

The Software Toolworks

Walt Bilofsky, Prop.

14478 GLORIETTA DRIVE
SHERMAN OAKS, CALIFORNIA 91423

TELEPHONE
(213) 986-4885

ELIZA
Conversation Writing Language
with
DOCTOR "Psychiatrist" Script
Version 1.0 - September 1981
Walt Bilofsky

Copyright (c) 1981 Walter Bilofsky. Sale of this software conveys a license for its use on a single computer owned or operated by the purchaser. Copying this software or documentation by any means whatsoever for any other purpose is expressly prohibited. ELIZA research paper and original DOCTOR script copyright (c) 1966 Association for Computing Machinery, Inc., reprinted by permission.

1. WHAT IS ELIZA?

ELIZA is a program which allows you to converse with your computer. It is a particularly good way to show your machine off to friends, since most people find it fascinating to "talk" with ELIZA.

ELIZA holds a conversation with you. You type a sentence to it, and it replies. A prepared script is used to determine its responses. The two scripts supplied on your disk are variations of DOCTOR, which makes ELIZA behave like a non-directive psychiatrist.

IMPORTANT: ELIZA with DOCTOR is not a "computer psychiatrist". It is not intended to provide psychiatric treatment or medical benefit of any kind.

ELIZA was originally written in 1965 by Dr. Joseph Weizenbaum at MIT. The program soon became famous for its ability to make people think they were talking not with a computer but with a person.

This version of ELIZA is a faithful recreation of Dr. Weizenbaum's original script interpreter, plus a few carefully selected new features. You can change ELIZA's behavior by modifying the scripts provided, or even writing completely new ones. So ELIZA is not only a computer game, but also a programming language.

Section 2 below tells how to run ELIZA. The remainder of this document describes the ELIZA script writing language.

To give you the opportunity to understand ELIZA in depth, we have also included a reprint of Weizenbaum's original paper on ELIZA, from the Communications of the Association for Computing Machinery, January 1966 (Vol. 9 No. 1).

As you converse with ELIZA, and become familiar with the scripts, you will gain insight into just how much the program really understands, and what methods it uses to give the appearance of understanding. You may decide that ELIZA is not so much "computer intelligence" as simply computer cleverness. Regardless, you and your friends will find it an absorbing and amazing program.

2. HOW TO RUN ELIZA.

Before doing anything else, you should copy all the files from the ELIZA disk onto another disk, and store the distribution disk in a safe place.

The ELIZA disk contains the following files:

ELIZA.COM (or .ABS)	The ELIZA program.
DOCTOR.SCR	The current DOCTOR script.
ODOCTOR.SCR	The original DOCTOR script.

To run ELIZA, the ELIZA program and DOCTOR.SCR files should be on the same disk. (ODOCTOR is the older, original DOCTOR script. You can use it if you wish but it probably will not carry on as good a conversation.)

ELIZA can run under CP/M and Heath/Zenith HDOS, and requires at least 40K of RAM. If ELIZA complains of insufficient memory, and you are using CP/M, be sure your system is configured for the full memory capacity of your computer.

If you are using CP/M, set the default disk to the disk containing the ELIZA files. For example, if they are on disk B:, type the CP/M command

```
B:
```

Then execute ELIZA by typing the command

```
ELIZA
```

If you are using HDOS, and the ELIZA files are on SY0:, all you need to do is give the ELIZA command. If the files are on another disk, such as SY1:, type

```
SY1:ELIZA
```

ELIZA will load from disk, and then greet you. You should talk to ELIZA using complete English sentences. Press the RETURN key

when you have finished typing your response. Since ELIZA replies in capitals you may find the dialogue easier to follow if you type in lower case.

To complete your conversation politely, type the word GOODBYE. (You may also leave abruptly by typing ctrl-C: that is, hold down the CTRL shift key and press C.)

If you want to run a script file other than DOCTOR.SCR, use the command

```
ELIZA scriptname
```

If no extension is specified for the script file, ELIZA assumes the extension .SCR.

If you want to record your conversation, when you first run ELIZA use the command

```
ELIZA >filename
```

where "filename" is any legal file name. On CP/M, "filename" can also be the device name LST: or PUN:. Or you can use the ctrl-P feature of CP/M to record the conversation on the printer. On HDOS, "filename" can be a disk file name, or a device such as LP:.

Sometimes you will be in the middle of a conversation and begin to wish you had saved it. ELIZA remembers as much of the conversation as can fit in the computer memory, so you can still save much or all of it. To start saving while in the middle of an ELIZA conversation, type the one word sentence

```
SAVE
```

ELIZA will ask you for a file name, on which it will save as much of the conversation as it has stored, and everything from now on. If your computer has a lot of memory, or if you have not been talking long, you will probably be able to save the entire conversation. Otherwise, only the most recent part of the conversation will appear on the file.

3. THE SECRETS OF ELIZA

"It is said that to explain is to explain away," Weizenbaum says of ELIZA. Although its behavior is remarkable, its effects are produced by an inherently simple mechanism. As long as ELIZA remains a mystery, you will marvel at its cleverness. Once we explain its secrets to you, however, you may feel it is just a gimmick.

Before reading further, therefore, we suggest you play with ELIZA as much as you like. Once you have studied its workings, you may

possibly find it less fun - or at least, a different kind of fun.

By showing you how it works, we provide the opportunity for you to learn how easy it can be for a programmer to make computers do marvelous things. In particular, you will be able to change ELIZA's conversational habits, and even write completely new scripts on whatever subject you like.

Once you decide to proceed, we recommend you first read the introduction to Weizenbaum's paper, which is included at the end of this manual. Although we have tried to explain the script language simply and clearly in the following sections, it contains some details that might be confusing. Weizenbaum's paper, although written mainly for computer scientists, presents the same material from a broader viewpoint. So you may find the paper helpful after you have read the rest of this manual.

4. THE ELIZA PROGRAMMING LANGUAGE.

The behavior of the ELIZA program is based on a simple mechanism. Each sentence you type is processed, using a prepared script, according to the following four steps:

- Step 1. Look at the sentence and pick out any keywords specified in the script.
- Step 2. Replace each keyword by a replacement word, if the script provides one.
- Step 3. See if the sentence matches any of the patterns specified in the script.
- Step 4. If a pattern match is found, prepare a reply using one of the rewriting rules associated in the script with that pattern. If no match is found, use one of the replies for "no match".

For example, suppose you tell ELIZA

YOU ARE REALLY HELPING ME.

The keywords in this sentence are YOU and ME. The script gives the replacements YOU = I and ME = YOU, so the sentence is rewritten

I ARE REALLY HELPING YOU.

One purpose of replacements is to "turn around" the sentences you type to ELIZA, so it can repeat parts of them back to you. Notice, though, that not all keywords have replacements. So at this stage of the process, the rewritten sentence does not always make sense, although the eventual reply probably will.

Under the keyword YOU, the script contains the pattern

(0 I ARE 0)

"0" matches zero or more words, so this pattern matches your sentence, with the first 0 matching nothing (zero words), and the second 0 matching "REALLY HELPING YOU".

One of the rewriting rules for this pattern is

(WHAT MAKES YOU THINK I AM 4 ?)

The "4" means to insert in that place the 4th element of the pattern: the word(s) matching the second 0. So ELIZA would generate the response

WHAT MAKES YOU THINK I AM REALLY HELPING YOU?

This four step process, although simple in concept, can be confusing when you first try to understand it in detail. ELIZA itself can explain the process to you. If you invoke ELIZA using the command

ELIZA -D

the program will print all sorts of information as it loads and executes. In particular, it will print the keywords, patterns and reassembly rules it uses as it generates each response. So you may wish to first study the description below of ELIZA's script language, and then, armed with a listing of the DOCTOR script, run ELIZA -D, type sentences to it, and try to follow along as it explains what it is doing. Remember that your operating system allows you to use ctrl-S and ctrl-Q to suspend printing so that you can read the output before it shoots off the screen.

5. LISTS.

An ELIZA script is made up of lists. A list is simply a set of words, numbers, and punctuation characters, enclosed in parentheses. For instance,

(THIS LIST CONTAINS 5 THINGS)

But a list can also contain another list:

(THIS CONTAINS 4 (AND THIS IS ONE OF THEM))

When writing lists, you must be extremely careful to get all the parentheses correct. The ELIZA -D command can help find parenthesis errors; see "Debugging Hints" in Section 9.

