

FOCUSING YOUR RST: A STEP TOWARD GENERATING COHERENT MULTISENTENTIAL TEXT

Eduard H. Hovy¹
Information Sciences Institute of USC

Kathleen F. McCoy²
University of Delaware

Abstract

In multisentence texts, the order and interrelationships of sentence topics is of crucial importance if the reader is to understand easily. What makes paragraphs coherent? What strategies do people employ to control the presentation order and linking of material? Without a theory of coherence, text generation systems have little hope of producing acceptable texts. While various theories of text coherence have been developed, no single theory captures all the phenomena of human-generated paragraphs. In this paper we argue that the coherence of a paragraph does not result from the application of a single theory, but instead results from the cooperation of a number of different coherence strategies. We illustrate this claim by showing how two very different theories about the planning of coherent text — 1) Rhetorical Structure Theory, based on structural and semantic relationships that hold between pieces of the text, and 2) Focus Trees, based on how the focus of attention shifts during discourse — can be used within a single system to complement each other to best advantage.

1 Introduction

In multisentence texts, the order and interrelationships of sentence topics is of crucial importance if the reader is to understand easily. But what makes a paragraph, and by extension, a text, coherent? By what strategies do people control the presentation of material so as to develop their ideas intelligibly? Without answers to these questions, text generation systems have little hope of producing acceptable paragraphs. Various theories of text coherence have been developed (e.g., see [Hobbs 78, Reichman 78, Cohen 83, Mann & Thompson 88]), each based on valid but quite different considerations. Unfortunately, no single theory suffices to define coherence well enough to compare with paragraphs written by people. An adequate theory of text planning must incorporate several coherence strategies under one framework. Care must be taken that the text planner give proper preference to the most effective coherence strategy at each step in the process, lest its effect be precluded by an independent decision.

In this paper we discuss two strategies that contribute greatly to the planning of coherent texts: Rhetorical Structure Theory and Focus Trees. Rhetorical Structure Theory (RST) [Mann & Thompson 88] provides coherence to a text based on structural considerations about the rhetorical relationships that hold between adjacent pieces of a paragraph. Focus Trees [McCoy & Cheng 88] provide coherence to a text based on anticipated shifts in the focus of

¹This author was supported in part by the Advanced Research Projects Agency under DARPA contract MDA903-87-C-641, and by AFOSR contract F49620-87-C-0005.

²This author was partially supported by UDRF Grant #LTR870112.

attention of the participants as the text proceeds. We first introduce these notions and illustrate how each is individually useful yet underconstraining to a generation system whose aim is to produce coherent multisentential text. Finally, we show how these two theories can be combined in a single text planning system that uses both strategies effectively. While we do not argue that combining these two methods into one system fully solves the coherence problem, we believe that it is a step toward a complete solution. We expect that ultimately other sources of guidance will have to be incorporated, and that the text planning methodology advocated here is general enough to support them.

2 Rhetorical Structure Theory and Coherence

The Penman project at ISI has been investigating the planning of coherent multisentential paragraphs of text by computer. The planner, a top-down hierarchical expansion planning system patterned on NOAH [Sacerdoti 77], uses plans which are operationalizations of some RST relations from Rhetorical Structure Theory [Mann & Thompson 88], which posits that approximately 20 relations suffice to relate adjacent blocks of text in the ways English speakers consider coherent. The planner is described in [Hovy 88a, Hovy 88b]. It operates antecedent to the natural language generator Penman [Mann & Matthiessen 83].

The text structure planner plans coherent paragraphs to achieve communicative goals to affect the hearer's knowledge in some way. It accepts one or more communicative goals along with a set of clause-sized inputs from the domain of discourse to be generated as an English paragraph. The planner assembles the input entities into a tree that expresses the paragraph structure. The nonterminals in the tree are RST relations while the terminal nodes contain the clause-sized inputs. Finally, the planner traverses the tree, dispatching the leaves (the input entities) to be generated by Penman.

