

Clause-internal coherence

Jerry R. Hobbs
University of Southern California

1. Introduction

About twenty years ago I was in a jogging group and was known in that group for my eagerness to take shortcuts. Our standard route included a stretch on a dirt road that became deep mud when it rained, but we had a longer detour on pavement for those days. One day after a rain, as we approached this stretch, I said,

(1) Let's take the muddy way.

Everyone laughed.

What's funny about this?

If I had said, "Let's take the short way," it would have been understood that the shortness used in my description also provided a motivation for taking that route. Since I used the word "muddy" instead, it seemed as if I wanted to run that way *because* it was muddy, as if I enjoyed sloshing through the mud.

Very often different parts of a clause are connected inferentially in this way even though there is no direct syntactic connection between them and even though there is no explicit lexical signal of the relation.

Among the relations we find are the kinds of coherence relations we find between successive clauses or larger stretches of discourse. In Hobbs (1985) and other papers, I present an account of discourse structure in which the coherence relations are seen as possible interpretations of the adjacency of segments in the discourse. For example, in

(2) Chris and Pat are studying. Don't bother them.

we not only need to figure out who "Chris", "Pat", and "them" refer to, what it would mean for them to study, what it would mean to bother them and to not bother them. We also need to figure out what is conveyed by the blank space between "." and "D". We need to explain what relation makes these two sentences a part of the same discourse.

The typical relations that are conveyed by adjacency of clauses, as argued in Hobbs (1985) include causal, interlocking change-of-state, similarity, and contrast

relations. Clauses describe events or situations, and these are some of the most common kinds of relations we express among events and situations. We see that the same set of relations can occur between different parts of a single clause, in a way that goes beyond the predicate-argument relations conveyed by syntax.

Consider another example:

- (3) A jogger was hit by a car last night in Marina del Rey.

Our assumption is that the person was jogging when the accident occurred, and that somehow played a role in the accident. Contrast this with

- (4) A professor was hit by a car last night in Marina del Rey.

Here we don't assume the professor's being a professor had anything to do with the accident. Professing is not an outdoor, on-street activity. In the sentence

- (5) We should listen to the warnings of scientists.

there is an implicit causal relation between the listening and the scientists. It is *because* they are scientists, and therefore know what they are talking about, that we should listen to them. In

- (6) Kids sometimes show great insight.

there is a violated expectation relation between the kids and insight. Normally children do not show great insight, and part of the message of this sentence is this violation of expected causality.

I will refer to events, situations, conditions, states, and so on as "eventualities". I have argued that nearly every morpheme in a sentence conveys some eventuality. For example, the sentence

- (7) Sugar is sweet.

describes three eventualities, first, that some entity x is sugar; second, that x is sweet; and third, that x 's sweetness holds at the present time.

Clause-internal coherence is the phenomenon of coherence relations holding between eventualities conveyed by morphemes in the clause, beyond the predicate-argument relations that are signalled explicitly by syntax.

Kronfeld (1989) discusses something like this problem under the heading of "conversationally relevant descriptions". In his account, in a sentence like

- (8) The city with the world's largest Jewish community welcomes Israel's Prime Minister.

the reader recognizes that the definite description is not as brief as it could be, decides it must be conversationally relevant, and interprets it as containing an

implicit universal—any city with the world’s largest Jewish community must welcome Israel’s Prime Minister.

Kehler et al. (2008) use the expectation of clause-internal explanations to manipulate the resolution of attachment ambiguities in sentences whose main verb tends to invite an explanation. In a sentence that begins

(9) John detested the servant of the actress who ...

the reader expects an explanation of John’s detesting the servant, and hence expects the relative clause to modify “the servant” and to provide the explanation, rather than expecting it to attach low to “the actress” as the ordinary syntactic default would indicate.¹

In Section 2 of this Chapter I briefly describe the “Interpretation as Abduction” framework that provides a unified approach to a large number of linguistic phenomena. It contains a position on what counts as an adequate interpretation of a text, which clause-internal coherence challenges. In Section 3 I show how some cases of clause-internal coherence can be handled in the course of validating explicit linguistic signals or as examples of coreference. However, there are other examples, mostly involving similarity and violated expectation relations, that cannot be handled by these means. In Section 4 I consider a number of other examples from several different genres. Section 5 summarizes what we can learn from this investigation.

2. The Abduction framework

The key idea behind the “Interpretation as Abduction” framework (Hobbs et al. 1993) is that we interpret the world by coming up with the best explanation for the observables in our environment. The process of abduction is proving the thing to be explained deductively where possible, making assumptions where necessary, and deciding among alternative possible proofs by means of some measure of economy in proofs. When applied to interpreting language, the observable is the text itself. We need to explain why this text occurred. This divides into two subquestions: what information does the text conventionally convey (the informational perspective) and why did the speaker or writer want to convey this information (the intentional perspective).