6. SCRIPTS AND RULE LISTS.

An ELIZA script is a text file containing lots of rule lists. You can prepare a script using any text editor program, such as PIE, ED or EDIT. You may use upper and lower case in the script; ELIZA translates everything to upper case.

The DOCTOR script provided with ELIZA contains many examples of rule lists. You may want to study it while reading the following sections.

First, we are going to show the most general form of rule list. We will use braces {} to enclose items that may be either included or left out. Ellipses (...) indicate that you can include more than one of the preceding item. Then most rule lists are of the form

```
(keyword {= replacement} {priority}
      {DLIST (/ listname ...)}
      { (pattern reassembly ...) ... } )
```

An example of a rule list containing all the above elements is

```
(YOU = I 3 DLIST (/ PERSON)
  ((0 I 1 YOU 0) (OH, YOU 3 ME?))
  ((0) (WHAT ABOUT ME?) (TELL ME ABOUT YOU)))
```

This is complicated, but we will go through each part of the rule list and explain it. As we do, you should keep in mind the four steps (Section 4) which ELIZA uses to generate its response to a user's sentence.

The first thing in a rule list is the keyword. This is the way the script defines the keywords used in Step 1 of the ELIZA conversation process. In the above example, the keyword is YOU.

If the next thing is the character "=" and a replacement word, then every time the keyword is seen in a typed input, it is immediately replaced by the replacement word (Step 2). In the example, the replacement is I. Every YOU typed by the user is replaced by I. If there is no "= replacement", the keyword is left as it was typed.

The next thing in the rule list is an optional priority for the keyword. This is a number between 0 and 127. If no priority is specified, 0 is assumed.

The priority is used in deciding which keyword in a sentence to consider first. As ELIZA scans an input sentence, it places each keyword it sees on a stack. If the keyword has a higher priority than any of the ones already seen, it is placed on top of the stack. Otherwise, it is placed on the bottom.

When the sentence has been read, ELIZA considers the keywords in the order they appear on the stack. Thus, the highest priority

keyword in a sentence is used first. If two keywords both have the highest priority, the one on the left is used first.

If the next thing in the rule list is the word `DLIST`, then the keyword is placed on one or more word lists. In the example, `YOU` is placed on the word list called `PERSON`. Word lists provide a way to define a category of words for use in a pattern, as explained in the next section.

7. PATTERNS.

The last thing in a rule list is one or more transformation rules. Each transformation rule consists of a pattern followed by one or more reassembly rules. Once the highest priority keyword has been selected from the user's sentence, an attempt is made to apply each transformation rule for that keyword, in turn. This is Step 3 of the response generation process.

The first transformation rule in our example is

```
((0 I 1 YOU 0) (OH, I 3 ?))
```

This consists of the pattern

```
(0 I 1 YOU 0)
```

and the single reassembly rule

```
(OH, I 3 ?)
```

The items in the pattern are used, left to right, to try to match the user's sentence. Here are the things you can use in a pattern:

- 0 The digit 0 matches any sequence of words, including the "sequence" containing no words. If 0 is not the last item in the pattern, it will match the smallest sequence before a match for the next item.
- number Any other number matches any sequence containing exactly that many words. Thus, 3 will match any three words.
- word Any word will match just that word. Remember that keyword replacement (Step 2) is performed before matching (Step 3), so to match a keyword that is replaced, you must use its replacement.

(* word ...)

A list beginning with "*" will match any of the words in the list. For example, in the pattern

```
(0 YOU (* WANT NEED) 0)
```

the third item in the pattern will match either WANT or NEED.

(/list) This will match any element on the word list named LIST. For example, in the DOCTOR script, the keywords MOTHER, FATHER, SISTER, etc., are all placed on the FAMILY word list using the DLIST feature. Then the pattern element (/FAMILY) will match any of these words.

Thus, the rule

(0 I 1 YOU 0)

will match any sentence containing the phrase "I (one word) YOU", with any number of words (including none) before and after. Notice that, if the keyword replacements YOU = I and ME = YOU were in this script, the user would actually have to type "YOU (one word) ME" in order to match this pattern after replacement.

If the first element in a pattern is not 0, that pattern will only match starting at the beginning of a sentence. If the last element is not 0, the end of the pattern must come at the end of a sentence in order to match. Thus,

(I 0 YOU)

will match any sentence starting with I and ending with YOU. (Note that ELIZA treats commas just like periods: what comes after the comma is a new sentence.)

The pattern (0) is used often. It will match any sentence.

8. REASSEMBLY RULES.

Recall that a transformation rule consists of a pattern followed by one or more reassembly rules. If the pattern is found to match the user's input sentence, ELIZA applies the first reassembly rule in order to generate a new sentence, which it prints as its reply. The reassembly rule is then moved to the end of the list. So if there are several reassembly rules for one pattern, they are used in rotation as the conversation progresses.

Most reassembly rules are simply lists of words and numbers. To create the reply, ELIZA replaces each number in the rule with the part of the input sentence that matched the corresponding pattern element. For example, if the input was

SOMETIMES YOU BOTHER ME

and the transformation rule was our old friend

((0 I 1 YOU 0) (OH, I 3 YOU?))

then, after rewriting the input sentence to

SOMETIMES I BOTHER YOU

the pattern gives a match, with the "1" matching BOTHER. In this case, the reassembly rule would replace the 3 by the match, BOTHER, for the third element in the pattern, "1". Then ELIZA would reply

OH, I BOTHER YOU?

9. SPECIAL FEATURES.

The special transformation rule

(= keyword)

is like a GOTO in other programming languages. It tells ELIZA to immediately begin trying to apply the transformations associated with that keyword. For instance, suppose there is a rulelist for the keyword WAS, and you want to treat WERE the same as WAS. You could use the following rulelist for WERE:

(WERE = WAS (=WAS))

This replaces WERE by WAS in the input, and then starts checking the patterns for WAS.

(=keyword) can also be used as a reassembly rule. Note that it can be one of several reassembly rules following a pattern, or the last of several transformation rules for a keyword. (It should be the last because you'll never get past it to the others.)

NEWKEY.

There are several other special reassembly rules. The rule

(NEWKEY)

tells ELIZA to give up on the current keyword and try finding another one in the input sentence. It often appears as part of the last transformation rule for a keyword:

((0) (NEWKEY))

PRE.

The reassembly rule

```
(PRE (reassembly) (=keyword))
```

is used to reassemble the input sentence, and then go back and try pattern matching again. The input sentence is reassembled according to the reassembly rule, and control is then transferred to the rule list for the specified keyword, for attempted matching on the transformed input.

For example, DOCTOR contains the transformation

```
(I'M = YOU'RE ((0 YOU'RE 0) (PRE (YOU ARE 3) (=I))))
```

If the input sentence is

```
SOMETIMES I'M VERY HAPPY.
```

it is rewritten to

```
SOMETIMES YOU'RE VERY HAPPY.
```

If I'M is the highest priority keyword in the sentence, the pattern (0 YOU'RE 0) will always match. The PRE rule rewrites the input as

```
YOU ARE VERY HAPPY
```

and the rule list for the keyword I is then used to try to match this sentence.

NONE.

ELIZA breaks up the user's typed input into one or more sentences, separated by any punctuation such as period, comma, dash, and so on. Each of these sentences in turn is searched for keywords, and a pattern match sought.

The reserved keyword

```
NONE
```

is invoked whenever ELIZA is unable to find a keyword in the input, or to match any pattern for the keywords that do appear. Every script must contain a rulelist for NONE.

EXIT.

If a reassembly rule begins with =EXIT, it has a special effect. The remainder of the rule is printed as the output, without any substitution, and ELIZA terminates its execution. (This feature was not in Weizenbaum's original program.)

MEMORY.

A script may contain one or more rulelists of the form

(MEMORY keyword transformations ...)

These rulelists look just like any other, except for the addition of the word MEMORY. The keyword must have a rulelist of its own elsewhere in the script.

When the keyword is selected for transformation, before invoking the keyword's own rulelist, an attempt is made to match the patterns in the MEMORY rulelist. If a match is found, the resulting output is not printed on the terminal, but is stored away.

When five responses have passed since the last time a MEMORY response was either stored or recalled, and no keyword match is found, instead of using a NONE reply ELIZA will print the oldest stored MEMORY response.

DEBUGGING HELP.

As mentioned in Section 4, running ELIZA with the command

```
ELIZA -D
```

produces debugging printout of various sorts. You can also use

```
ELIZA -D >file
```

to record the output on a disk file. Besides explaining the pattern match process, ELIZA provides several other kinds of output that may be useful in writing or modifying scripts.