The planner embodies a limited top-down hierarchical expansion planning framework. Figure 1 shows a typical relation/plan in this formalism. Each relation/plan has two parts, a *nucleus* and a *satellite*, and recursively relates some unit(s) of the input or another relation (cast as nucleus) to other unit(s) of the input or another relation (cast as satellite). In order to admit only properly formed relations, nuclei and satellites contain requirements that must be matched by characteristics of the input. In addition, nuclei and satellites contain growth points: collections of goals that suggest the inclusion of additional input material. On finding (an) RST relation/plan(s) whose effects include achieving (one of) the system's communicative goal(s), the planner searches for input entities that match the requirements holding for each of its parts. If fulfilled, the planner then considers the growth points of each part of the relation/plan. It tries to achieve each newly instantiated growth point goal by again searching for appropriate relation/plans and matching them to the input, recursively, adding successfully achieved goals to the paragraph tree structure. The planning process bottoms out when either all of the input entities have been incorporated into the tree or no extant goals can be satisfied by the remaining input entities. The tree is then traversed in a depth-first left to right manner, and the relation/plans' characteristic cue words or phrases are added to the appropriate input entities and transmitted to Penman to be generated as English clauses.

Up until now, the paragraphs produced by the system relied on a very important assumption about the growth points: their presence and order were treated as injunctions. That is to say, the structure planner always tried to achieve every growth point goal in the order given. As shown in [Hovy 88b], treating growth points this way is equivalent to using the relation/plans as schemas — structures that mandate the content of a paragraph-sized block of text [McKeown 85]. Though useful for many constrained applications, schemas do not support well systems that seek to exhibit dynamic and adaptive behavior. Since we are attempting to build such systems, we have

Figure 1: The RST relation/plan SEQUENCE

```

Name: SEQUENCE

Results:
  ((BMB SPEAKER HEARER (SEQUENCE-OF ?PART ?NEXT)))

Nucleus requirements/subgoals:
  ((AND (BMB SPEAKER HEARER (MAINTOPIC ?PART))
        (BMB SPEAKER HEARER (NEXT-ACTION ?PART ?NEXT))))

Nucleus growth points:
  ((BMB SPEAKER HEARER (CIRCUMSTANCE-OF ?PART ?CIR))
   (BMB SPEAKER HEARER (ATTRIBUTE-OF ?PART ?VAL))
   (BMB SPEAKER HEARER (PURPOSE-OF ?PART ?PURP)))

Satellite requirements/subgoals:
  ((BMB SPEAKER HEARER (MAINTOPIC ?NEXT)))

Satellite growth points:
  ((BMB SPEAKER HEARER (ATTRIBUTE-OF ?NEXT ?VAL))
   (BMB SPEAKER HEARER (DETAILS-OF ?NEXT ?DETS))
   (BMB SPEAKER HEARER (SEQUENCE-OF ?NEXT ?FOLL)))

Order: (NUCLEUS SATELLITE)
Relation-phrases: (" "then" "next")
Activation-question:
  "Could ~A be presented as start-point, mid-point, or end-point
   of some succession of items along some dimension? -- that is,
   should the hearer know that ~A is part of a sequence?"

```


reconsidered the interpretation of growth point goals by treating them merely as (unordered) suggestions for additional paragraph growth.

Under the new interpretation, the paragraph structurer produces many more paragraph trees, some of which do not seem coherent. We take here an example from one of the three domains to which the paragraph planner has been applied, the Integrated Interfaces domain, a multimodal system which satisfies user requests for information from a Navy database [Arens et al. 88]. The Integrated Interface display planner furnishes a set of 6 related entities along with the goal of describing the sequence of events starting from the first event, including as much of the given information as possible. Using the RST relation/plan that achieves the goal to express a sequence given in Figure 1, but taking the nucleus growth points out of order, the paragraph structure planner produces the tree in Figure 2, from which Penman generates the text shown. While this text is well-structured according to RST, it lacks the coherence found, for example, in the following rendition of the same propositional content:

With readiness C4, Knox is en route to Sasebo. It is at 79N 18E heading SSW. It will arrive 4/24 and will load for four days.

RST relation/plans do constrain the planner to produce only coherent paragraphs. Its valuable contribution to the coherence of a text must be supplemented with other coherence factors. One of these is focus.

Figure 2: RST Generated Navy text



Knox is en route to Sasebo. It is at 79N 18E heading SSW.
It is C4. It will arrive 4/24, and will load for four days.