Texts convey meanings via words (or morphemes) and via adjacency relations. Within individual sentences, the adjacency of words or larger stretches of text convey predicate-argument relations, and the recognition of syntactic structure

1. I am indebted to Andrew Kehler for drawing my attention to both his and Kronfeld’s work.

is precisely the discovery of these relations. Thus, in the sentence “Pat works,” the word “Pat” tells us that there is an entity x_1 named “Pat”. The word “works” tells us that an entity x_2 works. The interpretation of the adjacency of the two words as conveying the Subject-Verb Phrase relation tells us that x_1 and x_2 are the same entity; Pat is the one who works. Thus, conceptually, the first step in the informational analysis of the sentences in a text is the discovery of the syntactic structure and the corresponding logical form.

One then tries to find the best abductive proof of the logical form of the sentences in the text, where the criteria for what is best include the following:

- Short proofs are better than long proofs.
- Salient axioms are better than nonsalient axioms.
- The fewer assumptions the better.
- A proof is better if it exploits implicit redundancies in the text.

The last of these is important in coreference resolution. One often gets a better proof by assuming that two things mentioned in different parts of a text are in fact the same, or are inferentially related. In the sentence

(10) I just bought a used car. It gets good gas mileage.

the most economical explanation happens if we assume “it” to be the same as “a used car”. This is an example of *direct coreference*. The two references are to the same entity. In the sentence

(11) I just bought a used car. The tires are worn.

we get the best interpretation if we assume the tires are the tires of the used car, and not some random unrelated tires. The existence of the car proves the existence of the tires. This is an example of *indirect coreference*.

Sometimes we get the best interpretation by recognizing only a partial match. In

(12) I talked to my mechanic. She said the engine is in good shape.

both “mechanic” and “she” refer to persons, but both carry properties that are not contained in the other—occupation and gender. Assuming the persons implicit in each are the same gives us the best proof of the logical form and resolves this direct coreference by the partial overlap in their meaning. In

(13) My tires are worn, but the engine is okay.

neither the tires or the engine implies the existence of the other, but if we assume a common core of “car”, then the existence of both can be inferred. Again there is a partial overlap in meaning, and recognizing this allows us to resolve the indirect coreference.

Because abduction finds a proof of the logical form by backchaining, the interpretation is more specific than the explicit content of the sentence. The disambiguation of word sense ambiguities is one example of this. In the sentence

- (14) The plane taxied to the terminal.

the plane could be an airplane or a wood smoother, taxiing could be a plane moving on the ground or someone riding in a taxi, and a terminal could be an airport terminal or a computer terminal. But by identifying the common core of the first reading of each, we get a much more economical interpretation of the whole sentence, thereby disambiguating the word sense ambiguities as a by-product. Where we refer to “explicit signals” in Sections 3 and 4, we are talking about lexical items or punctuation that convey very general meanings but get more specific meanings as a result of abductive interpretation.

I said above that within a sentence, adjacency conveys a predicate-argument relation; in fact, we can define sentences as that region of a text in which adjacency is conventionally interpreted as conveying a predicate-argument relation. Beyond the sentence, adjacency also conveys information, but the constraints are much looser. Adjacency says that the two clauses or larger stretches of discourse are somehow related. Clauses and larger stretches of discourse generally describe states and events, or eventualities. Thus, the relations most frequently conveyed by adjacency will be the kinds of relations that normally occur in the world among states and events. Overwhelmingly, these tend to be positive or negative versions of interlocking change of state, or what I have called the occasion relation; enablement, causality, and implication, and the negation of these in the violated expectation relation; the figure-ground relation; similarity or parallelism, and its negation, contrast; and a limiting case of similarity, the elaboration relation, or event coreference.

The occasion relation occurs when the first segment of a pair describes the initial state or a change into the initial state of a change of state described by the second segment.

- (15) I flew to Paris. I traveled through France, Spain, and Italy.

The final state of flying to Paris is being in Paris, which is the initial state of the traveling.

The causal relation occurs when one of the segments describes a cause and the other an effect.

- (16) I went to Europe. I needed a vacation.

Needing a vacation is the cause of going to Europe.

The violated expectation relation occurs when an expected causal or implicational relation does not hold. It is normally, but not always, signalled explicitly, e.g., by “but”.

- (17) I needed a vacation, but I spent the summer working.

Normally, needing a vacation causes one to not work.

The similarity or parallelism relation occurs when the same predication is made of entities that are similar in that, roughly, they share the same properties. A more precise definition of similarity is given in Hobbs and Kehler (1997).

- (18) I flew to Paris. My wife flew to London.

The predication is flying. I and my wife are similar in that we are both people. Paris and London are similar in that both are major European capitals.

The contrast relation occurs when a property is predicated of one of two similar entities and its negation is predicated of the other.

- (19) I took a taxi to the Louvre. My wife walked.

Riding in a vehicle implies that one did not walk, so a property and its negation are asserted of the similar entities, me and my wife.

The elaboration relation occurs when the two segments describe the same eventuality, perhaps from a different perspective.

