

A CONCEPTUAL THEORY OF QUESTION ANSWERING

Wendy G. Lehnert
Department of Computer Science
Yale University
New Haven, Ct. 06520

ABSTRACT

A theory of Q/A has been proposed from the perspective of natural language processing that relies on ideas in conceptual information processing and theories of human memory organization. This theory of Q/A has been implemented in a computer program, QUALM. QUALM is currently used by two story understanding systems (SAM and PAM) to complete a natural language processing system that reads stories and answers questions about what was read.

Keywords: natural language processing, computational question answering, conceptual information

1. INTRODUCTION

If a computer is going to answer questions in a manner which is natural for human interaction, the computer must have knowledge of how people ask questions and what kinds of answers are expected in return. A competent question answering system must be based on a theory of human question answering that describes:

- (1) what it means to understand a question
- (2) how context affects understanding
- (3) what kind of responses are appropriate
- (4) how to extract answers from memory

A theory of conceptual question answering has been developed which addresses these four problems [Lehnert '77]. This theory has been implemented in a computer program (QUALM) which runs in conjunction with two story understanding systems, SAM [Cullingford '77] and PAM [Wilensky '76], enabling these systems to answer questions about the stories they read.

The theory of question answering proposed by QUALM is a theory of natural language processing. This distinguishes QUALM from many other question answering systems which are oriented towards information retrieval. Many systems which attempt to answer questions phrased in natural language have been designed in two pieces: (1) a memory retrieval system, and (2) a natural language interface. Very often the interface problem is considered secondary to the retrieval system and the two subsystems are designed as if they were

This work was supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored under the Office of Naval Research under contract N00014-75-C-1111.

theoretically independent of each other [Shortliffe '74, Woods '72].

The theory behind QUALM extends theories of memory processing which originated with the study of parsing [Riesbeck & Schank '77] and generation [Goldman '75]. Conceptual Dependency [Schank '75] has proven to be a strong representational system for the task of question answering. Parsing and generation strategies based on Conceptual Dependency were naturally adopted for question answering without significant alterations. This approach to question answering which utilizes existing theories of natural language processing constitutes a major departure from the information retrieval viewpoint where natural language is considered to be merely a "front end" for a question answering system.

In order to understand questions, QUALM must interface with a conceptual analysis program that parses an English question into its Conceptual Dependency representation [Schank '75]. In SAM and PAM, QUALM interfaces with a parser designed by Christopher Riesbeck [Riesbeck & Schank '76]. In order to produce answers in English, QUALM also needs a generator that can translate Conceptual Dependency representations into English. The generator used by SAM and PAM is based on a generator designed by Neil Goldman [Goldman '75]. All of the processing specific to answering questions occurs on a conceptual level that is language independent. If QUALM interfaced with a Russian parser and a Chinese generator, it would be able to understand questions stated in Russian and produce answers to these questions in Chinese. No changes in QUALM are required to accommodate different languages since the question answering processes are independent of language.

2. CONCEPTUAL CATEGORIES FOR QUESTIONS

When QUALM initially receives a question from the parser, the question is represented as a Conceptual Dependency conceptualization. This conceptualization must then be categorized into one of thirteen possible Conceptual Categories. The Conceptual Categories for questions are:

- | | |
|-----------------------|-----------------------------|
| (1) Causal Antecedent | (7) Instrumental/Procedural |
| (2) Goal Orientation | (8) Concept Completion |
| (3) Enablement | (9) Expectational |
| (4) Causal Consequent | (10) Judgemental |
| (5) Verification | (11) Quantification |
| (6) Disjunctive | (12) Feature Specification |
| | (13) Request |

The conceptual parse of a question represents a very literal or naive understanding of the question. Conceptual Categorization constitutes a higher level of interpretation which is designed to determine exactly what the questioner really means. For example, if a stranger walks up to John on the street and asks:

Q1: Do you have a light?

John would parse this question into a conceptualization equivalent to asking:

Q2: Do you have in your immediate possession an object capable of producing a flame?

If John does not interpret the question any further than this, he could answer:

A1: Yes, I just got a new lighter yesterday.

and then walk away. This sort of response indicates that John did not have a complete understanding of the question. He understood it on a preliminary level, but he did not understand it in terms of what the questioner had intended. His misinterpretation can be explained as faulty Conceptual Categorization. What John understood to be an inquiry deserving a yes or no answer, should have been understood as a request deserving a performative action. The person asking Q1 didn't just want to know if John had a light; he wanted John to offer him a light (flame). In terms of Conceptual Categories, we would say that the question should have been interpreted as a Functional Request rather than a Verification Inquiry.

