a BBS. He introduced me to *Eliza*.

Eliza Today

Eliza—or more properly, *Eliza/Doctor*—is a ground breaking system created by computer science researcher Joseph Weizenbaum at MIT in the mid-1960s. In the two decades between when Weizenbaum created the system and I experienced it at my friend's house, it had become one of the world's most famous demonstrations of the potential of computing. First unveiled two years before HAL 9000's screen debut in *2001: A Space Odyssey,* it seemed that *Eliza* made it possible to have a real conversation with a computer.[2]

In the computer science literature, under the name *Eliza,* Weizenbaum's system is a contribution to the field of natural language processing. On the other hand when *Eliza* plays *Doctor* it is a well-known computer character, famous far beyond computer science, often also known by the name Eliza. And *Eliza* has a third common usage in the computer world as well: *the Eliza effect*. This has generally been a term used to describe the not-uncommon illusion that an interactive computer system is more "intelligent" (or substantially more complex and capable) than it actually is. One of my purposes in this chapter is to revisit the *Eliza* effect and give it a further nuance so that it names not only this initial illusion but also the authorial choice that comes with it: severely restricted interaction (on the one hand) or eventual breakdown that takes a form based on the actual underlying processes (on the other).

In the next chapter, with an examination of the *Eliza* effect as background, I will look at the options selected by today's authors of digital fictions—particularly for computer games. These generally put aside the *Eliza* effect in favor of systems that more clearly communicate their structures to audiences. Yet there are two problems with these that I will consider. Some of them employ processes that, while legible, tend toward a non-*Eliza* form of breakdown. Others, while avoiding breakdown, have low ambitions in their use of computational processes. More ambitious routes will be the focus of the remaining chapters of this book.

But for now I'll start with the illusion.

The *Eliza* Effect

I don't remember the specifics of my first conversation with *Eliza,* but here is *Eliza*'s most famous exchange (with the computer's responses in all capital letters):

Men are all alike.
IN WHAT WAY

They're always bugging us about something or other.
CAN YOU THINK OF A SPECIFIC EXAMPLE
Well, my boyfriend made me come here.
YOUR BOYFRIEND MADE YOU COME HERE
He says I'm depressed much of the time.
I AM SORRY TO HEAR YOU ARE DEPRESSED
It's true. I am unhappy.
DO YOU THINK COMING HERE WILL
HELP YOU NOT TO BE UNHAPPY
I need some help, that much seems certain.
WHAT WOULD IT MEAN TO YOU
IF YOU GOT SOME HELP
Perhaps I could learn to get along with my mother.
TELL ME MORE ABOUT YOUR FAMILY
My mother takes care of me.
WHO ELSE IN YOUR FAMILY TAKES CARE OF YOU
My father.
YOUR FATHER
You are like my father in some ways.
WHAT RESEMBLANCE DO YOU SEE
You are not very aggressive but I think
you don't want me to notice that.
WHAT MAKES YOU THINK I AM
NOT VERY AGGRESSIVE
You don't argue with me.
WHY DO YOU THINK I DON'T ARGUE WITH YOU
You are afraid of me.
DOES IT PLEASE YOU TO BELIEVE
I AM AFRAID OF YOU
My father is afraid of everybody.
WHAT ELSE COMES TO MIND WHEN
YOU THINK OF YOUR FATHER
Bullies.

> DOES THAT HAVE ANYTHING TO DO WITH THE FACT
> THAT YOUR BOYFRIEND MADE YOU COME HERE
> (Weizenbaum 1966, 37)[3]

As you might imagine, I was impressed. And I certainly wasn't the first. Weizenbaum's employers at MIT had seized on *Eliza* as an accessible demonstration of the potential of computing, exciting to those without the specialized knowledge needed to appreciate much of the ongoing research. The computer world agreed. In fact, when the first demonstrations of the ARPANET (the predecessor of the Internet) were carried out, some of the terminals ran *Eliza* from a remote location (Leonard 1997, 52). Given the excitement produced by *Eliza,* astronomer and science popularizer Carl Sagan suggested that the future might hold "a network of computer psychotherapeutic terminals, something like arrays of large telephone booths" (Weizenbaum 1976, 5).