- (20) The Louvre is France's biggest tourist site. Millions go there every year.

Being the biggest tourist site and having millions visit are two descriptions of the same situation. It should be pointed out that this elaboration relation is not the same as the elaboration relation in Rhetorical Structure Theory (Mann & Thompson, 1986). Their elaboration relation is no more than coreference; two successive segments of text predicate properties of the same entity. In this elaboration relation, the eventuality whose existence is asserted or claimed by the two segments have to be the same. Thus, the pair

- (21) France's biggest tourist site is the Louvre. It was built in 1202.

would be an elaboration in RST, but it would not be in my account. The Louvre's being the biggest tourist site and its being built in 1202 are not the same eventuality.

The “Interpretation as Abduction” framework can be seen as an answer to the question, “What inferences should we draw from a text?” The answer is those that contribute to the best explanation of the fact that the text occurred, which in part is the best proof of its logical form. The two aspects of this most relevant to this paper are that direct and indirect coreference resolution and the discovery of specific interpretation of general predicates both fall out of the interpretation as a by-product of the process of finding the most economical proof.

3. Classes of clause-internal coherence

3.1 The data

In the paper as a whole, I will examine texts from four diverse sources to see what instances of clause-internal coherence we find and how they can be recognized in the abduction framework. The four sources are

- A *Science* article on AIDS.
- An article from the business section of the San Jose *Mercury-News*.
- The first paragraph of Carson McCullers' *Ballad of the Sad Cafe*.
- Shakespeare's 64th sonnet.

In this section, I will examine three classes of cases: clause-internal coherence indicated by explicit signals; clause-internal coherence that falls out in the same way as coreference falls out in interpretation by abduction; and a residue of cases that remain problematic for the abduction framework. The three cases are illustrated in the following sentence opening the business news article:²

- (22) In a stunning reversal for one of Silicon Valley's fastest growing companies, Media Vision Technology Inc. said Thursday it will report a sharp decline in sales and a "substantial loss" in the quarter ending March 31—a jolt that cut its stock price in half.

3.2 Clause-internal coherence from explicit signals

Recognizing clause-*external* coherence is a matter of interpreting adjacency. That is, the coherence relation is the best explanation of why the two discourse segments are next to each other. But many of these relations are in addition explicitly signalled by a conjunction or discourse adverbial. Similarly, many instances of clause-*internal* coherence are also explicitly signalled.

In Example (22), there is a similarity between "sharp" and "substantial", and a similarity between "decline" and "loss". But the phrases "a sharp decline in sales" and "a 'substantial loss'" are conjoined by "and".

The word "and" rarely just means logical conjunction. It has two principal specializations: "and then" and "and similarly". The latter is probably more common. Specialization of the information in the logical form of a text is precisely what abduction does. Finding this specialization of "and" is equivalent to discovering the similarity. The similarity rests on the fact that a decline in sales and a

2. From the San Jose *Mercury News*, March 25, 1994, p. 12E.

loss (in profits) are both drops in positive measures for a business, and “sharp” and “substantial” both indicate the high region of the scales measuring these drops.

We may not normally think of “and” as conveying much information. But here it is the explicit signal that drives the recognition of the clause-internal coherence relation of similarity.

3.3 Intra-clausal coherence as coreference

There is an elaboration relation between the word “stunning” and the phrase “a jolt that cut its stock price in half”. We can view the condition of being stunning and the condition of being a jolt as two descriptions of the same situation. That is, by unpacking their meaning into a common core of something like “a sudden and surprising event”, and assuming these are the same, as a way of getting the most economical abductive proof, we thereby recognize the elaboration relation and in a sense see the two conditions, the jolt and the “stunning-ness”, as coreferential.

There is a violated expectation relation between “fastest growing” and “reversal”; normally, fast growth leads to more growth. Suppose for the moment we could recognize this. Then that violated expectation proposition becomes part of the interpretation of the text and is itself available as a source of inferences. Violated expectations cause surprise, and this relation is thus a partial explanation of the existence of the condition of being stunning and the condition of being a jolt. This is a kind of indirect coreference; the violated expectation exists, and therefore the “stunning-ness” exists.

In both of these cases there is an implicit redundancy in the situations described by different words or phrases in the text. We get a better abductive proof by assuming these situations are identical. Where the result is the resolution of direct coreference, we have discovered an elaboration relation. Where the result is the resolution of indirect coreference, we have discovered either an implication relation, or something more specific encoded in the axiom we use, such as causality or enablement.

3.4 Problematic residue

Let us return to the violated expectation relation between “fastest growing” and “reversal”. If we saw the text

- (23) Media Vision Technology Inc. had been growing fast. There was a reversal today.

we would have to interpret the adjacency of the two sentences by finding a coherence relation between their claims, i.e., the fast growth and the reversal. This would drive the recognition of a clause-external violated expectation relation.

Within clause boundaries, however, adjacency only conveys predicate-argument relations, and in any case, the phrases are not adjacent.