If a question is not categorized correctly, it will be impossible to produce an appropriate response.

RIGHT: Q3: How could John take the exam?
(an Enablement question)

A3a: He crammed the night before.
(an Enablement answer)

WRONG: Q3: How could John take the exam?
(an Enablement question)

A3b: He took it with a pen.
(Instrumental/Procedural answer)

Q3 is asking about the enabling conditions for taking an exam. In order to take an exam, one has to be prepared for it, presumably be a student, and so forth. Q3 suggests that the questioner does not believe John satisfied some necessary enabling condition. An appropriate answer to Q3 will address this questioned enablement (He crammed the night before, or he bribed an administrator). A3b does not address the Enablement conditions at all. A3b answers the question on a much lower level of instrumentality, indicating that the question was understood to be an Instrumental/Procedural question instead of an Enablement question.

RIGHT: Q4: How did John die?
(Causal Antecedent question)

A4: He caught the swine flu.
(a Causal Antecedent answer)

WRONG: Q4: How did John die?
(Causal Antecedent question)

A4b: Well, first he was alive.
(an Enablement answer)

This time A4b indicates that Q4 was understood to be an Enablement question. A necessary enablement for dying is being alive. But Q4 should not have been interpreted to be asking about the enabling conditions for dying. Q4 is more reasonably understood to be asking about the cause of John's death: Was it an accident? Was he ill? Did he kill himself?

RIGHT: Q5: How did John get to Spain?
(Instrumental/Procedural question)

A5a: He went by plane.
(Instrumental/Procedural answer)

WRONG: Q5: How did John get to Spain?
(Instrumental/Procedural question)

A5b: He wanted to see Madrid.
(a Causal Antecedent Answer)

An appropriate answer to Q5 would specify the transportation means which was instrumental to John's getting to Spain (he took a cruise, he flew, etc.) But A5b tells us what caused John to go to Spain. A5b answers a Causal Antecedent question instead of an Instrumental/Procedural question.

When Q3-5 are represented in Conceptual Dependency, it is easy to see which Conceptual Category should be assigned to these questions. In QUALM, parsed conceptualizations are run through a discrimination net which assigns a Conceptual Category to each question. But Conceptual Categorization does not constitute complete understanding of a question. Each conceptual question is subject to further interpretive processing before a memory search for an answer can begin.

3. INFERENTIAL ANALYSIS

Complete understanding of a question often involves inferences in addition to Conceptual Categorization. When interpretation of a question does not include analysis by inference, answers may be produced which are technically correct, but completely useless. Suppose John is mixing cake batter and he asks his wife:

Q6: Now what haven't I added?

A6: A pound of dog hair and an air filter.

She's probably right; he probably didn't add a pound of dog hair and an oil filter. But her answer is inappropriate because John was "obviously" asking for what he hadn't added that he should have added. The intent of this question is obvious only when an interpretive inference mechanism can be invoked to supply an implicit constraint. There is an entire class of questions that require the same type of inferential analysis:

Q7: Who isn't here?
(Who isn't here who should be here?)

Q8: What did I forget to buy?
(What didn't I buy that I should have?)

In each of these questions, an inference must be made that specifies appropriate constraints for potential answers. When Q7 is asked by a professor upon entering his class, appropriate answers refer to members of the class. When Q8 is asked in the context of shopping for a dinner party, appropriate answers refer to those things that are needed for dinner.

The Universal Set Inference, a general inference mechanism, is needed for questions of this class. This mechanism examines the context of a question and determines appropriate constraining factors. But before this mechanism can be invoked, some process must be responsible for recognizing which questions require this particular inference mechanism. The Universal Set Inference should not be summoned for questions like:

Q9: Who is coming to your party?

Q10: Isn't this the book you wanted?