3 Focus Trees and Coherence

Discourse Focus Trees were introduced in [McCoy & Cheng 88] to capture the shifts in the focus of attention of discourse participants as a discourse progresses. They are an attempt to integrate into one unified approach the kinds of focusing phenomena noticed by researchers in specialized kinds of discourse

[Schank & Abelson 77, Garrod & Sanford 83, Grosz 77, Carberry 83, McCoy 85]. During the discourse, a focus tree is constructed and traversed, one node being visited at a time. Based on the position of the visited node in the tree, entities from the knowledge base are highlighted. In deciding what to say next, the generation system must either choose an element from the highlighted set or make one of a small set of legal moves to another node in the focus tree. If another focus move must be made, the shift must be marked in the text lest it seem incoherent.

Each node in the Focus Tree points to an entity from the knowledge base. Thus the nodes in the focus tree are of different types, depending on the ontology of the domain. In general, Focus Tree nodes belong to one of five types: object, attribute (property), setting, action, and event. Each type of node causes highlighting of a particular set of knowledge base entities, and in so doing, furnishes different candidates for what may next be coherently included in the text. Figure 3 lists the permissible focus shifts for three node types.

The node type alone is not sufficient to determine knowledge base highlighting. In different conversations, for example, the very same object can support various focus moves. Therefore the highlighting depends also on the position of the currently visited node with respect to its ancestors (sibling nodes also play a role) in the focus tree. The focused knowledge can be thought of as the intersection of the knowledge related to information about the currently visited node and the knowledge related to each of its ancestor nodes in the tree. In this way, the parent nodes lend a perspective through which the currently visited node, and its children, are viewed.

Associated with this inheritance of perspective is the following rule: Focus candidates are always interpreted in the most specific context. That is to say, they migrate down the tree if they can: when in a topic shift a child node becomes the new currently visited node, and it has focus candidates (potential child nodes) in common with any of its ancestors, then those candidates migrate down the tree to appear *only* under the child node. Then, when the parent node is later returned to, a subsequent shift of focus to one of the migrated candidates would involve the

Figure 3: Candidate Focus Shifts for Selected Node Types

NODE TYPE	FOCUS SHIFT CANDIDATES
OBJECT:	attributes of the object, actions the object plays a prominent role in (e.g., is actor of)
ACTION:	actor, object, etc., of the action (any participant role; see [Fillmore 77]), purpose (goal) of action, next action in some sequence, subactions, specializations of the action
ATTRIBUTE:	objects which have the attribute, more specific attribute

revisiting of the intermediate node, a move that is incoherent unless linguistically marked.

As the discourse proceeds, a Focus Tree is built up and traversed. A node may be added to the tree either by explicit reference in the discourse or by inference which can happen in both top-down and bottom-up fashion. By top-down we mean that each particular node in the tree, when visited, furnishes candidate nodes to which the focus may later progress, based on its type and its position in the tree. Each progression causes an appropriate child node to be grown in the tree. Bottom up inferencing may also be necessary to connect several seemingly unrelated nodes under a unifying theme.

The traversal of the tree (shifting of focus) normally proceeds depth-first: what is said next is either an expansion of the currently visited node, or an expansion of one of its previously unexpanded children, or a new expansion of one of its ancestors. A major difference between this focusing theory and other theories of focus (e.g., [Grosz 77, Sidner 79, Grosz & Sidner 86]) is that a depth first walk of the tree is only expected, not required. Other focus moves are indeed possible, though they require explicit marking in the text (by the use of such phrases as “to go back to. . .”). The further away from the standard depth-first traversal the move is, the stronger the marking must be.

Focus Trees can be used in the generation of paragraphs by constraining what is said next to respect legal focus shift moves in the tree. However, the use of the Focus Tree for this purpose does not preclude the generation of incoherent text in all cases. For instance, using the above rules, the following text could be generated from the Navy text input entities, assuming the initial focus is on Knox:

Knox, which is C4, will arrive 4/24 and load for 4 days. It is heading SSW and is at 79N 18E. It is en route to Sasebo.

The above text is not a coherent rendering of the text produced in the previous section. Prohibiting such text requires additional information not contained in Focus Trees.

4 Integration of the Two Methods

The preceding sections described the inability of either method alone — RST or Focus Trees — to fully control the generation of coherent paragraphs. In this section we describe a way of integrating these theories that uses each to best advantage.

