

want it to understand language, we must give it knowledge about the specific subject we want to talk about.

For our experiment, we pretended that we were talking to a simple robot, with a hand and an eye and the ability to manipulate toy blocks on a table. We can say, "Pick up a block which is bigger than the one you are holding and put it in the box.", or ask a sequence of questions like "Had you touched any pyramid before you put the green one on the little cube?" "When did you pick it up?" "Why?", or we can give it new information like "I like blocks which are not red, but I don't like anything which supports a pyramid." The "robot" responds by carrying out the commands (in a simulated scene on a display screen attached to the computer), typing out answers to the questions, and accepting the information to use in reasoning later on.

The dialog is carried out by typing on a terminal attached to the computer time-sharing system. There are a number of hard technical problems in getting a computer to communicate by voice, and it has not been attempted.

We had three main kinds of goals in writing such a program. The first is the practical goal of having a language-understanding system. Even though we used the robot as our test area, the language programs do not depend on any special subject matter, and they have been adapted to other uses.

The second goal is gaining a better understanding of what language is and how it is put together. To write a program we need to make all of our knowledge about language very explicit, and we have to be concerned with the entire language process, not just one area such as syntax. We need the most advanced theories which linguists and others have developed, and we must fit them together to get the program working. This provides a rigid test for linguistic theories, and leads us into making new