The successful application of an interpretive inference mechanism relies on the ability to know when that mechanism is needed. This is one way Conceptual Categorization is exploited. One of the thirteen Conceptual Categories is the class of Concept Completion questions. These correspond roughly to fill-in-the-blank questions. During the interpretation of a question, the Universal Set Inference is applied if and only if:

- (1) the question is categorized as a Concept Completion question, and
- (2) the conceptual question has MODE = NEG

Q6-8 each satisfy these requirements. While the lexical statement of Q8 does not appear to be negated, the conceptual representation for Q8 is equivalent to asking "What didn't I remember to buy?" which is encoded as an MTRANS with negative MODE. Q9 is a Concept Completion question but it fails to meet the criteria because it has a non-negative MODE. Q10 fails because it is a Verification question instead of a Concept Completion question.

A useful system of categorization will provide simple test criteria for inference mechanisms of the sort just described. Different questions require different processing. A strong categorization system can recognize which processes are required for a given question and dictate subsequent processing accordingly.

4. CONTEXT-SENSITIVE INTERPRETATION

In the last section we claimed that one general inference mechanism, the Universal Set Inference, could be invoked to establish appropriate constraints on Concept Completion questions with MODE = NEG. This inference mechanism relies on the context in which questions are asked.

Q7: Who isn't here?

requires contextual information that implicitly specifies who should be here. If this question is asked by a professor in a class, it means "Which of my students aren't here?" If it is asked by a host at a party, it means "Who isn't here who was invited?" Without contextual information, it is impossible to know what implicit constraints are appropriate.

Specific constraints on questions can be derived from whatever scripts [Schunk & Abelson '77] are actively operating in a given context. When a script is active, its script-defined roles and role instantiations [Cullingford '77] delineate the universal set these questions implicitly reference.

The Universal Set Inference

Question Category: Concept Completion

Question Criteria: MODE value = NEG

Contextual Criteria: there is an active script

If these test criteria are satisfied, interpretive constraints are imposed by those roles defined in the active script(s).

<u>question</u>	<u>scriptal context</u>	<u>constraining roles</u>
Who isn't here?	party	guests
What didn't I add?	cooking	ingredients
Who hasn't bid?	bridge game	bridge players

Many questions must be understood in terms of their surrounding context. It is therefore crucial to be able to characterize contextual information in terms of general knowledge structures (Schunk & Abelson '77) so that general interpretive inference mechanisms can be designed which are sensitive to context without being context-specific. That is, a contextually sensitive processing mechanism should be applicable in different contexts. A theory of question answering that needs to propose a new set of processing strategies for each new context encountered is not much of a theory.

5. CONTENT SPECIFICATION

Once a question has been sufficiently understood, retrieval processes can begin to look for an answer. The first part of the retrieval process decides how much of an answer is needed. Consider the following story:

John went to a restaurant and the hostess gave him a menu. When he ordered a hot dog the waitress said they didn't have any. So John ordered a hamburger instead. But when the hamburger came, it was so burnt that John left.

If asked:

Q14: Did John eat a hot dog?

There are many possible answers. When SAM reads this story, SAM can answer Q14 three different ways:

A14a: No.

A14b: No, the waitress told John they didn't have any hot dogs.

A14c: No, the waitress told John they didn't have any hot dogs and so John ordered a hamburger.

These answers are all different in terms of the amount of information they convey. In fact, answers can vary not only in terms of their relative content, but in terms of the kind of content they communicate. For example, if Q14 had been answered "Yes," in the context of our story where John didn't eat a hot dog, then the content of this answer would be described as wrong.

The decision-making processes that determine what kind of an answer should be returned are part of Content Specification. Content Specification takes into account the Conceptual Category of each question and intentionality factors that describe the "attitudinal" mode of the entire system in order to determine how a question should be answered. A system of descriptive instructions are produced by Content Specification to instruct and guide memory retrieval processes as they look for an answer.

The primary challenge involved in Content Specification is precisely how these instructions to memory retrieval are formalized. It is not enough to say "give a minimally correct answer," or "bring in everything you can find that's relevant." The instructions generated by Content Specification must tell the retrieval heuristics exactly how to produce a minimally correct answer and exactly what has to be done to do to come up with everything that's relevant.