The explicit signal that comes closest to relating the reversal and the fast growth is the preposition “for”. But “for” expresses a relation between the reversal and the company, which just happened to have been growing fast. There is no direct syntactic relation between the reversal and the growth, so we cannot hope to discover the coherence relation by means of an explicit signal.

If things of type X cause or imply things of type Y, then we can normally at least partially infer one from the other. But the violated expectation relation happens precisely when that expected causal or implicational relation does *not* hold. So we can’t expect to infer the reversal from the fast growth or vice versa, as a kind of indirect coreference.

It is true that a reversal requires some kind of directed motion to be reversed, and growth is just such a motion. Assuming the reversal allows us to partially prove the existence of the growth. However, this would only give us the occasion relation. The growth sets up the occasion for the reversal to happen. It does not yield the recognition of the violated expectation relation.

This example illustrates the problematic residue of cases, for which it is hard to see how the discovery of the clause-internal coherence relation would happen in the abductive framework.

In the next section we analyze a number of other examples of clause-internal coherence in our sample texts, and classify them into one of these three classes: explicit signal, coreference, or residue.

4. Further examples of clause-internal coherence

4.1 *Science* article

In the *Science* article on AIDS we find the sentence,³

- (24) For a short but variable period—a few weeks to a few months— after an individual is infected with HIV-1, virus is typically found in the blood (viremia), and high levels of virus replication can be observed.

Explicit Signal: There is a contrast between “short” and “variable”. It is explicitly signalled by “but”. In fact, this is an example of a common pattern that might be

3. From p. 964 of “Antigenic Diversity Thresholds and the Development of AIDS”, by Martin A. Nowak, Roy M. Anderson, Angela R. McLean, Tom F. W. Wolfs, Jaap Goudsmit, and Robert M. May, *Science*, November 15, 1991, pp. 963–969.

called a “Second-Order Refinement”. The pattern is “X but not completely X”. The first part makes a first-order approximation to the intended state, and the second part makes corrections at a finer granularity. Other examples are “Pat is tall, but he stoops a lot,” and “Chris is an A student, but sometimes she makes big mistakes.” This pattern is the source of the “Yes, but ...” construction. It expresses a contrast, but it blends into the Violated Expectation coherence relation since the first part defeasibly implies the state without the correction.

The phrase “a few weeks to a few months” is an elaboration of “short” (and of “variable”). There are often explicit signals of coherence relations that we might overlook in the analysis of a clause. In this case the explicit signal is the dash, “–”. This normally signals an appositive construction and hence an elaboration. The appositive is providing a different description of the same situation. In this case, we reason that weeks and months are both periods. The word “to” indicates a range of values, as does “variable”. The word “short” applied to periods and the word “few” applied to weeks and months describe the same situation. Hence, the elaboration relation signalled by the dash is validated.

There is a similarity relation between “virus is typically found in the blood” and “high levels of virus replication can be observed”; this is signalled by the word “and”. As noted above, similarity is one of the most likely specializations of “and”. Here the similarity is based on the fact that finding and observing are both acts of perception, the implicit agents of the perception in both cases are medical personnel, and what is found is the presence and activity of virus. In both cases, what is asserted is a perceiving of diagnostic properties of the virus.

Finally, there is an enablement relation between the infection, on the one hand, and the finding of HIV-1 in the blood and the observation of high levels of virus replication, on the other. Infection, precisely, is the physical transfer of the virus from outside the body to inside the body and its establishment there as a replicating entity. Its presence in the body enables the finding, and its replicating enables the observation of the replicating. Enablement implies a temporal relation, and this relation can be expressed by the word “after”. In the abduction framework, one tries to prove the “after” relation, along with the rest of the logical form, and the enablement relation between the infection and the diagnosis is what proves it. Enablement is the more specific interpretation of “after” found by the abduction process.

The explicit signals in this example are “but”, “—”, “and”, and “after”.

Later in the *Science* article on AIDS we find the sentence,⁴

4. P. 694.

- (25) Antibodies then appear in blood serum, after which it becomes difficult to isolate the virus; viral antigens are often undetectable during the long but variable asymptomatic or incubation period between primary HIV-1 infection and the occurrence of AIDS.

Explicit Signal: Three clause-internal coherence relations are explicitly signalled.

There is a causal relation between the appearance of antibodies in the blood serum and the difficulty of isolating the virus; it is the job of the antibodies to destroy the invading virus. This causal relation is explicitly signalled by the word “after”, so the analysis is similar to Example (24) above.

The clause “viral antigens are often undetectable ...” is an elaboration on the difficulty. Here the explicit signal is the semicolon. A semicolon “;” can represent a number of possible relations, but elaboration is one of them. Recognizing the elaboration rests on lining up the implicit negative in “difficult” with the “un-” of “undetectable”, and recognizing that what is negated in the two cases, the isolating and the detectability, are the same in that both are a matter of discovering the virus. Strictly speaking, this is an *inter*-clausal coherence relation, and would be recognized in the process of explaining the adjacency of the two clauses, regardless of the punctuation.

