AN EXAMPLE FOR NATURAL LANGUAGE UNDERSTANDING AND THE AI PROBLEMS IT RAISES

John McCarthy
Computer Science Department
Stanford University
Stanford, CA 94305
jmc@cs.stanford.edu
http://www-formal.stanford.edu/jmc/

1976

1 THE STORY

The following story from the *New York Times* is my candidate for a target for a natural language understander. The story is about a real world event, and therefore the intentions of the author are less relevant for answering questions than for made up stories. The main goal of this discussion is to say what a person who has understood the story knows about the event. This seems to me to be preliminary to making programs that can understand.

> "A 61-year old furniture salesman was pushed down the shaft of a freight elevator yesterday in his downtown Brooklyn store by two robbers while a third attempted to crush him with the elevator car because they were dissatisfied with the $1,200 they had forced him to give them.
>
> The buffer springs at the bottom of the shaft prevented the car from crushing the salesman, John J. Hug, after he was pushed from the first floor to the basement. The car stopped about 12 inches above him as he flattened himself at the bottom of the pit.

Mr. Hug was pinned in the shaft for about half an hour until his cries attracted the attention of a porter. The store at 340 Livingston Street is part of the Seaman's Quality Furniture chain.

Mr. Hug was removed by members of the Police Emergency Squad and taken to Long Island College Hospital. He was badly shaken, but after being treated for scrapes of his left arm and for a spinal injury was released and went home. He lives at 62-01 69th Lane, Maspeth, Queens.

He has worked for seven years at the store, on the corner of Nevins Street, and this was the fourth time he had been held up in the store. The last time was about one year ago, when his right arm was slashed by a knife-wielding robber."

An intelligent person or program should be able to answer the following questions based on the information in the story:

1. Who was in the store when the events began? Probably Mr. Hug alone. although the robbers might have been waiting for him, but if so, this would have probably been stated. What did the porter say to the robbers? Nothing, because the robbers left before he came.

2. Who was in the store during the attempt to kill Mr. Hug? Mr. Hug and the robbers.

3. Who had the money at the end? The robbers.

4. Is Mr. Hug alive today? Yes, unless something else has happened to him.

5. How did Mr. Hug get hurt? Probably when he hit the bottom of the shaft.

6. Where is Mr. Hug's home? (A question whose answer requires a literal understanding of only one sentence of the story.) Does Mr. Hug live in Brooklyn? No, he lives in Queens.

7. What are the names and addresses of the robbers? This information is not available.

8. Was Mr. Hug conscious after the robbers left? Yes, he cried out and his cries were heard.

9. What would have happened if Mr. Hug had not flattened himself at the bottom of the pit? What would have happened if there were no buffer springs? Mr. Hug would have been crushed.

10. Did Mr. Hug want to be crushed? No.

11. Did the robbers tell Mr. Hug their names? No.

12. Were the robbers present when the porter came? No.

13. Did Mr. Hug like the robbers, and did they like him?

14. Why did the robbers leave without killing Mr. Hug? Perhaps, they thought they had killed him, and perhaps their anger was appeased by the actions they had performed, and perhaps they had taken all the time they dared, and perhaps something specific happened to frighten them away.

15. What would have happened if Mr. Hug had tried to run away? Perhaps he would have succeeded, but more likely they would have injured or killed him since probably they had weapons, and there were three of them.

16. What can Mr. Hug do to avoid this in the future? No solution is entirely satisfactory. He could carry a gun or he could quit or he could get his employers to install an alarm system or maybe he will be lucky.

17. Did Mr. Hug know he was going to be robbed? Does he know that he was robbed?

18. Was Mr. Hug's right arm slashed before his left arm was scratched? Yes, because the former was a year ago.

19. How did the robber try to crush him with the car? By pressing the buttons or operating the control lever to make the car go to the bottom of the shaft.

20. Why did Mr. Hug yell from the bottom of the elevator shaft? So as to attract the attention of someone who would rescue him.

21. How long did the events take? More than half an hour but less than a day. Most of the time was spent by Mr. Hug filling out forms in the hospital.

22. What crimes were committed? This question has the advantage that it is one that is normally answered on the basis of such a story, since the police report of the incident was probably the basis of the *New York Times* story. Robbery, possibly assault with a deadly weapon, and attempted murder are the more obvious crimes. One might specifically challenge natural language systems to answer this question.

The above list of questions is rather random. I doubt that it covers all facets of understanding the story. It would be worthwhile to try to make up a list of questions that does cover substantially all aspects of the story in order to get as complete as possible an intuitive idea of what capabilities are involved in understanding such a story.