The input rules are printed as they are read from the script file. This can be helpful in finding missing or extra parentheses.

After loading the script, ELIZA prints out the amount of memory remaining. This memory is used to store the conversation for the SAVE feature, and its size is also a clue to how much larger your script can become before it can no longer fit in memory.

APPENDIX I

Following is Dr. Weizenbaum's research paper on ELIZA, reprinted by permission from the January 1966 Communications of the Association for Computing Machinery.

ELIZA—A Computer Program For the Study of Natural Language Communication Between Man And Machine

JOSEPH WEIZENBAUM

Massachusetts Institute of Technology, Cambridge, Mass.*

ELIZA is a program operating within the MAC time-sharing system at MIT which makes certain kinds of natural language conversation between man and computer possible. Input sentences are analyzed on the basis of decomposition rules which are triggered by key words appearing in the input text. Responses are generated by reassembly rules associated with selected decomposition rules. The fundamental technical problems with which ELIZA is concerned are: (1) the identification of key words, (2) the discovery of minimal context, (3) the choice of appropriate transformations, (4) generation of responses in the absence of key words, and (5) the provision of an editing capability for ELIZA "scripts". A discussion of some psychological issues relevant to the ELIZA approach as well as of future developments concludes the paper.

Introduction

It is said that to explain is to explain away. This maxim is nowhere so well fulfilled as in the area of computer programming, especially in what is called heuristic programming and artificial intelligence. For in those realms machines are made to behave in wondrous ways, often sufficient to dazzle even the most experienced observer. But 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. The observer says to himself "I could have written that". With that thought he moves the program in question from the shelf marked "intelligent", to that reserved for curios, fit to be discussed only with people less enlightened than he.

Work reported herein was supported (in part) by Project MAC, an MIT research program sponsored by the Advanced Research Projects Agency, Department of Defense, under Office of Naval Research Contract Number Nonr-4102(01).

* Department of Electrical Engineering.

The object of this paper is to cause just such a re-evaluation of the program about to be "explained". Few programs ever needed it more.

ELIZA Program

ELIZA is a program which makes natural language conversation with a computer possible. Its present implementation is on the MAC time-sharing system at MIT. It is written in MAD-SLIP [4] for the IBM 7094. Its name was chosen to emphasize that it may be incrementally improved by its users, since its language abilities may be continually improved by a "teacher". Like the Eliza of Pygmalion fame, it can be made to appear even more civilized, the relation of appearance to reality, however, remaining in the domain of the playwright.

For the present purpose it is sufficient to characterize the MAC system as one which permits an individual to operate a full scale computer from a remotely located typewriter. The individual operator has the illusion that he is the sole user of the computer complex, while in fact others may be "time-sharing" the system with him. What is important here is that the computer can read messages typed on the typewriter and respond by writing on the same instrument. The time between the computer's receipt of a message and the appearance of its response is a function of the program controlling the dialogue and of such MAC system parameters as the number of users currently corresponding with the system. These latter parameters generally contribute so little to the overall response time that conversational interaction with the computer need never involve truly intolerable delays.

When in conversation with ELIZA, the user types in some statement or set of statements in natural language using normal punctuation and sentence structures. Only the question mark may not be used, because it is interpreted as a line delete character by the MAC system. The user's statement is terminated by a double carriage return which serves to turn control over to ELIZA. ELIZA then analyzes the user's statement and generates some response which it types out. Control is then again with the user.

A typical conversation is the following:

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

The capitalized lines are the machine responses.

The gross procedure of the program is quite simple; the text is read and inspected for the presence of a *keyword*. If such a word is found, the sentence is transformed according to a *rule* associated with the keyword, if not a content-free remark or, under certain conditions, an earlier transformation is retrieved. The text so computed or retrieved is then printed out.

In detail, of course, the procedure sketched above is considerably more complex. Keywords, for example, may have a RANK or precedence number. The procedure is sensitive to such numbers in that it will abandon a keyword already found in the left-to-right scan of the text in favor of one having a higher rank. Also, the procedure recognizes a comma or a period as a delimiter. Whenever either one is encountered and a keyword has already been found, all subsequent text is deleted from the input message. If no key had yet been found the phrase or sentence to the left of the delimiter (as well as the delimiter itself) is deleted. As a result, only single phrases or sentences are ever transformed.

Keywords and their associated transformation¹ rules constitute the SCRIPT for a particular class of conversation. An important property of ELIZA is that a script is data; i.e., it is not part of the program itself. Hence, ELIZA is not restricted to a particular set of recognition patterns or responses, indeed not even to any specific language. ELIZA scripts exist (at this writing) in Welsh and German as well as in English.

The fundamental technical problems with which ELIZA must be preoccupied are the following:

- (1) The identification of the "most important" keyword

¹ The word "transformation" is used in its generic sense rather than that given it by Harris and Chomsky in linguistic contexts.

occurring in the input message.

(2) The identification of some minimal context within which the chosen keyword appears; e.g., if the keyword is "you", is it followed by the word "are" (in which case an assertion is probably being made).

(3) The choice of an appropriate transformation rule and, of course, the making of the transformation itself.

(4) The provision of mechanism that will permit ELIZA to respond "intelligently" when the input text contained no keywords.

(5) The provision of machinery that facilitates editing, particularly extension, of the script on the script writing level.

There are, of course, the usual constraints dictated by the need to be economical in the use of computer time and storage space.

The central issue is clearly one of text manipulation, and at the heart of that issue is the concept of the *transformation rule* which has been said to be associated with certain keywords. The mechanisms subsumed under the slogan "transformation rule" are a number of SLIP functions which serve to (1) decompose a data string according to certain criteria, hence to test the string as to whether it satisfies these criteria or not, and (2) reassemble a decomposed string according to certain assembly specifications.

While this is not the place to discuss these functions in all their detail (or even to reveal their full power and generality), it is important to the understanding of the operation of ELIZA to describe them in *some* detail.

Consider the sentence "I am very unhappy these days". Suppose a foreigner with only a limited knowledge of English but with a very good ear heard that sentence spoken but understood only the first two words "I am". Wishing to appear interested, perhaps even sympathetic, he may reply "How long have you been very unhappy these days?" What he must have done is to apply a kind of template to the original sentence, one part of which matched the two words "I am" and the remainder isolated the words "very unhappy these days". He must also have a reassembly kit specifically associated with that template, one that specifies that any sentence of the form "I am BLAH" can be transformed to "How long have you been BLAH", independently of the meaning of BLAH. A somewhat more complicated example is given by the sentence "It seems that you hate me". Here the foreigner understands only the words "you" and "me"; i.e., he applies a template that decomposes the sentence into the four parts:

- (1) It seems that (2) you (3) hate (4) me

of which only the second and fourth parts are understood. The reassembly rule might then be "What makes you think I hate you?"; i.e., it might throw away the first component, translate the two known words ("you" to "I" and "me" to "you") and tack on a stock phrase (What makes you think) to the front of the reconstruction.

A formal notation in which to represent the decomposition template is:

(0 YOU 0 ME)

and the reassembly rule

(WHAT MAKES YOU THINK I 3 YOU).

The "0" in the decomposition rule stands for "an indefinite number of words" (analogous to the indefinite dollar sign of COMIT) [6] while the "3" in the reassembly rule indicates that the third component of the subject decomposition is to be inserted in its place. The decomposition rule

(0 YOU 1 ME)

would have worked just as well in this specific example. A nonzero integer "n" appearing in a decomposition rule indicates that the component in question should consist of exactly "n" words. However, of the two rules shown, only the first would have *matched* the sentence, "It seems you hate and love me," the second failing because there is more than one word between "you" and "me".

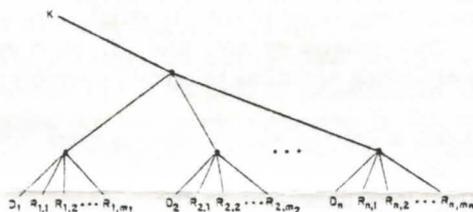


FIG. 1. Keyword and rule list structure

In ELIZA the question of which decomposition rules to apply to an input text is of course a crucial one. The input sentence might have been, for example, "It seems that you hate," in which case the decomposition rule (0 YOU 0 ME) would have failed in that the word "ME" would not have been found at all, let alone in its assigned place. Some other decomposition rule would then have to be tried and, failing that, still another until a match could be made or a total failure reported. ELIZA must therefore have a mechanism to sharply delimit the set of decomposition rules which are potentially applicable to a currently active input sentence. This is the keyword mechanism.

