DOCUMENT RESUME

ED 225 544 IR 010 549

AUTHOR	Lewis, Clayton; Mack, Robert
TITLE	The Role of Abduction in Learning to Use a Computer System.
PUB DATE	82
NOTE	11p.; Paper presented at the Annual Meeting of the American Educational Research Association (Los Angeles, CA, April 13-17, 1982).
PUB TYPE	Viewpoints (120) -- Speeches/Conference Papers (150)
EDRS PRICE	MF01/PC01 Plus Postage.
DESCRIPTORS	*Autoinstructional Aids; *Cognitive Processes; *Feedback; *Learning Activities; Learning Theories; Problems; *Word Processing
IDENTIFIERS	*Abductive Reasoning; Peirce (Charles S)

ABSTRACT

New users learning to use text-processing systems, especially those who are using self-study materials without expert supervision, must often make sense of actions and events in situations where they have little basis, in prior knowledge or current information, for inferring what is happening or why. Many cases have been observed where learners were able to generate an explanation to account for some particular fact or event, despite this limited knowledge. This process of explanation generation resembles abductive reasoning, as C. S. Peirce describes it: adopting a hypothesis when it, along with other assumptions, allows one to account for some fact or observation. This paper presents examples of how new users try to account for their experiences and the resemblance of these processes to abductive reasoning. Observations are also made about the possible role and implications of abductive reasoning for complex learning situations, such as learning to use a text-processor. Twelve references are listed. (Author/LMM)

The role of abduction in learning to use a computer system

Clayton Lewis
IBM UK Laboratories
Hursley House, Hursley Park
Winchester, Hampshire ENGLAND
and
Computer Science Department
IBM Watson Research Center
Yorktown Heights, NY 10598

Robert Mack
Computer Science Department
IBM Watson Research Center
Yorktown Heights, NY 10598

Abstract: New users learning to use text-processing systems must often make sense of actions and events in situations where they have little basis--in prior knowledge or current information--for inferring what is happening or why. Yet we have observed many cases where learners were able to generate an explanation to account for some particular fact or event despite this limited knowledge. This process of explanation generation resembles abductive reasoning, as C. S. Peirce describes it: adopting a hypothesis, when it, along with other assumptions one might have, allows one to account for some fact or observation. In this paper we present examples of how new users try to account for their experiences, and the resemblence of these processes to abductive reasoning. We also make observations about the possible role and implications of abductive reasoning for complex learning situations (like learning to use a text-processor).

Introduction

We have been studying how people learn to use text-processing systems. We have been especially interested in the problems people have when they are not familiar with computers, and must learn using self-study materials without expert supervision. In these studies we have asked people to "think aloud": to verbalize what they are thinking about as they learn. We find that new users are active learners, constructing interpretations of their training experience, and that these interpretations, right or wrong, form a large part of what they learn. In this paper we present a preliminary analysis of how these interpretations are formed. (See Lewis & Mack, 1982; Mack, Lewis & Carroll, 1982; Carroll & Mack, 1982 for further discussion of this research.)

An example of reasoning.

> E1: A learner was following instructions for a practice exercise which explained how to type in a letter. Part of the instructions required her to position the cursor at the beginning of various parts of the letter, such as date, heading, greeting, and closing. Because of flaws in the manual she skipped the instructions which covered the actual typing, and carried out only those which moved the typing point around to these various locations on the display screen. At some point she tried to to make sense of the fact that she had not really typed anything: she concluded that she was creating an invisible template into which the letter would later be typed.

What seems to be happening here is that the learner is constructing a story to explain what she believes she has been asked to do. The story consists of new assumptions which when added to the learner's existing belief allow the occurence of the happening to be deduced.

Forms of reasoning.

How does the learner develop the story about the template? What kind of reasoning is involved? Not induction, because the reasoner is not making a judgment about the truth (or adequancy) of a hypothesis by examining a lot of evidence and evaluating how well it supports the hypothesis. There is initially no hypothesis, and there is only one "observation" that might be used to support one.

Nor is the story a deduction: there is not some other premise from which it follows. Indeed, rather than being deduced from something else, the story seems designed to permit an observation, the contentless letter, to be derived from it.

The story seems to come from what Peirce (1958) calls abductive reasoning. In abductive reasoning, if some observation O is implied by an assumption A, together with beliefs already held, then A is adopted. In the example E1, for instance, the observation to be accounted for is the activity of moving the cursor around an empty typing display and not typing anything. The abductive inference is that the learner is typing a letter template which will be filled in later on: if one adopts this assumption then the observed activity is explained.

