Selection: Salience, Relevance and the Coupling between Domain-Level Tasks and Text Planning

T. Pattabhiraman
Nick Cercone

Centre for Systems Science
Simon Fraser University
Burnaby, B.C., Canada V5A 1S6
patta@cs.sfu.ca and *nick@cs.sfu.ca*

Abstract

In this paper we examine some issues pertaining to the task of selection in text planning. We attempt to distinguish salience and relevance, and characterize their role as important fundamental notions governing selection. We also formulate the problem of selection of text content in terms of the coupling between domain-level tasks and text planning tasks. We describe our research on generating bus route descriptions.

Keywords: Natural Language Generation, Text Planning, Selection, Salience, Relevance, Coupling, Route Descriptions

1. Introduction

Most models of natural language generation, be they computational or psychological, recognize that the task of text planning (also called conceptualization [14], [10]), comprises the following essential subtasks: (1) *content selection* and (2) *content organization.* The task of content selection (hereafter, selection) is concerned with computing, or retrieving from a knowledge base, the primary content of texts. Selection may also include a phase of *content expansion*, also called additional topic inclusion [8]. Content expansion consists in inferentially generating or selecting additional material for expression in text, once the primary material is selected. The task of content organization, also called topic organization [8], content ordering or linearization [13], [14] deals with ordering the chosen content into a sequence of propositional (linear) representations, appropriate for realization into coherent text.

Many text planning models concentrate chiefly on content organization and perhaps content expansion. Such planners accept as input a pre-selected, and to some extent, pre-sequenced collection of representations that serve as the raw material of text content. The knowledge base from which these input elements are selected, as well as the selection process itself, are treated to be external to the text planner. In such systems, whatever selection is performed by the text planner is confined to the task of choosing a subset of the input elements according to control exercised by knowledge sources resident within the text planner. Other examples of systems in which pre-selected content is input to the text planner include natural language generation front-ends to expert systems, database systems or other application programs- in other words, a problem-solver or a host system which is devoted to domain-level non-linguistic activities.

In some other models, such as those of Paris [23], the knowledge base from which much of the text content is drawn is resident within the text generation system itself. Selection in such systems is totally a responsibility of the text planner. In the TAILOR system of [23], for instance, facts describing objects are stored in a *knowledge base*, and the *textual component* selects the content of the description from the knowledge base under the regulatory influence of a *user model.* The urge to generate text is input to the generator in the form of a request for the definition of an object.

In all cases, the beholder of the natural language output sees the text as coming from a single program which can perform domain-level tasks as well as text production tasks: *the speaker and the thinker are one and the same.* If the beholder were a text planning researcher, she could be inclined to pose questions about the origin of text content. Our recent research has been motivated and directed by the adoption of such a role. The problem of selection in texts (in our case, multisentential descriptive texts) led us to examine the nature of input to the text planner, and the boundary between the domain-level program and the text planner. When we regarded the thinker and the speaker as a unified whole, we were led to search for very general factors that influenced selection in diverse domains of discourse production. When we viewed domain-level activities and text planning activities as distinct tasks, we examined the division of duties between the text planner and its underlying program in the task of selection, and explored the conditions under which the modular boundary between the domain-level and the text-planning level could be kept intact and those under which it might break down.

This paper is devoted to a presentation of our research on some issues pertinent to selection in text planning. We found that the fundamental notions that were crucial to understanding selection in text planning were *salience* and *relevance.* These are not altogether unfamiliar notions. Conklin and McDonald [2] and Waltz [33] describe some effects of salience in generating scene descriptions. Researchers in natural language *understanding* have studied salience and relevance in some detail and have used them profitably in their accounts. However, in the natural language generation research community, the terms are used in their literal sense, and often interchangeably, whereas in fact they are distinct notions. In this paper we present the notions of salience and relevance as they pertain to natural language generation, in particular, to (content) selection.

Our presentation in section 2 is a synthesis of several analyses of salience and relevance in the disciplines of language understanding, psycholinguistics and communication. In section 3 we consider the coupling between domain-level tasks and text planning tasks from the point of view of selection. Our work in the domain of route description generation is presented in section 4. In this domain, interesting questions emerge regarding mode of knowledge representation, connections between text planning and domain-level problem solving, and selection. In the concluding section we briefly state our current work and research plans for the near future.

2. Salience and Relevance

Salience and relevance are theoretical notions which are influential in accounting for how or why certain objects, concepts, properties or actions are highlighted or preferred in natural language processing. Of particular interest to the subject of this paper is the role of these important notions in controlling selection (and omission) in text planning. What is common to salience and relevance is their role as determinants of decisions in selection. However, salience is connected with speaker-external objects or properties, while relevance is related to speaker-internal factors such as goals and motivation.

