

AI Working Paper 156

November 15, 1977
Massachusetts Institute of Technology
Artificial Intelligence Laboratory

Reporter: An Intelligent Noticer

by

Steven Rosenberg

Abstract

Some researchers, notably Schank and Abelson, (1975) have argued for the existence of large numbers of scripts as a representation for complex events. This paper adopts a different viewpoint. I consider complex events to have no fixed definition. Instead they are defined by a set of target components. At any given time an arbitrarily complex description which contains the target components can be generated from semantic memory. This description provides evidence for a complex event containing the target components. It can be as complex or as simple as the task demands.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. It was supported in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643.

Working Papers are informal papers intended for internal use.

Introduction

This paper develops a proposal for the design of an intelligent noticing program, REPORTER, capable of recognizing complex events from partial descriptions. It takes the point of view that complex events are composed of simple events, which in turn can be represented by at least one frame. A complex event will then consist of a set of linked frames. For example, the great Russian Wheat Fiasco of 1973 can be considered a complex event consisting of at least 1) a poor Russian wheat harvest; 2) the sale of U.S. wheat to Russia at low prices, and 3) a consequent rise in domestic wheat prices.

Complex events don't have precise definitions. Did the Wheat Fiasco start with the poor Russian harvest? Or with the introduction of forced collectivization which led to a history of marginal wheat crops? The answer will depend on our goals, and the extent to which we are willing to search for complex causes.

We define complex events as a set of linked frames representing the various component events, with at least one target event specified. For example, we might ask Reporter whether there is a complex event involving a rise in the price of wheat. Or we might ask for an event which involves both higher grain prices and a poor Russian wheat harvest.

The set of links among the components of the event provide supporting evidence for each component. Typically, we are interested in the occurrence of the target event, and the job of Reporter is to generate a complex event description linking the target event to known events. This complex event can be evaluated to provide evidence for the occurrence of the target event. Since not all the facts necessary to generate a complete description may be available, Reporter will make assumptions about missing data in order to generate links to other events which can be used as evidence.

While simple events are most easily represented as instances of canonical forms such as

Frames, complex events have variable descriptions. The specification of the target events, together with the web of links, provide a flexible descriptive mechanism for complex events which can be used as evidence for the unknown components of that event. As a trivial example, we assume it has rained if we notice the ground is wet, using a known causal link, together with the known fact, as evidence for the unknown rain event.

For any description of a complex event, many of the component events may be missing. Thus a theory of evidence based on partial event descriptions is required. If Reporter finds enough evidence among existing events and their links, we conclude the target event is likely. The model of understanding which underlies this theory of noticing is one in which new facts are linked to what is already known through the application of extrapolation rules. These rules support new data by generating a labeled link between new and existing knowledge. The label provides a connection between the linked pieces of knowledge and the underlying rules which generate these relations. Since these rules embody the causality of the domain, they provide a justification for the links they generate.

Theory

Intelligent noticing is useful for domains which are:

- A) Incomplete.
- B) New knowledge is continually being added, while current knowledge can become outdated.
- C) In assimilating new information not all necessary facts may be available at the time new information is added.
- D) Decisions may be deferred unless rules can generate assumptions about the missing data's form.

Suppose we wish to notice whenever a complex event happens in which the price of wheat will rise. Such an event might consist of a catastrophic Russian harvest, a plague of rats, a season of low rainfall, a sudden increase in the average amount of bread consumed, and so on. We

cannot limit the number of possible events which might provide evidence by using a canonical description or script. Instead, we must be able to examine the relation between any relevant new knowledge and the target event. This relation provides evidence. We would like to be able to quiz Reporter about target events, and get an evaluation of the target event in terms of the current state of knowledge. This complex event description will outline how currently known events and states of the world contribute to the target event, and what currently unknown facts are important in assessing the situation.

As new facts are provided, validating some assumptions, invalidating others, Reporter must be able to change its assessments. By leaving expectations active in the database when needed facts are missing, and when assumptions are made, Reporter will be able to re-evaluate supporting evidence when new data becomes available. When gaps in an incomplete database are encountered, the rules of the domain suggest what missing facts they require. Using these assumptions may lead to conflicts with other rules, requiring Reporter to decide if there is enough evidence for these assumptions to use them. When a known fact is changed by a later report, any links using this fact will notice the change, alerting Reporter to reassess any evidence based on this fact.