An input sentence is scanned from left to right. Each word is looked up in a dictionary of keywords. If a word is identified as a keyword, then (apart from the issue of precedence of keywords) only decomposition rules containing that keyword need to be tried. The trial sequence can even be partially ordered. For example, the decomposition rule (0 YOU 0) associated with the keyword "YOU" (and decomposing an input sentence into (1) all words in front of "YOU", (2) the word "YOU", and (3) all words following "YOU") should be the last one tried since it is bound to succeed.

Two problems now arise. One stems from the fact that

almost none of the words in any given sentence are represented in the keyword dictionary. The other is that of "associating" both decomposition and reassembly rules with keywords. The first is serious in that the determination that a word is not in a dictionary may well require more computation (i.e., time) than the location of a word which is represented. The attack on both problems begins by placing both a keyword and its associated rules on a list. The basic format of a typical key list is the following:

$$\begin{aligned} &(K ((D_1) (R_{1,1}) (R_{1,2}) \cdots (R_{1,m_1})) \\ &((D_2) (R_{2,1}) (R_{2,2}) \cdots (R_{2,m_2})) \\ &\vdots \\ &((D_n) (R_{n,1}) (R_{n,2}) \cdots (R_{n,m_n})) \end{aligned}$$

where K is the keyword, D_i the i th decomposition rule associated with K and $R_{i,j}$ the j th reassembly rule associated with the i th decomposition rule.

A common pictorial representation of such a structure is the tree diagram shown in Figure 1. The top level of this structure contains the keyword followed by the names of lists; each one of which is again a list structure beginning with a decomposition rule and followed by reassembly rules. Since list structures of this type have no predetermined dimensionality limitations, any number of decomposition rules may be associated with a given keyword and any number of reassembly rules with any specific decomposition rule. SLIP is rich in functions that sequence over structures of this type efficiently. Hence programming problems are minimized.

An ELIZA script consists mainly of a set of list structures of the type shown. The actual keyword dictionary is constructed when such a script is first read into the hitherto empty program. The basic structural component of the keyword dictionary is a vector KEY of (currently) 128 contiguous computer words. As a particular key list structure is read the keyword K at its top is randomized (hashed) by a procedure that produces (currently) a 7 bit integer "i". The word "always", for example, yields the integer 14. $KEY(i)$, i.e., the i th word of the vector KEY, is then examined to determine whether it contains a list name. If it does not, then an empty list is created, its name placed in $KEY(i)$, and the key list structure in question placed on that list. If $KEY(i)$ already contains a list name, then the name of the key list structure is placed on the bottom of the list named in $KEY(i)$. The largest dictionary so far attempted contains about 50 keywords. No list named in any of the words of the KEY vector contains more than two key list structures.

Every word encountered in the scan of an input text, i.e., during the actual operations of ELIZA, is randomized by the same hashing algorithm as was originally applied to the incoming keywords, hence yields an integer which points to the only possible list structure which could potentially contain that word as a keyword. Even then, only the tops of any key list structures that may be found there need be interrogated to determine whether or not a keyword has been found. By virtue of the various list

sequencing operations that SLIP makes available, the actual identification of a keyword leaves as its principal product a pointer to the list of decomposition (and hence reassembly) rules associated with the identified keyword. One result of this strategy is that often less time is required to discover that a given word is not in the keyword dictionary than to locate it if it is there. However, the location of a keyword yields pointers to all information associated with that word.

Some conversational protocols require that certain transformations be made on certain words of the input text independently of any contextual considerations. The first conversation displayed in this paper, for example, requires that first person pronouns be exchanged for second person pronouns and vice versa throughout the input text. There may be further transformations but these minimal substitutions are unconditional. Simple substitution rules ought not to be elevated to the level of transformations, nor should the words involved be forced to carry with them all the structure required for the fully complex case. Furthermore, unconditional substitutions of single words for single words can be accomplished during the text scan itself, not as a transformation of the entire text subsequent to scanning. To facilitate the realization of these desiderata, any word in the key dictionary, i.e., at the top of a key list structure, may be followed by an equal sign followed by whatever word is to be its substitute. Transformation rules may, but need not, follow. If none do follow such a substitution rule, then the substitution is made on the fly, i.e., during text scanning, but the word in question is not identified as a keyword for subsequent purposes. Of course, a word may be both substituted for and be a keyword as well. An example of a simple substitution is

(YOURSELF = MYSELF).

Neither "yourself" nor "myself" are keywords in the particular script from which this example was chosen.

The fact that keywords can have ranks or precedences has already been mentioned. The need of a ranking mechanism may be established by an example. Suppose an input sentence is "I know everybody laughed at me." A script may tag the word "I" as well as the word "everybody" as a keyword. Without differential ranking, "I" occurring first would determine the transformation to be applied. A typical response might be "You say you know everybody laughed at you." But the important message in the input sentence begins with the word "everybody". It is very often true that when a person speaks in terms of universals such as "everybody", "always" and "nobody" he is really referring to some quite specific event or person. By giving "everybody" a higher rank than "I", the response "Who in particular are you thinking of" may be generated.

The specific mechanism employed in ranking is that the rank of every keyword encountered (absence of rank implies rank equals 0) is compared with the rank of the highest ranked keyword already seen. If the rank of the

new word is higher than that of any previously encountered word, the pointer to the transformation rules associated with the new word is placed on top of a list called the keystack, otherwise it is placed on the bottom of the keystack. When the text scan terminates, the keystack has at its top a pointer associated with the highest ranked keyword encountered in the scan. The remaining pointers in the stack may not be monotonically ordered with respect to the ranks of the words from which they were derived, but they are nearly so—in any event they are in a useful and interesting order. Figure 2 is a simplified

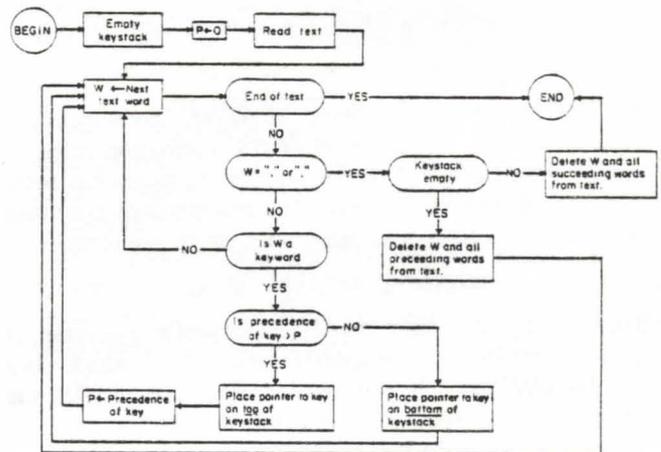


FIG. 2. Basic flow diagram of keyword detection

flow diagram of keyword detection. The rank of a keyword must, of course, also be associated with the keyword. Therefore it must appear on the keyword list structure. It may be found, if at all, just in front of the list of transformation rules associated with the keyword. As an example consider the word "MY" in a particular script. Its keyword list may be as follows:

(MY = YOUR 5 (transformation rules)).

Such a list would mean that whenever the word "MY" is encountered in any text, it would be replaced by the word "YOUR". Its rank would be 5.

Upon completion of a given text scan, the keystack is either empty or contains pointers derived from the keywords found in the text. Each of such pointers is actually a sequence reader—a SLIP mechanism which facilitates scanning of lists—pointing into its particular key list in such a way that one sequencing operation to the right (SEQLR) will sequence it to the first set of transformation rules associated with its keyword, i.e., to the list

$((D_1) (R_{1,1}) (R_{1,2}) \dots (R_1, R_m))$.

The top of *that* list, of course, is a list which serves a decomposition rule for the subject text. The top of the keystack contains the first pointer to be activated.

The decomposition rule D_1 associated with the keyword K , i.e., $\{(D_1), K\}$, is now tried. It may fail however. For example, suppose the input text was:

You are very helpful.

The keyword, say, is "you", and $\{(D_1), \text{you}\}$ is

(0 I remind you of 0).