2.1 Salience

Salience explains the prioritization or foregrounding of objects or information. Such prioritization may arise from direct perception of the world or context, or from conceptual knowledge, built up by shared perceptual experiences of the speaker and the hearer. According to Sridhar [30] and Levelt [14], salience guides the direction of the speaker's attention by making certain aspects of a situation stand out relative to others. By virtue of being ascribed to factors of context, situation or *background*, salience is a relatively fixed notion, invariant with respect to specific perceivers. Salience information can therefore be lexicalized, stored as an invariant component of the knowledge base, incorporated as a structural property of the design of the knowledge base, and/or built into the control processes that select text content from the knowledge base.

Perceptual salience arises from the prominence of *external* characteristics such as size (bigness), amplitude of sound (loudness), colour (brightness), etc. Osgood and Bock [22] call this factor *vividness*. Conceptual salience originates in (shared) experiential knowledge. For example, the slyness of foxes, the strength of gorillas, and other such distinguishing characteristics are conceptually salient [4]. The head of a queue of people is a conceptually salient part of the queue with respect to its function [7]. Conceptual salience also arises from the degree of unexpectedness, unusualness or *deviation from the norm* of an object or event, giving the object or event of description its *information value* and thus motivating its selection in the planned text. Salience is attributed to properties of objects as well as to components of complex objects.

An important implication of salience for objects that possess it is the ease of their availability in the knowledge base or knowledge structure [34]. The higher the level of salience of an object, the higher is its level of activation in the speaker's mind, and therefore, the greater is its probability of being selected in the speaker's description, and the earlier is it likely to be mentioned in the description. This fact has been recognized by Conklin and McDonald [2], Osgood [21], Sridhar [30] and Levelt [14].

Salience and Text Planning. The significance of visual salience in text planning has been pointed out by Conklin and McDonald [2] in their empirical study involving scene descriptions (of suburban houses) and the GENARO program that generates similar descriptions. In GENARO, the chosen representation for the set of objects is a list, ordered according to decreasing levels of salience of the objects. Non-salient objects are not represented. This list mediates between the deep generation component and the domain data base in regulating the order of description of objects.

Other Work on Salience. Sridhar [30] amplifies upon earlier psycholinguistic treatments of the role of salience in single sentence production. He examines the effects of salience on intrasentential syntactic ordering in phenomena like passivization, topicalization and constituent order shift. His results are of significance to the realization component of natural language generation. On the subject of content selection in the description of states and events, Sridhar hypothesizes that actions are more salient than state changes, which in turn are more salient than constant states. Ortony [20] invokes salience in explaining the interpretation of metaphor. Herskovits [7] uses salience in her account of prepositional semantics and the construction of mental models from prepositional phrases.

2.2 Relevance

Relevance constrains participants' utterances and interpretations in communicative contexts. It connects texts and their generators through goals: speakers are assumed to generate utterances which they believe to be relevant to their goals [1]. It connects the context of text generation with the generator and the text by the fact that communication as well as the goals themselves are situated in context, and by the assumption that the conversational (or communicative) goals are relevant to the generator's personal goals. The speaker's communicative goals are characterized by the speaker wanting the goals to be made known to the hearer, and intending to reach them by the hearer knowing and adopting them [25]. Relevance is thus an important notion that connects texts, contexts, speakers and hearers. Selecting relevant information in a text planning scenario is thus an issue of selecting information that is pertinent to the generator's communicative goals.

Leech [12] defines relevance as follows:

An utterance is relevant to a speech situation if it can be interpreted as contributing to the conversational goals of the speaker or hearer.

Haslett [6] modifies this definition to account for the communicative actions and knowledge. Her definition of relevance is stated as follows:

An utterance, action or unit of knowledge is relevant to a *speech situation if it can be interpreted as contributing to the communicative goals of the speaker or listener.*

Relevance and Coherence. Haslett [6] emphasizes relevance over coherence as the essential property of texts emerging in communicative contexts. While coherence is essential in interpreting texts, it is not sufficient in accounting for how texts are planned and generated. Texts make sense not only in the general way of being interpretable, but also in the particular way of contributing to conversational goals. Generators try to make their texts both relevant and coherent. A relevant text must also be coherent in the sense of being meaningful to the conversational goals. However, not all coherent texts need be relevant (cf. retorts like *yes, but that's not relevant to the issue*). In their work on generating expert system explanations, Moore and Swartout [19] implicitly affirm the precedence of relevance over coherence. Although the finer details of how they generate explanations are not directly relevant to the concerns of this paper, their stress on the role of the generator's goals in accounting for selected text content is of significance. A similar view is evident in Paris [24].