One type of Content Specification mechanism that guides retrieval heuristics are Elaboration Options. Each Elaboration Option has four parts: an Intentionality Threshold, a Question Criterion, an Initial Answer Criterion, and Elaboration Instructions. Intentionality refers to variables within the system that are set with suggestively-named values like "talkative," "cooperative," "minimally responsive," etc. The Intentionality Threshold specifies what sort of Intentionality must be assigned to the system in order for an Elaboration Option to be attached to the question. The Question Criterion describes what Conceptual Category must be assigned to the question in order for it to receive the Elaboration Option. If either the Intentionality of the system or the Conceptual Category of the question fail to meet the specifications of the Intentionality Threshold and the Question Criterion, then the Elaboration Option is not used. The Answer Criterion specifies the type of conceptual answer which the memory search must initially return in order for the Elaboration Option to be

executed. And the Elaboration Instructions specify exactly how an elaboration is to be extracted from memory and integrated into the conceptual answer.

To see exactly how an Elaboration Option works, we will discuss one of the simpler Elaboration Options that has been implemented in QUALM. Consider the following story:

John went to New York by bus. On the bus he talked to an old lady. When he left the bus, he thanked the driver. He took the subway to Leone's. On the subway his pocket was picked. He got off the train and entered Leone's. He had some lasagna. When the check came, he discovered he couldn't pay. The management told him he would have to wash dishes. When he left, he caught a bus to New Haven.

After reading this story SAM answers:

Q15: Did John go to New York?

A15: Yes, John went to New York by bus.

Q16: Did John eat?

A16: Yes, John ate lasagna.

Q17: Did someone pick John's pocket?

A17: Yes, a thief picked John's pocket.

Q18: Did John pay the check?

A18: Yes, John paid the bill.

These are answers SAM gives when it has a talkative Intentionality. If SAM were running with a less than talkative Intentionality, each of these questions would have been answered with a simple "Yes." The longer answers (A15-18) are the result of the Verification Option. This is a very simple Elaboration Option which is defined as follows:

The Verification Option

Intentionality Threshold: Talkative
Question Category: Verification
Answer Criterion: initial answer is Yes
Elaboration Instructions:

final conceptual answer is "Yes, *X*" where *X* is the conceptualization found in the story representation that matches the question concept.

The retrieval heuristics for a Verification question search the story representation for a conceptualization matching the conceptual question. If it finds a match, the initial answer is Yes. A conceptualization from the story representation doesn't have to correspond to the question concept exactly in order to match it; it may contain more information than the question concept. This is why A15-17 appear to volunteer information. A15 tells how John went to New York, A16 says what John ate, and A17 asserts who stole John's wallet.

6. FINDING AN ANSWER

Retrieval heuristics vary for each Conceptual Category of questions. A number of interesting problems arise in designing processes that extract information from memory. We will briefly outline three such problems which reflect the scope and depth of the difficulties involved.

6.1 Integrative Memory Processing

Expectational questions are interesting because they cannot be answered on the basis of a story representation alone. Expectational questions correspond roughly to why-not questions. These questions require "integrative" memory processing. The term integration is very often used in the context of adding new information to memory. A single unit of information is "integrated" into a larger memory structure. But in the context of retrieving information from memory, an integrative process is one which combines information from different sources to produce new information.

After reading the burnt-hamburger story, SAM answers:


- Q19: Why didn't John eat a hot dog?
 A19: Because the waitress told John they didn't have any hot dogs.
 Q20: Why didn't John eat the hamburger?
 A20: Because the hamburger was burnt.

These questions are answered by an integrative process that combines the story representation with predictive mechanisms in order to reconstruct expectations that were alive at some time during the understanding process. When John ordered a hot dog we had an expectation that he would eat a hot dog until we heard there were none. When John ordered a hamburger we expected him to eat a hamburger until we heard that the hamburger was burnt and John just left. Expectational questions ask about expectations which were aroused at some point during the understanding process and then subsequently violated by an unexpected turn of events. Had we asked "Why didn't John swim across the lake?" the question would seem unreasonable since we never had any expectations about John going swimming or crossing a lake.

The theories of memory representation implemented in SAM and PAM adhere to the premise that a story representation should encode information about things that happened in the story. This includes inferences about things that probably occurred (but weren't explicitly mentioned) as well as conceptualizations for events that were explicitly described in the input story. But Expectational questions ask about things that didn't happen. To answer an Expectational question, we must use the same predictive processes used during story understanding to reconstruct failed expectations which were alive at some time during understanding. The reconstruction of failed expectations is achieved by an integrative

memory process called ghost path generation.