In its most self-conscious and systematic form, abduction is a companion process to deduction and induction, involved in the logic of discovery in science or other intellectual disciplines (see also Kinneavy, 1971). Abduction is the process by which a scientist generates a hypothesis to account for one or more observations. In the ideal case, new consequences are deduced from the hypothesis and tested against additional data. Inductive reasoning is used to make a judgment about the likelihood of the hypothesis being true given the accumulating evidence. But abduction is the first phase of reasoning in this development of knowledge because it generates the hypothesis (or new assumptions) from which further reasoning can take place.

<u>Additional Examples of Abductive Reasoning</u>.

We can make this description of abduction more concrete by considering three more examples of abductive reasoning in the context of learning text-processing.

E2: A learner was attempting to enter a password when a typing mistake caused the system to halt awaiting a correction. An indicator light marked "input inhibited" came on. The learner attributed both the delay and the light to a heavy work load on the system.

E3: A learner had made a mistake in issuing a "file" command, and wondered whether her work had been filed. Seeing a message on the screen which said

INPUT MODE 1 FILE

she concluded that it had been.

E4: A learner interpreted the sentence

Type four "blank lines" by typing in one space (hit SPACE BAR once) and pressing ENTER.

to mean

What I should do now is type four blank lines by pressing CARRIAGE RETURN four times and then type a space.

In E2, as in E1, the interpretation of the incidents in terms of abduction is straighforward: something has happened, and the learner develops a plausible (though in fact incorrect) explanation to deduce the occurence. In each case the explanation involves an elaboration of what might be called the <u>space of discourse</u>, introducing the idea of a template in E1 and ideas about system workload in E2. These new elements now form part of the learners' ideas about the system they are dealing with.

E3 illustrates a use of abduction that differs a bit from these. In E3 the thing that is being explained is the presence of a message that has the word FILE in it. But the real work that's going on is deciding whether the FILE command has filed the document. Here's how abduction can be used to do that.

The learner constructs a collection of assumptions, one of which is that the document was filed, and others of which indicate that the message refers to the outcome of the FILE command. Taken together, these assumptions "explain" the presence of the message. That is, if one made these assumptions, one would deduce the occurence of the message that did occur. Since one of the assumptions used was that the document was filed, the question of real interest has been answered.

E4 is another application of abduction. Here it is used to replace ordinary reading comprehension. The learner's task was to find an interpretation of the sentence

> Type four "blank lines" by typing in one space (hit SPACE BAR once) and pressing ENTER.

Her first attempt was something like this:

> What I should do now is type a space and press ENTER. This will produce four blank lines.

When she tried this, however, only one blank line resulted. She then tried other ways of producing blank lines, and discovered that pressing CARRIAGE RETURN would work. That leads to this interpretation:

> What I should do now is type four blank lines by pressing CARRIAGE RETURN four times.

But there is a problem here. This leaves the last part of the sentence unexplained: why does it say "by typing one space (hit SPACE BAR once) and pressing ENTER"? This can be accounted for, at least roughly, by the interpretation the learner adopted:

> What I should do now is type four blank lines by hitting CARRIAGE RETURN four times and then type a space (hit the SPACE BAR).

The argument here is that this is not an interpretation of the sentence that can be developed by ordinary linguistic analysis of the form of the sentence. The "by" just can't mean "and then". But the interpretation is a reasonable abduction from the occurence of the sentence in the manual. That is, if one assumes that what one should do is what the interpretation says, then the occurence of a sentence something like the one that is observed can be explained.

Carroll and Lasher (1982) have reported another kind of case in which the comprehension of text seems dominated by abduction. Their subject was learning to use a calculator from an instruction book. Having skipped some introductory material, the subject had formed misconceptions about the function of many of the keys. When she then encountered material that contradicted these misconceptions she concluded first that the offending material was misprinted, and finally that she had the wrong book of instructions for the calculator. In this case the reader accounted for the presence of material by invoking an explanation that avoided interpreting it linguistically at all.

Implications of abduction for the learner.

As the examples show, the learner who uses abduction suffers from it in a number of ways. Here are some that have appeared often in our study.

Not realizing anything's wrong. As E1 and E2 show, learners are often able to explain away the consequences of errors they have made. This is an especially serious problem for self-study learning, since the learner can't seek help for undetected problems.

Use of irrelevant information. New users do not always know what is relevant to some problem that they have. In this case they can be influenced by superficial resemblences between what they think they need and what they see or do. Example E3 is a good example. In fact, the message containing the word "file" had nothing to do with the FILE command, and only the coincidental use of the word "file" linked the message to the command. This kind of spurious appearance of relevance is very common among computer messages, displays, and instructions, and learners are often misled by it.

Partial interpretation of information. Learners in our situation carried out abductive inferences only as far as needed to account for something, without testing the abductions further. In E4, for example, abduction allows text to interpreted without requiring careful analysis of what it actually says.