Relevance and Selection. Haslett's definition is suited to our purposes, since it explicitly recognizes the role of the knowledge sources (representations) as well as the processes involved in selecting text content. Knowledge could be represented in any number of ways. However the definition of Haslett has implications for the design of particular text planners and specific models of text planning: a very natural design (for the knowledge base that stores much of what is eventually expressed in text, as well as for the processes that select from the knowledge base) will be one which aids, and is in harmony with, the goal structures present in the given setting or domain of text generation. Mann and Moore [15], in their work on the KDS system of text generation, point out the importance of relevance in judging whether an item of knowledge should be selected for expression in text. Their method, however, has numerous deficiencies, a discussion of which is presented in [15]. Paris [24] also recognizes explicitly the relation between selection and relevance through goals.

Other Work on Relevance. Sperber and Wilson [29] stress the role of relevance in language comprehension. Their work is partly a response to the shortcomings of the semiotic code model of verbal communication. They emphasize the place of intended inference in verbal communication, over and above the encoding and decoding of messages. The reader is referred to [6] for a critique of the theory of Sperber and Wilson from the point of view of relevance. van Dijk and Kintsch [32] and Schank [28] treat the more *local*, restricted issue of topical relevance in sentences.

2.3 Relevance and Salience Contrasted

Salience is a speaker-external, contextual principle, while relevance is a speaker-internal principle related to the speaker's goals. Waltz [33] states that in generating scene descriptions, what the speaker notices is a function of external factors (scene characteristics) and internal factors (like goals). Thus he effectively recognizes the distinction between salience and relevance, though he doesn't use the term *relevance* to describe the external factors. Similarly, psycholinguists like Sridhar [30] use the term *salience* to describe *relevance* as presented here. Notions derived from salience and relevance include *focus* and *attention.*

Not all settings of text production may admit characterization of salience, but all settings involve relevance, to the extent that text production in the given setting is purposive. However, relevance can be a nebulous matter in many text production settings. Relevance and salience could sometimes reinforce each other and at other times, conflict with each other. When they do conflict, relevance, given Haslett's definition, would have greater weight in its dictates on selection decisions in particular and language generation decisions in general. This observation is in accord with Sridhar's psycholinguistic findings on single sentence production.

3. Domain-Level Tasks and Text Planning Tasks

We have noted relevance and salience as fundamental notions that guide selection of text content. The effects of relevance and salience can be brought out in text planning systems in several ways, of which one is sketched below. We assume that the primary content of texts is drawn from a *knowledge base*, which may consist of several component modules in the implementation.

A large body of surely-irrelevant knowledge is not represented. This corresponds to a *closed world assumption* for knowledge bases with respect to relevance and salience. The knowledge base only contains potentially relevant and salient items for all likely instances of text production. Selection is still responsible for choosing only the portion of the knowledge base that is pertinent to a *particular instance* of text planning. The principles of relevance and salience, as may be embodied in a text planner, are thus a matter of *representation* as well as *control.* The control could be exercised by a variety of knowledge modules, such as the user model (as in the case of Paris [23]), the hearer model (in effect, the user model) (as in the case of Dale [3]), etc. What needs to be said is selected; at the same time, what should be omitted is skipped, by the search process turning a blind eye to irrelevant knowledge, as it were. Selection processes mediate between the controlling knowledge sources and the knowledge base from which content is selected.

The texts produced by such systems may be perceived to be emerging from communicative goals, though communicative goals themselves are not explicitly represented. In the text planning models of [3] and [23], communicative goals need not be explicitly represented, since relevance information is implicit in the semantics and the content of the user model. When do the goals themselves require explicit representation and reasoning in the selection process? This question directs our analysis towards two aspects:

1. the nature of the text produced
2. the connection (or, *coupling*) between domain-level tasks and text planning tasks

3.1 Explanation Generation

The text could be self-referential, in that (for instance) the generator might allude to, justify or explain the prior content of text that it produced or the actions that were performed at the non-linguistic, domain-level. (The identity of meta-language with object-language is a quintessential characteristic of natural languages). An exemplary case in point is that of explanation generation in expert systems. Explanation systems explicitly represent goals so as to explain their domain-level actions, or to explain their own prior utterances, for instance in answering follow-up questions in dialogues with the user, as in the works of Moore and Paris [18], Paris [24] and Moore and Swartout [19]. The explanation system of McKeown [16] identifies its goals with the user's goals, which are hence represented in detail.

3.2 Coupling between Domain-Level Tasks and Text Planning Tasks

We observed in the introduction that the speaker and the thinker are one and the same, in that a single system is perceived to perform domain-level activities such as problem-solving, action planning, etc, as well as text production tasks. In this context, we could speak of *loose coupling* and *tight coupling* between domain-level and text planning tasks. Loose coupling implies modularization between the two tasks in process models and implementations. On the one hand we could ask whether the domain-level tasks are in some sense *naturally* coupled to text planning tasks, and on the other, we could ask what the chosen theoretical paradigm, modelling method, implementation or application has to do with the coupling.