3. Erik Loyer, in reimplementing the *Eliza* system for a project on which we were collaborating for the digital media journal *Vectors,* noted that this conversation appears to be edited. Unless the *Eliza* system had undocumented …

Notes continued at end of this chapter.

Originally *Eliza* ran on a computer less powerful than the one in my mobile phone, though at the time it was one of the most advanced at MIT. Rather than being hooked up to a monitor, keyboard, and mouse—or HAL's microphone, speaker, and camera—*Eliza* was experienced through something like a typewriter, allowing the computer to type to the user and the user to type back. People could type anything they wanted to *Eliza,* and the software would respond in different ways depending on the currently active script.

This last word, "script," is important. There is nothing magical about *Eliza*—it is simply a bundle of data and processes, and pretty simple processes at that. Each time that *Eliza* runs, it uses a particular script to guide its behavior. The example conversation given was created using *Eliza*'s most famous script, *Doctor,* which causes the software to parody the conversational patterns of a nondirective therapist during an initial visit. All of this is described in Weizenbaum's 1966 article in *Communications of the ACM*.[4]

How *Eliza* Works

A session with *Eliza* can begin with a greeting. Weizenbaum's *Doctor* script starts with: "How do you do. Please tell me your problem."[5] After this point *Eliza* will not take the initiative again—only respond.

Each time an audience member types something, *Eliza* examines it, looking for words that have entries in the currently active script. Some of the words are *keywords*, which I will discuss further below. Some words are marked for simple substitution. For example, when *Eliza* runs the *Doctor* script, these substitutions switch all the first-person pronouns for second-person pronouns ("I" becomes "you") and vice versa ("yourself" becomes "myself"). "Well, you are very helpful," for instance, would become "Well, I are very helpful." A word can be both substituted and used as a keyword.

Periods and commas are treated as delimiters. If a period or comma is encountered, *Eliza* checks to see if a keyword has already been found. If one has, then everything that the audience member typed after the delimiter is discarded. If no keyword has yet been found, everything before the delimiter is discarded. For example, "Well, I are very helpful" would become "I are very helpful."

Each keyword has a priority level or rank. When the first keyword is found in a text, it is added to a "keystack." Each time another keyword is found, the rank of the

4. This article is not only a good, clear source of explanation for *Eliza*'s processes (and the most-cited publication about *Eliza* in the computer science literature). It also served as the basis for many homegrown versions of *Eliza* created at …

5. While the original script text is in all capitals, I am regularizing it here. Also, this description focuses on the core processes at work in *Eliza;* to describe them all would make this section as long as Weizenbaum's paper.

Notes continued at end of this chapter.

new keyword is compared with that of the highest-rank keyword yet found. If the new word has a higher rank it is added to the top of the stack; otherwise it is added to the bottom of the stack. The result, at the end of scanning a text for keywords, is that the highest-rank keyword is at the top of the stack.

After keyword scanning, the next step is to find a "decomposition rule" that matches the postsubstitution version of what the audience member typed (minus any parts discarded because of commas or periods). Decomposition rules are associated with keywords, so this search begins by popping the top keyword off the stack. (I'll soon discuss what happens if there is nothing on the stack.) *Eliza* tries out the decomposition rules associated with that keyword, looking for one with an appropriate pattern of words in combination with the wild card "0" (which is equivalent to "any number of words, including zero"). For example, if the keyword is "you," the first decomposition rule is:

```
(0 I remind you of 0)
```

This wouldn't match the sample phrase from above, "I are very helpful." The first rule would only match a postsubstitution phrase like, "With that last statement I remind you that I are a machine," or (because 0 can represent zero words), "I remind you of the happier days of your childhood." The next decomposition rule for the "you" keyword, though, would match the example phrase used earlier:

```
(0 I are 0)
```

Once a match is found, the decomposition rule does what it is named for: it decomposes the phrase into parts. This decomposition is based on the number of words and the number of 0s in the rule. For instance, "I are very helpful" would become:

(1) *empty* (2) I (3) are (4) very helpful.

All that remains, at this point, is to decide what to say back to the audience. Each decomposition rule can have a set of reassembly rules associated with it. An example reassembly rule (for the decomposition "0 I are 0") is:

(What makes you think I am 4)

This would result in the response, "What makes you think I am very helpful."[6] Once they were used during any run of *Eliza*, the reassembly rules were marked so that all the reassemblies associated with a particular decomposition would be cycled through before any were repeated.

