

Strategies for Sequencing as a Planning Task

Daniel D. Suthers

Learning Research and Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh PA 15260
suthers+@pitt.edu

Abstract. The paper summarizes an ongoing investigation of the discourse planning tasks concerned with the sequencing of utterances and their parts. Content selection provides some important constraints on sequencing, most notably those derived from the preconditions and effects of planning operators. However these operators underconstrain sequencing, especially below the granularity at which they interface with a domain knowledge source. Further constraints are available from the integrative processes of the hearer or reader and from working memory limits. Application of these constraints is a matter for discourse planning because the choices relate to one's communicative goals. The planning task is one of translating functionally relevant relationships between units to be ordered into ordering constraints. A collection of strategies for this task are presented. Some of the strategies were used in an earlier implemented system; many are justified by prior psycholinguistic research. Also discussed include current efforts to extend the work to focus structure in general, and to address the handling of conflicts between strategies.

Introduction

The extent to which the components of a text or utterance succeed in carrying out their intended function depends in part on the sequence in which they are realized. For example, a critical aspect of understanding an explanation is to integrate the concepts and propositions in the explanation with existing knowledge. Sequencing decisions should attempt to facilitate this integration and otherwise enhance the intended functionality of the segments of the explanation. Superficially, the "sequential structure" of discourse is simply the order in which its elements are positioned in a linear medium. However, some of the ordering may be arbitrary. In a theoretical analysis, it is more useful to define the sequential structure of discourse as a *partial* ordering that has specific justifications. This paper provides a collection of such justifications in the form of strategies for translating functionally significant relationships between discourse elements into pairwise ordering¹ constraints between those elements.

We begin with a discussion of the nature of the se-

quencing task and the advantages of explicit operators for this task. A number of strategies for the coherent ordering of an explanation are then presented, some of which were used in an earlier implemented system [Suthers 1993a], and many of which are justified by prior psycholinguistic research. Directions for further research are also discussed, including current efforts to extend the work to focus structure in general (i.e., subordination structure as well as sequential structure), and to address the handling of conflicts between strategies (e.g., between centering theory and McKeown's focus preferences).

Approaches to the Sequencing Task

Previous work in generation has handled sequencing decisions in a number of ways. Schematic approaches specify allowable orderings implicitly in terms of the transitions of finite state automata [McKeown 1985]. Non-determinism in these automata has been addressed using focus preferences for selecting from a content pool, these preferences being embodied in a selection mechanism. Other approaches exploited the structure of domain knowledge with mechanisms for traversing data structures representing this knowledge [Paris & McKeown 1986, Sibun 1992]. Planning approaches initially utilized more local yet still schematic specifications of ordering, expressed as preconditions or optional satellites for plan operators [Cawsey 1989, Hovy 1988, Moore 1989]. More recently, partial order causal link (POCL) planning is being applied to discourse planning, with partial ordering derived in a principled manner from the relationships between preconditions and postconditions of plan steps [Young *et al.* 1994].

Content selection processes provide some important constraints on sequential structure, most notably in the form of satisfaction-precedence relations derived from the preconditions and effects of discourse planning operators. However these processes underconstrain sequential structure. This is especially true below the granularity at which the operators interface with a domain knowledge source, because the latter is partially responsible for providing collections of related content that can't be specified by domain-independent operators. For example, a specification that some distinguish-

¹"Ordering" and "sequencing" are used interchangeably.

ing attributes of an entity should be expressed might bind to several nonexclusive alternatives, or a specification for a description of the constituents of an object or process might result in retrieval of a collection of propositions. Some sequencing decisions above the granularity of access to domain knowledge may be underconstrained as well, for example the order in which to express a list of multiple reasons for a conclusion. Further constraints are available, for example from reasoning about the integrative processes of the hearer or reader and the impact of working memory limits on these processes. Application of these constraints is a matter for discourse planning because the choices relate to one's communicative goals, as discussed below. Sequential structure is too important to treat arbitrarily and too context-sensitive to treat in a schematic manner.