Assuredly, the knowledge of relevance and salience that is necessary in selecting text content is *conceptually* distributed between the non-linguistic, domain-level and the text production level. Therefore, the knowledge and the processes involved in selecting are, in principle, distributed at both levels. In actual implementations or computational process models, if the coupling between domain-level tasks and text planning tasks is loose, then additional work will have to be done to make the coupling appear tight (i.e., to make the thinker identical with the speaker). This issue arises in the construction of explanation facilities for expert systems and front-end natural language generators for various application programs. *Virtual* tight coupling is achieved by building interfaces, or by augmenting the design of the expert system and/or the text planner.

In the Explainable Expert System architecture described in Paris [24], this is achieved by designing expert systems with explanation in mind, and by using for explanations the support knowledge applied in deriving the expert system. Problem-solving knowledge is kept insulated in the implementation from the knowledge necessary for generating explanations. Selection in the explanation task is decided by the text planner and mediated by the use of a rich text plan language, described in Moore and Paris [18]. The design of powerful plan languages (for selecting text content from the knowledge base) is an interesting research problem [9].

In the EPICURE recipe-generation system of Dale [3], tight coupling is achieved by modelling discourse plans to be isomorphic to domain-level plans. There is no domain-level activity proceeding independently of (separately from) text planning. In the work of Appelt [1], tight coupling is ensured in the theoretical paradigm of *action planning* in which both domain-level tasks and utterance production (linguistic) tasks are uniformly viewed as *goal-directed actions*. Content selection in Appelt's system is distributed in a complex way throughout the generation process by commitment to the view that the task of *what to say* is inseparable from the task of *how to say*. What about complex domains of language generation in which domain-level activities are naturally coupled to text planning activities, irrespective of the theoretical paradigm that may be used? In the next section we turn our attention to one such domain, viz., that of route communication.

4. Generating Route Descriptions

In this section we describe the domain of route description generation, and report some aspects of our research on generating descriptions of bus route directions from a given source to an intended destination within a city. Our work includes an implementation in CProlog of a prototype system that generates descriptions of bus route directions in Vancouver. A representation of the city or region in which route description generation occurs is available in the form of a *cognitive map*. The map serves as the knowledge base in the domain-level activity of finding a route, as well as in providing the primary content of route description texts.

4.1 Route Communication

The verbal protocol that consists of requesting route directions and giving them is termed *route communication* by Klein [11]. Klein uses route communications that occurred naturally on the streets of Frankfurt in exploring the relation between the *cognitive* (domain-level) task of route finding and the *linguistic* task of generating a route description as manifested in the choice of local deictic terms like *here*, *there*, *left* and *right*. Wunderlich and Reinelt [35] deal with the discourse structure of route communication. Although full route communication occurs as a dialogue, the roles of the participants are not symmetric as they would be, say, in a casual conversation. All varieties of route communication involve selecting from maps. However, the nature of information selected, and in particular the kind of spatial information included in the route description depend on such factors as the mode of transportation used, whether the questioner herself will be navigating or whether she will be using a public transport system, etc.

Bus Route Descriptions. In our research on generating bus route directions we emulated the methods of Klein [11] and Wunderlich and Reinelt [35] to gather and analyze natural text. Our primary linguistic data consist of over 40 bus route descriptions to various destinations in Vancouver, and include written (printed) as well as spoken descriptions. We concentrate on the *route description* phase of the dialogue, which emerges more or less as monological text. Route descriptions involve knowledge of different *granularities*. For instance, when the full journey involves taking two or more buses, the connecting buses may be available at the point of disembarkation (from the previous bus) itself, or just across the street. But often, one may have to walk a fair distance before transferring to the next bus. In such cases, the bus route description also includes walk-route information, as illustrated by the italicized portion of the following text:

from Kingsway and Edmonds you want to catch bus number 106 called Metrotown ... take it down to the

Edmonds skytrain station ... take the skytrain as far as Burrard station ... *and at Burrard station you want to walk north on Burrard one block ... as far as West Pender* ... and on West Pender going westbound it's bus number 19 and it's called Stanley Park

Route Descriptions vs Route Sketches. Route descriptions are also interesting in that they can also be communicated non-verbally (graphically) in the form of a route sketch. The following spoken text describes how to get from *Lougheed Mall* to *Grouse Mountain* in Vancouver. The "..."s correspond to pauses or confirmatory expressions like *ok* or *yeah* by the questioner.

from Lougheed Mall? ... ok, you could catch a 151 or a 152 called Vancouver ... get off at Hastings and Kootenay right by the Kootenay Loop ... on Hastings in front of the loop transfer to bus number 28 called Phibbs Exchange ... will cross Second Narrows Bridge to Phibbs Exchange on the other side ... and that's where you could catch the 232 Grouse Mountain bus ... and it goes right up to Grouse Mountain.