This is the bulk of *Eliza*'s approach. There are a few refinements, such as techniques that allow scripts to share rules between keywords (e.g., mother and father). But from this the dominant operational logic at work in *Eliza* should be clear: transformation. Each statement by an *Eliza* script is the result of a multistep transformation of the most recent audience statement.

Many of these transformations include a verbatim portion of the statement being transformed, as outlined above. Others are more indirect transformations, as when "Perhaps I could learn to get along with my mother" is translated into "Tell me more about your family." The real transformation difficulty comes, however, when no keywords—a central aspect of *Eliza*'s transformation logic—are found in the audience's most recent text. For example, this is the situation when *Eliza/Doctor* asks, "What else comes to mind when you think of your father," and receives the response, "Bullies."

One set of possible transformations, in these cases, can be defined in the script based on a special keyword

6. Question marks were not used because they were interpreted as "line delete" on the computer system that Weizenbaum was using.

called "none." In *Doctor* these are content-free phrases such as "Please go on" and "I see." Yet some of the most surprising moments with *Eliza* come from another special structure called "memory." This structure works in combination with a particular keyword. In the case of *Doctor* it works with the keyword "my" and the decomposition "0 your 0." When this keyword is selected, a response is generated in the usual way. In addition, another reassembly may be created and stored in a queue for later use when no keyword is found. In *Doctor* these reassemblies are "Lets discuss further why your 3," "Earlier you said your 3," and "Does this have anything to do with the fact that your 3"—the last of these we saw above, with 3 being "boyfriend made you come here." In other words, the surprising punch line in the above conversation was nothing but a fallback, a default move that *Eliza* can make when the most recent response has no identifiable keywords.

Here we can see what moved Weizenbaum to write that "once a particular program is unmasked, once its inner workings are explained in language sufficiently plain to induce understanding, its magic crumbles away; it stands revealed as a mere collection of procedures, each quite comprehensible" (1966, 36). But when we interact with a piece of software, we don't necessarily achieve such understanding, we don't get a clear picture of how it actually operates internally. And Dr. Sagan and I were not alone regarding *Eliza/Doctor*; its audiences often developed mistaken ideas about its internals. They assumed that since the surface appearance of an interaction with the program could resemble something like a coherent dialogue, internally the software must be complex. Some thought it must be

7. This is the version of the story reported in Janet Murray's *Hamlet on the Holodeck*. In addition to this version, many clearly apocryphal versions circulate, including one in which Weizenbaum is a participant in the events. This version itself cannot be a …

Notes continued at end of this chapter.

something close to the fictional HAL: a computer program intelligent enough to understand and produce arbitrary human language. This happened so frequently, and was so striking, that in computer science circles people began to use a specific term for this kind of misunderstanding, one mentioned earlier in this chapter: the *Eliza* effect.

Why *Eliza* Worked

The stories of *Eliza/Doctor*'s effectiveness are many. Weizenbaum himself tells us that "my secretary, who had watched me work on the program for many months and therefore surely knew it to be merely a computer program, started conversing with it. After only a few interchanges with it, she asked me to leave the room" (1976, 6). There is also the oft-repeated tale, which exists in many versions, of the boss thinking he is communicating with a subordinate while actually conversing with *Eliza/Doctor*:

> *Vice president*: Do you think I can use the computer this morning?
> *Teletype*: WHY DO YOU ASK?
> *Vice president*: I might be able to make some additional telcomp (computer system) sales.
> *Teletype*: WHY AREN'T YOU SURE?
> *Vice president*: My prospective customers have never seen a demonstration of the system.
> *Teletype*: DOES THIS MEAN VERY MUCH TO YOU?
> *Vice president*: Of course it does.
> *Teletype*: YOU SEEM VERY POSITIVE ABOUT THAT.
> *Vice president*: This is ridiculous.
> *Teletype*: WHAT DO YOU MEAN THIS IS RIDICULOUS?
> (Murray 1997, 7)[7]