(Recall that the "you" in the original sentence has already been replaced by "I" in the text now analyzed.) This decomposition rule obviously fails to match the input sentence. Should $\{(D_1), K\}$ fail to find a match, then $\{(D_2), K\}$ is tried. Should that too fail, $\{(D_3), K\}$ is attempted, and so on. Of course, the set of transformation rules can be guaranteed to terminate with a decomposition rule which must match. The decomposition rule

(0 K 0)

will match any text in which the word K appears while

(0)

will match any text whatever. However, there are other ways to leave a particular set of transformation rules, as will be shown below. For the present, suppose that some particular decomposition rule (D_i) has matched the input text. (D_i) , of course, was found on a list of the form

$((D_i)(R_{i,1})(R_{i,2}) \cdots (R_{i,m_i}))$.

Sequencing the reader which is presently pointing at (D_i) will retrieve the reassembly rule $(R_{i,1})$ which may then be applied to the decomposed input text to yield the output message.

Consider again the input text

You are very helpful . . .

in which "you" is the only key word. The sentence is transformed during scanning to

I are very helpful

$\{(D_1), \text{you}\}$ is "(0 I remind your of 0)" and fails to match as already discussed. However, $\{(D_2), \text{you}\}$ is "(0 I are 0)" and obviously matches the text, decomposing it into the constituents

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

$\{(R_{2,1}), \text{you}\}$ is

(What makes you think I am 4)

Hence it produces the output text

What makes you think I am very helpful.

Having produced it, the integer 1 is put in front of $(R_{2,1})$ so that the transformation rule list in question now appears as

$((D_2)1(R_{2,1})(R_{2,2}) \cdots (R_{2,m_2}))$.

Next time $\{(D_2), K\}$ matches an input text, the reassembly rule $(R_{2,2})$ will be applied and the integer 2 will replace the 1. After (R_{2,m_2}) has been exercised, $(R_{2,1})$ will again be invoked. Thus, after the system has been in use for a time, every decomposition rule which has matched some input text has associated with it an integer which corresponds to the last reassembly rule used in connection with

that decomposition rule. This mechanism insures that the complete set of reassembly rules associated with a given decomposition rule is cycled through before any repetitions occur.

The system described so far is essentially one which selects a decomposition rule for the highest ranking keyword found in an input text, attempts to match that text according to that decomposition rule and, failing to make a match, selects the next reassembly rule associated with the matching decomposition rule and applies it to generate an output text. It is, in other words, a system which, for the highest ranking keyword of a text, selects a specific decomposition and reassembly rule to be used in forming the output message.

Were the system to remain that simple, then keywords that required identical sets of transformation rules would each require that a copy of these transformation rules be associated with them. This would be logically sound but would complicate the task of script writing and would also make unnecessary storage demands. There are therefore special types of decomposition and assembly rules characterized by the appearance of "=" at the top of the rule list. The word following the equal sign indicates which new set of transformation rules is to be applied. For example, the keyword "what" may have associated with it a transformation rule set of the form

((0) (Why do you ask) (Is that an important question) . . .)

which would apply equally well to the keywords "how" and "when". The entire keyword list for "how" may therefore be

(How (=What))

The keywords "how", "what" and "when" may thus be made to form an equivalence class with respect to the transformation rules which are to apply to them.

In the above example the rule "(=what)" is in the place of a decomposition rule, although it causes no decomposition of the relevant text. It may also appear, however, in the place of a reassembly rule. For example, the keyword "am" may have among others the following transformation rule set associated with it:

((0 are you 0) (Do you believe you are 4) . . . (=what) . . .)

(It is here assumed that "are" has been substituted for "am" and "you" for "I" in the initial text scan.) Then, the input text

Am I sick

would elicit either

Do you believe you are sick

or

Why do you ask

depending on how many times the general form had already occurred.

Under still other conditions it may be desirable to

perform a preliminary transformation on the input text before subjecting it to the decompositions and reassemblies which finally yield the output text. For example, the keyword "you're" should lead to the transformation rules associated with "you" but should first be replaced by a word pair. The dictionary entry for "you're" is therefore:

```
(you're = I'm ((0 I'm 0) (PRE (I AM 3) (=YOU))))
```

which has the following effect:

(1) Wherever "you're" is found in the input text, it is replaced by "I'm".

(2) If "you're" is actually selected as the regnant keyword, then the input text is decomposed into three constituent parts, namely, all text in front of the first occurrence of "I'm", the word "I'm" itself, and all text following the first occurrence of "I'm".

(3) The reassembly rule beginning with the code "PRE" is encountered and the decomposed text reassembled such that the words "I AM" appear in front of the third constituent determined by the earlier decomposition.

(4) Control is transferred, so to speak, to the transformation rules associated with the keyword "you", where further decompositions etc. are attempted.

It is to be noted that the set

```
(PRE (I AM 3) (=YOU))
```

is logically in the place of a reassembly rule and may therefore be one of many reassembly rules associated with the given decomposition.

Another form of reassembly rule is

```
(NEWKEY)
```

which serves the case in which attempts to match on the currently regnant keyword are to be given up and the entire decomposition and reassembly process is to start again on the basis of the keyword to be found in the keystack. Whenever this rule is invoked, the top of the keystack is "popped up" once, i.e., the new regnant keyword recovered and removed from the keystack, and the entire process reinitiated as if the initial text scan had just terminated. This mechanism makes it possible to, in effect, test on key phrases as opposed to single key words.

A serious problem which remains to be discussed is the reaction of the system in case no keywords remain to serve as transformation triggers. This can arise either in case the keystack is empty when NEWKEY is invoked or when the input text contained no keywords initially.

The simplest mechanism supplied is in the form of the special reserved keyword "NONE" which must be part of any script. The script writer must associate the universally matching decomposition rule (0) with it and follow this by as many content-free remarks in the form of transformation rules as he pleases. (Examples are: "Please go on", "That's very interesting" and "I see".)

There is, however, another mechanism which causes the system to respond more spectacularly in the absence of a key. The word "MEMORY" is another reserved pseudo-keyword. The key list structure associated with it differs

from the ordinary one in some respects. An example illuminates this point.

Consider the following structure:

```
(MEMORY MY
 (0 YOUR 0 = LETS DISCUSS FURTHER WHY YOUR 3)
 (0 YOUR 0 = EARLIER YOU SAID YOUR 3)
 :)
```

The word "MY" (which must be an ordinary keyword as well) has been selected to serve a special function. Whenever it is the highest ranking keyword of a text one of the transformations on the MEMORY list is randomly selected, and a copy of the text is transformed accordingly. This transformation is stored on a first-in-first-out stack for later use. The ordinary processes already described are then carried out. When a text without keywords is encountered later and a certain counting mechanism is in a particular state and the stack in question is not empty, then the transformed text is printed out as the reply. It is, of course, also deleted from the stack of such transformations.

The current version of ELIZA requires that one keyword be associated with MEMORY and that exactly four transformations accompany that word in that context. (An application of a transformation rule of the form

```
(LEFT HAND SIDE = RIGHT HAND SIDE)
```

is equivalent to the successive application of the two forms

```
(LEFT HAND SIDE), (RIGHT HAND SIDE).)
```

Three more details will complete the formal description of the ELIZA program.

The transformation rule mechanism of SLIP is such that it permits tagging of words in a text and their subsequent recovery on the basis of one of their tags. The keyword "MOTHER" in ELIZA, for example, may be identified as a noun and as a member of the class "family" as follows:

```
(MOTHER DLIST (/NOUN FAMILY)).
```

Such tagging in no way interferes with other information (e.g., rank or transformation rules) which may be associated with the given tag word. A decomposition rule may contain a matching constituent of the form (/TAG1 TAG2 ...) which will match and isolate a word in the subject text having any one of the mentioned tags. If, for example, "MOTHER" is tagged as indicated and the input text

```
"CONSIDER MY AGED MOTHER AS WELL AS ME"
```

subjected to the decomposition rule

```
(0 YOUR 0 (/FAMILY) 0)
```

(remembering that "MY" has been replaced by "YOUR"), then the decomposition would be

```
(1) CONSIDER (2) YOUR (3) AGED (4) MOTHER
 (5) AS WELL AS ME.
```

Another flexibility inherent in the SLIP text manipulation mechanism underlying ELIZA is the ordering of matching criteria is permitted in decomposition rules. The above input text would have been decomposed

precisely as stated above by the decomposition rule:

(0 YOUR 0 (*FATHER MOTHER) 0)

which, by virtue of the presence of "*" in the sublist structure seen above, would have isolated either the word "FATHER" or "MOTHER" (in that order) in the input text, whichever occurred first after the first appearance of the word "YOUR".