The words “asymptomatic” and “incubation” are in a kind of contrast relation; in fact, they stand in a function-structure relation, with “incubation” describing what is going on structurally or internally, and “asymptomatic” describing what is visible to the exterior in terms of the functioning of the entity. This contrast is explicitly signalled by the word “or”, which here is used in a kind of speech act sense. It could be paraphrased as “the period which could be called ‘asymptomatic’ or could be called ‘incubation’”. That is, “asymptomatic” and “incubation” are embedded in a kind of metonymic operator— “could be called ...”. The contrast rests on the fact that “asymptomatic” means that nothing is happening, and “incubation” means that something is about to happen.

The explicit signals in this example are “after”, “;”, and “or”.

Coreference: There is an elaboration relation between the propositions conveyed by “undetectable” and “asymptomatic”. This can be recognized in the same way that coreference resolution happens, as a by-product of abduction. Symptoms are what enable someone to detect something. The words “undetectable” and “asymptomatic” both mean that there are no outward signs of the condition. They describe the same situation. The lowest-cost proof of the logical form of the sentence will be one that assumes this absence is the same in both cases. The identification of these two situations is equivalent to recognizing the elaboration relation.

There is an occasion relation between the situations conveyed by the word “incubation” and the phrase “the occurrence of AIDS”. The incubation period of

a disease is the time between infection with the vector and the occurrence of the disease. If the incubation happens, the occurrence will happen. Thus, the occurrence need not be assumed; it can be proved after assuming or proving the incubation, yielding a lower-cost proof and explicating the inferential relation between them. Because the incubation of a disease is a change of state into the state in which the disease occurs, the relation between them is the occasion relation.

In the first of these two cases, we can say that the undetectability and the asymptomaticity are directly coreferential. In the second case we can say the incubation and the occurrence of the disease are indirectly coreferential. The existence of the first implies the existence of the second.

The next sentence in the article on AIDS is the following:

- (26) The incubation period is characterized by low viral replication (interspersed with minor and short-lived upsurges of viremia in some patients), and by constant or slowly decreasing numbers of CD4⁺ cells.

Explicit Signal: The two issues in this sentence are the level of viral replication and time. These are similar in that they are both parameters. They are contrasting in that time is the independent parameter and the level is the dependent parameter. Recognizing that they are both parameters and that the level depends on time requires us to interpret the word “characterized” correctly. If X characterizes Y, then there is a functional relation from Y to X. We can think of the statement “Red hair characterizes Irishmen” as positing a function from people to hair color that maps Irishmen into red hair. Here, different levels characterize different periods of time.

The word “with” also conveys that functional relation. Its first argument is “interspersed”, which indicates a temporal aggregate, a set of temporal intervals. For these intervals to be “with” upsurges means that there is a mapping from each of the elements of the aggregate to the upsurge that characterizes it.

If we had in our knowledge base an axiom that said that elements of the domain of a function are independent parameters while elements in the range are dependent parameters, then we could in principle recognize the contrast relation.

Residue: The word “period” refers to a temporal interval. The word “interspersed” describes an aggregate of intervals all contained within a longer interval. The phrase “short-lived” describes the length of an interval. The phrase “slowly decreasing” says something about the length of the interval occupied by the decreasing event. Thus, all of the eventualities described by these words and phrases are in a similarity relation, by virtue of their reference to a temporal interval.

The word “low” describes the level of viral activity. The word “upsurges” describes episodes in which the level is higher than usual. The word “minor” moderates that, in a kind of Second-Order Correction contrastive relation with “upsurges”, so it also refers to the level of viral activity. The phrase “numbers of

CD4⁺ cells” describes a measure of the level of viral activity, and consequently the words “constant” and “decreasing” describe such levels. Thus, all of the eventualities described by these words and phrases are in a similarity relation, by virtue of their reference to level of viral activity.

These clause-internal similarity relations exist. But it is not clear how they would be recognized in the abduction framework. We can axiomatize similarity in a way that allows it to be validated when it is explicitly signalled. But here it is not. For example, in “interspersed with short-lived upsurges”, the fact that “interspersed” and “short-lived” both make reference to time is, in some sense, accidental. We could have as felicitously said “interspersed by unexpected upsurges”, where there is no similarity relation.

We cannot recognize the similarity as a kind of direct anaphora. They are not the *same* intervals or the *same* levels, so we would not want to identify them, on the way to finding the lowest-cost proof.

To recognize them as a kind of *indirect* anaphora, we would need an axiom that said for any entity, there is a similar, nonidentical entity. This axiom seems to be much too powerful.

Nevertheless, people are very very good at spotting similarities wherever they occur.

4.2 Business news

In the article from the San Jose *Mercury-News* business section, the two sentences immediately after Example (22) are the following:

- (27) Media Vision plummeted to 11, down 10 1/2 in frantic NASDAQ trading as 14.2 million shares were traded, more than 25 times normal volume. Thursday’s decline continues a precipitous two-month slide from a peak of 45 1/4 Jan. 20 that has wiped out \$480 million in market value.