The information conveyed by this text could also be expressed graphically as illustrated in figure 1. There are interesting parallels between natural language route descriptions and route sketches as bearers of information. A one-to-one correspondence may be seen between certain aspects of route sketches and natural language route descriptions. However, route descriptions are linear, and the use of certain mechanisms such as connectives, pronominals and deictic terms is exclusively associated with the use of natural language as the medium of description.

4.2 Planning and Discourse Units in Route Description Generation

A journey consists of a series of connected *route sections*. However, not all route sections need be described in a route description. For instance, when a public transportation system is used, all the details of the turns taken by the conveyance used are not relevant for the description. However, when one has to drive or walk, information on each route section and turn should be communicated. Thus, route planning units at the domain level are in general more detailed than route description units at the discourse level.

Route descriptions consist of a sequence of units, each specifying, *at the very least*, a source, a destination and a route label. In the case of bus route descriptions, such a unit consists of a specification of a boarding point, a disembarkation point and the bus label for a single bus ride. Such units provide the *skeletal plans* for route description texts. In figure 1, the skeletal plans correspond to the connected sequence of arrows marked with the bus labels. The skeletal plans are computed by the route finder using the knowledge base (map), and are input to the text planner. The text planner may augment the skeletal plans by selecting from the knowledge base additional descriptions of landmarks, location, orientation and so on.

4.3 Knowledge Base and Route Finder

We have described the design of the knowledge base and the route finder in [26]. For discussions of issues on the form, content, function and structure of spatial knowledge in the domain of route description generation, the reader is referred to Habel [5], and to experimental works such as those of McNamara [17] and Thorndyke [31].

As noted earlier, the cognitive (problem-solving) task of route finding and the linguistic task of description generation are simultaneously manifested in route communication. We do not make specific cognitive claims regarding how the route finding and description generation tasks are interleaved. In instances of human route communication, one may become aware of the final sections of the route even before the initial sections are realized into descriptions. However, a route section will have to be determined before it is realized. Accordingly, we separate route finding and route description generation, and feed the text planner with the sequence of skeletal plans (described in section 4.2) furnished by the route finder.

The route finder is *specialized* for the description generation task, in the sense that its output units correspond to the discourse units (skeletal plans) of the descriptions. The sequence of skeletal plans processed by the text planner preserves the spatiotemporal connectivity of the journey described. The route finder thus contributes to the discourse structure, coherence and gross linearization of the route description. Therefore the text planner need not maintain any explicit paragraph-level representation for the overall text in its discourse model.

The input to the route finder is a top-level goal represented in Prolog as in

?- rfind(pats_house,sfu).

It outputs the following skeletal plans in sequence. We use the Prolog list notation to indicate that the three elements of the skeletal plan have no positional significance in the representation.

1.[source(pats_house),dest(lougheed_mall), rlabel(busid(lm134))]

2.[source(lougheed_mall),dest(sfu), rlabel(busid(sfu145))]

4.4 Selection

Input to the text planner is a sequence of skeletal plans generated by the route finder. The planner consists of two major submodules that correspond to two ordered stages of processing: (1) selector and (2) realization-specification (r-spec) synthesizer. The details of the r-spec synthesizer and other aspects of the route description generation system will be presented in [27]. The selector augments the skeletal plan with additional information retrieved (or computed) from the knowledge base (map), possibly under the control of various knowledge sources. The output of the selector is an *expanded plan*. The expanded plan is input to the r-spec synthesizer, which accommodates the information in the plan into one or more r-specs. It attends to such tasks as forming predicates, choosing the utterance type and certain language-oriented tasks like topicalization, choice of verb modality and tense. The r-specs are input to the realization module which generates surface sentences. The text planner thus deals with two distinct sorts of representations, one of which is closer to the domain (knowledge-base) and one, to the language.