Overall, ordering heuristics have not typically been made explicit as causally efficacious plan operators. Explicit ordering operators facilitate the expression and study of alternate theories of sequential structure, and have benefits for planning as well. For example, they enable a planner to tell whether it can achieve a communicative goal by ordering its utterances in a certain way, and they provide a handle for changing strategies in different discourse situations. Suthers [1993a] treated sequencing as a distinct planning task and used sequencing operators that derived ordering constraints from relevant relationships between the elements being ordered. However, prior planning approaches (including that just cited) have not made the *effects* of ordering decisions explicit. Effects can help choose between conflicting operators, and must be considered in reasoning about what communicative goals are achieved by particular orderings.

Choice Between Conflicting Operators. When ordering operators conflict, the operators express principled reasons for choosing between alternate sequential structures. The choice is a matter for discourse planning because it relates to one's goals. The intended effects of conflicting ordering operators can be used to select between them provided that these effects can be related to contextual factors such as superordinate goals and stylistic preferences [DiMarco & Hirst 1993, McCoy & Cheng 1991, Hovy 1990]. For example, a common conflict is between a sequential structure that makes a single entity salient and a sequential structure that flows smoothly from one entity to another ("dovetailing"). Entity salience might be preferred when the entity in question is a topic of the current segment, while dovetailing might be preferred when communicating a relational structure among equally important entities or to make a transition to a new topic. (A definition of

"topic" is forthcoming.) Or consider the ordering of examples with respect to the generality they exemplify. If the examples are presented first, the reader has an opportunity to engage in inductive inference towards the generality, yet may fail to see the relevance of the examples. Presented after a generality, examples provide concrete instances under which the generality may be indexed. The choice depends on whether the speaker or writer is trying to get the hearer or reader to engage in active induction or trying to ease the comprehension process.

Achieving Communicative Goals by Sequential Structure. Some planners (e.g., POCL planners) can notice when a goal can be achieved by actions that have already been planned. However, these planners can only take advantage of existing actions when the explicit effects of a single action meets the goal in question. It should be possible to extend these planners to use ordering operators in an opportunistic manner to identify ways in which goals can be achieved by felicitous ordering of multiple actions, e.g., to achieve communicative goals by implicit relations [Lascarides & Oberlander 1992, Mann & Thompson 1983]. If an operator's effect matches to an active unsatisfied goal and the operator's constraints match existing utterance components or planning can satisfy these constraints with new utterances, then mere installation of the ordering constraint can be used to achieve the communicative goal. For example, suppose a discourse planner has the goals of describing a number of events and the causal relations between them. Communication of the causal relations might be achieved implicitly by describing the events in their causal order.

Sequencing as Exploiting Relationships. The sequencing of two elements of a discourse can only be decided if there exists some relationship between the elements which has implications for their order of expression. The sequencing task is relational, not merely selective. Approaches that treat sequencing as a selective task, for example by using predicates that select the most preferred element out of a set of remaining candidates, are forced to generate a full ordering (the sequence of selected items). The structural aspects of sequencing are not a natural consequence of the model. When sequencing is treated as the accumulation of explicit constraints between elements, partial orderings can be constructed. Derivation of these constraints from relationships between the elements to be ordered addresses the structural aspect of sequencing directly.

The remainder of this paper will present a number of sequencing strategies that illustrate how relationships

between elements yield sequencing constraints. But first some notational preliminaries are required.

Notation for Ordering Strategies

Two aspects of sequential structure are distinguished: precedence and juxtaposition. **Precedence** indicates that one segment should occur *sometime before* another segment. Precedence is significant when prior communication of the contents of one segment facilitate the intended functionality of the contents of another segment. For example, in technical writing definitions of terms usually precede their use. **Juxtaposition** indicates that one segment should occur *next to* another segment in the sequential realization of the explanation. Juxtaposition is significant when the contents of both segments must be in focus of attention at the same time in order for the segments to fulfill their intended communicative function. For example, statements of similarity and difference are usually juxtaposed when making a comparison so that the relative significance of the similarities and differences can be weighed. Any constraint on sequential structure must involve one of precedence or juxtaposition. **Succession** indicates the simultaneous presence of both constraints (i.e., that one segment should occur *immediately before and next to* another segment).

