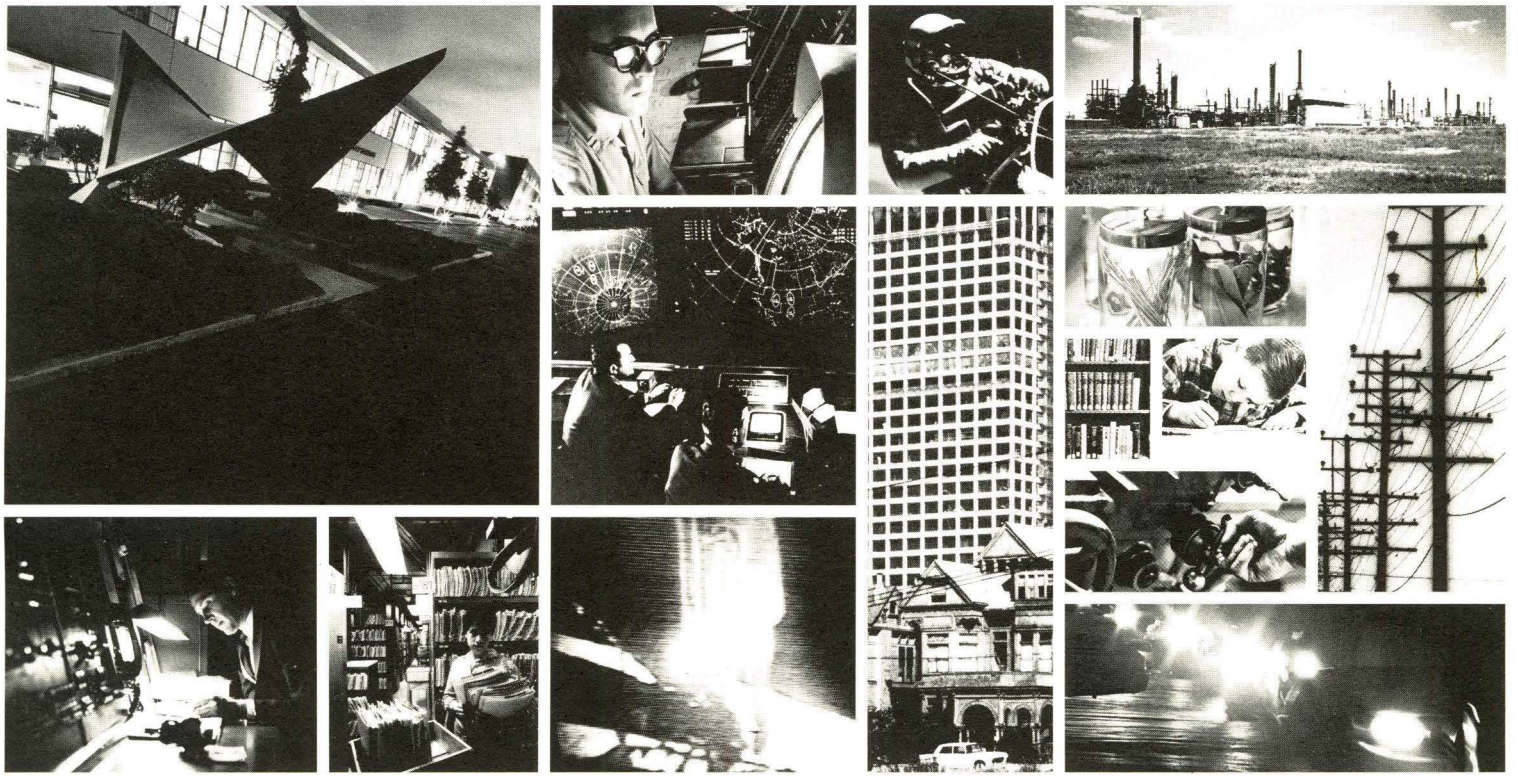

SP-3622

A Computer's View of the World

Antonio Leal

25 August 1971

SP-3622

SP *a professional paper*

A Computer's View of the World

by

Antonio Leal

25 August 1971

SYSTEM
DEVELOPMENT
CORPORATION
2500 COLORADO AVE.
SANTA MONICA
CALIFORNIA
90406

A-1162 (9/68)

Computers are becoming more than giant calculating machines. For example, we can program them to accept and store information (knowledge) and answer questions about that knowledge. Indeed, we can program them in any way that we wish, but this does not help us to decide exactly how to program them. What attitude should the computer have toward the outside world? What should the computer assume if insufficient knowledge is presented to it? I would like to partially answer these questions.

Let us assume that we have a computer like the one described above. We will assume, for convenience, that the computer understands English. I would first like to consider the following question: What should the computer's view of the world be? Should it be (1) that nothing exists except the information to which it has access, or (2) that it only has a partial view of the world and that certain things exist which it does not know. In my estimation, the latter view is far superior to the former. But let us assume the former view to be the case and look at an example.