Evidence: A complex event description supports its component events through the web of links which form the description. The links can be used to evaluate the likelihood of the target event. If the target event is a rise in the price of wheat, and the new data is that the domestic crop is very small, we have evidence for a price rise, if we believe the rule that reduced supply causes a price increase until demand is once more in balance. The justification for a link is the rule which was applied to generate the link. These rules express our understanding of the underlying relationships in our domain. For instance, a poor Russian wheat crop supports the notion of higher wheat prices, through a rule which states that when the Russians do not have enough

wheat, they buy more American wheat. Another rule states that reduced supplies drive up prices. In order for a chain of links between a poor Russian harvest and higher domestic prices to exist, we assume that the intervening event of greater Russian wheat purchases postulated by the rule occurs, although this fact may not be available. Support may be provided through a chain of links of arbitrary depth.

Support: Links in event descriptions provide support through the evidence they supply that the linked parts are consistent with the causality of the world embedded in the rules. New data may conflict with this model. The price of wheat may not rise; instead people may eat more soybeans. In that case, a new event provides feedback which will allow us to revise our strategy for selecting rules. In some cases a rule may simply be wrong, and can be discarded.

Assumptions: such as the one involving greater Russian wheat purchases are open to confirmation or disproof by subsequent data. The greater the number of assumptions, and the greater the number of links in a support chain, the less sure we can be that this particular piece of data actually provides support. The chain of support is grounded in the actual piece of data which generated the chain. Notice that if a piece of data later occurs which grounds an intermediate assumption, the chain is shortened, and our confidence in the evidence is correspondingly strengthened.

Caveats: A poor Russian wheat crop only drives up prices if the domestic supply does not increase enough to fill the shortfall. Typically, one piece of data supports another conditional on certain assumptions about the rest of the world. When applying a rule to generate a link, we must be careful to check that the assumptions this rule is based on aren't violated. I will call such assumptions caveats. Caveats are annotations associated with rules. Thus a caveat of the rule that "wheat price increases are caused by increases in demand" compares the change in supply to the increase in demand, to determine if this difference has increased. If supply has also increased,

then the rule cannot be used.

A rule can have many forms, with some forms requiring less complete information, but having correspondingly more complex caveats. The most rigorous form of the rule in the above example might simply state that the difference between Supply and Demand must increase for prices to rise. This form has no caveats. (Unless we wish to consider the possibility of price controls!) Since we often do not know both these amounts, a less rigorous form of this rule would simply state "if Demand increases, Price rises". However, in this case it is necessary to have a caveat which checks that supply has not also increased. (Note that people will generally make such assumptions. For instance in the above example, few people would bother to check for changes in supply, on hearing of the poor Russian crop, before deciding prices will rise.) Finally, we retain the option, when using a particular form of a rule, of deciding it is not necessary to check its caveats.

Caveats check for violations of assumptions on which a rule is based. Detecting such a violation removes any support the application of the rule has given. Therefore caveats which cannot be applied due to a lack of pertinent facts must leave expectations which when triggered, modify the event description by withdrawing support. Caveats are useful precisely because they can leave expectations. This allows different forms of a rule to run under varying conditions of incomplete information. The caveats can also be used to generate reasons why a rule is not applicable.

Degrees of Evidence: Some evidence is "better" than other evidence. For example, a small Russian harvest supports the notion of an increase in wheat prices. Evidence that the Russians are chartering large numbers of grain cargo ships can also support this idea. Such evidence is less compelling, however, particularly if the ships might be useful for other purposes. The evaluation of evidence is complex, and will depend on many factors, which together form a strategy. Some

evidence is compelling because it is seen as causally necessary (e.g. a poor crop). Other evidence is compelling because it is typically the case. (e.g. the ship charters). Still another reason for accepting evidence is that to reject it means changing our rules, with consequences for future events. For example, in the past, at least one Russian leader, (Stalin), has let Russians starve rather than purchase wheat abroad. However, we now consider such an event to be unlikely, and instead use the rules introduced above, which do involve foreign wheat purchases. Changing these rules will have general consequences for the system, as they may be used in recognizing other complex event sequences. To the degree we have confidence in the assumptions the rules embody, the evidence is strong. Changing these assumptions can be difficult unless they are made explicit. Consequently we may wish to encode them as rationales associated with the rule, as in Goldstein and Grimson's annotated production systems (Goldstein & Grimson, 1977).