Sequential constraints are placed between text plan elements at three granularities. **Inter-intentional** constraints are placed between intentions to perform rhetorical and communicative acts [Suthers 1993a], and constrain the ordering of the utterance segments that achieve these intentions. For example, an inter-intentional constraint would be used to ensure that a description of the structure of a device precede an account of how the device carries out its function. **Inter-propositional** constraints are placed between propositions and constrain the ordering of sentences, adjectives, relative clauses, and other surface realizations of the propositions. For example, an inter-propositional precedence constraint between (Parallel plates-1) and (Made-of plates-1 metal-1) would allow any of "parallel metal plates," "parallel plates made of metal" or "The plates are parallel. The plates are made of metal." but not "metal parallel plates," "metal plates that are parallel," or "The plates are made of metal. The plates are parallel." (Clearly, this constraint leaves other realization decisions open.) **Intra-propositional** constraints, placed between roles of a proposition, controls voice (i.e., which role filler is expressed as the subject of a clause).² For example, an intra-propositional precedence constraint

could select between "The plates are made of metal." vs. "Metal is what the plates are made of." Juxtaposition constraints between role fillers discourage the insertion of a subordinate clause between the realization of the role fillers.

The ordering strategies are expressed in the form of rules for translating other kinds of relationships into ordering relations. The general form of the rules is:

If S_1 bears relation R to S_2 then $S_1 \xrightarrow{ord} S_2$

where S_1 and S_2 are segments and $S_1 \xrightarrow{ord} S_2$ is one of the following:

$S_1 \xrightarrow{prec} S_2$ for precedence (S_1 occurs sometime before S_2),

$S_1 \xrightarrow{juxt} S_2$ for juxtaposition (S_1 and S_2 are next to each other in either order), and

$S_1 \xrightarrow{succ} S_2$ for succession (S_1 and S_2 are juxtaposed and S_1 precedes S_2).

Propositions are denoted by $(P \tau_1 \tau_2 *)$ where P is a predicate, the τ_i are role fillers, and $*$ denotes 0 or more additional role fillers. No ordering of τ_1 , τ_2 , and $*$ is implied by this notation; in particular $*$ may represent other role fillers that can be expressed before or between τ_1 and τ_2 as well as after them. The notation for intra-propositional ordering is

$(P \tau_1 \xrightarrow{ord} \tau_2 *)$.

This constrains the realizations of the fillers of τ_1 and τ_2 to be sequenced in a manner respecting relation *ord*. For example, if $(P \tau_1 \xrightarrow{prec} \tau_2)$ is expressed as a clause, then τ_1 will be the surface-initial subject.³

Two predicates and a function are needed to express some of the ordering strategies:

End-p(P), true when P is a proposition at the end of a chain of ordered propositions:

$P_1 \xrightarrow{ord} \dots P_n \xrightarrow{ord} P$.

Topic-p(t), true when t is the argument of or a constituent of the argument of a rhetorical goal scoping over the propositions to be ordered. For example, in Describe((Structure capacitor-1)) both capacitor-1 and (Structure capacitor-1) are "topics."

Familiarity(c) = l when a possibly fallible oracle indicates that concept c is assumed to be familiar to the questioner at level l , a member of a *partially ordered* set of levels, perhaps using categories such as in Prince's [1981] taxonomy. $l = \text{false}$ for unfamiliar c .

²In the implemented system of Suthers [1993a], subjects were always surface-initial. Subject and surface-initial are not conflated in forthcoming revisions.

³The notation fails to constrain the realization of the predicate relative to its role fillers, a deficiency that will be addressed in a future revision.