As a high school student, I found that *Eliza/Doctor* created an effective initial illusion for the same reasons it worked for the iconic "emotional" secretary and "clueless" boss evoked by the above stories. First, all of us were accustomed to text-only computing and to having conversations with other people within that environment. Second, *Eliza/Doctor* makes a remarkably good match between process and data. The situation of the initial visit to the therapist, the clever writing in the reassemblies and nonresponses, and the well-chosen keywords do the most possible to leverage the simple linguistic tricks available via *Eliza*'s transformation processes. (In fact, there were other *Eliza* scripts created besides *Doctor,* but none of them became nearly as well known or widespread.) Third, for myself, the secretary, and the boss, this was one of our first experiences with computer characters. But all three of these reasons are only *Eliza/Doctor*'s specific nuances on a much more general phenomenon: when a system is presented as intelligent and appears to exhibit intelligent behavior, people have a disturbingly strong tendency to regard it as such.[8]

This phenomenon derailed Weizenbaum's career. He came to focus his work on the conceptual mismatch that gives the *Eliza* effect its name and specifically on how it could "induce powerful delusional thinking in quite normal people" (1976, 7). Weizenbaum wrote a book dedicated to demonstrating that the internals of computers aren't magical, and that we do ourselves a disservice when we assume that human beings are so mechanical that we could or should have our intelligence matched by computational machines. In a sense, he moved from being a computer scientist to being one of the first knowledgeable critics to interrogate the cultures

8. Mark J. Nelson, in the blog-based peer review of this book, urged me to clarify the fact that this isn't only true of computer systems—as demonstrated by an example later in this chapter: Harold Garfinkel's yes/no therapy experiment.

of computing and artificial intelligence.

Following Weizenbaum, a number of other authors saw the *Eliza* effect as important to address in understanding our relationship with computers and our culture more generally. A decade after Weizenbaum's book, Lucy Suchman published *Plans and Situated Actions* (1987), in which she sees *Eliza/Doctor* as an iconic example in human-computer interaction of the broad phenomenon of treating systems as intelligent based on limited evidence. Specifically, she discusses what ethnomethodologist Harold Garfinkel (citing Karl Mannheim) has called the documentary method of interpretation.

Suchman presents one of Garfinkel's experiments as a demonstration of the idea that people tend to "take appearances as evidence for, or the document of, an ascribed underlying reality, while taking the reality so ascribed as a resource for the interpretation of the appearance" (23). In this experiment student subjects were introduced to a new kind of therapy in which they asked yes/no questions about their personal problems. These were answered by "counselors" who were not visible to the subjects. Unbeknownst to the subjects, the counselors answered each question randomly.

After the experiment, the students were found to have constructed stories that made sense of each string of answers as a coherent exchange and set of advice. This happened even when, as would almost inevitably happen in such a circumstance, the answers given were self-contradictory. The apparent contradictions were explained away by the experimental subjects as revised views based on further information, evidence of a deeper agenda on the part of the counselor or something else that fit with the frame of the therapeutic situation.

Yet another decade later, Janet Murray revisited *Eliza/Doctor* in *Hamlet on the Holodeck* (1997). Here she makes a crucial turn for my purposes, seeing the project in terms of media. She views *Eliza/Doctor* as a dramatic character, Weizenbaum as a kind of playwright, and the source of the *Eliza* effect in such cases as "the human propensity to suspend disbelief in the presence of a persuasive dramatic presence" (224). Murray makes visible what wasn't clear in previous accounts: that people enjoy *Eliza*. Initial interactions are often quite fun and engaging, and create the sense of open-ended possibilities with an interesting character.

But that first impression doesn't necessarily last.

Revisiting the *Eliza* Effect

Being a teenager, after my initial interaction with *Eliza/Doctor,* I did the natural thing: I started to play with it. (Software, of course, doesn't need to be a game to inspire play.) I don't remember the exact exchange, but I know that I soon discovered the reason Murray doesn't go far into talking about ongoing interaction with *Eliza/Doctor,* just as Suchman and Weizenbaum don't.[9] If they did, it would complicate their discussions. This is because, while the initial experience of *Eliza/Doctor* can create the surface impression of an incredibly complex internal system, sustained interaction with the system, the verbal back-and-forth, invites play—and linguistic play with *Eliza/Doctor* quickly begins to destroy the illusion. In other words, precisely the open-ended textual interaction that helped foster the illusion of internal complexity and intelligence enables play that draws attention to the system's rote

9. To be fair, at the time of Weizenbaum's initial observations, almost no one could experience ongoing interaction with *Eliza/Doctor,* due to the limited availability of computing resources. As Weizenbaum notes, "Since the subject cannot …

Notes continued at end of this chapter.