Explicit Signal: There is an occasion relation between the two-month slide and Thursday’s decline. The slide is a change of state into a final state that was the initial state in the decline. But this is explicitly signalled by the verb “continues”.

In addition, when X continues Y, the implication is that X and Y are the same sort of eventuality. That is, X and Y are similar. In this case, they are similar because both are downward motion.

Coreference: The word “plummeted” indicates a rapid movement downward. The word “down” of course indicates a downward direction. The word “decline” also refers to a downward movement, as does slide. Money is a metaphorically vertical scale, so that decreasing a measure on that scale is a downward movement; the phrase “wiped out \$480 million” thus indicates a downward movement. All of the eventualities conveyed by these words and phrases at least stand in a similarity relation by virtue of the downward movement they all indicate.

But in fact there are coreference relations involved here. The word “down” describes the same as the downward motion implicit in the plummeting. Thursday’s decline is the same as the plummeting event. Thus, assuming a downward motion at least partially accounts for all three eventualities. Moreover, the wiping out and slide are both descriptions of the same event, and can similarly be recognized by identifying the two downward motions. The statement that the decline continues the slide should block coreference resolution between the two sets. Thus, we have an elaboration relation among the eventualities in each of the sets. Because of the word “continues”, we have an occasion and a similarity relation between the first set and the second.

There is also a similarity among the various indicators of the intensity of that decline, including “plummeted”, “frantic”, “14.2 million shares”, “more than 25 times normal volume”, “precipitous”, “wiped out”, and “\$480 million”. Of course, for the numeric indicators to be recognized as signals of intensity, one has to know the normal range of values. There seem to be three situations being described here. For a process to be frantic is for its subevents to occur in rapid succession driven by fear in the agents. To plummet is to drop rapidly, so there is an identity in the implicit rapidity of the plummeting and the “frantic-ness”. But the two situations themselves are not identical. The plummeting does not necessarily have to be driven by fear and the “frantic-ness” does not necessarily involve downward motion. The identity of the trading of 14.2 million shares and the volume of 25 times normal is indicated by the fact that the latter phrase acts as an appositive on the former phrase, once we coerce the latter from the verb “traded” to its subject (cf. Hobbs, 2001). Recognizing the other identities of eventualities depend on fairly complex reasoning about rates. For 25 times normal trading to occur in one day must mean that the trading was rapid, linking with the rapidity implicit in “plummeted” and “frantic”.

The eventuality conveyed by “precipitous” is downward motion with a steep slope. This situation is described again in the reference to wiping out \$480 million. To recognize this we must know that \$480 million is a lot of money and is a loss that occurred in a two-month period, and thus is a quantitative measure of a slope.

The precipitousness of the slide is distinct from the rapidity of the plummeting simply because we can establish that the slide and the plummeting are not the same.

4.3 The novelette

The first sentence of the Carson McCullers novelette is as follows:⁵

5. Carson McCullers, “Ballad of the Sad Cafe”, Houghton-Mifflin Co., Boston, MA, 1943.

- (28) The town itself is dreary; not much is there except the cotton mill, the two-room houses where the workers live, a few peach trees, a church with two colored windows, and a miserable main street only 100 yards long.

Coreference: There is a causal or implicational relation between the cotton mill and the workers. A cotton mill is a factory and factories have workers, so we can see this inferential relation as a kind of indirect coreference.

There is also an implicational relation between the houses and the people who live there. Houses are where people live, and workers are people. There is thus a partial proof from “houses” of the existence of the workers, and a full proof of the existence of the living.

Explicit Signal: There is an exemplification, or more properly, “bad whole - bad part” relation, between the “dreary” town and the “miserable main street”. There is also a similarity relation among the various indicators of quantity in the sentence—they are all small—including “not much”, “two-room (houses)”, “few (peach trees)”, “two (colored windows)”, and “only (100 yards long)”.

It is possible that all of these should fall out from a recognition of the coherence relation between the two clauses, signalled by a semicolon, and an interpretation of “and” as “and similarly”. The first clause describes the town as a whole. The second clause elaborates on this by describing its various parts—its factory, houses, trees, church, and street. These are all similar in that they *are* all parts of a town. Furthermore, all of these items except the cotton mill have quantitative descriptors. This adds to the similarity. But it is more important that the quantities are all small. This is explicit in “a few peach trees” and “only 100 yards long”. To recognize “two-room houses” and “a church with two colored windows” as indicating a small quantity we need to know the normal range of quantities. The interpretation of “and” as “and similarly” spreads across the entire conjoined noun phrase object of the preposition “except”, and the similarity of these items is established since all are small examples of things found in town. The preposition “except” conveys a Second-Order Correction. The phrase “not much” gives us a baseline pretty close to nothing, and the “except” phrase gives us a more detailed accounting of how it differs from nothing.