Table 1: Rules for Ordering Strategies

<i>Supplemental Constraints</i>		
Antithesis	If A is an <i>antithesis</i> of T	then $T \xrightarrow{juxt} A$.
Background	If B is <i>background</i> for F	then $B \xrightarrow{prec} F$.
Enrichment	If E is an <i>enrichment</i> of S	then $S \xrightarrow{succ} E$.
Evidence	If E provides <i>evidence</i> for assertions in H	
a	(<i>data-driven task</i>)	then $E \xrightarrow{succ} H$
b	(<i>goal-driven task</i>)	or $H \xrightarrow{succ} E$.
Exemplification	If E provides an <i>example</i> of G	
a	(<i>illustrative use</i>)	then $G \xrightarrow{succ} E$
b	(<i>inductive use</i>)	or $E \xrightarrow{prec} G$.
Motivation	If M provides a <i>motivation</i> for S	then $M \xrightarrow{succ} S$.
Preview	If P provides a <i>preview</i> of B	then $P \xrightarrow{succ} B$.
Summary	If S is a <i>summary</i> of B	then $B \xrightarrow{prec} S$.
<i>Memory and Processing Constraints</i>		
Topic Initial	If Topic-p(t), ($P_1 t *$) is unordered and ($P_2 *$) does not contain t	then ($P_1 t *$) \xrightarrow{prec} ($P_2 *$).
Topic as Subject	If Topic-p(t) in ($P_1 t *$)	then ($P_1 t \xrightarrow{prec} *$)
Dovetailing	If $R = (P_1 x \xrightarrow{ord} y *)$, End-p(R), and ($P_2 y z *$) is unordered	then ($P_1 x y *$) \xrightarrow{succ} ($P_2 y \xrightarrow{prec} z *$).
<i>Domain Knowledge Constraints</i>		
Natural Ordering	If ($N x y$) where N is a Natural-Ordering and x is in the <i>predecessor</i> role of N ,	then ($N x \xrightarrow{prec} y$)
Differentia	If (Subsumption $c s$) and s is differentiated within c by ($P s *$)	then (Subsumption $c s$) \xrightarrow{succ} ($P s *$).
Context	If C is context in which S holds	then $C \xrightarrow{succ} S$
<i>Epistemic Context Constraints</i>		
Familiar First	If ($P f n *$), Familiarity(P), and Familiarity(f) > Familiarity(n)	then ($P f \xrightarrow{prec} n *$).
<i>(See Suthers [1993a,b] for exchange-level constraints on model choice.)</i>		

Table 1 lists the strategies. Selected strategies are discussed below.

Supplemental Constraints

The most obvious constraints on sequential structure are those derived from the inclusion of supplemental material. Supplemental material facilitates the understanding or acceptance of other segments of an explanation in specific ways, the success of which is often affected by order of presentation. Suthers [1993a] used a collection of *supplemental relations*, these being “rhetorical” relations that are primarily intentional rather than informational [Moore & Pollack 1992] and in which one can pre-identify a “nucleus” that is more essential to the goals of the discourse than the other *relata*. Ordering strategies are associated with each sup-

plemental relation in Table 1. Some relations give rise to an unambiguous ordering, and thus have only one rule in the table. The ordering implications of others are complicated by possible differences in tutorial strategy and individual differences in learning style. One advantage of separating ordering decisions from supplemental relationships is the ability to model stylistic differences by changing ordering strategies independently of supplemental strategies.

Background. Background material is that which functions to enable the *comprehension* of nuclear foreground material. Thus background functions best if it precedes the foreground. Succession is not necessary as long as the delay between background and foreground is small enough that the background will not have been forgotten when the foreground is encountered.

Exemplification. Strategic variation is possible in the ordering of examples and illustrations. An explainer can encourage a questioner to engage in inductive inference by giving examples before the generalizations or concepts that they exemplify, as expressed by the $E \xrightarrow{prec} G$ constraint of version b of the rule. Alternately, the example can be given immediately after the concept or generalization being exemplified, as expressed by the $G \xrightarrow{succ} E$ constraint of version a. Under this strategy, the questioner does not have to guess the generalization and will appreciate why the example was introduced. When multiple examples are present other ordering strategies are available for inter-exemplar ordering: see Mittal & Paris [1993] and Rissland [1978].

Motivation. A motivation segment is intended to point out the utility of another segment of an explanation so that the hearer will appreciate the relevance of the motivated segment enough to take it seriously. Motivation to attend doesn't work retroactively, so the motivating segment should occur prior to the motivated segment. Succession is preferred, but not necessary.

Previews and Summaries. By definition, a preview precedes the main body of an explanation. To serve the function of preparing the questioner for the sequence of utterances to follow, a previewed text should be the immediate successor of the preview, because a preview sets up an expectation that the subsequent segments will be those mentioned in the preview. Violation of this expectation with intervening material can cause confusion. A summary is similar to a preview in that both provide skeletal characterization of the main body of an explanation, though summaries can refer back to content that was not available at the time of a preview. The pedagogical utility of a summary is in repetition and consolidation. Succession is not as important for summaries. In fact, a summary might be used *because* there is some extra material between the segments related by the summary relation: the summary functions to refocus on the main points after the digression.

