THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Microplanning with Communicative Intentions: The SPUD System

Citation for published version:

Stone, M, Doran, C, Webber, B, Blears, T & Palmer, M 2003, 'Microplanning with Communicative Intentions: The SPUD System', *Computational Intelligence*, vol. 19, no. 4, pp. 311-381. <https://doi.org/10.1046/j.0824-7935.2003.00221.x>

Digital Object Identifier (DOI):

[10.1046/j.0824-7935.2003.00221.x](https://doi.org/10.1046/j.0824-7935.2003.00221.x)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Computational Intelligence

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Microplanning with Communicative Intentions: The SPUD System

Matthew Stone	Christine Doran	Bonnie Webber	Tonia Bleam	Martha Palmer
Rutgers	MITRE	Edinburgh	Pennsylvania	Pennsylvania

Abstract

The process of microplanning encompasses a range of problems in Natural Language Generation (NLG), such as referring expression generation, lexical choice, and aggregation, problems in which a generator must bridge underlying domain-specific representations and general linguistic representations. In this paper, we describe a uniform approach to microplanning based on declarative representations of a generator's communicative intent. These representations describe the RESULTS of NLG: communicative intent associates the concrete linguistic structure planned by the generator with inferences that show how the meaning of that structure communicates needed information about some application domain in the current discourse context. Our approach, implemented in the SPUD (sentence planning using description) microplanner, uses the lexicalized tree-adjoining grammar formalism (LTAG) to connect structure to meaning and uses modal logic programming to connect meaning to context. At the same time, communicative intent representations provide a RESOURCE for the PROCESS of NLG. Using representations of communicative intent, a generator can augment the syntax, semantics and pragmatics of an incomplete sentence simultaneously, and can assess its progress on the various problems of microplanning incrementally. The declarative formulation of communicative intent translates into a well-defined methodology for designing grammatical and conceptual resources which the generator can use to achieve desired microplanning behavior in a specified domain.

Contents

1	Motivation	3
2	Introduction to Microplanning Based on Communicative Intent	5
2.1	Representing Communicative Intent	5
2.2	Reasoning with Communicative Intent in Conversation	13
2.3	Communicative-Intent-Based Microplanning in SPUD	18
3	Grammar Organization	19
3.1	Syntactic Constructions	19
3.2	Lexical Entries	21
3.3	Lexico-grammar	21
3.4	Morphological rules	22
4	Grammatical Derivation and Communicative Intent	23
4.1	Grammatical Inference	23
4.2	Contextual Inference	26
4.3	Concrete Representations of Communicative Intent	30
4.4	Recognition of Communicative Intent	34

5	Microplanning as a Search Task	35
5.1	A Formal Search Problem	36
5.2	A Greedy Search Algorithm	37
6	Solving NLG tasks with SPUD	39
6.1	Referring Expressions	39
6.2	Syntactic Choice	41
6.3	Lexical Choice	42
6.4	Aggregation	43
6.5	Interactions in Microplanning	45
7	Building specifications	47
7.1	Syntax	48
7.2	Semantic Arguments and Compositional Semantics	50
7.3	Lexical Semantics	53
8	Previous Work	56
9	Conclusion	57
A	Instruction Grammar Fragment	65
A.1	Syntactic Constructions	65
A.2	Lexical Entries	66
B	Motion Verb Entries	68
B.1	Pure Motion Verbs	68
B.2	Pure Change-of-state Verbs	69
B.3	Near-motion Verbs	70
B.4	<i>Put</i> Verbs	71

Contact Address
 Matthew Stone
 Department of Computer Science
 and Center for Cognitive Science
 Rutgers, the State University of New Jersey
 110 Frelinghuysen Road
 Piscataway NJ 08854-8019
 mdstone@cs.rutgers.edu
 April 30, 2001

Figure 1: Microplanning in the NLG pipeline.

1 Motivation

Success in Natural Language Generation (NLG) requires connecting domain knowledge and linguistic representations. After all, an agent must have substantive and correct knowledge for others to benefit from the information it provides. And an agent must communicate this information in a concise and natural form, if people are to understand it. The instruction in (1) from an aircraft maintenance manual suggests the challenge involved in reconciling these two kinds of representation.

(1) Reposition coupling nut.

The domain knowledge behind (1) must specify a definite location where the coupling nut goes, and a definite function in an overall repair that the nut fulfills there. However, the linguistic form does not indicate this location or function explicitly; instead, its precise vocabulary and structure allows one to draw on one's existing understanding of the repair to fill in these details for oneself.

In the architecture typical of most NLG systems, and in many psycholinguistic models of speaking, a distinctive process of MICROPLANNING is responsible for making the connection between domain knowledge and linguistic representations.¹ Microplanning intervenes between a process of CONTENT PLANNING, in which the agent assembles information to provide in conversation by drawing on knowledge and conventions from a particular domain, and the domain-independent process of REALIZATION through which a concrete presentation is actually delivered to a conversational partner. These processes are frequently implemented in a pipeline architecture, as shown in Figure 1. Concretely, the content planner is typically responsible for responding to the information goals of the conversation by identifying a body of domain facts to present, and by organizing those facts into a rhetorical structure that represents a coherent and potentially convincing argument. Microplanning takes these domain facts and recodes them in suitable linguistic terms. Finally, realization is responsible for a variety of low-level linguistic tasks (including certain syntactic and morphological processes), as well as such formatting tasks as laying out a presentation on a page or a screen or performing speech synthesis. See Reiter and Dale for a thorough overview of these different stages in NLG systems (Reiter and Dale, 2000).