Strategies: Different people may have partially different rule sets. These differences express different strategies. For example, in a jury trial, one juror may have a rule that decides that the testimony of a witness should be disregarded if there is evidence that he is unreliable. Another juror may have a rule for such cases which involves looking for independent verification of the testimony. Still another juror may try to find evidence for a compelling motive for the witness to be unreliable in this instance. Such rules embody a particular strategy for the generation of evidence. These strategies can be encoded as sets of rules of the same form as any other rule. These rules will be triggered by specific information, such as evidence of a witness's unreliability. They will then operate on the functioning of other rules, and change the links previously generated by these modified rules. They allow Reporter to perform according to different strategic goals.

Control: When we wish to generate a complex event which contains a great deal of links, we must be able to evaluate the evidence produced. There are several ways of evaluating evidence: a)

the depth of the support chain; b) the type of rule; c) the number of chains of support; d) the number of assumptions made; e) the support for these assumptions; f) whether applying the rule required many caveats; g) the total number of caveats; h) partial evidence; and i) degrees of evidence, such as 1) causally strong evidence; 2) typical evidence.

Each piece of new knowledge is considered a target event for the purposes of linking it to the database. However, we do not want to generate very complex descriptions for target events we are not interested in. Hence we will consider these as "target events" with which we will be satisfied when one link is created.

Cost: There is a cost to each piece of support. If a chain of support is extended to a great depth it may be too expensive relative to the degree of evidence provided. For example, to support the target event of wheat prices rising, Reporter tries to find evidence for an increase in demand. To support this, it looks at increased (foreign) purchases. Not finding direct evidence for this, it tries to find support, in turn, for foreign purchases, by looking for evidence of the Russians accumulating the hard currencies necessary to purchase wheat. If no direct data on this exists, support is given by the fact that the price of gold has dropped recently. This, in turn, supports the notion of increased gold sales by the Russians. Thus, the drop in the price of gold supports, through several intermediate links and assumptions, the notion that the price of wheat will rise. The cost of generating such a complex event description may outweigh the importance of the support it gives.

Active data search: The ability to decide which missing information is most useful for current goals can be used to create requests for information from various sources. The goal of completing partial event descriptions allows purposeful requests for knowledge from the system.

Expectations: When assumptions are made, or data is found to be missing, expectations can be left which, when the necessary facts do occur, will be able to fill in the gaps. Such expectations

maintain consistency and allow for partial evidence to be completed. Given a target event, we would like Reporter to be able to automatically adjust the evidence as the facts available change.

Payoff function: A classic problem in psychology consists of deciding what level of false alarms we are willing to tolerate to ensure always getting alerted whenever the real event occurs. This has been a problem in training radar operators. They often notice blips which then disappear. Is it worth staging an alert, if the blip is a gremlin? Obviously, if we are worried about nuclear attack, we will tolerate a high false alarm rate. This issue manifests itself when we need to decide how many assumptions, and what level of support, we are willing to accept in a complex event description. Reporter will be adjustable to different sensitivities through an evaluation function which examines each piece of generated evidence before linking it to the rest of the event description.

An Example Scenario

In the following scenario, I shall develop the wheat price increase example, and use it to outline a simple system, capable of propagating chains of support and using expectations. If we give Reporter the following sentences as target events we might get:

(S1) Wheat prices rose.

Alert: wheat supply is probably not responsible for the price rise, unless we assume very poor rain fall. It is more likely a rise in wheat demand has caused the price rise, but the evidence is based on many assumptions.

(S2) There was no change in weather conditions.