Memory and Processing Constraints

The ordering operators of this section rely on a few processing assumptions that have been supported in the psycholinguistic literature. Memory (retention and retrieval) is better for "integrated" items, i.e. those that the subject can relate to other prior knowledge [Keenan *et al.* 1984, Kintsch & van Dijk 1978]. New integrative links are constructed in a limited working memory during comprehension [Kintsch & van Dijk 1978]. There is a cost (allocation of attention, probability of error) associated with changing the contents of this working mem-

ory, and the longer an item is kept in working memory the more likely it is to be encoded in long term memory. Finally, subjects attempt to identify thematically central entities and use these as the default locus of integration when attempting to integrate new material [Carpenter & Just 1977].

Comments on "topic" and "theme" may be helpful at this point. Lavid & Hovy (unpublished working paper) define "theme" as "that element that informs the listener as discourse unfolds how to relate the incoming information to what is already known." The topic predicate is *not* intended as a generally applicable definition of thematic elements outside its use in the ordering rules. It merely provides candidates for being made thematic elements though appropriate ordering and other devices.

When a questioner asks a question, the topics of the query are brought into the focus of attention. Their focal status motivates the relevance of assertions made and concepts introduced by the explainer in the response. If assertions or concepts that had no apparent relation to the topics were introduced, the questioner might be unable to integrate them and could become confused due to the conversational implicature of the apparent change in subject. For example:

"What killed the dinosaurs?"

"Many rocks at the KT-boundary have an unusual concentration of iridium ...

(The iridium poisoned them? The speaker doesn't want to talk about dinosaur demise?)"

In contrast, the following explanation changes focus of attention from the question's topic to other concepts and propositions in a well connected manner:

"The dinosaur extinctions may have been caused by a huge meteorite. Evidence for such a meteorite is provided by an unusual concentration of iridium found in KT-boundary rocks. ..."

This example illustrates the next three rules: the response starts with a proposition about the topic, the topic is in the subject position, and new concepts are introduced by their relation to prior concepts.

Topic Initial and Topic as Subject. The "topic initial" rule specifies that propositions involving a topic entity *t* are to be expressed before propositions not involving a topic. As illustrated above, in situations where the topic has been pre-identified, this rule ensures continuity with the recipient's expectations. In other situations, first-mention aids the subject's identification of the topic [Kieras & Boviar 1981, Gernsbacher & Hargreaves 1988], helping to ensure that facts about a particular entity among the many mentioned are remembered.

The “topic as subject” rule marks the topic as the surface-initial subject of any clause it occurs in. Repetition in the subject position is another way to facilitate the hearer’s or reader’s identification of the topic [Kieras & Boviar 1981]. This rule can also help smooth over violations of the following “dovetailing” rule by returning to a concept that is presumably easy to reactivate.

Dovetailing. This strategy relies on a combination of “argument overlap” [Kintsch & van Dijk 1978] and “given-new” [Haviland & Clark 1974]. The inter-propositional $P_1 \xrightarrow{succ} P_2$ constraint introduces a proposition P_2 when it involves a concept or proposition that has been brought into immediate focus of attention by another ordering decision to express P_1 (i.e., there is argument overlap). The intra-propositional $y \xrightarrow{prec} z$ constraint makes the surface-initial subject of each proposition be the role filler by which it was introduced (i.e., proceeds from given to new). Dovetailing is intended to minimize working memory changes and maintain the connectedness of the subject’s evolving conceptual model or “text base.”