Microplanning often looks like a grab-bag of idiosyncratic tasks, each of which calls for its own representations and algorithms. For example, consider the three microplanning tasks that Reiter and Dale survey: referring expression generation, lexical choice, and aggregation.

- In referring expression generation, the task is to derive an identifying description to take the place of the internal representation of some discourse referent. To carry out this task, generators often execute rules to elaborate an incomplete semantic specification of an utterance (*the*

¹The name microplanning originates in Levelt's psycholinguistic model of language production (Levelt, 1989), and is adopted in Reiter and Dale's overview of NLG systems (Reiter and Dale, 2000). The process has also been termed SENTENCE PLANNING, beginning with (Rambow and Korelsky, 1992).

rabbit, say) by incorporating additional descriptive concepts (for instance *white*, to yield *the white rabbit*) (Dale and Haddock, 1991; Dale, 1992; Dale and Reiter, 1995).

- In lexical choice, the task is to select a word from among the many that describe an object or event. To perform lexical choice, generators often invoke a pattern-matching process that rewrites domain information (that there is a caused event of motion along a surface, say) in terms of available language-specific meanings (to recognize that there is *sliding*, for example) (Nogier and Zock, 1991; Elhadad et al., 1997; Stede, 1998).
- In aggregation, the task is to use modifiers, conjoined phrases, and other linguistic constructions to pack information concisely into fewer (but more complex) sentences. Aggregation depends on applying operators that detect relationships within the information to be expressed, such as repeated reference to common participants (that Doe is a patient and that Doe is female, say), and then reorganize related semantic material into a nested structure (to obtain *Doe is a female patient*, for example) (Dalianis, 1996; Shaw, 1998).

But tasks like referring expression generation, lexical choice and aggregation interact in systematic and intricate ways (Wanner and Hovy, 1996). These interactions represent a major challenge to integrating heterogeneous microplanning processes—all the more so in that NLG systems adopt widely divergent, often application-specific methods for sequencing these operations and combining their results (Cahill and Reape, 1999).

In contrast to this heterogeneity, we advocate a UNIFORM approach to microplanning. Our generator, called SPUD (for sentence planning using description), maintains a common representation of its provisional utterance during microplanning and carries out a single decision-making strategy using this representation. In what follows, we draw on and extend our preliminary presentations of SPUD in (Stone and Doran, 1996; Stone and Doran, 1997; Stone and Webber, 1998; Stone et al., 2000) to describe this approach in more detail.

The key to our framework is our generator’s representation of the INTERPRETATION of its provisional utterances. We call this representation COMMUNICATIVE INTENT. In doing so, we emphasize that language use involves a LADDER OF RELATED INTENTIONS (Clark, 1996), from uttering particular words, through referring to shared individuals from the context and contributing new information, to answering open questions in the conversation. (Clark’s ladder metaphor particularly suits the graphical presentation of communicative intent that we introduce in Section 2.) Since many of these intentions are adopted during the course of microplanning, communicative intent represents the RESULTS of generation. At the same time, we emphasize that microplanning is a deliberative process like any other, in which the provisional intentions that an agent is committed to can guide and constrain further reasoning (Bratman, 1987; Pollack, 1992). Thus, communicative intent also serves as a key resource for the PROCESS of generation.

Our specific representation of communicative intent, described in Sections 2–4, associates a concrete linguistic structure with inferences about its meaning that show how, in the current discourse context, that structure describes a variety of generalized individuals² and thereby communicates specific information about the application domain. As argued in Sections 5–6, this representation has all the information required to make decisions in microplanning. For example, it

²meaning not only objects but also actions, events, and any other constituents of a rich ontology for natural language, as described in (Bach, 1989) and advocated in (Hobbs, 1985)

records progress towards unambiguous formulation of referring expressions; it shows how alternative choices of words and syntactic constructions suit an ongoing generation task to different degrees because they encapsulate different constellations of domain information or set up different links with the context; and it indicates how given structure and meaning may be elaborated with modifiers so that multiple pieces of information can be organized for expression in a single sentence. Thus, with a model of communicative intent, SPUD can augment the syntax, semantics and pragmatics of an incomplete sentence simultaneously, and can assess its progress on the various interacting subproblems of microplanning incrementally.

In communicative intent, the pairing between structure and meaning is specified by a grammar which describes linguistic analyses in formal terms. Likewise, links between domain knowledge and linguistic meanings are formalized in terms of logical relationships among concepts. To construct communicative intent, we draw conclusions about interpretation by reasoning from these specifications. Thus, communicative intent is a DECLARATIVE representation; it enjoys the numerous advantages of declarative programming in Natural Language Processing (Pereira and Shieber, 1987). In particular, as we discuss in Section 7, the declarative use of grammatical resources leads to a concrete methodology for designing grammars that allow SPUD to achieve desired behavior in a specified domain.

Performing microplanning using communicative intent means searching through derivations of a grammar to construct an utterance and its interpretation simultaneously. This search is facilitated with a grammar formalism that packages meaningful decisions together and allows those decisions to be assessed incrementally; SPUD uses the lexicalized tree-adjoining grammar formalism. Meanwhile, the use of techniques such as logic programming and constraint satisfaction leads to efficient methods to determine the communicative intent for a given linguistic form and evaluate progress on a microplanning problem. These design decisions, combined for the first time in SPUD, lend considerable promise to communicative-intent-based microplanning as an efficient and manageable framework for practical NLG.