It seems to be important to the use of abduction that learners do not require an abduction to explain everything that is observed. Rather, it has to explain something that is observed that doesn't have some other obvious explanation. In the case of the re-interpreted instruction the new interpretation is acceptable because it is able to explain a good deal of the observed sentence. An interpretation based on existing misconceptions may account for enough of even a flatly contradictory statement to appear acceptable. Carroll and Lasher's subject gives a clear though extreme illustration of this.

This characteristic of abductive reasoning is related to the observation that people tend to look for evidence that supports their beliefs, and not evidence that would disprove them (see, e.g., Nisbett & Ross, 1980; Wason & Johnson-Laird 1974.)

Abductions are often wrong.

Abduction is the generation of an implication, with the observation to be accounted for as the consequent and the hypothesis generated to account for it as the antecedent premise. When the hypothesis is affirmed learners are committing the deductive fallacy of affirming the consequent.

This is a common error in deductive reasoning (Johnson, 1972) and the fallacy occurs even in scientific reasoning (Kinneavy, 1971). What saves abductive reasoning in the latter case is that additional deductions are drawn from the abduction in order to test it. Unfortunately this testing was not applied to the abductions we observed, and as a result many of them are just wrong.

The value of abduction?

Despite these drawbacks, abduction has powerful advantages. Indeed, these are so strong that we believe the use of abduction is indispensable in complex learning. Abduction is a way of forming generalizations on the basis of very little evidence. Put another way, abduction allows learners to convert episodic information instantly into semantic information (Tulving, 1972). When the abduction is correct, this is a cheap way to generate new knowledge.

Consider the occurence of a message on the computer screen. On the face of it, all one can make of this is that at such and such a time this message appeared: one can code the episode that was observed. But that is in itself worthless: it has no value as a guide to the future, which is what learning is for. Using abduction, one does not bother to store the fact that the message occurred at such and such time, but rather produces an explanation of the occurence and stores that. For example, one may guess something like "when I issue a file command the message "INPUT MODE 1 FILE" appears. Again, such abductions may be wrong, but if correct they have real value in interpreting future events, in a way that simply remembering the event would not.

Often there appears to be no alternative to abduction. Many complex learning situations are characterized by incomplete and ambiguous information on the basis of which people are expected to form useful concepts. In such situations learners do not have a body of evidence on which to base systematic inductions, nor premises from which to deduce the structure of the task domain. If learners are ever to build up ideas about the processes that might lie behind what they experience in these situations, they have to use abduction.

How is abduction done?

Our learners developed their explanations of what was happening with amazing fluency, and yet with considerable aptness. That is, they could quickly produce an explanation that accounted for several different, and previously unrelated, aspects of a situation. A fuller description of the case sketched in E1 illustrates this.

The learner's idea of template not only explained the otherwise aimless activity of the botched exercise, but also related it to an earlier discussion of office tasks, in which filling information into a form letter was described. This connection makes the explanation of the exercise that much more attractive, since not only the exercise but the earlier discussion can be "explained". But how did the learner find that linking explanation?

Semantic network representations provide one way to think about how abductive inferences can be made: they provide a framework for representing people's beliefs about some knowledge domain, and a mechanism for searching out and creating new links. In this framework, the abduction about the template is linked both to the aimless cursor movement and to the earlier discussion about form letters.

How does one find a concept like this that links two existing points in the network? One mechanism is spreading activation, proceeding in parallel from both of the concepts to be linked. The intersection of activation from both sources outlines the available linking concepts. This will work, but it assumes that the points to be linked have to be identified at the start of the search. How can the learner determine what points are candidates for linking together?

This question can be avoided by positing a simpler search process, that starts from just one point in the network. One can trace out a plausible path from the botched exercise that gets close enough to the content of the earlier discussion that the link might be built from just one end, as follows.

The instructions that the learner was following were things like

Space over for the date.

The oddity is that the date was not then typed in. The wording "for the date" suggests that the date will be typed, so one can imagine that it will be typed in in the future. The instruction also says where the date will be typed, so one has the idea of a place where the date has not been but will be filled in. This is now very close to the content of the earlier discussion of which the learner was reminded.

Abstractly, then, an explanation might be built up by starting with an explanation of some one observed element, but then elaborating this by adding elements that are linked to those already included. If the structure grows in this way to include other elements that are actually observed in the existing situation (such as the recent discussion, in the example) the value of the explanation is increased. That is, search for relevant material would move outward from one starting point, but any encounter with other points that can be explained will be noted. The resulting explanations will therefore often tie together several different observations.