10. Turkle comments that "some people embark on an all-out effort to 'psych out' the program, to understand its structure in order to trick it and expose it as a 'mere machine.' Many more do the opposite. . . . They didn't ask questions that they knew …

11. Suchman argues that Garfinkel's experiment lends support to Weizenbaum's view that the feeling of intelligence in conversations with *Eliza/Doctor* comes from the work of the audience. Further, she contends that …

Notes continued at end of this chapter.

simplicity, its distance from human interaction. Before long, many exchanges begin to look like this:

You are going to repeat what I say
in the form of a question
WHAT MAKES YOU THINK I AM GOING TO REPEAT
WHAT YOU SAY IN THE FORM OF A QUESTION

Sherry Turkle, in *The Second Self*, notes this phenomenon, writing, "As one becomes experienced with the ways of *Eliza,* one can direct one's remarks either to 'help' the program make seemingly pertinent responses or to provoke nonsense" (1984, 39). Turkle's book is based on interviews and observations—and she found that some of her subjects worked to keep the illusion going, while others sought to maximize *Eliza*'s nonsense.[10] Both strategies, though, depended on coming to understand something of the internal processes at work and shaping surface interaction accordingly. Even working to maintain the illusion required a type of seeing past it, something that those who discuss the *Eliza* effect rarely acknowledge.

The *Eliza* Breakdown

From my point of view, what Turkle describes points toward a further lesson of Garfinkel's yes/no therapy experiment. For Suchman, this experiment demonstrates the importance of ethnomethodology and the documentary hypothesis for understanding *Eliza/Doctor* and human-computer interaction.[11] And certainly it is essential to understand that *Eliza/Doctor* succeeds, to the extent that it does, because it plays on the interpretative expectations brought to each interaction by audience members. But for my purposes here, Garfinkel's experiment also serves to illustrate something rather different: the *Eliza* effect can be shielded from breakdown by severely restricting interaction. The experiment allowed the subjects to maintain the illusion that something much more complex was going on inside the system (a human considering her problems seriously and answering questions thoughtfully, rather than random yes/no answers) because the scope of possible responses was so limited. If it had been expanded only slightly—say, to random choice between the responses available in a "magic eight ball"—almost any period of sustained interaction would have shattered the illusion through too many inappropriate responses.

When breakdown in the *Eliza* effect occurs, its shape is often determined by that of the underlying processes. If the output is of a legible form, the audience can then begin to develop a model of the processes. This is what Turkle observes in those interacting with *Eliza/Doctor*: from the shape of the breakdown they begin to understand something of the processes of the system—and then employ that knowledge to help maintain or further compromise the illusion.

In this context, it is interesting to note that most systems of control that are meant to appear intelligent have extremely restricted methods of interaction. In some cases the reasons for this are quite obvious. If the public were allowed playful interaction with software that identifies possible targets for financial surveillance, the shape of the underlying system would become relatively apparent, making it possible to "game" the system. At the same time, this restricted interaction also serves to maintain the *Eliza* effect for software that is not nearly as intelligent as the public has been asked to believe.

Further, within a rather different community, this

12. Except for that limited number of fictions that might want to explore one of these effects.

choice—between severely restricted interaction and the boom/bust of illusion followed by breakdown—presents no good options to those with an interest in creating digital fictions.[12] So while some have argued that it is best to capitalize on the *Eliza* effect, depending on temporary illusion and the willful suspension of disbelief to carry the day, most digital fiction authors employ a different approach: exposing important elements of the structures of their processes to the audience from the outset. This allows for interaction that matches the process employed, and avoids the *Eliza* illusion and breakdown. Still, as I will discuss next, the most common of these approaches suffer from limitations of their own.

Finally, I should mention that some authors—such as Jeremy Douglass (2007)—assert that breakdown can be an interesting mode for digital fictions. And certainly breakdowns can be fascinating. On a linguistic level, for example, we're attracted to study every form of breakdown from occasional slips of the tongue to hemmorage-induced aphasia.