Dovetailing can be iterated on z in P_2 to produce chains. However, overapplication of dovetailing risks obtaining stream-of-consciousness texts which lack thematic development. Topic salience and dovetailing can be in conflict. One manifestation of this conflict is the conflict between McKeown’s [1985] preference of “change” over “maintain” over “return” in order to avoid having to reintroduce entities one has more to say about (based on Sidner [1979]) and Gordon *et al.*’s [1993] preference of “continuing” over “retaining” over “shifting” in order to maintain local coherence. Further work is required to identify how the choice of one strategy over the other depends on register and task demands, and to examine the interaction with other factors such as location in the discourse. For example, the author expects that topic salience will be preferred in contexts where the theme is being established (e.g., early in a document, paragraph, or other new discourse segment) while dovetailing might be preferred once the theme is established and a transition to a subordinate theme is needed.

Domain Knowledge Constraints

Now we consider constraints derived from relationships in the domain knowledge being expressed.

Natural Orderings. Temporal and causal relations are normally experienced in a particular direction, for example from prior to posterior events or from cause to effect. The assumption that our cognitive apparatus is adapted to more easily use these relations in the “for-

ward” direction suggests that predicates categorized as “natural orderings” [Bienkowski 1986] be expressed with the prior event or cause as the subject, for example “ X caused y ” is preferred to “ y is caused by x ,” all other things being equal. This strategy is consistent with psycholinguistic evidence indicating that reverse causal and temporal ordering inhibits comprehension [Irwin 1980] and disrupts thematic processing [Townsend 1983]. It also facilitates the hearer’s identification of implicitly expressed temporal relations [Lascarides & Oberlander 1992]. Related strategies are available for spatial descriptions [Linde 1974, Sibun 1992].

Differentia and Context. The “differentia” relation holds between two propositions when one proposition ($P s *$) differentiates a subclass s from other subclasses of a class c . An explainer chooses the statement ($P s *$) from amongst all the possible predicates one could apply to s because P distinguishes s from the other subdivisions of c that the questioner might know about. The questioner cannot recognize or assess this significance of ($P s *$) unless he or she has been informed of the “contrast class” against which the claim ($P s *$) is being made. This rule suggests that the contrast class c be introduced first, for example:

“An electric field is a kind of force field that applies a force to a charged object.”

is preferred over

“An electric field applies a force to a charged object and is a kind of force field.”

Mere precedence is insufficient because the genus provides the context in which the differentia is meaningful. Succession places them both in focus of attention at once. This strategy can be generalized to the level of sibling communicative acts or rhetorical intentions [Suthers 1993a], and to other context/statement relationships, as suggested by the “Context” rule in Table 1.

“Epistemic Context” Constraints

Suthers [1993a,b] discusses how the “epistemic context” (the knowledge available to the explainer and questioner and the knowledge shared in prior dialogue) influences the choice between alternate domain models on which to base an explanation. Some of the “preferences” presented in these publications address sequential concerns. For example, when preferences to “say something new,” “minimize new propositions,” and “elaborate on focal models” are applied together in a dialogue about some phenomenon, incremental construction of increasingly elaborate domain models of the phenomenon will result [Suthers *et al.* 1992]. In general, the epistemic context

provides important constraints on sequential structure across *multiple* exchanges but has less impact on sequential structure within a *single* utterance. This paper does not discuss constraints across exchanges. Only one constraint originating in assumptions about hearer familiarity is discussed.

Familiar First. New concepts can be introduced in relation to familiar ones using any domain relation. Suppose concept f is familiar and n is new. Then any proposition ($P f n *$) will do the job, provided the predicate P itself is familiar. (An unfamiliar predicate won't be much help in integrating an unfamiliar concept.) The strategy assumes an ordering of familiarity levels. If f is more familiar than n , it installs an intra-propositional constraint that f should be expressed surface-initial, yielding expressions of form " F is P -related to n ." (not " n is P^{-1} -related to f "). Once familiar-first has been applied, dovetailing can be used to introduce other unfamiliar concepts.

The assumption behind familiar-first is that it is easier to retrieve a known concept and integrate a new concept in relation to it than it is to construct a new concept from scratch and subsequently retrieve a known concept to which it can be integrated. However, at and below the clause level the delay between the introduction of n and f may be so small that this strategy does not have a discernible effect. Also, if n is a topic, "topic as subject" may be more relevant to the speaker's goals. These are questions for empirical work.