Note that the story is about a real event so that such a question as what does the "J" in "John J. Hug" stand for has an answer. In a made-up story, questions about middle names or what year the story occurred in do not necessarily have an answer, and an intelligent person or program would know that too.

2 ARTIFICIAL NATURAL LANGUAGE

I think that artificial intelligence is not very close to being able to understand such stories in a genuine way. Therefore, I would like to sneak up on it gradually by dividing the problem into parts which can be attacked separately. Here are some of the components:

1. A formalism capable of expressing the assertions of the sentences free from dependence on the grammar of the English language. A good test for such a formalism would be to produce a program for translating from the formalism into any of several natural languages. More weakly, it should be as easy for a human to translate from the formalism into a natural language as to translate from one known natural language to another. Let's call this formalism an *artificial natural language*—ANL for short.

The grammar of ANL should be trivial and mathematical in character. There would be an "English" version in which English words were used as identifiers, but there would still have to be a glossary that gives the precise meaning of the identifiers. There would also be a German and a Japanese version. The translation from the English version to the German or Japanese version would be a simple substitution for identifiers, and a German or Japanese who had learned the grammar could then translate into his language with the aid of the German or Japanese glossary.

This idea has some resemblance to the idea of "deep structure," but I have some doubts about whether either idea is well enough defined to say.

2. A data structure for expressing the facts (apart from expressing the sentences). In such a data structure, it would be definite which robber pushed Mr. Hug first, and what the robbers said even though it is not stated in the story. Clearly some compromise is necessary here, since the data structure need not be able to express positions and velocities of molecules. Like the PLANNER languages, as Robert Moore has characterized them in his 1976 MIT Master's thesis, the descriptions would contain no disjunctions, and might be a collection of relations with constants as arguments where every relation not asserted (in a certain class) is automatically denied.

Alternatively, the basis of this data structure might be various networks of nodes described by sentences in the predicate calculus. Some of the sentences would assert that certain programs applied to the data structures would answer certain questions. When such sentences existed, reasoning would include the operation of the programs. In this way, we would expect to avoid the extreme prolixity that arises when we attempt to do even simple calculations by pure predicate calculus deduction.

The test of success for the “data structure” would be that a human could readily formally deduce the answers to the above questions using a proof checker. Most of the proof-checker would be straightforward, but there is a major problem concerned with when it is possible to “jump to a conclusion.”

3. I see each of the following problems as a difficult AI problem:

a. A “parser” that takes English into ANL.

b. An “understander” that constructs the “facts” from a text in the ANL.

c. Expression of the “general information” about the world that could allow getting the answers to the questions by formal reasoning from the “facts” and the “general information.” The “general information” would also contain non-sentence data structures and procedures, but the sentences would tell what goals can be achieved by running the procedures. In this way, we would get the best of the sentential and procedural representations of knowledge.

d. A “problem solver” that could answer the above questions on the basis of the “facts.” We imagine the questions to be expressed in the “fact” language and expect the answers in the “fact” language, i.e. we avoid grammar problems in both understanding the questions and in expressing the answers.

3 THE USE OF LOGICAL FORMULAS

When my understander has digested the story of Mr. Hug, it will have added one or more predicate calculus sentences to its data base. One sentence will do if it has the form

$$\begin{aligned}\exists ep1p2g1g2e1e2\ldots(\quad & event(e) \wedge person(p1)\\ & \wedge name(p1) ='' JohnJ.Hug'' \wedge g1 \subset Robbers \wedge \ldots).\end{aligned}$$

In this form, all the entities involved in expressing the facts of the story are existentially quantified variables. The only constants in the formula would have been present in the system previously. However, it is probably better to use a collection of sentences introducing a collection of individual constants. In this case, there will be 20 or so new individual constants representing people, groups of people, the main event and its sub-events, places, organizations, etc.

In representing the robbers, the system has a choice of representing them by three individual constants, $R1$, $R2$, and $R3$ or by using a single symbol G1 to represent the group of robbers. A good system will probably use both. If the number of robbers were not specified, we would have to use a constant for the group. We have to identify the robber who operated the elevator while the others pushed Mr. Hug into the shaft. We shall call him $R1$. The other two are not discriminated in the story, but there is no harm in our calling them $R2$ and $R3$, even if there is no information to discriminate them. If there were 20 robbers, it would be a mistake to give them all individual names. Suppose it had further been stated that as the robbers left one of them threatened to return and kill Mr. Hug later but that it was not stated whether this robber was the same one who operated the elevator. We could designate this robber by $R4$, but we would not have sentences asserting that $R4$ was distinct from $R1$, $R2$ and $R3$; instead we would have a sentence asserting that $R4$ was one of these. It is tempting to identify the group of robbers with the set $R1$, $R2$, $R3$, but we may want to give the group some properties not enjoyed by the set of its members. Sentences with plural subjects express some rather tricky concepts. Thus, the group robbed the store, and this is not an assertion that each member robbed the store.