The generation of ghost paths cannot be fully understood without a fundamental understanding of script application [see Cullingford '77]. But some sense of what goes on should be apparent from the following diagram. In this diagram, the chain of events in the center corresponds very roughly to information in the story representation that SAM generated at the time it read the burnt-hamburger story. The two chains on either side correspond to the two ghost paths needed to answer Q19 and Q20.



6.2 Answer Selection

While Expectational questions are interesting because they cannot be answered on the basis of a story representation alone, there are many questions that do not need information outside of the story representation which are still difficult to answer. Causal Antecedent questions are complicated in this respect. A Causal Antecedent question is one which asks for the reason behind an event. After reading the Leone's story, consider the following answers:

- Q21: Why did John wash dishes?
 A21a: Because he couldn't pay the check.
 A21b: Because he had no money.
 A21c: Because he had been pickpocketed on the subway.

SAM answers Q21 "Because he had no money." But is this the best answer of the three? What factors determine the superiority of one answer over another?

Effective answer selection entails making assumptions about what the questioner knows. Anyone who asks Q21 can be assumed to know that John washed dishes. If we go on to assume that the questioner knows two more things: (1) John washed dishes in a restaurant, and (2) washing dishes in a restaurant is classically what happens when someone eats and then can't pay, then the questioner can infer: (3) John couldn't pay the check. If the questioner can figure out for himself that John couldn't pay the check then A21a does not tell him anything he doesn't know to begin with. A good answer must take into account what the questioner does and doesn't know, and address the knowledge state of the questioner by telling him something new.

A21b is a weak answer for the same reasons that A21a was weak. If someone knows that John couldn't pay a check, they can reasonably infer that John didn't have (enough) money. Both inferences:

- 1) John couldn't pay the check.
- 2) John didn't have any money.

can be made by the questioner on the basis of general world knowledge and knowing that John washed dishes in a restaurant. But there is no way the questioner can infer that John was pick-pocketed on the subway without additional knowledge of the story. Therefore A21c is the best answer to Q21 as long as we assume the questioner has knowledge about the world and can make inferences on the basis of that knowledge.

If we assumed that the questioner knew nothing about restaurants, A21a would be the best answer. If we assumed that the questioner knew about restaurants but didn't understand about paying for things, A21b would be the best answer. It is impossible to judge various answers to a question without knowing (or assuming) something about the person being addressed.

6.3 Conceptual Organization of Knowledge

When people answer questions, their answers sometimes tell us something about the form and organization of conceptual information in human memory. For example, consider the following story:

John was sitting in a dining car. When the train jerked, the soup spilled.

Suppose we ask:

Q22: Where was the soup?

This is a specification question that can be answered a number of ways. Two common answers are:

- A22a: In a bowl.
A22b: On the table.

A much less natural answer would be:

A22c: On a plate.

A22c seems to be very odd answer which conjures up an image of a soup puddle on a plate. This is not the scene most people envision when hearing the story. Most people imagine the soup in a bowl on a plate on a table.

The acceptable and unacceptable answers to Q22 tell us something about human memory organization. It never occurs to people to answer "On a plate." Furthermore, when this answer is given it provokes a wrong image of soup resting directly on a plate. But "On the table," is a natural answer. Why is it that "On a plate," is a bad answer but "On the table," is perfectly reasonable? The soup does not rest directly on the table any more than it rests directly on a plate. Why is it acceptable in one case but not the other? This phenomenon must be accounted for in terms of memory organization.

When people hear this story they assume a causality between the train jerking and the soup spilling. (If asked "Why did the soup spill?" people will answer "Because the train moved.") This causality relies on the fact that the soup is physically connected to the train in some way. This physical connection can only be recognized by constructing a path of physical objects between the soup and the train. This path of connections must be accessed in order to answer Q22. If a path is constructed the same way people build one, it will be easy to retrieve answers to Q22 that seem natural. If the path is built differently, we may end up with an answer like A22c.

Suppose we construct a path like the following:

A BAD PATH: soup (inside-of)
 bowl (on-top-of)
 plate (on-top-of)
 tablecloth (on-top-of)
 table (on-top-of)
 floor (inside-of)
 dining car (part-of)
 train

With this memory representation it is not clear how we can extract the answers A22a and A22b without also getting answers like "On a plate," or "On a tablecloth." There is nothing in this memory representation that tells us where the good answers are. What we need is a memory representation that makes it easy to find a bowl and a table but hard to retrieve a plate.