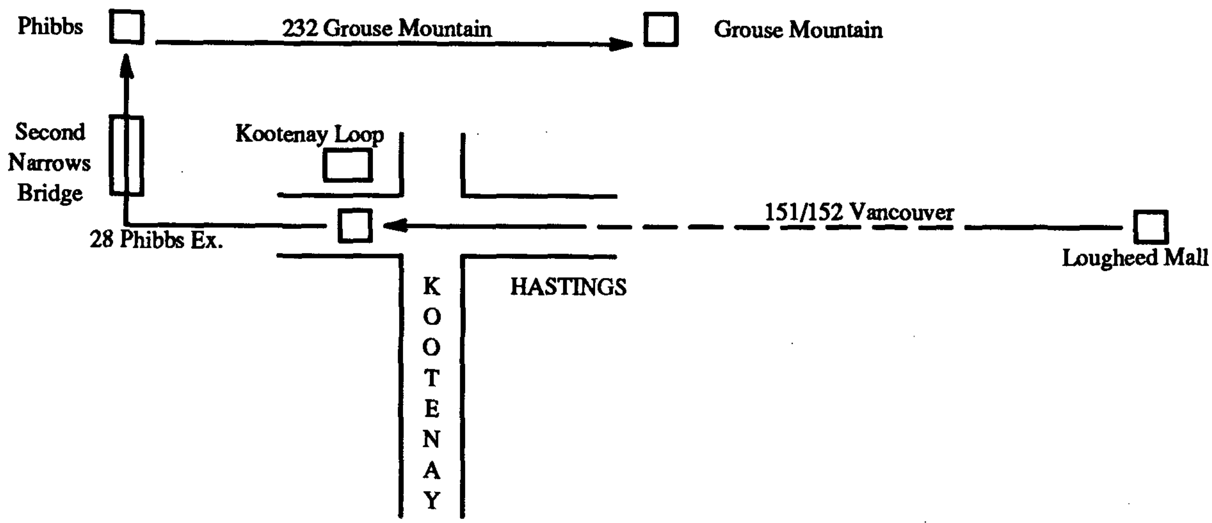

Figure 1: From *Lougheed Mall* to *Grouse Mountain*

What kind of extra information is selected? The following are some of the kinds of information selected at this stage.

- **Direction** or **orientation**, as in the descriptions *On West Pender going* **westbound** *it's bus number 19. Walk* **north** *on Renfrew.*
- **Distance** information, for realizing such descriptions as *You walk north on Burrard* **one block** *as far as West Pender.*
- **Landmarks** and **location descriptions** for sources and destinations of route sections, as included in the descriptions *Get off at Hastings and Kootenay* **right by the Kootenay Loop**. *On Hastings* **in front of the loop** *transfer to bus number 28 called Phibbs Exchange.*
- Additional descriptions of route sections involving *landmarks*, as in *Will cross* **Second Narrows Bridge** *to Phibbs Exchange* **on the other side.**
- Straightforward matters like retrieving names corresponding to route labels in the skeletal plan (**bus number, bus name**, etc).

How is the information derived? Computation of direction (orientation) information requires availability of a co-ordinate system and additional inference modules. The text planner (selector) uses additional procedural domain knowledge for retrieving or computing orientation information. Information so computed is expressed linguistically as *north, left, right,* and so on, and incorporated into the text plan. The search and computation of orientation from a map using qualitative or quantitative information forms a separate subject matter of study by itself.

Distance is expressed in route descriptions in terms of various units, as, for example, in *one block, 200 yards* or *three stoplights*, but may not be represented in the map in the same terms. In our system we simplify the task of distance estimation/evaluation by directly representing distances in terms required by the text planner.

Retrieving landmarks for describing locations of sources and destinations requires examining finer-grainer spatial layout information at the neighbourhood of these points. *Salience* plays a key role in the choice of landmarks chosen for location description. Along with the task of choosing the landmark comes the task of computing the locative relation between the two objects (for example, *adjacence*), and expressing the relation in language (as in *right by*).

Salience also influences the choice of landmarks chosen for describing long route sections (as in *will cross Second Narrows Bridge...*). In our implementation we have represented certain landmarks along the bus route stages. Given the end-points of a route section, the landmarks along the way can be selected for description. It is to be noted that landmarks like the *Second Narrows Bridge* aren't quite points, but are idealized as points and represented in the knowledge base as such. Our implementation at present doesn't include the detailed representations and procedures necessary for deriving orientation, distance and location information. We sidestep the task by incorporating such information in the form required by the text planner.

What is the rationale for selecting such information? The kinds of information described in this section are common in route descriptions, but are nevertheless secondary to the essential information contained in the skeletal plans. When a route section should be covered by a walk or a drive, some kind of distance information is required. Location information is crucial at transfer points (intermediate destinations). Additional descriptions of routes in terms of landmarks crossed give the prospective traveller a feel for how long she needs to be on the bus, train, etc, and assure her that she is still on the right track. Orientation information is given in bus route descriptions when the listener has to take a bus after a spell of walk (self-navigation), as evident in the description

> ...you want to walk north on Burrard one block as far as West Pender... and on West Pender *going westbound* it's bus number 19...

But one cannot always attribute motives for including additional material in route descriptions. At present we include in the text plan all additional relevant material that can be extracted from the knowledge base.