Alert: since rainfall has not decreased, wheat supply has probably not changed. Expect therefore that the wheat price rise is due to rising wheat demand, but consider that this still involves several assumptions.

The system is first given (S1). Events are encoded as Frames using FRL (Goldstein and

Roberts, 1977). The representation for (S1) is:

```

Price-increase
ako      =      thing
from     =      wheat-price1
          $require :from less-than :to
to       =      wheat-price2
          $require :from less-than :to
rule     =      Price-increase1-rule
              Price-increase2-rule

```

```

Wheat-price
ako      = price
price    = $x
date     = previous

```

```

Wheat-price
ako      = price
price    = $y
date     = now

```

There are no existing expectations waiting to be triggered, since Reporter has not previously attempted to notice any target events. Nor does this sentence directly conflict with any known facts, as we have no other values for the current price of wheat. The task is to generate a complex event description which supports this target event, allowing it to be added to the database.

To support (S1) we must apply at least one appropriate rule. Frames will have a rule slot which contains rules relevant to supporting instances of that frame. Rules are themselves represented as frames. The rules associated with this slot of the Price-increase frame are the Price-increase1-rule and the Price-increase2-rule. The Price-increase1-rule illustrates all the slots of a rule frame. For brevity, other rules are shown only with those slots which are relevant to the example, generally the support, alternate and alternate slots.

Price-increase1-rule

ako = rule
 test = increase-difference-test
 caveat = nil
 support = assume-increase-difference
 alternate = supply-decrease-rule
 demand-increase-rule
 exists = increase-difference-args
 rationale = \uparrow (supply - demand) \Rightarrow \uparrow price

Price-increase1-rule

ako = rule
 alternate = Less-selling-rule
 More-buying-rule
 rationale = \uparrow (amount-offered-for-sale - amount-offered-buy) \Rightarrow \uparrow price

Rule frames will have the following structure:

1) A test slot. The value of this slot is a function which computes the relation embodied in the rationale. There are three possible results. First, the test may succeed. In this case the test supports the new knowledge, and we link the new knowledge through a link frame which also specifies the rule involved. (The test that Price-increase1-rule specifies, "increase-difference", tests for an increase in the difference between the amount of supply and the amount of demand.) The test may fail. In this case we have evidence against the desired event. The link for this indicates the negative nature of the evidence. Lastly, we may not be able to apply the test because the all data it requires has not yet been added to the database. In this case we wish to leave expectations for this data, which are triggered by its occurrence, and can resume running the test. In the meantime, however, we wish to be able to seek further evidence by using an alternate form of the rule, or by making the necessary assumptions required to apply the test.

2) A caveat slot. The value of this slot is a predicate which checks that the assumptions of the test are not violated. For instance, suppose that a test (in a rule which is being used to support an increase in the price of wheat) determines if Russia has recently chartered a large number of ships. If the test succeeds, the price rise is supported, since the ships are necessary to transport the

wheat. The caveat for this rule checks that there is no evidence the ships will be used to transport another commodity. Caveats are a type of rule, and like rules may leave expectations when there is not sufficient data. For instance, if supply of a material decreases, a price increase is supported, with the caveat that demand does not also decrease. While demand is unknown, the decrease in supply supports the price increase. When demand is known, the caveat determines whether the test still applies. Whenever a test is applied, the caveat, if it exists, can also be run. Caveats are themselves rules. Consequently the value of the caveat slot is a set of rules whose success provides evidence against the evoking rule's assumptions. These assumptions provide a context within which a rule's test is valid evidence. Caveats in Reporter perform only this limited subset of the caveat functions in APS (Goldstein & Grimson, 1977).

Caveats represent the causal assumptions on which our rules are based. Some researchers (Reiger, 1975) have attempted to classify the types of caveats which exist, such as sustaining causes, enabling causes, and so on. I prefer to allow the user the discretion to encode whichever pre-requisites are relevant to a particular rule.