2 Introduction to Microplanning Based on Communicative Intent

We begin with an extended illustration of communicative intent and motivation for its use in microplanning. In Section 2.1, we situate representations of communicative intent more broadly within research on the cognitive science of contributing to conversation, and we use a high-level case-study of communicative intent to discuss more precisely how such representations may be constructed from linguistic and domain knowledge. In Section 2.2, we show how such representations could be used to guide reasoning in conversational systems, particularly to support microplanning decisions. Finally, in Section 2.3, we identify the key assumptions that we have made in SPUD, in order to construct an effective NLG system that implements a model of communicative intent.

2.1 Representing Communicative Intent

Communicative intent responds to a view of contributing to conversation whose antecedents are Grice's description of communication in terms of intention recognition ((Grice, 1957), as updated by Thomason (Thomason, 1990)) and Clark's approach to language use as joint activity (Clark, 1996).

According to this view, conversation consists of joint activity undertaken in support of common goals. Participants take actions publicly; they coordinate so that all agree on how each action is in-

Figure 2: Carrying out instruction (2) in an aircraft fuel system.

tended to advance the goals of the conversation, and so that all agree on whether the action succeeds in its intended effects. This joint activity defines a CONVERSATIONAL PROCESS which people engage in intentionally and collaboratively, and, we might even suppose, rationally. The fundamental component of conversational process is the coordination by which speakers manifest and hearers recognize communicative intentions carried by linguistic actions. But there are many other aspects of conversational process: acknowledgment, grounding and backchannels; clarification and repair; and even regulation of turn-taking. (See (Clark, 1996) and references therein.) Dialogue systems increasingly implement rich models of conversational process; see e.g. (Cassell, 2000). This makes it vital that a sentence planning module interface with and support a system's conversational process.

Like any deliberative process (Pollack, 1992), this conversational process depends on plans, which provide resources for decision-making. In conversation, these plans map out how the respondent might use certain words to convey certain information: they describe the utterances of words and linguistic constructions, spell out the meanings of those utterances, and show how these utterances, with these meanings, could contribute structure, representing propositions and intentions, to the CONVERSATIONAL RECORD, an evolving abstract model of the the dialogue (Thomason, 1990). In other words, a communicative plan is a structure, built by reasoning from a grammar, which summarizes the interpretation of an utterance in context. Such plans constitute our ABSTRACT LEVEL OF REPRESENTATION OF COMMUNICATIVE INTENT. Note that this level of representation presupposes, and thereby suppresses, the specific collaborative activity that determines how meaning is actually recognized and ratified. Communicative intent is a resource for these processes in conversation, not a description of them.

We can develop these ideas in more detail by considering an illustrative example.

(2) Slide coupling nut onto elbow to uncover fuel-line sealing ring.

We draw (2) from an aircraft maintenance domain that we have studied in detail and report on fully in Section 7; Figure 2 shows the effect of the action on the aircraft fuel system. In this system, pipes are joined together using sealing rings that fit snugly over the ends of adjacent pipes. Sometimes these joints are secured by positioning a coupling nut around the seal to keep it tight and then installing a retainer to keep the coupling nut and seal in place. In checking (and, if necessary, replacing) such sealing rings, personnel must gain access to them by first removing the retainer and then

sliding the coupling nut away. Figure 2 illustrates part of this process for a case where an instructor could use (2) to direct an actor to perform the step of sliding the coupling nut clear.

We draw on this account of the domain in which (2) is used, to describe the communicative intent with which we represent (2). We consider three components of communicative intent in turn.

- The first derives from the update to the conversational record that the instruction is meant to achieve. This update includes the fact that the actor is to carry out a motion specified in terms of the given objects and landmarks—namely, the actor is to move the coupling nut smoothly along the surface of the fuel line from its current position onto the elbow. But the update also spells out the intended purpose of this action: the action is to uncover the sealing ring and, we may presume, thereby enable subsequent maintenance steps. So communicative intent must show how the meanings of the words in (2) are intended to put on the record this characterization of movement and purpose.
- The second component relates to the set of referents that the instruction describes and evokes: the elbow, the coupling nut adjacent to the elbow, the fuel line, and the sealing ring on the fuel line. The actor is expected to be familiar with these referents; this familiarity might come from the actor's general experience with the aircraft, from a diagram accompanying a block of instructions, or just from the physical surroundings as the actor carries out the instructions. In any case, the expectation of familiarity corresponds to a constraint on the idealized conversational record: the specified referents with the specified properties must be found there. Indeed, in understanding (2), the actor can and should use this constraint together with the shared information from the conversational record to identify the intended objects and landmarks. Thus, communicative intent must represent this constraint on the conversational record and anticipate the actor's use of it to resolve the instructor's references.
- The third component accounts for the collection of constructions by which the instructor frames the instruction. The instruction is an imperative; that choice shows (among other things) that the instructor's relationship with the actor empowers the instructor to impose obligations for action on the actor. (In our domain, maintenance instructions are in fact military orders.) Meanwhile, the use of definite noun phrases that omit the article *the* reflects the distinctive telegraphic style adopted in these instructions. Of course, the relationship of instructor and actor and the distinctive linguistic style of the domain are both part of the conversational record, and the instructor anticipates that the actor will make connections with these shared representations in interpreting the constructions in (2). Thus communicative intent must also represent these connections.

To represent communicative intent, then, we will need to associate a formal representation of the utterance in (2) with a model of interpretation that describes these three components: how the utterance adds information that links up with the goals of communication; how it imposes constraints that link up with shared characterizations of objects; and how it establishes specific connections to the status of participants and referents in the discourse.