The insight that focus and structural considerations can be combined to produce coherent text is not new. McKeown implemented a combined scheme in [McKeown 85]. Her TEXT system generated paragraphs by first partitioning off the portion of the knowledge base that might be

included in the text using simple rules. This pool of relevant knowledge lent a “global focus” [Grosz 77] to the text. Once the potentially relevant knowledge was identified, the assembly of the paragraph was controlled by structural rules encoded in a schema. Variability, primarily in the inclusion of optional material, was controlled by a focusing mechanism based on [Sidner 79].

While it is clear that an RST planner would greatly benefit from the incorporation of a focusing mechanism, the mechanism used by McKeown is not sufficient because of the recursive method of planning employed by the RST planner. McKeown’s algorithm always controls what should be said immediately following the text planned so far. Using the RST planner, two pieces of text may be planned under a particular RST operator, but then growth points in either the nucleus or the satellite may cause additional text to be inserted *between* the already planned parts. McKeown’s algorithm provides no way to handle focus dependencies over discontinuous pieces of text. In addition, since her algorithm uses a stack, it does not maintain a record of popped entities. As a result, the algorithm would allow returning to a previously focused entity which had been popped off the stack without reference to its previous mention.

The use of focus trees avoids these problems: 1) inserting additional text corresponds to introducing new nodes into the Focus Tree, which is a routine operation, and, 2) a tree is precisely a stack that records its history.

We integrate the two methods as follows: While the RST planner constructs the paragraph structure tree, a focus tree is constructed in tandem. During the expansion of a node in the RST tree, the structurizer applies all the growth point goals active at that point and collects the resulting candidate relations and their associated clause-sized input entities. Each candidate growth entity is then checked against the currently allowed focus shifts in the Focus Tree, and invalid candidates are simply removed from consideration. One of three possibilities ensues:

1. Only one candidate remains and growth proceeds straightforwardly.
2. More than one candidate remains. In this case all candidates are coherent based on rhetorical structure and focus but additional measures, still to be developed, must be employed to select the best of these. (As an interim practical solution, the growth points in the RST relation/plan can be ordered by typical occurrence, and the tree can be grown in this default order.)
3. No candidates remain. In this case the overall stylistic goals of the system may dictate to either stop tree growth at this point, or continue tree growth in the default order as above but linguistically mark the text to indicate a focus shift.

One further subtlety remains: Some RST trees are unacceptable to the Focus Tree criterion in their initial form, but can be made acceptable by reordering their parts (which may involve generating appropriate linguistic focus words in the text). Consider the planning of the RST tree in Figure 2, which is such a case, under the additional control of the Focus Tree. The initial goal to express a sequence starting with *enroute* and focusing on *Knox* generates the RST and Focus trees in Figure 4. Next, using the growth point calling for an ATTRIBUTIVE relation, the RST planner finds the C4 readiness attribute of *Knox*. However, the Focus Tree requires that the C4 clause precede the *enroute* clause in the text — otherwise, as is clear from Figure 4, generating C4 causes a shift up the Focus Tree away from *enroute*, a shift that must be undone directly in order to generate the subsequent *arrive* and *load* clauses. The RST planner handles this requirement by inverting the ATTRIBUTIVE relation nucleus and satellite in the RST tree. After subsequent planning, the final result is the RST tree in Figure 5, which would give rise to the text shown. Note that simple reordering of the attributive information makes the text more coherent, and prevents both the text generated in Figure 2 and the unacceptable text allowed by the Focus Trees alone.

Figure 4: Initial Trees





Figure 5: Joint RST and Focus Generated Navy text



With readiness C4, Knox is en route to Sasebo. It is at 79N 18E heading SSW. It will arrive 4/24 and will load for four days.

5 Conclusion

In this paper we illustrated how a text planner which relies on a single coherence method will not generate coherent paragraphs in all circumstances. We presented two coherence theories, RST and Focus Trees, and showed how they may be integrated into a single planning methodology to overcome problems that neither addresses alone.

Though a step in the right direction, this combination is not yet sufficient to guarantee coherent text in all cases. We envision that other aspects of text coherence will give rise to other theories which must ultimately be integrated into this framework. We believe the top-down hierarchical expansion method is powerful enough to support such additions. We hope to continue this investigation by identifying other coherence techniques useful for the generation task and integrating them into this framework.