3) A support slot. If a test cannot be applied because of incomplete data, the value of the support slot can be used to determine if we can assume the existence of the missing information. The value of this slot is a function which will determine which information is missing, and attempt to add this missing knowledge to the database. This causes rules to be run to support these new assertions. If the missing knowledge can be supported by these rules running successfully, the support functions will re-invoke the test. They will also leave expectations which, should the assumed knowledge occur, remove the assumed evidence, and determine any changes in the support caused by the actual knowledge.

The process of supporting assumed data is identical to the process of supporting the original data. This process may in turn require further assumptions. We are in effect creating a series of

sub-goals in order to build a support chain. Successive uses of support slot values build chains of support with each link requiring further assumptions. Each rule, if it cannot directly support or contradict by running a test function, can pass the buck and set the goal of finding enough evidence to achieve its own goal of applying a test. These support chains can be arbitrarily deep. New data may satisfy a test anywhere in such a chain, shortening its depth, and the number of assumptions which must be made.

4) The Alternate slot: The value of this slot are less powerful forms of the rule, which require less information, but more complex caveats. For example, the Price-increase-rule has a test which looks for an increase in the difference between supply and demand. Less rigorous forms of this rule are contained in the alternate slot. The Supply-decrease-rule will look for support by checking for a decreased supply, while the Demand-increase-rule checks for an increased demand. These rules implement a notion of cost and degree of evidence. The secondary rules are those which do not provide as good support, since they require further caveats, but are more likely to run, requiring less information. These alternate rules may also have an alternate slot value. The alternate forms of a rule are ordered in terms of the increasing number of assumptions required. The most rigorous forms are the first to be applied. When we use the alternate forms, the more rigorous form may still have its expectations fulfilled, allow it to run at a later date. When this happens support generate by the alternate forms is removed. Note that if either alternate rule failed, this would not provide conclusive evidence against the likelihood of the price rise, since the other rule may provide support. The strong form of this relation, expressed in the original rule, provides this kind of powerful evidence. There are degrees of negative evidence, depending on the existence of alternative positive support.

Supply-decrease-rule

ako = rule
 test = decrease-wheat-supply-test
 caveat = decrease-wheat-exports
 support = assume-supply-decrease
 rationale = ↓ supply => ↑ price

Demand-increase-rule

ako = rule
 test = increased-exports-test
 caveat = increase-wheat-supply
 support = assume-export-increase
 rationale = ↑ demand => ↑ price

5) The Exists slot contains a function which checks for those frames whose existence is required by the test. This allows us to distinguish cases where a test fails because of missing data from those cases where the test ran, but failed.

6) The Rationale slot provides an explicit description of the rule which can be used in generating explanations.

Links: Support is indicated in the database through a link. A link is a frame which points to the frames supported, the rule used, and the frames support depends on. It has slots to indicate whether the test supported or unsupported, whether assumptions were made, and whether the caveat applied. These links provide the only permanent record of support chains. Removing or modifying support requires removing or modifying these frames.

Let's pursue our example, and consider how the support chain might continue using the two alternate rules. Both increased demand and decreased supply provide support. Consider decreased supply. Supply-decrease-rule first tests for the existence of a supply-decrease frame for wheat. The caveat on this rule tests that there is not a decrease in wheat exports. If we cannot find a decrease in wheat-supply, we try to assume its existence. The support slot value will try and find enough evidence to assume the existence of a Supply-decrease frame for wheat.

Supply-decrease

ako = thing
 domain = wheat
 rule = Decrease-soil-moisture-rule

To support a decrease in the supply of wheat, the Decrease-soil-moisture-rule looks for a decrease in soil moisture.

Decrease-soil-moisture-rule

ako = rule
test = decrease-soil-moisture-test
caveat = increase-irrigation
support = assume-decrease-rainfall
rationale = ↓ soil-moisture => ↓ supply

Since our data base is very simple, we do not know the soil moisture. If this knowledge is not available, we must support this new assumption. A decrease in soil moisture is supported by a decrease in rainfall. Such weather events are not supported by any other data, but are observational primitives. We cannot continue to find antecedents for support. However, we can ask what would be some of the consequences of less rain having fallen, and place expectations on these, since we may learn this information before learning about the weather. One consequence is a reduced crop, which was the assumption in support of which we looked for reduced rain. Consequently, in mixing antecedent and consequent rules, we must be careful not to get into loops.