Closing Comments

The sequencing of expository texts and speech should be chosen to enhance the intended functionality of each textual unit and to facilitate the questioner's integration of the communicated information. This aspect of the planning task is usefully seen as one of translating functionally relevant relationships between textual units into ordering constraints. The strategies presented in this paper were derived from examination of example explanations and found to be necessary to enable an automated explanation generator to produce coherently sequenced explanations [Suthers 1993a]. (Space constraints necessitate leaving a full example to the work just cited.) Many of the strategies were subsequently found to correspond to results in the psycholinguistic literature.

As discussed previously, the strategies can conflict. Suthers [1993a] handled conflicts with a simple preference ordering. This approach is inadequate because conflict resolution is expected to depend on contextual factors such as the speaker or writer's goals and the relationships of the elements to be ordered with respect

to discourse segment boundaries. The author is currently planning psycholinguistic experiments to test the impact of the strategies through reading time and recall studies, with particular concern for how the resolution of conflicts between topic salience and other heuristics should be sensitive to discourse context. Another question for future work is the extent to which the sequencing task fits top-down models of planning such as DPOCL [Young *et al.* 1994] vs. requiring a distinct mechanism for the application of conflicting heuristics. The modeling of working memory limits and memory decay in DPOCL operators would require generalizing the POCL notion of a "threat" to be a matter of degree rather than absolute. The author suspects that important generalities will be easier to capture if factored out and expressed as explicit sequencing operators rather than manifested in variations of multiple decompositional operators.

The foregoing work is being extended to include the subordination structure of an utterance as well as its sequential structure. "Focus trees" [McCoy & Cheng 1991, Hovy & McCoy 1989] will be used to represent the combined sequential and subordination structure, with partial orderings placed between siblings at each level of the hierarchy. The nodes of a focus tree represent units of a text at various granularities ranging from intentionally defined segments of the sort discussed by Grosz & Sidner [1986] down through clausal propositions to predicates and entity references. "Focusing operators," including versions of the sequencing heuristics of this paper rewritten in the focus tree notation, will manipulate the tree structure by subordinating one or more subtree to another or by installing cross-links between siblings to constrain possible traversals. This reformulation is expected to be an improvement for several reasons. It allows the expression of heuristics for the subordination aspects of focus structure, not just sequencing, and thematic development can take place at multiple granularities. Only one ordering relation, precedence, is required. Juxtaposition is handled by grouping within a subordinate structure rather than by a different ordering relation that requires special interpretation. An availability metric can be defined in terms of distance to search back through the tree. Finally, the notation can be used in the analysis of texts, and promises to support application of the sequencing heuristics to text revision as well as text generation.

Acknowledgements

This research was conducted while supported by grant MDR-9155715 from the National Science Foundation Applications of Advanced Technology program. The author thanks the anonymous reviewers for thoughtful comments.