How is the selected information incorporated in the plan? Additional information selected at this stage is incorporated in a skeletal plan like

> [source(lougheed_mall),dest(hastings@kootenay), rlabel(busid(van151))]

by adding to the list or by argument-adjunction and replacement. For example, the skeletal plan given above may be expanded as:

> [source(lougheed_mall),dest(d(hastings@kootenay), loc(next(kootenay_loop))), rlabel(busid(busnumber(151),busname(vancouver))), via(brentwood_mall)]

Discussion. We have used the skeletal plan as a substrate on which further information is deposited by the selector in the text planner. At the same time, the information available in the skeletal plan enables fast access to the information needed by the selector, by providing the route-label, source and destination names. This corresponds to our intuition regarding route description generation using visual search over maps: once we spot the end-points of a route section, we can easily spot the additional relevant objects around the end-points and over the route-section connecting them. The selector derives the information from the same knowledge base (map) consulted by the route finder.

The text plan thus has a dual function: (1) for selection from the knowledge base, in its skeletal form and (2) for conversion into text eventually, in its expanded form. The expanded plan also has a third important function: a copy of the expanded plan is retained in the discourse record for consultation when the corresponding r-spec and sentences, and the next expanded plan are generated. It thus provides a record of the objects that have been introduced in the discourse, and forms the basis for decisions on pronominalization, referential expression generation and other factors that govern text coherence.

It is to be noted that the domain of route description generation involves the use of a knowledge base which is primarily object-oriented. The knowledge base consists of points, landmarks, bus routes, streets and so on. The route finder works with the knowledge base and gives some objects the status of *source*, *destination* and *route-label* in particular instances of route finding. The selector in the text planner picks up some more objects, endowing them with such attributes as landmarks, distance, direction. As the expanded plan stands now, it can just as well be expressed pictorially as a route-sketch.

5. Work in Progress and Future Work

We have examined the roles of salience and relevance as determinants of selection decisions, and presented their distinction mostly in intuitive terms. The notions should be formulated in more precise terms to be usable in specific models or implementations. We are continuing our investigation on salience and relevance by attempting to cast them in concrete terms in the domain we are currently examining, viz., route description generation. As in the case of Grice's maxims and many other pragmatic principles, the formulation of formal principles appears to be a difficult task. Herskovits [7] has encountered similar difficulties in attempting to formalize salience and relevance rigorously (in her case, for the purpose of preposition understanding).

We are also continuing our research on the coupling between the *thinker* and the *speaker*. Examining the kinds and degrees of coupling will help us understand the nature of text planning tasks, design better text planners and compare various models of natural language generation.

Our system is intended to generate bus route descriptions. However, there are other route descriptions that are more demanding in the kind of knowledge representation and processing that they require, and more complex in their syntax. The representations for text plans and r-specs we use are also limited in variety and are tailored to the domain under examination. We are at present completing the implementation of the system.

Route description generation is in itself a very complex process, and this research can proceed along several lines of further work: introduction of various knowledge sources to regulate selection from the knowledge base, representation of more detailed spatial knowledge at several levels, and so on. Upon completion of the current implementation effort, our primary plans, however, are to examine the implications of the system as a process model for speech-like monological multisentential text generation, refine its details and examine its applicability to other domains.

Acknowledgments

Our thanks are due to Ed Hovy, Fred Popowich and Dan Fass for help with some of the literature, and to the three anonymous reviewers for their comments. This work has been supported in part by a Simon Fraser University Graduate Research Fellowship and a British Columbia Advanced Systems Institute Graduate Student Scholarship.

References

1. Appelt, D.E. *Planning English Sentences.* Cambridge University Press, 1985.

2. Conklin, E.J. and McDonald, D.D. Salience: The Key to the Selection Problem in Natural Language Generation. Proceedings of the 20th ACL Conference, Toronto, 1982, pp. 129-135.

3. Dale, R. Generating Recipes: An Overview of EPICURE. The Second European Workshop on Natural Language Generation, Edinburgh, 1989, Extended abstract.

4. Fass, D. Lexical Semantic Constraints. CSS/LCCR TR 89-11, Centre for Systems Science, Simon Fraser University,1989.

5. Habel, C. Cognitive Linguistics: The Processing of Spatial Concepts. Lilog-Report 45, IBM Deutschland GmbH,1988.

6. Haslett, B.J. *Communication: Strategic Action in Context.* Lawrence Erlbaum Associates, 1987.

7. Herskovits, A. *Language and Spatial Cognition: An Interdisciplinary Study of the Prepositions in English.* Cambridge University Press, 1986.

8. Hovy, E.H. Planning Coherent Multisentential Text. Proceedings,26th Conference of the ACL, Buffalo, 1988.

9. Hovy, E.H. Unresolved Issues in Paragraph Planning. Selected Papers from the Second European Workshop on Language Generation, Edinburgh, 1989. Mellish, C. and Dale, R., Editors, to appear, 1990. Preprint obtained from the author.