Suppose the computer contains information about the population of certain cities. We will assume that the concept of "city" is known and that "population" is a function that maps elements of the set of cities into the non-negative integers. A user may ask the computer for information as follows:

"What is the population of Los Angeles?"

and the computer would apply the function "population" to the city "Los Angeles" and reply with the numeric result. Suppose, however, that no value for the population of Los Angeles had been given to the computer although "Los Angeles" had previously been defined as a city. What should the computer answer to the above question? The answer could not be "I don't know.", since this would imply that there was some knowledge that the computer did not know and would thus violate our assumption that the computer's world consists solely of what it knows. The computer must answer:

"The question is meaningless."

or

"Population does not apply to Los Angeles."

The same answer must be given even if it were the case that Los Angeles had never been defined as a city. If the user were to ask:

"Does Los Angeles have a population?"

the answer would be either "yes", if a value were present or "no", if either no value were present or Los Angeles were not a city.

Suppose, now, the user wishes to tell the computer that population should be applicable to all cities. He might say:

"All cities have population."

This produces an internal contradiction. The user is telling the computer that population can be applied to all cities, and, yet, there are certain cities for which no population figure has been given. As far as the computer's view is concerned, non-applicability is identical to no value. The computer cannot even ask the user a question, since this would imply that there is information it does not know. To carry this line of reasoning to its end would be to say that no new information at all could be given to the computer, since that again implies there is something it does not know. I believe you will agree that such a view of the world is too restricted to be of any use.

Assume now that the computer has only a partial view of the world. That is, it realizes that information exists which is unknown to it. If the user asks:

"What is the population of Los Angeles?"

The computer should answer with one of the following:

(a) A number, if Los Angeles is a city and if a value for population is available.

(b) "I don't know." if Los Angeles is a city and if no population figure is available.

(c) "Los Angeles has no population." if Los Angeles is a city and if population does not apply to it. We know from our experience and knowledge that there is no city for which population has no meaning. (Ghost towns have a population of zero.) But the computer does not know this, and, besides, this situation may occur in other examples. The computer can be made to believe this simply by telling it: "Some cities have population."

(d) "The question has no meaning." if Los Angeles is not a city.

If the user asks:

"Does Los Angeles have a population?"

the computer should answer "yes" in cases (a) and (b), "no" in case (c), and "I don't know." in case (d). It should be clear that the partial view of the world is the only one that is tenable and acceptable.

I would now like to consider a second question. If a user makes an assertion to the computer that is lacking a clearly needed quantifier, which quantifier should the computer assume: "all" or "some"? For example, the user says:

"Cities have population."

Does he mean "all cities" or "some cities"? Either choice is logically acceptable, but which one would be implied in normal conversation? Which would a user expect the computer to assume? Consider a different example. Suppose someone came up to you and said: "Dogs are mean." Now, you, on the basis of your experience and knowledge, know that only some dogs are mean. It is certainly not the case that all dogs are mean. But this is not the issue. The issue is, what should you assume the person meant when he

said that? My immediate reaction would be either to ask why the person hated dogs or to begin giving counter examples by naming dogs I knew to be friendly. Either of these reactions implies that I assumed he meant "all" dogs. Otherwise I would have said: "You're right!" *Sentences with missing quantifiers imply "all" as their quantifier.* Correct?

My second argument has just been given. What was your reaction to the last sentence in the above paragraph? Did you immediately assume I meant *all* sentences with missing quantifiers and did you try to come up with some counter example? Or did you agree with my statement because I had just shown you one case where it was true? I would be willing to bet that you assumed I meant all sentences with missing quantifiers.

But let us look at some arguments for assuming "some". One might be that it is well and good for humans to assume "all", but computers must assume the weakest position possible, since they have no information upon which to base a decision. Accepting "some" is certainly weaker than accepting "all". But this is circular reasoning. It is obvious that assuming "some" implies accepting the weaker position. Therefore, one cannot use "accept the weaker position" as an argument for assuming "some". As a point of fact, they are one and the same thing!

Another argument may stem from missing data. For example, since we may not have population figures for all cities, we cannot say that population necessarily applies to all cities. But this is just not true. The fact that values of a given function are not known for certain points in its domain does not imply that the function is defined only on a subset of the assumed domain. That is to say, the question of whether or not population is applicable to all cities or just some cities cannot be determined by how many actual population values are known unless all of the values are known, in which case population applies to all cities.

We must be careful not to confuse the function with the values for that function. In the former example, "population" was the function, "cities" was the domain, and the non-negative integers were the range. But, consider the following statement:

"Cities have smog."

In this case, "smog" is the function, "cities" is the domain, and the set {true, false} is the range. Such functions with two-valued ranges are called "properties" or "attributes" and their use refers to the values of the function as well as the applicability. The above assertion is much stronger than:

"Cities have population."

It asserts not only that the function smog applies to all cities, but additionally, that the value for every city with respect to the function smog is "true". Similarly, the statement:

"Some cities have smog."

is stronger than:

"Some cities have population."

It asserts not only that smog applies to at least one city, but additionally, that at least one of these cities has as its value "true".

By answering questions such as these, with both logical and philosophical arguments, we should gain deeper insight not only into how best to utilize the computer, but also, and more importantly, into how man's own mind functions.


SDC