Schematically, we can represent the form of the utterance using a dependency tree as shown in (3).

(3)

This tree analyzes the utterance as being made up of ELEMENTS bearing specific content and realized in specific syntactic constructions; these elements form the nodes in the tree. Thus, the left-most leaf, labeled *coupling-nut (zero-def)*, represents the fact that the noun *coupling nut* is used here, in construction with the zero definite determiner characteristic of this genre, to contribute a noun phrase to the sentence. Generally, these elements include lexical items, as *coupling nut* does; but in cases such as the *<purpose>* element, we may simply find some distinctive syntax associated with meaning that could otherwise be realized by a construction with explicit lexical material (*in order*, for a purpose relation). Edges in the tree represent operations of syntactic combination; the child node may either supply a required COMPLEMENT to the parent node (as the node for *coupling-nut* does for its parent *slide*) or may provide an optional MODIFIER that supplements the parent's interpretation (as the node for *<purpose>* does for its parent *slide*).

We pair (3) with a record of interpretation by taking into account two sources of information: the GRAMMATICAL CONVENTIONS that associate meaningful conditions with an utterance across contexts, in a public representation accessible to speaker and hearer; and the SPEAKER'S PRESUMPTIONS which describe specific instantiations for these conditions in the current context, and determine the precise communicative effects of the utterance in context.

We assume that grammatical conventions associate each of the elements in (3) with an ASSERTION that contributes to the update intended for the utterance; a PRESUPPOSITION intended to ground the utterance in shared knowledge about the domain; and a PRAGMATIC condition intended to reflect the status of participants and referents in the discourse. There is a long history in computational linguistics for the assumption that utterance meaning is a conjunction of atomic contributions made by words (in constructions); see particularly (Hobbs, 1985). Our use of assertion and presupposition reflects the increasingly important role of this distinction in linguistic semantics, in such works as (van der Sandt, 1992; Kamp and Rossdeutscher, 1994); the particular assertions and presuppositions we use draw not only on linguistic theory but also on research in connecting linguistic meanings with independently-motivated domain representations, such as those required for animating human avatars (Badler et al., 1999; Badler et al., 2000). Our further specification of pragmatic conditions is inspired by accounts of constructions in discourse in terms of contextual requirements, such as (Hirschberg, 1985; Ward, 1985; Prince, 1986; Birner, 1992; Gundel et al., 1993).

As an illustration of these threefold conventions, consider the item *slide* as used in (2) and rep-

resented in (3). Here *slide* introduces an event *a1* in which *H* (the hearer) will *move* *N* (the coupling nut) along a path *P* (from its current location along the surface of the pipe to the elbow); this event is to occur *next* in the maintenance procedure.

At the same time, *slide* provides a presupposed constraint that *P* *start at* the current location of the nut *N* and that *P* lie along the *surface* of an object. This constraint helps specify what it means for the event to be a sliding, but also helps identify both the nut *N* and the elbow *E*. As an imperative, *slide* carries a presupposed constraint on who the *participants* in the conversation are, which helps identify the agent *H* as the hearer, and at the same time introduces a variable for the speaker *S*. Moreover, *slide* carries the pragmatic constraint that *S* be capable of imposing *obligations* for physical action on *H*.

These conditions can be schematized as in (4)³:

- (4) a Assertion: $move(a1, H, N, P) \wedge next(a1)$
- b Presupposition: $partic(S, H) \wedge start-at(P, N) \wedge surf(P)$
- c Pragmatics: $obl(S, H)$

Note that these conditions take the form of constraints on the values of variables; this helps explain why we see DESCRIPTION as central to the problem of sentence planning. We call the variables that appear in such constraints the DISCOURSE ANAPHORS of an element; we call the values those variables take, the element's DISCOURSE REFERENTS. Our terminology follows that of (Webber, 1988), where a discourse anaphor specifies an entity by relation (perhaps by an inferential relation) to a referent represented in an evolving model of the discourse. (Throughout, we follow the Prolog convention with anaphors—variables in upper case and referents—constants in lower case.)

When elements are combined by syntactic operations, the grammar describes both syntactic and semantic relationships among them. Semantic relationships are represented by requiring coreference between discourse anaphors of combined elements. We illustrate this by considering the element *coupling-nut*, which appears in combination with *slide*. The grammar determines that the element presupposes a coupling nut (*cn*) represented by some discourse anaphor *R*. The pragmatics of the element is the condition that the genre supports the zero definite construction (*zero-genre*) and that the referent for *R* has definite status in the conversational record. The element carries no assertion. Thus, this use of *coupling-nut* carries the conditions schematized in (5).

- (5) a Assertion: —
- b Presupposition: $cn(R)$
- c Pragmatics: $def(R) \wedge zero-genre$

Now, when this element serves as the direct object of the element *slide* as specified in (3), the coreference constraints of the grammar kick in to specify that what is slid must be the coupling nut; formally, in this case, the *N* of (4) must be the same as the *R* of (5). Applying this constraint, we would represent the conditions imposed jointly by *slide* and *coupling-nut* in combination as in (6).

- (6) a Assertion: $move(a1, H, N, P) \wedge next(a1)$
- b Presupposition: $partic(S, H) \wedge start-at(P, N) \wedge surf(P) \wedge cn(N)$
- c Pragmatics: $obl(S, H) \wedge def(N) \wedge zero-genre$

³From here on, we adopt the abbreviations *partic* for *participants*, *surf* for *surface*, and *obl* for *obligations*.