What breakdowns can do—in the case of *Eliza/Doctor,* linguistic slips, and neurological problems alike—is give us some insight into the shape of the underlying system processes. This fact points to the reason why I still talk with people online (even if I no longer dial in to a BBS to do so) but I no longer play with *Eliza* in my spare time: a system prone to breakdown is only as interesting as the shape of the processes that the breakdowns partially reveal. And as shown earlier in this chapter, the *Eliza* system processes are mostly a relatively uninteresting set of substitutions. We can do better.

Notes

1. For more on BBS culture, I highly recommend Jason Scott's *BBS: The Documentary* (2005).

2. *Eliza* was not the first system to give audiences the impression of meaningful exchange with a computer. Matthew Kirschenbaum's *Mechanisms* (2008) offers the intriguing example of "Professor RAMAC"—a four-ton IBM machine anthropomorphized at the 1958 World's Fair—which used a then-revolutionary hard drive to provide question-answering access to a collection of historical facts.

3. Erik Loyer, in reimplementing the *Eliza* system for a project on which we were collaborating for the digital media journal *Vectors,* noted that this conversation appears to be edited. Unless the *Eliza* system had undocumented features, its response to "You are not very aggressive but I think you don't want me to notice that" would have been "WHAT MAKES YOU THINK I AM NOT VERY AGGRESSIVE BUT YOU THINK I DON'T WANT YOU TO NOTICE THAT" rather than the snappier "WHAT MAKES YOU THINK I AM NOT VERY AGGRESSIVE."

4. This article is not only a good, clear source of explanation for *Eliza*'s processes (and the most-cited publication about *Eliza* in the computer science literature). It also served as the basis for many homegrown versions of *Eliza* created at computing centers across the country in the years that followed. This was still decades away from when researchers and hobbyists could easily transport software to one another over computer networks—so paper publications such as *Communications of the ACM* (for researchers) and *Byte* (for hobbyists) often included all the information needed (including, in the case of Weizenbaum's article, the complete *Doctor* script) for reimplementing the software locally.

5. While the original script text is in all capitals, I am regularizing it here. Also, this description focuses on the core processes at work in *Eliza;* to describe them all would make this section as long as Weizenbaum's paper.

6. Question marks were not used because they were interpreted as "line delete" on the computer system that Weizenbaum was using.

7. This is the version of the story reported in Janet Murray's *Hamlet on the Holodeck.* In addition to this version, many clearly apocryphal versions circulate, including one in which Weizenbaum is a participant in the events. This version itself cannot be a verbatim conversation with *Eliza,* at least not as the system existed at the time of Weizenbaum's 1966 paper. That paper's Doctor script, for example, contains the responses "YOU AREN'T SURE" and "WHY THE UNCERTAIN TONE"—but not "WHY AREN'T YOU SURE." Beyond differences in wording, it is also worth remembering the previous note: the system used by Weizenbaum did not support question marks, which appear throughout this transcript.

8. Mark J. Nelson, in the blog-based peer review of this book, urged me to clarify the fact that this isn't only true of computer systems—as demonstrated by an example later in this chapter: Harold Garfinkel's yes/no therapy experiment.

9. To be fair, at the time of Weizenbaum's initial observations, almost no one could experience ongoing interaction with *Eliza/Doctor,* due to the limited availability of computing resources. As Weizenbaum notes,

"Since the subject cannot probe the true limits of *Eliza*'s capabilities (he has, after all, only a limited time to play with it, and it is constantly getting new material from him), he cannot help but attribute more power to it than it actually has" (1976, 191).

10. Turkle comments that "some people embark on an all-out effort to 'psych out' the program, to understand its structure in order to trick it and expose it as a 'mere machine.' Many more do the opposite. . . . They didn't ask questions that they knew would 'confuse' the program, that would make it 'talk nonsense'" (40). Turkle attributes this to a desire to "maintain the illusion that *Eliza* was able to respond to them." It is also entirely in line with Murray's interpretation of *Eliza* as a media experience, however, with the audience shaping their interaction to help maintain the willful suspension of disbelief.

11. Suchman argues that Garfinkel's experiment lends support to Weizenbaum's view that the feeling of intelligence in conversations with *Eliza/Doctor* comes from the work of the audience. Further, she contends that the strongly situated understandings of the students (they interpreted the random series of yes/no answers based on assumed context) is a challenge not only to the strong structure-oriented assumptions of the social sciences but also those of cognitive science.

12. Except for that limited number of fictions that might want to explore one of these effects.