Finally, the script writer must begin his script with a list, i.e., a message enclosed in parentheses, which contains the statement he wishes ELIZA to type when the system is first loaded. This list may be empty.

Editing of an ELIZA script is achieved via appeal to a contextual editing program (ED) which is part of the MAC library. This program is called whenever the input text to ELIZA consists of the single word "EDIT". ELIZA then puts itself in a so-called dormant state and presents the then stored script for editing. Detailed description of ED is out of place here. Suffice it to say that changes, additions and deletions of the script may be made with considerable efficiency and on the basis of entirely contextual cues, i.e., without resort to line numbers or any other artificial devices. When editing is completed, ED is given the command to FILE the revised script. The new script is then stored on the disk and read into ELIZA. ELIZA then types the word "START" to signal that the conversation may resume under control of the new script.

An important consequence of the editing facility built into ELIZA is that a given ELIZA script need not start out to be a large, full-blown scenario. On the contrary, it should begin as a quite modest set of keywords and transformation rules and permitted to be grown and molded as experience with it builds up. This appears to be the best way to use a truly interactive man-machine facility—i.e., not as a device for rapidly debugging a code representing a fully thought out solution to a problem, but rather as an aid for the exploration of problem solving strategies.

Discussion

At this writing, the only serious ELIZA scripts which exist are some which cause ELIZA to respond roughly as would certain psychotherapists (Rogerians). ELIZA performs best when its human correspondent is initially instructed to "talk" to it, via the typewriter of course, just as one would to a psychiatrist. This mode of conversation was chosen because the psychiatric interview is one of the few examples of categorized dyadic natural language communication in which one of the participating pair is free to assume the pose of knowing almost nothing of the real world. If, for example, one were to tell a psychiatrist "I went for a long boat ride" and he responded "Tell me about boats", one would not assume that he knew nothing about boats, but that he had some purpose in so directing the subsequent conversation. It is important to note that this assumption is one made by the speaker. Whether it is realistic or not is an altogether separate question. In any case, it has a crucial psychological utility

in that it serves the speaker to maintain his sense of being heard and understood. The speaker further defends his impression (which even in real life may be illusory) by attributing to his conversational partner all sorts of background knowledge, insights and reasoning ability. But again, these are the *speaker's* contribution to the conversation. They manifest themselves inferentially in the *interpretations* he makes of the offered responses. From the purely technical programming point of view then, the psychiatric interview form of an ELIZA script has the advantage that it eliminates the need of storing *explicit* information about the real world.

The human speaker will, as has been said, contribute much to clothe ELIZA'S responses in vestments of plausibility. But he will not defend his illusion (that he is being understood) against all odds. In human conversation a speaker will make certain (perhaps generous) assumptions about his conversational partner. As long as it remains possible to interpret the latter's responses consistently with those assumptions, the speaker's image of his partner remains unchanged, in particular, undamaged. Responses which are difficult to so interpret may well result in an enhancement of the image of the partner, in additional rationalizations which then make more complicated interpretations of his responses reasonable. When, however, such rationalizations become too massive and even self-contradictory, the entire image may crumble and be replaced by another ("He is not, after all, as smart as I thought he was"). When the conversational partner is a machine (the distinction between machine and program is here not useful) then the idea of *credibility* may well be substituted for that of *plausibility* in the above.

With ELIZA as the basic vehicle, experiments may be set up in which the subjects find it credible to believe that the responses which appear on his typewriter are generated by a human sitting at a similar instrument in another room. How must the script be written in order to maintain the credibility of this idea over a long period of time? How can the performance of ELIZA be systematically degraded in order to achieve controlled and predictable thresholds of credibility in the subject? What, in all this, is the role of the initial instruction to the subject? On the other hand, suppose the subject is told he is communicating with a machine. What is he led to believe about the machine as a result of his conversational experience with it? Some subjects have been very hard to convince that ELIZA (with its present script) is *not* human. This is a striking form of Turing's test. What experimental design would make it more nearly rigorous and airtight?

The whole issue of the credibility (to humans) of machine output demands investigation. Important decisions increasingly tend to be made in response to computer output. The ultimately responsible human interpreter of "What the machine says" is, not unlike the correspondent with ELIZA, constantly faced with the need to make credibility judgments. ELIZA shows, if nothing else, how easy it is to create and maintain the illusion of understanding, hence perhaps of judgment

deserving of credibility. A certain danger lurks there.

The idea that the present ELIZA script contains *no* information about the real world is not entirely true. For example, the transformation rules which cause the input

Everybody hates me

to be transformed to

Can you think of anyone in particular

and other such are based on quite specific hypotheses about the world. The whole script constitutes, in a loose way, a model of certain aspects of the world. The act of writing a script is a kind of programming act and has all the advantages of programming, most particularly that it clearly shows where the programmer's understanding and command of his subject leaves off.

A large part of whatever elegance may be credited to ELIZA lies in the fact that ELIZA maintains the illusion of understanding with so little machinery. But there are bounds on the extendability of ELIZA's "understanding" power, which are a function of the ELIZA program itself and not a function of any script it may be given. The crucial test of understanding, as every teacher should know, is not the subject's ability to continue a conversation, but to draw valid conclusions from what he is being told. In order for a computer program to be able to do that, it must at least have the capacity to store *selected* parts of its inputs. ELIZA throws away each of its inputs, except for those few transformed by means of the MEMORY machinery. Of course, the problem is more than one of storage. A great part of it is, in fact, subsumed under the word "selected" used just above. ELIZA in its use so far has had as one of its principal objectives the *concealment* of its lack of understanding. But to encourage its conversational partner to offer inputs from which it can select remedial information, it must *reveal* its misunderstanding. A switch of objectives from the concealment to the revelation of misunderstanding is seen as a precondition to making an ELIZA-like program the basis for an effective natural language man-machine communication system.

One goal for an augmented ELIZA program is thus a system which already has access to a store of information about some aspects of the real world and which, by means of conversational interaction with people, can reveal both what it knows, i.e., behave as an information retrieval system, and where its knowledge ends and needs to be augmented. Hopefully the augmentation of its knowledge will also be a direct consequence of its conversational experience. It is precisely the prospect that such a program will converse with many people and learn something from each of them, which leads to the hope that it will prove an interesting and even useful conversational partner.

One way to state a slightly different intermediate goal is to say that ELIZA should be given the power to slowly build a model of the subject conversing with it. If the subject mentions that he is not married, for example, and later speaks of his wife, then ELIZA should be able to

make the tentative inference that he is either a widower or divorced. Of course, he could simply be confused. In the long run, ELIZA should be able to build up a *belief structure* (to use Abelson's phrase) of the subject and on that basis detect the subject's rationalizations, contradictions, etc. Conversations with such an ELIZA would often turn into arguments. Important steps in the realization of these goals have already been taken. Most notable among these is Abelson's and Carroll's work on simulation of belief structures [1].

The script that has formed the basis for most of this discussion happens to be one with an overwhelmingly psychological orientation. The reason for this has already been discussed. There is a danger, however, that the example will run away with what it is supposed to illustrate. It is useful to remember that the ELIZA program itself is merely a translating processor in the technical programming sense. Gorn [2] in a paper on language systems says:

Given a language which already possesses semantic content, then a translating processor, even if it operates only syntactically, generates corresponding expressions of another language to which we can attribute as "meanings" (possibly multiple—the translator may not be one to one) the "semantic intents" of the generating source expressions; whether we find the result consistent or useful or both is, of course, another problem. It is quite possible that by this method the same syntactic object language can be usefully assigned multiple meanings for each expression . . .

It is striking to note how well his words fit ELIZA. The "given language" is English as is the "other language", expressions of which are generated. In principle, the given language could as well be the kind of English in which "word problems" in algebra are given to high school students and the other language, a machine code allowing a particular computer to "solve" the stated problems. (See Bobrow's program STUDENT [3].)

The intent of the above remarks is to further rob ELIZA of the aura of magic to which its application to psychological subject matter has to some extent contributed. Seen in the coldest possible light, ELIZA is a translating processor in Gorn's sense; however, it is one which has been especially constructed to work well with natural language text.