Let us now return to instruction (2).

(2) Slide coupling nut onto elbow to uncover fuel-line sealing ring.

In all, our exposition in this paper represents the content of (2) with the three collections of constraints on discourse anaphors in (7); we associate (7) with (2) through the derivation of (2) as tree (3) in our grammar for English.

- (7) a Assertion: $move(a1, H, N, P) \wedge next(a1) \wedge purpose(a1, a2) \wedge uncover(a2, H, R)$
 b Presupposition: $partic(S, H) \wedge start-at(P, N) \wedge surf(P) \wedge cn(N) \wedge end-on(P, E) \wedge el(E) \wedge sr(R) \wedge fl(F) \wedge nn(R, F, X)$
 c Pragmatics: $obl(S, H) \wedge def(N) \wedge def(E) \wedge def(R) \wedge def(F) \wedge zero-genre$

Spelling out the example in more detail, we see that in addition to the asserted constraints *move* and *next* contributed by the element *slide*, we have a *purpose* constraint contributed by the bare infinitival adjunct and an *uncover* constraint contributed by the element *uncover*; in addition to the presupposed constraints *partic*, *start-at*, *surf* and *cn* contributed by *slide* and *coupling-nut*, we have an *end-on* constraint contributed by *onto*, an *el* constraint contributed by *elbow*, an *sr* constraint contributed by *sealing-ring* and *fl* and *nn* constraints contributed by the noun-noun modifier use of *fuel-line*; *nn* uses a variable *X* to abstract some close relationship between the fuel line *F* and the sealing ring *R* which grounds the noun-noun compound.

In any use of an utterance like (2), the speaker intends the presupposition and the pragmatics of the utterance to link up in a specific way with particular individuals and propositions from the conversational record; the speaker likewise intends the assertion to settle particular open questions in the discourse in virtue of the information it presents about particular individuals. These links constitute the PRESUMPTIONS the speaker makes with an utterance; these presumptions must be recorded in an interpretation over and above the shared conventions that we have already outlined. We assume that these presumptions take the form of INFERENCES that the speaker is committed to in generation and that the hearer must recover in understanding.

We return to the element *slide* of (3) to illustrate this ingredient of interpretation. We take the speaker of (2) to be a computer system (including an NLG component), which represents itself as a conversational participant *s0* and represents its user as a conversational participant *h0*. We suppose that the coupling nut to be moved here is identified as *n11* in the system's model of the aircraft, the fuel-line joint is identified as *j2* and the elbow is identified as *e2*. In order to describe paths, we use a function *l* whose arguments are a landmark and a spatial relation and whose result is the place so-related to the landmark. For example, $l(on, e2)$ is the place *on the elbow*. We also use a function *p* whose arguments are two places and whose result is the direct path between them. For example, $p(l(on, j2), l(on, e2))$ is the path that the coupling nut follows here. (For a similar spatial ontology, see (Jackendoff, 1990).) Then the system here intends the contribution that the next action, *a1*, is one where *h0* moves *n11* by path $p(l(on, j2), l(on, e2))$. This contribution follows by inference from the meaning of *slide* in general together with the speaker's commitments to pick out particular discourse referents from the conversational record and, where necessary, to rely on background knowledge about these referents and about aircraft maintenance in general.

Let's adopt the notation that a boxed expression represents an update to be made to the conversational record, while an underlined expression represents a feature already present in the conversational record; boxed and underlined expressions are DOMAIN representations and can be specialized, when appropriate, to application-specific ontologies and models. The other expressions we

have seen are LINGUISTIC representations, since they are associated with lexical items and syntactic constructions in a general way. An edge indicates an inferential connection between a linguistic representation and a domain representation. Then we can provide representations of the presumption associated with the assertion of *slide* in (2) by (8).

$$(8) \quad \begin{array}{ccc} \boxed{\text{move}(a1, h0, n11, p(l(on, j2), l(on, e2)))} & & \boxed{\text{next}(a1)} \\ | & & | \\ \text{move}(a1, H, N, P) & & \text{next}(a1) \end{array}$$

Given what we have supposed, in uttering (2), the system is also committed to inferences which establish instances of the presupposition and the pragmatics of *slide* for appropriate referents. Our conventions represent these further inferences as in (9).

$$(9) \quad \begin{array}{ccc} \text{partic}(S, H) & \text{start-at}(P, N) & \text{surf}(P) \\ | & | & | \\ \text{partic}(s0, h0) & \text{start-at}(p(l(on, j2), l(on, e2)), n11) & \text{surf}(p(l(on, j2), l(on, e2))) \\ \hline \text{obl}(S, H) & & \\ | & & \\ \text{obl}(s0, h0) & & \end{array}$$

In (9), we use the same predicates for domain and linguistic relationships, so the inferences required in all cases can be performed by simple unification. But our framework will enable more complicated (and more substantive) connections. For example, suppose we use a predicate $\text{loc}(L, O)$ to indicate that the place L is the location of object O . Then we would represent the fact that the nut is located on the joint as (10).

$$(10) \quad \text{loc}(l(on, j2), n11)$$

We know that if an object is in some place, then any path from that place begins at the object; (11) formalizes this generalization.

$$(11) \quad \forall loe(\text{loc}(l, o) \supset \text{start-at}(p(l, e), o))$$

Since they provide common background about this equipment and about spatial action in general, both of these facts belong in the conversational record.