References

- [Bienkowski 1986] M. Bienkowski. A Computational Model for Extemporaneous Elaborations, CSL Report 1, Cognitive Science Laboratory, Princeton University, Princeton, N.J., 1986.
- [Carpenter & Just 1977] P. A. Carpenter & M. A. Just. Integrative processes in comprehension. In D. Laberge & S. J. Samuels, *Basic Processes in Reading: Perception and Comprehension*, pp. 217-241. Hillsdale, N.J.: LEA, 1977.
- [Cawsey 1989] A. Cawsey. *Generating Explanatory Discourse: A Plan-Based, Interactive Approach*. Ph.D. Dissertation, University of Edinburgh, 1989.
- [DiMarco & Hirst 1993] C. DiMarco & G. Hirst. A computational theory of goal-directed style in syntax. *Computational Linguistics* 19:451-499, 1993.
- [Gordon et al. 1992] P. Gordon, B. Gross, & L. Gillion. Pronouns, names, and the centering of attention in discourse. *Cognitive Science* 17:311-347, 1992.
- [Gross & Sidner 1986] B. Gross & C. Sidner. Attention, intention, and the structure of discourse. *Computational Linguistics* 12:175-204, 1986.
- [Haviland & Clark 1974] S. E. Haviland & H. H. Clark 1974. What's new? Acquiring new information as a process in comprehension. *Journal of Verbal Learning and Verbal Behavior* 13:512-521, 1974.
- [Hovy 1988] E. Hovy. Approaches to the planning of coherent text. Presented at the 4th International Workshop on Text Generation. Catalina Island, California, July 1988.
- [Hovy 1990] E. Hovy. Pragmatics and natural language generation. *Artificial Intelligence* 43:153-197, 1990.
- [Hovy & McCoy 1989] E. Hovy & K. McCoy. Focusing your RST: A step toward generating coherent multi-sentential text. *Proc. 11th Annual Conference of the Cognitive Science Society*, pp. 667-674. Ann Arbor, Michigan, 1989.
- [Keenan et al. 1984] J. Keenan, S. Baillet, & P. Brown. The effects of causal cohesion on comprehension and memory. *Journal of Verbal Learning and Verbal Behavior* 23:115-126, 1984.
- [Kintsch & van Dijk 1978] W. Kintsch & T. A. van Dijk. Toward a model of text comprehension and production. *Psychological Review* 85:363-394, 1978.
- [Lascarides & Oberlander 1992] A. Lascarides & J. Oberlander. Abducing Temporal Discourse. In R. Dale, E. Hovy, D. Rösner, & O. Stock (Eds.) *Aspects of Automated Natural Language Generation*, pp. 167-182. Berlin: Springer-Verlag, 1992.
- [Linde 1974] C. Linde. *The Linguistic Encoding of Spatial Information*. Doctoral Dissertation, Columbia University, 1974.
- [Mann & Thompson 1983] W. C. Mann & S. A. Thompson. Relational Propositions in Discourse. *ISI/RR-83-115*. Information Sciences Institute, University of Southern California, Marina del Rey, California, 1983.
- [McCoy & Cheng 1991] K. McCoy & J. Cheng. Focus of attention: Constraining what can be said next. In C. Paris, W. Swartout, & W. Mann (Eds.) *Natural Language Generation in Artificial Intelligence and Computational Linguistics*, pp. 103-124. Boston: Kluwer Academic, 1991.
- [McKeown 1985] K. McKeown. Discourse strategies for generating natural language text. *Artificial Intelligence* 27:1-41, 1985.
- [Mittal & Paris 1993] V. Mittal & C. Paris. Automatic document generation: The interaction of text and examples. *Proc. Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, pp. 1158-1163. Chambéry, France, 1993.
- [Moore 1989] J. Moore. *A Reactive Approach to Explanation in Expert and Advice-giving Systems*. Ph.D. Dissertation, University of California, Los Angeles, 1989.
- [Paris & McKeown 1986] C. Paris & K. McKeown. Discourse strategies for describing complex physical objects; *Proc. NATO Advanced Research Workshop on Natural Language Generation*. Nijmegen, The Netherlands, August 19-23, 1986.
- [Prince 1981] E. F. Prince. Toward a taxonomy of given-new information. In Peter Cole (Ed.), *Radical Pragmatics*, pp. 223-255. NY: Academic Press, 1981.
- [Rissland 1978] Published under E. Rissland Michener. Understanding Understanding Mathematics. *Cognitive Science* 2:361-383, 1978.
- [Sibun 1992] P. Sibun. Generating text without trees. *Computational Intelligence* 8, 1992.
- [Sidner 1979] C. Sidner. Focusing in the comprehension of definite anaphora. In M. Brady & R. Berwick (Eds.), *Computational Models of Discourse*, pp. 267-330. Cambridge, Mass: MIT Press, 1979.
- [Suthers 1993a] D. Suthers. *An Analysis of Explanation and its Implications for the Design of Explanation Planners*. Ph.D. Dissertation, Dept. of Computer Science, University of Massachusetts, 1993.
- [Suthers 1993b] D. Suthers. Preferences for Model Selection in Explanation. *Proc. Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, pp. 1208-1213. Chambéry, France, 1993.
- [Suthers et al. 1992] D. Suthers, M. Cornell, & B. Woolf. Steps from explanation planning to model construction dialogues. *Proc. Tenth National Conference on Artificial Intelligence (AAAI-92)*, pp. 24-30. San Jose, CA, 1992.
- [Young et al. 1994] R. Young, J. Moore, & M. Pollack. Towards a principled representation for discourse plans. To appear in *Proc. Sixteenth Annual Conference of the Cognitive Science Society*, Atlanta, GA, 1994.