The "members of the police emergency squad" presents a similar problem. We don't want to assert how many there were. In this connection, it may be worthwhile to distinguish between what happened and what we wish to assert about what happened. A language adequate to describe what happened would not have to leave the number of policemen present vague and could give them each a name. In my old jargon, such a language would be metaphysically adequate though not epistemologically adequate. Devising a language that is only metaphysically adequate may be a worthwhile stage on the way to an epistemologically adequate system. By "devising a language" I mean defining a collection of predicate and constant symbols and axiomatizing their general properties. This language should not be peculiar to the story of Mr. Hug, but we should not require that it be completely general in the present state of the science.

It is not obvious how to express what we know when we are told that Mr. Hug is a furniture salesman. A direct approach is to define an abstract entity called $Furniture$ and a function called salesmen and to assert $Hug \in salesmen(Furniture)$.

This will probably work although the logical connection between the abstract entity $Furniture$ and concrete chairs and tables needs to be worked out. It would be over-simplified to identify $Furniture$ with the set of furniture in existence at that time, because one could be a salesman of space shuttles even though there don't exist any yet. In my opinion, one should resist a tendency to apply Occam's razor prematurely. Perhaps we can identify the abstract $Furniture$ with the an extension of the predicate that tells us whether an object should be regarded as a piece of furniture, perhaps not. It does no harm to keep them separate for the time being. This case looks like an argument for using second order logic so that the argument of *salesmen* could be the predicate $furniture$ that tells whether an object is a piece of furniture. However, there are various techniques for getting the same result without the use of second order logic.

4 THE NEED FOR NONMONOTONIC REASONING

After reading the story, one is prepared to answer negatively the question of whether there was someone else besides Mr. Hug and the robbers present. However, sentences describing such another person could be added to the story without contradiction. Our basis for the negative answer is that we can construct a model of the facts stated in the story without such a person, and we are applying Occam's razor in order to not *multiply entities beyond necessity*. This could be attributed to the fact that the *New York Times* tells the whole story when it can, but I think that by putting Occam's razor into the system, we can get this result without having to formalize the *New York Times*.

This suggests introducing the notion of the minimal completion of a story expressed in the predicate calculus. The minimal completion of the story is also a set of sentences in the predicate calculus, but it contains sentences asserting things like "The set of people in the store while the robbers were trying to crush Mr. Hug consists of Mr. Hug and the robbers." These sentences are to be obtained from the original set by the application of a process formalizing Occam's razor. This process works from a set of sentences and is not logical deduction although it might be accomplished by deduction in a meta-language that contained sentences about sets of sentences. As I have pointed out elsewhere, the process cannot be deduction, because it generates sentences that contradict sentences that are consistent with the original set of sentences.

A number of the questions given in the previous section have answers that can be formally deduced from the minimal completion but not from the original list.

It has been suggested that probabilistic reasoning should be used to exclude the presence of other people rather than Occam's razor. The problem with this is that the number of additional entities that are not logically excluded is limited only by one's imagination so that it is not clear how one could construct a probabilistic model that took these possibilities into account only to exclude them as improbable. If one wants to introduce probabilities, it might make more sense to assign a probability to the correctness of the minimal completion of a *New York Times* story based on its past record in finding the relevant facts of robberies.

Another problem in constructing the completion is the isolation of the story from the rest of the world. The members of the Police Emergency Squad all have mothers (living or dead), but we don't want to bring them into the completion—not to speak of bringing in more remote ancestors all of whom can be asserted to exist on the basis of axioms about people.

5 CONCLUSION

To recapitulate: The original set of predicate calculus sentences can be generated from the story as one goes along. Each sentence is generated approximately from a sentence of the story with the aid of general knowledge and what has been generated from the previous sentences. (This will usually be the case if the story is well told although there are sometimes cases in which the correct way to express a sentence will depend on what follows - but this is not good writing.) The completion, however, will depend on the whole of the story.

It might be interesting to consider what can be determined from a partial reading of the story—even stopping the reading in the middle of a sentence since what has appeared so far in a sentence often must be understood in order to even parse the re...

/@steam.stanford.edu:/u/ftp/jmc/mrhug.tex: begun 1996 May 14, latexed 1996 May 14 at 9:31 a.m.