From (10) and (11) we can infer that the path on the joint starts at the nut; that leads to a record of inference as in (12).

$$(12) \quad \begin{array}{c} \text{start-at}(P, N) \\ | \\ \text{loc}(l(on, j2), n11) \end{array}$$

That is, the understanding behind (12) is that $\text{loc}(l(on, j2), n11)$ is a fact from the conversational record intended to be linked with the linguistic presupposition $\text{start-at}(P, N)$ by appeal to the premise (11) from the conversational record.

Similarly, we propose to analyze the modifier *fuel-line* in keeping with the inferential account of noun-noun compounds proposed in (Hobbs et al., 1988; Hobbs et al., 1993). This item carries a very general linguistic presupposition. There must be a fuel line F and some close relationship X between F and the object R that the modifier applies to. In the context of this aircraft, this presupposition is met because of the fact that the particular ring intended here is designed *for* the fuel

Structure:
<i>dependency representation of the utterance</i>
Assert:
<i>links for utterance assertion</i>
Presuppose:
<i>links for utterance presupposition</i>
Pragmatics:
<i>links for utterance pragmatic conditions</i>

Figure 3: General form of communicative intent representation.

line: $X = \text{for}$. This link exploits a domain-specific inference rule to the effect that one thing's being designed for another counts as the right kind of close relationship for noun-noun modification. Concretely, we might use this structure to abstract the inference:

$$(13) \quad \begin{array}{c} nn(R, F, X) \\ | \\ \underline{\text{for}(r11, f4)} \end{array}$$

As with (12), (13) represents that $\text{for}(r11, f4)$ is a shared fact linked with the linguistic presupposition $nn(R, F, X)$ by appeal to a shared rule, here (14).

$$(14) \quad \forall ab(\text{for}(a, b) \supset nn(a, b, \text{for}))$$

In general, then, the communicative intent behind an utterance must include three inferential records. The first collection of inferences links the assertions contributed by utterance elements to updates to the conversational record that the instruction is intended to achieve; in the case of (8), we add instances of the assertion identified by the speaker. The second collection of inferences links the presuppositions contributed by the utterance elements to intended instances in the conversational record. The final collection of inferences links the pragmatic constraints of the utterance elements to intended instances in the conversational record. We will represent these inferences in the format of Figure 3. Reading Figure 3 from bottom to top, we find a version of Clark's ladder of intentions, with higher links dependent on lower ones: that is, the inference to pragmatics and presupposition are prerequisites for successful interpretation, while the inferences from the assertion contingently determine the contribution of interpretation. Such diagrams constitute a complete record of communicative intent, since they include the linguistic structure of the utterance and lay out the conventional meanings assigned to this structure as well as the presumed inferences linking these meanings to context. For example, Figure 4 displays the communicative intent associated with the utterance of *slide*.

Figure 5 schematizes the full communicative intent for (2) using the notational conventions articulated thus far. As a whole, the utterance carries the syntactic structure of (3); in Figure 5 this structure is paired with inferential representations that simply group together the inferences involved in interpreting the individual words in their specific syntactic constructions.

Structure:
slide-[onto] (imperative)
Assert:
<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;"> <div style="border: 1px solid black; padding: 2px;"> $move(a1, h0, n11, p(l(on, j2), l(on, e2)))$ </div> <div style="border-bottom: 1px solid black; width: 100%;"></div> $move(a1, H, N, P)$ </div> <div style="text-align: center;"> <div style="border: 1px solid black; padding: 2px;"> $next(a1)$ </div> <div style="border-bottom: 1px solid black; width: 100%;"></div> $next(a1)$ </div> </div>
Presuppose:
<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;"> $partic(S, H)$ <div style="border-bottom: 1px solid black; width: 100%;"></div> $partic(s0, h0)$ </div> <div style="text-align: center;"> $start-at(P, N)$ <div style="border-bottom: 1px solid black; width: 100%;"></div> $loc(l(on, j2), n11)$ </div> <div style="text-align: center;"> $surf(P)$ <div style="border-bottom: 1px solid black; width: 100%;"></div> $surf(p(l(on, j2), l(on, e2)))$ </div> </div>
Pragmatics:
<div style="text-align: center;"> $obl(S, H)$ <div style="border-bottom: 1px solid black; width: 100%;"></div> $obl(s0, h0)$ </div>

Figure 4: Interpretation of *slide* in (2). The speaker’s presumptions map out intended connections to discourse referents as follows: the speaker *S*, *s0*; the hearer *H*, *h0*; the nut *N*, *n11*; the path *P*, $p(l(on, j2), l(on, e2))$; the elbow *E*, *e2*. The fuel-line joint is *j2*.

2.2 Reasoning with Communicative Intent in Conversation

We now return to our initial characterization of conversation as a complex collaborative and deliberative process, guided by representations of communicative intent such as that of Figure 5. This characterization locates microplanning within the architecture depicted in Figure 6.

In Figure 6, content planning is one of a number of subtasks carried out by a general dialogue manager. The dialogue manager tracks the content of conversation through successive turns, through such functions as following up on an utterance (Moore and Paris, 1993; Moore, 1994), repairing an utterance (Heeman and Hirst, 1995), and updating a model of the ongoing collaboration (Rich et al., 2001). The dialogue manager also coordinates the interaction in the conversation, by managing turn-taking, acknowledgment and other conversational signals (Cassell, 2000).

Once content planning has derived some updates that need to be made to the conversational record, the dialogue manager passes these updates as input to the microplanning module. In response, the microplanner derives a communicative-intent representation that spells out a way to achieve this update using an utterance of concrete linguistic forms. To construct this representation, the microplanner consults both the grammar and a general KNOWLEDGE BASE. This knowledge base specifies the system’s private domain knowledge, as well as background information about the domain that all participants in the conversation are presumed to share. It maintains information conveyed in the conversation, thus including and extending the system’s model of the conversational record.