Starting from one point works here because the points that are to be linked aren't specified in advance. The learner is free to link up whatever is encountered, rather than having to develop an explanation of certain given facts. By contrast, in memory retrieval from a network, or pattern matching, the points that must be connected are specified, and more sophisticated processing is needed to find links between them.

<u>The role of feedback</u>. Feedback is relevant to abduction in at least two ways. First, abductions appear to be triggered by discrepancies between what learners expect and what happens. In all four cases discussed, the observation or fact to be accounted for corresponded to a discrepancy.

<u>What do people learn when they learn?</u>

The prevalence of abduction in learning suggests to us a different analysis of what people learn in a given situation than has been traditional. We want to distinguish an abduction-centered view from schema-centered views (e.g. Rumelhart & Norman, 1978; Rumelhart & Ortony, 1977) on the one hand, and older stimulus-centered views on the other.

Schema-centered views don't readily account for the ability of learners to develop explanations of events in a completely unfamiliar domain. Nor do they account for the cross-linking of different ideas that seem to characterize learners' explanations. Both these facts point to an active building process, not to identification of new material with existing structures, as the dominant process in learning of the sort we have observed.

On the other hand, it seems that the structures that are built aren't mainly representations of stimuli, even if "representation of stimuli" is interpreted broadly enough to encompass generalized forms like propositional representations of the meanings of sentences, labelling of objects, and the like.

Consider a learner who is given material to study that presents certain facts, say about a computer system. On our view, the learner develops and stores a tangle of explanations, including why the material has been presented, what there is about the computer that makes these facts pertinent, why the facts are important, and the like. In building this explanatory material the learner may hypothesize many things about the computer, its uses, and the learning situation that substantially extend the content of the material actually presented.

Somewhere in this tangle there may be some representation of the stimuli, that is, the facts actually presented, but it will be a minor component. Since the facts are imbedded in a tangle of explanations, we can expect accessing them later on to be complex, and to exploit the presence of deductive linkages in the tangle. A particular fact may not be stored at all, but may be reconstructed as an implication of the assumptions that were formed to explain and encode it. This may entail loss of information, because the explanation can explain lots of things besides the specific instance that led to its formulation. In effect, instead of facts being the building blocks of which more elaborate conceptualizations are formed, they may be rather special objects that are represented and retrieved only with difficulty, using machinery that works smoothly to store the general and the vague.

This view, along with many others in contemporary psychology, challenges our everyday conception of learning as a process of "taking in" some definite thing that is presented to us in the environment, say by a teacher or self-study manual. Instead, the thing that happens in the environment seems to act like the bit of debris that irritates the oyster into forming a pearl. The thing that we learn is the tangle of abductions, which may be very remote indeed from any officially sanctioned interpretation of what we were given.

References

Carroll, J. and Lasher, M. D. Getting to know a small computer. Manuscript, 1982

Carroll, J. M. and Mack, R. Learning to use a word processor: By doing, thinking and knowing. To appear in Thomas, J. & Schneider, M. (Eds.) Human Factors in computer systems. Norwood, N.J.:Ablex, in press.

Johnson, D. A systematic introduction to the psychology of thinking. New York; Harper and Row, 1972.

Kinneavy, J. A theory of discourse. New York; W. W. Norton, 1971.

Lewis, C. and Mack, R. Learning to use a text processing system: Evidence from "Thinking Aloud" protocols. Presented at the Conference in Human Factors in Computer Systems, March 15-17, 1982, Gaithersburg, Maryland.

Mack, R., Lewis, C. and Carroll, J. Learning to use office systems: Problems and prospects. Submitted to IBM Systems Journal, 1982.

Nisbett, R. and Ross, L. Human inference: Strategies and shortcomings of social judgment. Englewood Cliffs, New Jersey; Prentice-Hall, 1980.

Peirce, C.S. The logic of drawing history from ancient documents. In A. Burks (Ed.) Collected Papers of Charles Sanders Peirce. Cambridge, Mass.; Harvard University Press, 1958.

Rumelhart, D. and Ortony, A. The representation of knowledge in memory. In R. C. Anderson, R. Spiro & W. Montague (Eds.) Schooling and the acquisition of knowledge. Hillsdale, N.J.; Erlbaum, 1977.

Rumelhart, D. and Norman, D. Accretion, tuning, and restructuring: Three models of learning. In J. Cotton & R. Klatkzy (Eds.) Semantic factors in cognition. Hillsdale, N.J.; Erlbaum, 1978.

Tulving, E. Episodic and semantic memory. In E. Tulving and W. Donaldson (Eds.) Organization of Memory. New York: Academic Press, 1972.

Wason, P. and Johnson-Laird, P. The Psychology of Reasoning. Cambridge, Massachusetts; Harvard University Press, 1974.