

Michael Lebowitz

Columbia University, Department of Computer Science

406 Mudd Building, New York, NY 10027

ABSTRACT

The parts of a typical piece of text vary greatly in interest. Presented in this paper are three ways a concept can fail to be interesting — it can be irrelevant, reconstructible or overshadowed. The uses of interest in understanding are also discussed.

1. Introduction

Not everything we read is interesting. Even within a basically interesting piece of text, some concepts will seem much more important than others. Some stories will carry so little interest that they can be ignored altogether. A key issue in studying the variable nature of interest is to understand why certain pieces of text are not interesting. After all, a priori, we might expect anything included in a story to be worth our full attention.

As a simple example of varying interest, consider story SI.

SI - UPI, 28 July 1980, Lebanon

Unidentified gunmen Monday ambushed a pro-Iraqi politician, riddling him and his bodyguard with bullets, in the latest of the region's almost daily political assassinations.

Police sources said Mussa Shueib, a top official of the Lebanese chapter of Iraq's ruling Baath party, was driving from Beirut airport when his car was overtaken by armed men and sprayed with gunfire in a gangland-style hit.

Shueib and his bodyguard died on the spot.

In SI, the death of the bodyguard would usually not be considered to be as interesting as the killing of Shueib (it is overshadowed). Neither is the method of the killing as a shooting too interesting, as it can be reconstructed from general knowledge of terrorism. Several other points are simply irrelevant. Having dispensed with the non-interesting, we are left with the interesting concept of an important pro-Iraqi politician being killed in Lebanon.

It is important to recognize that interest depends upon the current state of memory. Interest is highly idiosyncratic from reader to reader and from time to time. Concepts interesting to one reader at a given time may not be to other readers or at different times. For example, a reader who knows Shueib's bodyguard personally would find the pieces of SI about him quite interesting.

A computer understanding system will require long-term memory, such as that of IPP (the Integrated Partial Parser) [2, 3], to be able to make use of interest. However, with this memory, it can use interest to help decide what to remember and to guide the understanding process. It was in the context of IPP that the ideas introduced here were developed. In Section 3 I will describe the uses that interest can be put to.

Generally, a testing process that concentrates on recognizing non-interest will be quite effective. For example, an algorithm that recognizes an instance of non-interest can immediately reduce or terminate its processing effort. The alternative, however, involves assuming that an item can be ignored only when it does not satisfy positive criteria for interest, a test that may produce an answer too late to be useful.

2. The Nature of (non) Interest

Interest might seem to be an arbitrary and hard-to-define idea. However, some confusion is simply due to the fact that there are several distinct ways that a concept can be (or can fail to be) interesting. Figure 1 lists three ways a concept can fail to be interesting.

Figure 1: Grounds for non-interest

1. Concepts not causally connected to the main point of the story.
2. Concepts that can be reconstructed.
3. Concepts that are overshadowed.

Lack of causal connection (see [4]) and reconstructibility (see [1]) explain a large number of situations where a piece of a text is not interesting. These are in some sense "logical" reasons for finding a concept non-interesting.

Here I will just present an example of each case.

52 - UPI, 21 June 1980, El Salvador

A right-wing terrorist group called the Squadron of Death killed 10 men Saturday including a labor leader and three others shot to death as they ate breakfast in a restaurant, police said.

In story S2, the fact that the victims of the terrorist shooting were eating breakfast is not connected to the main causal chain of the story, and hence not interesting. This part of the story is not important to our overall understanding, and will be quickly forgotten.

53 provides an illustration of reconstructive concepts•

S3 - UPI, 20 June 80, Guatemala

The terrorists who kidnapped a Nestle Corp executive said Friday he will be released only if the swiss food firm comes up with an undisclosed ransom and pays for the publication of a terrorist manifesto.

In this story, the terrorists asking for ransom is quite typical of kidnappings, and so can be easily reconstructed from the simple knowledge that there was a kidnapping. As with the irrelevant part of S2, the reconstructibility of the ransom demand of S3 makes it not interesting. The demand for ransom will not be long remembered, nor need it be processed in detail.

The third idea in Figure 1, overshadowing, lacks the obvious rationale of irrelevance and reconstructibility. Nonetheless it plays an important role in interest evaluation. Consider the following story.

S4 - UPI, 29 July 1980, Madrid

Leftist gunmen ambushed an army general in the center of Madrid today wounding him but killing a soldier in his escort, police said. Unofficial reports said one of the general's aides and his chauffeur also were wounded in the early morning attack.

Normally the killing of the unnamed soldier would seem less interesting than the wounding of the general and the various injuries less interesting still. However, this cannot be explained by either reconstructibility or non-causality. The death of the soldier and the wounding of the aides are certainly causal results of the attack. Furthermore, the results surely cannot be reconstructed from memory. So why is it that the death and other injuries are so much less interesting?

Stated simply, the attack on the general overshadows the other events. It is not that a soldier being killed is inherently non-interesting but that the attack on the general is more interesting. Even though an event might be interesting by itself, most processing will be devoted to an overshadowing event.