A BETTER PATH: soup (inside-of)
 bowl (part-of)
 placesetting (part-of)
 tablesetting (on-top-of)
 table (part-of)
 dining area (part-of)
 dining car (part-of)
 train

This path suggests a very simple retrieval heuristic to produce the answers A22a and A22b: trace the path looking for objects which are connected by either "inside-of" or "on-top-of" links.

The closer a memory representation is to human memory organization, the easier it will be to produce answers that make sense to people. A system of memory representation for physical objects has been proposed [Lehnert '77] which is designed to facilitate inference and retrieval problems of the sort just described. Conceptual descriptions of objects in this system are based on decompositions into a set of seven object primitives in much the same way that Conceptual Dependency [Schank '75] describes actions by decomposing them into a set of primitive acts.

7. CONCLUSIONS

The overall question answering process can be intuitively approached in two stages: understanding the question (interpretation) and finding an answer (memory retrieval). Each of these stages is likewise divided into two parts:

INTERPRETATION:

- [1] Conceptual Categorization
- [2] Inferential Analysis

MEMORY RETRIEVAL:

- [3] Content Specification
- [4] Searching Heuristics

[1] Conceptual Categorization guides the subsequent processing by dictating which specific inference mechanisms, elaboration options, and retrieval heuristics should be invoked in the course of answering a question.

[2] Inferential Analysis is responsible for understanding what the questioner really meant when a question should not be taken literally.

[3] Content Specification determines how much of an answer should be returned in terms of detail and elaborations.

[4] Searching Heuristics do the actual digging in order to extract an answer from memory.

All of the processes within these four phases are specific to question answering per se and are language-independent, operating within a conceptual representation system.

QUALM represents a theory of question answering which is motivated by theories of natural language processing. Within the context of story understanding, QUALM has provided a concrete criterion for judging the strengths and weaknesses of story representations generated by SAM and PAM. If a system understands a story, it should be able to answer questions about that story in the same way that people do. Although the computer implementation of QUALM is currently

limited to the application of answering questions about stories, the theoretical model [Lehnert '77] goes beyond this particular context. As a theoretical model QUALM is intended to describe general question answering, where question answering in its most general form is viewed as a verbal communication device between people.

While many of QUALM's question answering techniques are designed for answering questions about stories, QUALM is not limited to stories about a specific content domain. QUALM is applicable to any story that can be understood in terms of scripts and plans [Schank & Abelson '77]. This limitation is not content-specific; it is dependent on the general knowledge structures that are used in text understanding. When new scripts and plans are added to the knowledge base for SAM and PAM, questions can be answered about stories using this new knowledge without any additional alterations to QUALM.

BIBLIOGRAPHY

Cullingford, R. E. (1977). Organizing World Knowledge for Story Understanding by Computer. (thesis) Department of Engineering and Applied Science. Yale University, New Haven, Ct.

Goldman, N. (1975). Conceptual generation. In R. C. Schank, ed: Conceptual Information Processing. North Holland Amsterdam.

Lehnert, W. (1977). The Process of Question Answering. (thesis) Research Report #88. Department of Computer Science, Yale University, New Haven, Ct.

Riesbeck, C. and Schank, R. (1976). Comprehension by Computer: Expectation-Based Analysis of Sentences in Context. Research Report #78. Department of Computer Science, Yale University, New Haven, Ct.

Schank, R. C. (1975). Conceptual Information Processing. North Holland, Amsterdam.

Schank, R. C. and Abelson, R. P. (1977). Scripts, Plans, Goals and Understanding. Lawrence Erlbaum Associates. Hillsdale, N.J.

Shortliffe, E.H. (1974). MYCIN: A Rule-Based Computer Program for Advising Physicians Regarding Anti-Microbial Therapy Selection. (thesis) Stanford Artificial Intelligence Laboratory Memo-AIM251. Stanford University, Stanford, CA.

Wilensky, R. (1976). Using Plans to Understand Natural Language. Proceedings of the Annual Conference of the Association for Computing Machinery. Houston, Texas.

Woods, W.A., Kaplan, R.M., Nash-Webber, B. (1972). The Lunar Sciences Natural Language Information System: Final Report. BBN Report No. 2378, Bolt, Beranek and Newman, Inc. Cambridge, Mass.