REFERENCES

1. ABELSON, R. P., AND CARROLL, J. D. Computer simulation of individual belief systems. *Amer. Behav. Sci.* 9 (May 1965), 24-30.
2. GORN, S. Semiotic relationships in ambiguously stratified language systems. Paper presented at Int. Colloq. Algebraic Linguistics and Automatic Theory, Hebrew U. of Jerusalem, Aug. 1964.
3. BOBROW, D. G. Natural language input for a computer problem solving system. Doctoral thesis, Math. Dept., MIT. Cambridge, Mass., 1964.
4. WEIZENBAUM, J. Symmetric list processor. *Comm. ACM* 6, (Sept. 1963), 524-544.
5. ROGERS, C. *Client Centered Therapy: Current Practice, Implications and Theory*. Houghton Mifflin, Boston, 1951.
6. YNGVE, J. *COMIT Programming Manual*. MIT Press, Cambridge, Mass., 1961.

APPENDIX. An ELIZA Script

```
(HOW DO YOU DO. PLEASE TELL ME YOUR PROBLEM)
START
(SORRY ((0) (PLEASE DON'T APOLOGIZE)
(APOLOGIES ARE NOT NECESSARY) (WHAT FEELINGS
DO YOU HAVE WHEN YOU APOLOGIZE) (I'VE TOLD YOU
THAT APOLOGIES ARE NOT REQUIRED)))
(DONT = DON'T)
(CANT = CAN'T)
(WONT = WON'T)
(REMEMBER 5
((0 YOU REMEMBER 0) (DO YOU OFTEN THINK OF 4)
(DOES THINKING OF 4 BRING ANYTHING ELSE TO MIND)
(WHAT ELSE DO YOU REMEMBER)
(WHY DO YOU REMEMBER 4 JUST NOW)
(WHAT IN THE PRESENT SITUATION REMINDS YOU OF 4)
(WHAT IS THE CONNECTION BETWEEN ME AND 4))
((0 DO I REMEMBER 0) (DID YOU THINK I WOULD FORGET 5)
(WHY DO YOU THINK I SHOULD RECALL 5 NOW)
(WHAT ABOUT 5) (=WHAT) (YOU MENTIONED 5))
((0) (NEWKEY)))
(IF 3 ((0 IF 0) (DO YOU THINK ITS LIKELY THAT 3) (DO YOU WISH THAT 3)
(WHAT DO YOU THINK ABOUT 3) (REALLY, 2 3)))
(WHAT DO YOU THINK ABOUT 3) (REALLY, 2 3)))
(DREAMT 4 ((0 YOU DREAMT 0)
(REALLY, 4) (HAVE YOU EVER FANTASIED 4 WHILE YOU WERE AWAKE)
(HAVE YOU DREAMT 4 BEFORE) (=DREAM) (NEWKEY)))
(DREAMED = DREAMT 4 (=DREAMT))
(DREAM 3 ((0) (WHAT DOES THAT DREAM SUGGEST TO YOU)
(DO YOU DREAM OFTEN) (WHAT PERSONS APPEAR IN YOUR DREAMS)
(DON'T YOU BELIEVE THAT DREAM HAS SOMETHING TO DO WITH
YOUR PROBLEM) (NEWKEY)))
(DREAMS = DREAM 3 (=DREAM))
(HOW (=WHAT))
(WHEN (=WHAT))
(ALIKE 10 (=DIT))
(SAME 10 (=DIT))
(CERTAINLY (=YES))
(FEEL DLIST(/BELIEF))
(THINK DLIST(/BELIEF))
(BELIEVE DLIST(/BELIEF))
(WISH DLIST(/BELIEF))
(MEMORY MY (0 YOUR 0 = LETS DISCUSS FURTHER WHY YOUR 3)
(0 YOUR 0 = EARLIER YOU SAID YOUR 3)
(0 YOUR 0 = BUT YOUR 3)
(0 YOUR 0 = DOES THAT HAVE ANYTHING TO DO WITH THE FACT THAT YOUR 3))
(NONE ((0) (I AM NOT SURE I UNDERSTAND YOU FULLY)
(PLEASE GO ON)
(WHAT DOES THAT SUGGEST TO YOU)
(DO YOU FEEL STRONGLY ABOUT DISCUSSING SUCH THINGS)))
(PERHAPS ((0) (YOU DON'T SEEM QUITE CERTAIN)
(WHY THE UNCERTAIN TONE)
(CAN'T YOU BE MORE POSITIVE)
(YOU AREN'T SURE) (DON'T YOU KNOW)))
(MAYBE (=PERHAPS))
(NAME 15 ((0) (I AM NOT INTERESTED IN NAMES)
(I'VE TOLD YOU BEFORE, I DON'T CARE ABOUT NAMES -
PLEASE CONTINUE)) )
(DEUTSCH (=XFREMD))
(FRANCAIS (=XFREMD))
(ITALIANO (=XFREMD))
(ESPAÑOL (=XFREMD))
(XFREMD ((0) (I AM SORRY, I SPEAK ONLY ENGLISH)))
(HELLO ((0) (HOW DO YOU DO. PLEASE STATE YOUR PROBLEM)))
(COMPUTER 50 ((0) (DO COMPUTERS WORRY YOU)
(WHY DO YOU MENTION COMPUTERS) (WHAT DO YOU THINK MACHINES
HAVE TO DO WITH YOUR PROBLEM) (DON'T YOU THINK COMPUTERS CAN
HELP PEOPLE) (WHAT ABOUT MACHINES WORRIES YOU) (WHAT
DO YOU THINK ABOUT MACHINES)))
(MACHINE 50 (=COMPUTER))
(MACHINES 50 (=COMPUTER))
(COMPUTERS 50 (=COMPUTER))
(AM = ARE ((0 ARE YOU 0) (DO YOU BELIEVE YOU ARE 4)
```

```
(WOULD YOU WANT TO BE 4) (YOU WISH I WOULD TELL YOU YOU ARE 4)
(WHAT WOULD IT MEAN IF YOU WERE 4) (=WHAT))
((0) (WHY DO YOU SAY 'AM') (I DON'T UNDERSTAND THAT)))
(ARE ((0 ARE I 0)
(WHY ARE YOU INTERESTED IN WHETHER I AM 4 OR NOT)
(WOULD YOU PREFER IF I WEREN'T 4) (PERHAPS I AM 4 IN YOUR
FANTASIES) (DO YOU SOMETIMES THINK I AM 4) (=WHAT))
((0 ARE 0) (DID YOU THINK THEY MIGHT NOT BE 3)
(WOULD YOU LIKE IT IF THEY WERE NOT 3) (WHAT IF THEY WERE NOT 3)
(POSSIBLY THEY ARE 3)) )
(YOUR = MY ((0 MY 0) (WHY ARE YOU CONCERNED OVER MY 3)
(WHAT ABOUT YOUR OWN 3) (ARE YOU WORRIED ABOUT SOMEONE ELSE'S 3)
(REALLY, MY 3)))
(WAS 2 ((0 WAS YOU 0)
(WHAT IF YOU WERE 4) (DO YOU THINK YOU WERE 4)
(WERE YOU 4) (WHAT WOULD IT MEAN IF YOU WERE 4)
(WHAT DOES ' 4 ' SUGGEST TO YOU) (=WHAT))
((0 YOU WAS 0)
(WERE YOU REALLY) (WHY DO YOU TELL ME YOU WERE 4 NOW)
(PERHAPS I ALREADY KNEW YOU WERE 4) )
((0 WAS I 0) (WOULD YOU LIKE TO BELIEVE I WAS 4)
(WHAT SUGGESTS THAT I WAS 4)
(WHAT DO YOU THINK) (PERHAPS I WAS 4)
(WHAT IF I HAD BEEN 4))
((0) (NEWKEY)) )
(WERE = WAS (=WAS))
(ME = YOU)
(YOU'RE = I'M ((0 I'M 0) (PRE (I ARE 3) (=YOU)))
(I'M = YOU'RE ((0 YOU'RE 0) (PRE (YOU ARE 3) (=I))))
(MYSELF = YOURSELF)
(YOURSELF = MYSELF)
(MOTHER DLIST(/NOUN FAMILY))
(MOM = MOTHER DLIST(/ FAMILY))
(DAD = FATHER DLIST(/ FAMILY))
(FATHER DLIST(/NOUN FAMILY))
(SISTER DLIST(/FAMILY))
(BROTHER DLIST(/FAMILY))
(WIFE DLIST(/FAMILY))
(CHILDREN DLIST(/FAMILY))
(I = YOU
((0 YOU (= WANT NEED) 0) (WHAT WOULD IT MEAN TO YOU IF YOU GOT 4)
(WHY DO YOU WANT 4) (SUPPOSE YOU GOT 4 SOON) (WHAT
IF YOU NEVER GOT 4) (WHAT WOULD GETTING 4 MEAN TO
YOU) (WHAT DOES WANTING 4 HAVE TO DO WITH THIS DISCUSSION))
((0 YOU ARE 0 (=SAD UNHAPPY DEPRESSED SICK ) 0)
(I AM SORRY TO HEAR YOU ARE 5) (DO YOU THINK COMING HERE
WILL HELP YOU NOT TO BE 5) (I'M SURE ITS NOT PLEASANT TO
BE 5) (CAN YOU EXPLAIN WHAT MADE YOU 5))
((0 YOU ARE 0 (=HAPPY ELATED GLAD BETTER ) 0)
(HOW HAVE I HELPED YOU TO BE 5)
(HAS YOUR TREATMENT MADE YOU 5) (WHAT MAKES YOU 5 JUST
NOW) (CAN YOU EXPLAIN WHY YOU ARE SUDDENLY 5))
((0 YOU WAS 0) (=WAS))
((0 YOU (/BELIEF) YOU 0) (DO YOU REALLY THINK SO) (BUT YOU ARE
NOT SURE YOU 5) (DO YOU REALLY DOUBT YOU 5))
((0 YOU 0 (/BELIEF) 0 I 0) (=YOU))
((0 YOU ARE 0)
(IS IT BECAUSE YOU ARE 4 THAT YOU CAME TO ME)
(HOW LONG HAVE YOU BEEN 4)
(DO YOU BELIEVE IT NORMAL TO BE 4)
(DO YOU ENJOY BEING 4)
((0 YOU (= CAN'T CANNOT) 0) (HOW DO YOU KNOW YOU CAN'T 4)
(HAVE YOU TRIED)
(PERHAPS YOU COULD 4 NOW)
(DO YOU REALLY WANT TO BE ABLE TO 4))
((0 YOU DON'T 0)(DON'T YOU REALLY 4)(WHY DON'T YOU 4)
(DO YOU WISH TO BE ABLE TO 4) (DOES THAT TROUBLE YOU))
((0 YOU FEEL 0) (TELL ME MORE ABOUT SUCH FEELINGS)
(DO YOU OFTEN FEEL 4)
(DO YOU ENJOY FEELING 4)
(OH WHAT DOES FEELING 4 REMIND YOU))
((0 YOU 0 I 0) (PERHAPS IN YOUR FANTASY WE 3 EACH OTHER)
(DO YOU WISH TO 3 ME)
(YOU SEEM TO NEED TO 3 ME)
(DO YOU 3 ANYONE ELSE))
```