The output communicative intent constructed by the microplanner returns to the dialogue man-

Figure 5: Communicative intent for (2). The grammar specifies meanings as follows: For *slide*, assertions *move* and *next*; for the bare infinitival adjunct, *purpose*; for *uncover*, *uncover*. For *slide*, presuppositions *partic*, *start-at* and *surf*; for *coupling-nut*, *cn*; for *onto*, *end-on*; for *elbow*, *el*; for *sealing-ring*, *cn*; for *fuel-line*, *fl* and *nn*. For *slide*, pragmatics *obl*; for other nouns, pragmatics *def* and *zero-genre*. The speaker's presumptions map out intended connections to discourse referents as follows: the speaker *S*, *s0*; the hearer *H*, *h0*; the nut *N*, *n11*; the path *P*, $p(l(on, j2), l(on, e2))$; the elbow *E*, *e2*; the ring *R*, *r11*; the fuel-line *F*, *f4*; the relation *X*, *for*. The fuel-line joint is *j2*.

Figure 6: A conversational architecture for communicative-intent-based microplanning.

ager; the dialogue manager not only can forward this communicative intent for realization but also can use it as a general resource for collaboration. Thus, Figure 6 reproduces and extends the NLG pipeline of Figure 1. (Cassell et al., 2000) describes more fully the integration of dialogue management and communicative-intent-based microplanning in one implemented conversational agent.

In a communicative-intent representation, as illustrated in the structure of Figure 5, we find the resources required for a flexible dialogue manager to pursue instruction (2) with an engaged conversational partner. To start with, the structure is a self-contained record of what the system is doing with this utterance and how it is doing it. The structure maps out the contributions that the system wants on the record and the assertions that signal these contributions; it maps out the constraints presupposed by the utterance and the unique matches for these constraints that determine the referents the instruction has. Because the structure combines grammatical knowledge and information from the conversational record in this unambiguous way, the dialogue manager can utter it with the expectation that the utterance will be understood (provided the model of the conversational record is correct and provided the interpretation process does not demand more effort than the user is willing or able to devote to it).

More generally, we expect that communicative-intent representations offer a resource for the dialogue manager to respond to future utterances. Although we have yet to implement such deliberation, let us outline briefly how communicative intent may inform such responses; such considerations help to situate structures such as that of Figure 5 more tightly within our general characterization of conversation.

As a first illustration, suppose the user asks a clarification question about the instructed action, such as (15).

(15) So I want to get at the sealing-ring at the joint under the coupling-nut?

By connecting the communicative intent from (2) with the communicative intent recognized for (15), the dialogue manager can infer that the actor is uncertain about which sealing-ring the system intended to identify with *fuel-line sealing-ring*. In carrying out this inference and in formulating an appropriate answer (*that's right*, perhaps), the explicit links in communicative intent between pre-supposed content and the conversational record are central. In other words, the dialogue manager can use communicative intent as a data structure for plan recognition and plan revision in negotiating referring expressions, as in (Heeman and Hirst, 1995).

As a second illustration, suppose the user asks a follow-up question about the instructed action, perhaps (16).

(16) How does that uncover the sealing ring?

(16) refers to the sliding and the uncovering introduced by (2); in fact, (16) shares with (2) not only reference but also substantial vocabulary. Accordingly, by connecting the intent behind (16) to that for (2), the dialogue manager may infer that the intent for (2) was successfully recognized. At the same time, by comparing the intent for (16) with that for (2), the dialogue manager can discover that, because the actor needs to know how the sliding will achieve the current purpose, the actor has not fully accepted instruction (2). The information provided in (2) and (16) can serve as a starting point for repair: knowing what information the actor has narrows what information the user might need. More generally, if structures for communicative intent also record the inferential relationships that link communicative goals to one another, the dialogue manager may attempt the more nuanced responses to expressions of doubt and disagreement described in (Moore and Paris, 1993; Carberry and Lambert, 1999).

With this background, we can now present the key idea behind the SPUD system: The structure of Figure 5 provides a resource for deliberation not just for the dialogue manager but also for the microplanner itself. The microplanner starts with a task set by the dialogue manager: this utterance is to contribute, in a recognizable way, the updates that a *move* is *next* and its *purpose* is to *uncover*. The microplanner can see to it that its utterance satisfies these requirements by adding interpreted elements, such as the structure for *slide* of Figure 4, one at a time, to a provisional communicative-intent representation. In each of these steps, the microplanner can use its assessment of the overall interpretation of the utterance to make progress on the interrelated problems of lexical choice, aggregation and referring expression generation.

Figure 7 offers a schematic illustration of a few such steps: it tracks the addition first of *slide*, then of a purpose adjunct, then of *uncover*, and finally of *coupling-nut*, all to an initially empty structure. (Note that in Figure 7 we abbreviate inference structures and specified updates to the predicates they establish; we use the tag *recognition* as a mnemonic that the microplanner is responsible for making sure these structures can be recognized as intended.)