A concept can have sufficient interest to overshadow for at least two distinct reasons. It can be either important or unique. For example, the event, "Terrorists shot the President," is interesting because the victim of the shooting is important. This is contrasted with an event like, "Three clowns fired shots into a group of spectators," where nothing in the event is very important, but it is a relatively unique event.

Importance simply represents a concept for which the reader has a complex structure in memory. This is extremely idiosyncratic and dependent upon the current state of memory. Important concepts for a given person might be his best friend, the President of the United States, and the status of his favorite football team.

Uniqueness, in contrast, represents a concept that is different from others in memory. It can be used to explain the interest of a shooting by a child, an explosion aboard a space ship, or the US hockey team's upset of the Russians. A key problem to be studied involves figuring out just how different an event must be before it seems unique (after all, in some sense all events are unique).

The lack of either importance or uniqueness will quite likely result in a concept being overshadowed. It is the specific analysis of just what determines importance and uniqueness that will allow interest to be fruitfully applied to understanding.

3. The Uses of (non) Interest in understanding

Interest has effects on both long-term memory for stories read and the parsing process. We have already seen that interest interrelates with memory. Almost by definition, the parts of a story that are interesting are those that are remembered. This is an important control on the accumulation of information in memory. Particularly important is the ability to omit material that can be reconstructed from memory, thus drastically reducing the quantity of information that must be stored. (See [1, 2] for more on the storage of redundant information and the reconstruction process.)

The analysis of interest described in this paper does make one important addition to the relation of interest and memory. The idea of overshadowing will have an important effect on what is remembered. As I will mention below, resources are usually not expended on detailed processing of the less interesting aspects of a story. This will normally

result in such items being lost and not saved in memory. While this result is not clearly desirable for a computer memory system, it may be necessary for the same reason it occurs in human understanding — to allow the corresponding increase in efficiency of processing.

So we see that the basic effect of interest on memory is a simple one — interesting items are remembered, non-interesting ones are not. However, the effects on direct processing are more subtle. I will mention here some of the possibilities.

The primary role of interest in parsing is to help control processing and allow the optimal allocation of resources. This has been discussed to some degree in [6] and [5]. Here I will discuss one very powerful way interest can be used in an understanding system designed to process large numbers of texts.

The simplest use of interest in understanding is to make a decision to stop processing a sentence (or a story) that does not appear to be interesting. There is an obvious advantage in doing this. Stopping the processing of a sentence or story clearly gives the maximum possible gain in efficiency - processing time is reduced to zero for the rest of the item! And this is clearly sufficient, since as was noted before, low interest items usually need not be remembered at all. This use of interest will give computer understanders an ability, that people clearly have — the ability to stop reading something. Informally, a program employing the rules that I will describe below will have the ability to become bored.

The basic idea is simply to track interest at both a sentence and a story, beginning at uncertain and continuing until a judgement can be made as to whether the item is or is not interesting. This is done using the rules described in this paper. If the conclusion is drawn that a story or sentence is not going to be interesting, then processing of that item can be stopped.

It would seem at first glance as if this plan should not work at all. One of the rules that I have described differentiate between elements of a story that appear at the beginning and those that appear later. Thus it would seem possible that the items that make a sentence or story significant could occur at very end. Using the rule described in the last paragraph, these items would never be processed. Fortunately, in practice the key items do usually occur early on, presumably because authors know that rules such as the one above will be applied. As an example of this, we would be much more likely to find a story like, "The President was shot ..." rather than, "In an official limousine today, while on a motorcade through Dallas, a gunman in the Book Depository, who was later apprehended by police, shot the President."

Out interest heuristics will sometimes fail if the key items are idiosyncratic — not those the

author would consider significant. For example, in a story about a building takeover where there were many hostages, the identity of the "low level" people taken hostage would not be found early in the story, but it might make the story significant to a reader who knew one of these people. However, that is exactly the type of situation where a human reader would miss the key information. Some sacrifices must be made to achieve adequate performance.

4. Conclusion

In this paper I have just touched the surface of the problem of interest. However, it is an important problem that justifies further effort. In particular, an understanding of the purpose and mechanism behind overshadowing will allow better control of natural language understanding programs.

REFERENCE

1. Kolodner, J. L. Retrieval and organizational strategies in conceptual memory: A computer model. Tech. Rept. 187, Yale University Department of Computer Science, 1980.
2. Lebowitz, M. Generalization and memory in an integrated understanding system. Tech. Rept. 186, Yale University Department of Computer Science, 1980. PhD Thesis
3. Lebowitz, N. The nature of generalization in understanding. Proceedings of the Seventh International joint Conference on Artificial Intelligence, International Joint Conference on Artificial Intelligence, Vancouver, Canada, 1981.
4. Schank, R. C. The structure of episodes in memory. In D. Bobrow and A. Collins, Eds., Representation and Understanding: Studies in Cognitive science, Academic Press, New York, 1975.
5. Schank, R. C. "Interestingness: Controlling inference." Artificial Intelligence 12, 3 (1979), 273 - 297.
6. Schank, R. C, Lebowitz, M., and Birnbaum, L. "An integrated understander." American Journal of Computational Linguistics 6, 1 (1980), 13 - 30.