((0) (YOU SAY 1)
 (CAN YOU ELABORATE ON THAT) (DO YOU SAY 1 FOR SOME SPECIAL REASON)
 (THAT'S QUITE INTERESTING))
 (YOU = 1 ((0) I REMIND YOU OF 0) (=DIT))
 ((0) I ARE 0) (WHAT MAKES YOU THINK I AM 4)
 (DOES IT PLEASE YOU TO BELIEVE I AM 4)
 (DO YOU SOMETIMES WISH YOU WERE 4)
 (PERHAPS YOU WOULD LIKE TO BE 4))
 ((0) I 0 YOU) (WHY DO YOU THINK I 3 YOU)
 (YOU LIKE TO THINK I 3 YOU - DON'T YOU)
 (WHAT MAKES YOU THINK I 3 YOU)
 (REALLY, I 3 YOU) (DO YOU WISH TO BELIEVE I 3 YOU)

 (SUPPOSE I DID 3 YOU - WHAT WOULD THAT MEAN)
 (DOES SOMEONE ELSE BELIEVE I 3 YOU))
 ((0) I 0) (WE WERE DISCUSSING YOU - NOT ME)
 (OH, I 3) (YOU'RE NOT REALLY TALKING ABOUT ME - ARE YOU)
 (WHAT ARE YOUR FEELINGS NOW))
 (YES ((0) (YOU SEEM QUITE POSITIVE) (YOU ARE SURE)
 (I SEE) (I UNDERSTAND)))
 (NO ((0) (ARE YOU SAYING 'NO' JUST TO BE NEGATIVE)
 (YOU ARE BEING A BIT NEGATIVE) (WHY NOT) (WHY 'NO'))
 (MY = YOUR 2 ((0) YOUR 0 (/FAMILY) 0)
 (TELL ME MORE ABOUT YOUR FAMILY) (WHO ELSE IN YOUR FAMILY
 5) (YOUR 4) (WHAT ELSE COMES TO MIND WHEN YOU THINK OF
 YOUR 4))
 ((0) YOUR 0) (YOUR 3) (WHY DO YOU SAY YOUR 3) (DOES THAT
 SUGGEST ANYTHING ELSE WHICH BELONGS TO YOU) (IS IT
 IMPORTANT TO YOU THAT 2 3))
 (CAN ((0) CAN I 0) (YOU BELIEVE I CAN 4 DON'T YOU) (=WHAT)
 (YOU WANT ME TO BE ABLE TO 4) (PERHAPS YOU WOULD LIKE TO BE
 ABLE TO 4 YOURSELF)) ((0) CAN YOU 0) (WHETHER OR NOT YOU CAN 4
 DEPENDS ON YOU MORE THAN ON ME) (DO YOU WANT TO BE ABLE TO 4)
 (PERHAPS YOU DON'T WANT TO 4) (=WHAT))
 (WHAT ((0) (WHY DO YOU ASK) (DOES THAT QUESTION INTEREST YOU)
 (WHAT IS IT YOU REALLY WANT TO KNOW)
 (ARE SUCH QUESTIONS MUCH ON YOUR MIND)
 (WHAT ANSWER WOULD PLEASE YOU MOST)
 (WHAT DO YOU THINK) (WHAT COMES TO YOUR MIND WHEN YOU ASK THAT)
 (HAVE YOU ASKED SUCH QUESTION BEFORE)
 (HAVE YOU ASKED ANYONE ELSE))
 (BECAUSE ((0) (IS THAT THE REAL REASON) (DON'T ANY OTHER REASONS
 COME TO MIND) (DOES THAT REASON SEEM TO EXPLAIN ANYTHING ELSE)
 (WHAT OTHER REASONS MIGHT THERE BE))
 (WHY ((0) WHY DON'T I 0) (DO YOU BELIEVE I DON'T 5) (PERHAPS I
 WILL 5 IN GOOD TIME) (SHOULD YOU 5 YOURSELF) (YOU WANT ME TO 5)
 (=WHAT))

 ((0) WHY CAN'T YOU 0) (DO YOU THINK YOU SHOULD BE ABLE TO 5)
 (DO YOU WANT TO BE ABLE TO 5) (DO YOU BELIEVE THIS WILL HELP YOU
 TO 5) (HAVE YOU ANY IDEA WHY YOU CAN'T 5) (=WHAT))
 (= WHAT))
 (EVERYONE 2 ((0) (= EVERYONE EVERYBODY NOBODY NOONE) 0)
 (REALLY, 2) (SURELY NOT 2) (CAN YOU THINK OF
 ANYONE IN PARTICULAR) (WHO, FOR EXAMPLE) (YOU ARE THINKING OF
 A VERY SPECIAL PERSON)
 (WHO, MAY I ASK) (SOMEONE SPECIAL PERHAPS)
 (YOU HAVE A PARTICULAR PERSON IN MIND, DON'T YOU) (WHO DO YOU
 THINK YOU'RE TALKING ABOUT))
 (EVERYBODY 2 (= EVERYONE))
 (NOBODY 2 (=EVERYONE))
 (NOONE 2 (=EVERYONE))
 (ALWAYS 1 ((0) (CAN YOU THINK OF A SPECIFIC EXAMPLE) (WHEN)
 (WHAT INCIDENT ARE YOU THINKING OF) (REALLY, ALWAYS))
 (LIKE I0 ((0) (=AM IS ARE WAS) 0 LIKE 0) (=DIT))
 ((0) (NEWKEY)))
 (DIT ((0) (IN WHAT WAY) (WHAT RESEMBLANCE DO YOU SEE)
 (WHAT DOES THAT SIMILARITY SUGGEST TO YOU)
 (WHAT OTHER CONNECTIONS DO YOU SEE)
 (WHAT DO YOU SUPPOSE THAT RESEMBLANCE MEANS)
 (WHAT IS THE CONNECTION, DO YOU SUPPOSE)
 (COULD THERE REALLY BE SOME CONNECTION)
 (HOW)))
 ()

RECEIVED SEPTEMBER, 1965