Not all small quantities are bad. If the town had few crimes and few tornadoes, that would be good. But we know that small house size, small church size, and a small number of trees is usually not very good. Thus, the quantitative properties all share a badness property with the dreariness of the town. We thus get an elaboration relation between the first clause—the town is bad—and the second—its parts are bad. The adjective “miserable” modifying “main street” fits into this pattern.

Thus, the recognition of the clause-internal similarity relations is driven by the interpretation of “and” and by the recognition of the elaboration relation between the two clauses.

The second sentence in the novelette is

(29) On Saturday the tenants from near-by farms come in for a day of talk and trade.

Coreference: There is an implicational relation between “Saturday” and “day”. Saturday is a day, and the lowest cost proof results by assuming they are the same day.

There is also an implicational relation from the farms to the tenants. Farms have farmers who work them, and a tenant is a farmer who does not own the farm, so there is a partial proof of the existence of the tenant if we assume the existence of the farm.

The fact that the farms are “near-by” enables the tenants to come in, and this coming in in turn enables the talk and trade. Qualitative scalar concepts like “near-by” are generally associated with functional properties. If two points are literally or metaphorically near each other, then it is easy to traverse the distance between them. Thus, the nearness enables the coming in. Social interactions require the participants to be in the same location (which in the electronic age can be defined by a network of connectivities). So the coming into town enables the talk and trade. If we assume the talk and trade occurred, then we can assume the enabling identity of location occurred, so we can assume the coming together in that place occurred. If that is so, an enabling nearness to the starting points must hold as well. Thus, axioms with the structure “If X occurs, then Y defeasibly enabled it” allow us to infer the existence of the coming in and the nearness from the assumption of the talk and trade. These enabling relations are thus a kind of indirect coreference resolution.

4.4 Shakespeare’s sonnet

The next four examples come from Shakespeare’s 64th sonnet, and because it is sometimes hard to understand lines from it in isolation, the entire sonnet is presented here.

When I have seen by Time’s fell hand defaced
The rich, proud cost of outworn buried age,
When sometime lofty towers I see down-rased
And brass eternal slave to mortal rage;

When I have seen the hungry ocean gain
Advantage on the kingdom of the shore,
And the firm soil win of the wat’ry main,
Increasing store with loss and loss with store;

When I have seen such interchange of state,
Or state itself confounded to decay,
Ruin hath taught me thus to ruminare,
That Time will come and take my love away.

This thought is as a death, which cannot choose
But weep to have that which it fears to lose.

In this sonnet, the poet observes that time destroys everything, and that time will eventually destroy his love. It is a poem about entropy.

The first example from the sonnet is the first quatrain:

- (30) When I have seen by Time's fell hand defaced
The rich, proud cost of outworn buried age,
When sometime lofty towers I see down-rased
And brass eternal slave to mortal rage;

Coreference: Defacing is one instance of something being fell. The two words overlap in meaning; there are other ways of being fell, and defacement describes the process as well as merely the result. But this common core of meaning in the two words represents the same underlying situation. Assume this, and there is a partial proof of both the condition of being fell and the defacing event, thereby capturing the elaboration relation between them.

There is a causal relation between "Time" and the defacing. The passage of time causes things to be no longer intact. This is one aspect of being defaced, so the defacement can be seen partially as an indirectly coreference from Time. The passage of time exists, so possibly a defacement does too.

There is also a causal relation between the "time" of "sometime" and "down-rased". This can be analyzed in the same way as the relation between "Time" and "defaced", as indirect coreference.

Residue: The whole quatrain is built on the contrast between valuable and intact things—"rich, proud cost", "lofty towers", and "brass eternal"—and the condition of being broken—"fell", "defaced", "outworn buried", "down-rased", and "slave to mortal rage". The general pattern is "Valuable and intact things break." We can recognize an Occasion relation from the intactness to the breaking, since breaking is a change of state whose initial state is intactness. But as in our analysis of Example (22), this does not give us the contrast relation.

The second quatrain is

- (31) When I have seen the hungry ocean gain
Advantage on the kingdom of the shore,
And the firm soil win of the wat'ry main,
Increasing store with loss and loss with store;

Coreference: There is a metaphorical causal relation between “hungry” and “gain” (or “gain advantage”)—the ocean’s hunger causes it to eat the land. Eating is consuming, and consuming causes a gain. Thus, the causal relation can be discovered as a kind of indirect coreference resolution. The existence of the hunger causes the existence of the gain.

Residue: There is a clause-internal contrast between “firm” and “wat’ry”. To be firm is to be solid; to be watery is to be not solid. But neither the firmness nor the watery-ness imply each other, so a relation cannot be discovered by coreference resolution. There is no explicit signal of the contrast. The word “and” indicates a parallelism between the first two lines and the third line, but the parallelism rests on one domain consuming the other. It does not go down to the level of the descriptors of the domains. That is, the hunger of the ocean bears no relation to the firmness of the soil.