The other alternate rule, Demand-increase-rule, tests for an increase in wheat exports. It has a caveat that checks that wheat-supply does not also increase. If this test does not run due to lack of data, the support slot indicates supporting an increase in exports. In our model we have just one foreign purchaser of wheat, Russia. For Russian demand to increase, Russia must have a decrease in local supply. Thus the rules used in support of this assumption will mirror the ones just described, except that they will apply to Russia's production, and not domestic production.

Either alternate rule succeeding provides the same support for the evoking frame. Suppose none of the support chains built so far finds support, since none of the relevant data is available. These chains will have left many expectations around, but will be unable to provide support.

We now input the second sentence, (S2). (S2) states that the weather has not changed. This fact, when added to the database, will trigger a sentinel watching for decreased rainfall. This sentinel will test the new weather fact, and find that the test fails, since rainfall has not decreased.

Thus the chain of support for decreased supply cannot be maintained. This leaves only the alternative rule which is checking for increased demand still active. Because we have two alternatives, the failure of one does not provide evidence against the price rise. It merely removes one possible chain of support.

One result of chaining the forward consequences of the non-change in weather is that we replace the old expectations for decreases in soil moisture and wheat-supply with new ones expecting unchanged values. If these new values are changed, we then have a conflict generated by the rules which link two observed values. Real world feedback will cause us to improve rules. Some things, like weather, although they have consequences which can be searched for, will have no antecedents. Similarly, given limits on our semantics, some facts will have antecedents but no consequents.

Sub-goals: Propagating support chains implicitly involves goals and subgoals. We apply rules to achieve our support goals. A busy noticing system might have many things it wishes noticed. Over time, some of these might be dropped. How do we clean up the chains of expectations? It would be costly to garbage collect through the support chains each time a goal changed. An alternative is to propagate the top level goal down any chains created to support it. Each expectation on firing, would test that this goal still existed, if not it would abort, removing itself. Thus, only those expectations are removed which actually get triggered.

Conclusion To conclude, Reporter's central problem is coping with the potential explosion of theories for interpreting complex events in large domains. Consequently the ability of Reporter to control the computational explosion will determine its success. Understanding whether the theory of evidence formulated in this paper is adequate to control this explosion is the underlying issue. Our plan for development and testing Reporter is as follows:

Plan

1. Expected duration: September 1977 through November 1977.

Goals

- A) Designing a format for rules and adding these to the system.
- B) Developing a set of sentinels capable of leaving expectations for running tests, creating supports, and handling caveats.
- C) Developing a way of maintaining the database when tests fail, caveats succeed, and expectations run. This will involve being able to modify, remove and shorten chains of evidence.
- D) Designing a control module so that new data can be automatically linked to the database, and so that when requests for noticing occur, we are able to make use of the measures of cost such as number of assumptions, caveats, depth, and so on, discussed previously, to control noticing.
- E) Develop methods for handling negative support.

2. Expected duration: December 1977 through January 1978.

Goals

- A) Applying the system to the symbolic correlation testbed.
- B) Developing a more sophisticated control language which will allow different strategies of noticing corresponding to the different strategic roles in the symbolic correlation testbed.
- C) Design a simple way for a user to interrogate Reporter concerning the state of the world and particular events.

3. Expected duration: February 1978 through March 1978.

Goals

- A) Integrate Reporter with Stansfield's Strategist.
- B) Develop the communications channel between Reporter and Strategist.

4. Expected duration: April 1978 through July 1978.

Goals

- A) Test and evaluate.

References

Goldstein, Ira P. and Grimson, Eric. 1977. Annotated Production Systems: A Model for Skill Acquisition. MIT-AI Memo 407.

Reiger, Charles. 1975: The Commonsense Algorithm As A Basis For Computer Models Of Human Memory, Inference, Belief And Contextual Language Comprehension. Proc. TINLAP.

Schank, Roger, and Abelson, Robert p. 1975. Scripts, Plans, and Knowledge. in Proc. 4IJCAI