References

[Arens et al. 88]

Arens, Y., Miller, L., Shapiro, S.C. & Sondheimer, N.K. Automatic Construction of User-Interface Displays. In *Proceedings of the 7th AAAI Conference*, St. Paul, MN, 1988. Also appears as USC/Information Sciences Institute Research Report RR-88-218.

HOVY, MCCOY

- [Carberry 83] Carberry, S.M. Tracking user goals in an information-seeking environment. In *Proceedings of the 3rd AAAI Conference*, Washington, D.C., 1983 (59-63).
- [Cohen 83] Cohen, R. *A Computational Model for the Analysis of Arguments*. Ph.D. dissertation, University of Toronto, 1983. Also appears as University of Toronto Computer Systems Research Group Technical Report no. 151.
- [Fillmore 77] Fillmore, C.J. The Case for Case Reopened. In P. Cole and J.M. Sadock (editors), *Syntax and Semantics VIII: Grammatical Relations*. Academic Press, New York, 1977 (59-81).
- [Garrod & Sanford 83] Garrod, S. and Sanford, A. Topic dependent effects in language processing. In G.B. Flores d'Arcais and R.J. Jarvella (editors), *The Process of Language Understanding*. John Wiley & Sons Ltd., 1983 (271-296).
- [Grosz 77] Grosz, B.J. *The Representation and Use of Focus in Dialogue Understanding*. Technical Report no. 151, SRI International, Menlo Park CA, 1977.
- [Grosz & Sidner 86] Grosz, B.J. and Sidner, C.L. Attention, Intentions, and the Structure of Discourse. *Computational Linguistics Journal* 12(3), 1986 (175-204).
- [Hobbs 78] Hobbs, J.R. Why is discourse coherent? Technical Note no. 176, SRI International, Menlo Park CA, 1978.
- [Hovy 88a] Hovy, E.H. Planning coherent multisentential text. In *Proceedings of the 26th ACL Conference*, Buffalo, 1988 (163-169).
- [Hovy 88b] Hovy, E.H. Approaches to the planning of coherent text. Presented at the *4th International Workshop on Text Generation*, Los Angeles, 1988.
- [Mann & Matthiessen 83] Mann, W.C. and Matthiessen, C.M.I.M. Nigel: A systemic grammar for text generation. USC/Information Sciences Institute Research Report RR-83-105, 1983.
- [Mann & Thompson 88] Mann, W.C. and Thompson, S.A. Rhetorical structure theory: Toward a functional theory of text organization. In *Text* 8(3), 1988 (243-281). Also appears as USC/Information Sciences Institute Research Report RR-87-190.
- [McCoy 85] McCoy, K.F. *Correcting object-related misconceptions*. Ph.D. dissertation, University of Pennsylvania, 1985.
- [McCoy & Cheng 88] McCoy, K.F. and Cheng, J. Focus of attention: Constraining what can be said next. Presented at the *4th International Workshop on Text Generation*, Los Angeles, 1988.
- [McKeown 85] McKeown, K.R. *Text generation: Using discourse strategies and focus constraints to generate natural language text*. Cambridge University Press, Cambridge, 1985.
- [Moore & Swartout 88] Moore, J.D. and Swartout, W.R. A reactive approach to explanation. Presented at the *4th International Workshop on Text Generation*, Los Angeles, 1988.
- [Paris 88] Paris, C.L. Generation and explanation: Building an explanation facility for the Explainable Expert Systems framework. Presented at the *4th International Workshop on Text Generation*, Los Angeles, 1988.
- [Reichman 78] Reichman, R. Conversational coherency. *Cognitive Science* 2, 1978 (283-327).
- [Sacerdoti 77] Sacerdoti, E. *A structure for plans and behavior*. North-Holland Publishing Company, Amsterdam, 1975.
- [Schank & Abelson 77] Schank, R.C. and Abelson, R.P. *Scripts, plans, goals, and understanding*. Lawrence Erlbaum Associates, Hillsdale, 1977.
- [Sidner 79] Sidner, C.L. *Towards a Computational Theory of Definite Anaphora Comprehension in English Discourse*. Ph.D. dissertation, MIT, 1979.