The third quatrain is

- (32) When I have seen such interchange of state,
Or state itself confounded to decay,
Ruin hath taught me thus to ruminate,
That Time will come and take my love away.

Coreference: There is a clause-internal contrast between the second occurrence of “state” and the word “decay”. The sense of “state” here is “majesty, royalty, or splendor”, but the poem works just as well if we take “state” to mean the more modern sense of a politically organized body of people under a single government. In either case we have highly structured entity. Decay is the loss of internal structure in a structured entity, and confounding is a cognitive version of the process of losing orderly structure. The higher level of structure in the start state of a decay process is implicit in the word “state”. This is thus an example of partial indirect coreference.

There is an enablement relation between “come” and “take . . . away”. If we assume the existence of the taking away event, it must have had its locational enabling conditions hold, and that is the final state of the coming. Thus, we can see the recognition of this enablement relation as a partial solution to an indirect coreference problem.

The final couplet is

- (33) This thought is as a death, which cannot choose
But weep to have that which it fears to lose.

Residue: There is a clause-internal violation of expected causality between “weep” and “have”—normally it is a loss that would be wept, not a possession. We cannot

discover this as a variety of coreference, because neither the having nor the weeping imply the other. The relation between the two eventualities is, in more modern terms, “weep at having”, where “at” conveys causality. But there is no rule that says that having what we want causes us to weep. Quite the contrary. So there is no mechanism in the abduction framework that would force the discovery of this clause-internal coherence relation.

5. Summary

The dense clause-internal coherence structure we have seen might be expected in poetry, and perhaps in the novelette as well, but the examples from the business news and the *Science* article indicate that the phenomenon is quite pervasive in all *written* discourse, at least. (Examples are harder to find in conversational data.)

The interpretation of many of these examples falls out from ordinary interpretation by abduction. The lowest cost explanation of the text as a whole contains the coherence relations within it. There are two mechanisms by which this can occur.

The first mechanism is the process of finding more specific meanings in context for words and other explicit signals than they would convey in isolation. Many of these signals might often escape our notice as requiring interpretation, such as words like “and” and punctuation like dashes and semicolons. The situation is analogous to when an inter-clausal coherence relation is signalled not only by adjacency but also by a conjunction or a clause-level discourse adverbial. In both cases, the abduction framework dictates that we prove abductively the very general meaning conveyed by the signal, and the specific coherence relation falls out of that proof.

The second mechanism resembles coreference resolution. The lowest cost abductive proof results if we assume that entities or eventualities described or implied by different parts of a text are in fact identical. Where they are *described*, we have thereby recognized an elaboration relation, analogous to *direct* coreference. Where one or more are merely *implied*, we have thereby recognized at least an implicational relation, analogous to *indirect* coreference, and if there is causality or enablement encoded in the axioms we use, we have thereby recognized those relations as well. Example (3) is just such a case, although a somewhat complex one. If a car hits someone, they must be in the same place. The location of the victim enables the accident. Jogging is usually done outside, and often joggers are in the street, which is where cars usually are. The existence of the jogger partially implies being located in the street, which enables the accident. This is thus a case of partial indirect coreference resolution in which enablement is part of the supporting abductive proof.

However, problematic cases remain. When the relation is one of violated expectation, we cannot expect the same sort of coreference based on an inferential or causal relation, because that is precisely what is violated. When the relation is one of similarity, there is no direct coreference to discover, since the eventualities are not identical. There is no indirect coreference to discover, since the existence of an eventuality does not imply the existence of similar eventualities. These cases constitute a challenge to the “Interpretation as Abduction” picture of what counts as an interpretation.

References

- Hobbs, Jerry R. 1985. “On the coherence and structure of discourse.” Report No. CSLI-85-37, Center for the Study of Language and Information, Stanford University.
- Hobbs, Jerry R. 2001. “Syntax and metonymy.” In *The Language of Word Meaning*, Pierrette Bouillon & Federica Busa (eds), 290–311. Cambridge, United Kingdom: Cambridge University Press.
- Hobbs, Jerry R., & Kehler, Andrew. 1997. “A theory of parallelism and the case of VP ellipsis.” *Proceedings*, 35th Annual Meeting of the Association for Computational Linguistics, 394–401. Madrid, Spain, July 1997.
- Hobbs, Jerry R., Stickel, Mark, Appelt, Douglas, & Martin, Paul. 1993. “Interpretation as abduction.” *Artificial Intelligence*, 63 (1–2): 69–142.
- Kehler, Andrew, Kertz, Laura, Rohde, Hannah, & Elman, Jeffrey L. 2008. “Coherence and coreference revisited.” *Journal of Semantics*, 25: 1–44.
- Kronfeld, Amichai. 1989. “Con conversationally relevant descriptions.” *Proceedings*, 27th Annual Meeting of the Association for Computational Linguistics, 60–67. Vancouver, British Columbia, June 1989.
- Mann, William, & Thompson, Sandra. 1986. “Relational propositions in discourse.” *Discourse Processes*, 9 (1): 57–90.