10. Kempen, G. and Hoenkamp, E. "An Incremental Procedural Grammar for Sentence Formulation". *Cognitive Science 11* (1987), 201-258.

11. Klein, W. Local Deixis in Route Directions. In *Speech, Place and Action, Jarvella, R.J. and Klein, W., Editors,* John Wiley & Sons, Ltd., 1982, pp. 161-182.

12. Leech, G. *Principles of Pragmatics.* Longman, London, 1983.

13. Levelt, W. J. M. Linearization in Describing Spatial Networks. In *Processes, Beliefs, and Questions, Peters, S. and Saarinen, E., Editors,* Reidel, Dordrecht, 1982, pp. 199-220.

14. Levelt, W.J.M. *Speaking: From Intention to Articulation.* M.I.T.Press, 1989.

15. Mann W. C. and Moore, J. A. Computer as Author-Results and Prospects. ISI/RR-79-82, Information Sciences Institute,1980.

16. McKeown, K.R. "Generating Goal-Oriented Explanations". *International Journal of Expert Systems 1,* 4 (1988), 377-395.

17. McNamara, T.P. "Mental Representation of Spatial Relations". *Cognitive Psychology 18* (1986), 87-121.

18. Moore, J.D. and Paris, C.L. Planning Text for Advisory Dialogues. Proceedings of the 27th Annual Meeting of the ACL, Vancouver, 1989, pp. 203-211.

19. Moore, J.D. and Swartout, W.R. A Reactive Approach to Explanation. Presented at the Fourth International Workshop on Natural Language Generation, Preprint, 1988.

20. Ortony, A. "Beyond Literal Similarity". *Psychological Review 86,* 3 (1979), 161-179.

21. Osgood, C.E. Where do sentences come from? In *Semantics: An Interdisciplinary Reader in Philosophy, Linguistics and Psychology, Steinberg, D. D. and Jakobovits, L. A., Editors,* Cambridge University Press, 1971.

22. Osgood, C.E. and Bock, J.K. Salience and Sentencing: Some Production Principles. In *Sentence Production: Developments in Research and Theory, Rosenberg, S., Editor,* Erlbaum, Hilldale, N.J., 1977.

23. Paris, C.L. Responding to a User's Domain Knowledge: Incorporating the User's Level of Expertise into a Generation System. Proceedings of the International Computer Science Conference, Hong Kong, 1988, pp. 301-309.

24. Paris, C.L. Generation and Explanation: Building an Explanation Facility for the Explainable Expert Systems Framework. Submitted to the Fourth International Workshop on Natural Language Generation, draft version, 1988.

25. Parisi, D. and Castelfranchi, C. A Goal Analysis of Some Pragmatic Aspects of Language. In *Possibilities and Limitations of Pragmatics, Parrett, H., Sbisa, M. and Verscheuren, J., Editors,* John Benjamins B.V., Amsterdam, 1981, pp. 551-567.

26. Pattabhiraman, T. and Cercone, N. Representing and Using Protosemantic Information in Generating Bus Route Descriptions. In *Knowledge Based Computer Systems, Ramani, S. et al., Editors,* Narosa Publishing House, 1990, pp. 341-352. To appear in the Springer-Verlag LNCS Series.

27. Pattabhiraman, T. and Cercone, N. *Generation of Speech-Like Monologue: The Case of Bus Route Descriptions.* CSS, Simon Fraser University, 1990. Unpublished manuscript.

28. Schank, R. "Rules and Topics in Conversation". *Cognitive Science 1* (1977), 421-444.

29. Sperber, D. and Wilson, D. *Relevance: Communication and Cognition.* Basil Blackwell, 1986.

30. Sridhar, S.N. *Cognition and Sentence Production.* Springer-Verlag, 1988. Springer Series in Language and Communication, 22.

31. Thorndyke, P.W. "Distance Estimation from Cognitive Maps". *Cognitive Psychology 13* (1981), 526-550.

32. van Dijk, T.A. and Kintsch, W. *Strategies of Discourse Comprehension.* Academic Press, 1983.

33. Waltz, D. Generating and Understanding Scene Descriptions. In *Elements of Discourse Understanding, Joshi, A. et al., Editors,* Cambridge University Press, 1982, pp. 266-282.

34. Weiner, E.J. "A Knowledge Representation Approach to Understanding Metaphors". *American Journal of Computational Linguistics 10* (1984), 1-14.

35. Wunderlich, D. and Reinelt, R. How to Get There From Here. In *Speech, Place and Action, Jarvella, R.J. and Klein, W., Editors,* John Wiley & Sons, Ltd., 1982, pp. 183-201.