To start, the first transition in Figure 7, which results in a structure that repeats Figure 4, can be viewed as a description of the use of the particular word *slide* in a particular syntactic construction to achieve particular effects. We will see that a generator can create such descriptions by an inferential matching process that checks a pattern of lexical meaning against the discourse context

Figure 7: A schematic view of the initial stages of microplanning for (2). Each state includes a provisional communicative intent and an assessment of further work required, such as updates to achieve. Each transition represents the addition of a new interpreted element.

and against the specified updates. In particular, to be applicable at a specific stage of generation, a lexical item must have an interpretation to contribute: the item's assertion must hold; the item's presupposition and pragmatics must find links in the conversational record. Moreover, to be preferred over alternative options, use of the item should push the generation task forward: in general, the updates the item achieves should include as many as possible of those specified in the microplanning problem, and as few others as possible; in general, the links the item establishes to shared context should appeal to specific shared content that facilitates the hearer's plan-recognition interpretation process.

Thus, in deriving structures like that of Figure 4 from its grammatical inventory, the generator can implement a model of lexical and grammatical choice. The generator determines available options by inference and selects among alternatives by comparing interpretations.

Meanwhile, in extending provisional communicative intent as suggested in Figure 7, the generator's further lexical and syntactic choices can simultaneously reflect the its strategies for aggregation and for referring expression generation. Take the addition of an element like the bare infinitival purpose clause, in step two of Figure 7. As with *slide*, this entry represents a pattern of interpretation where linguistic meaning mediates between the current context and potential update to the context. In particular, the entry for a bare infinitival purpose clause depends on an event *a1* with an agent *h0* already described by the main verb of the provisional instruction (in this case *slide*). The entry relates *a1* to another event *a2* which *a1* should achieve and which also has *h0* as the agent; here *a2* is to be described as an *uncovering* by a subsequent step of lexical choice. Thus the syntax and semantics of the entry amount to a pattern for aggregation: the modifier provides a way of extending an utterance that the generator can use to include additional related information about referents already described in the ongoing utterance.

As another illustration, take the addition of a complement like *coupling nut*, as in step four of Figure 7, or a modifier like *fuel-line*. The contribution of these entries is to add constraints on the context that the hearer must match to interpret the utterance. With *coupling nut*, for example, the hearer learns that the referent for *N* must actually be a coupling nut; similarly, with *fuel-line*, the hearer learns that the referent for *R* must be for some fuel line *F*. Here we find the usual means for ensuring reference in NLG: augmenting the content of an utterance by additional presupposed relationships.

2.3 Communicative-Intent-Based Microplanning in SPUD

Sections 2.1–2.2 have characterized microplanning as a problem of constructing representations of communicative intent to realize communicative goals. Communicative intent is a detailed representation of an utterance that combines inferences from a declarative description of language, the grammar, and from a declarative description of context, the conversational record. This representation supports the reasoning required for a dialogue manager to produce, support and defend the generated utterance as part of a broader conversational process. At the same time, by setting up appropriate microplanning choices and providing the means to make them, this representation reconciles the decision-making required for microplanning tasks like lexical choice, referring expression generation and aggregation.

Our characterization of sentence planning is not so far from Appelt’s (Appelt, 1985). One difference is that Appelt takes a speech-act view of communicative action, so that communicative intent is not an abstract resource for conversational process but a veridical inference about the dynamics of agents’ mental state; this complicates Appelt’s representations and restricts the flexibility of his system. Closer still is the work of Thomason and colleagues (Thomason et al., 1996; Thomason and Hobbs, 1997) in the interpretation-as-abduction framework (Hobbs et al., 1993); they construct abductive interpretations as an abstract representation of communicative intent, by reasoning from a grammar and from domain knowledge.

A key contribution of our research, over and above these antecedents, is the integration of a suite of assumptions and techniques for effective implementation and development of communicative-intent-based microplanners.

- We use the feature-based lexicalized tree-adjoining grammar formalism (LTAG) to describe microplanning derivations (Joshi et al., 1975; Schabes, 1990). Each choice that arises in using this grammar for generation realizes a specified meaning by concrete material that could be added to an incomplete sentence, as advocated by (Joshi, 1987) and anticipated already in Section 2.1. In fact, LTAG offers this space of choices directly on the derivation of surface syntactic structures, eliminating any need for “abstract” linguistic structures or resources.
- We use a logic-programming strategy to link linguistic meanings with specifications of the conversational record and updates to it. We base our specification language on modal logic in order to describe the different states of information in the context explicitly (Stone, 1999; Stone, 2000b); however, the logic programming inference ensures that a designer can assess and improve the computational cost of the queries involved in constructing communicative intent.
- By treating presuppositions as anaphors (cf. (van der Sandt, 1992)), we carry over efficient

constraint-satisfaction techniques for managing ambiguity in referring expressions from prior generation research (Mellish, 1985; Haddock, 1989; Dale and Haddock, 1991).

- We associate grammatical entries with pragmatic constraints on context that model the different discourse functions of different constructions (Ward, 1985; Prince, 1986). This provides both a principled model of syntactic choice and a declarative language for controlling the output of the system to match the choices observed in a given corpus or sublanguage.
- We adopt a head-first, greedy search strategy. Our other principles are compatible with searching among all partial representations of communicative intent, in any order. But a head-first strategy allows for a particularly clean implementation of grammatical operations; and the modest effort required to design specifications for greedy search is easily repaid by improved system performance.

Although many of these techniques have seen success in recent generation systems, SPUD's distinctive focus on communicative intent results in basic and important divergences from other systems; we return to a more thorough review of previous work in Section 8.

In the remainder of this paper, we first describe the grammar formalism we have developed and the model of interpretation that associates grammatical structures declaratively with possible communicative intent. We then introduce the SPUD sentence planner as a program that searches (greedily) through grammatical structures to derive a communicative intent representation that describes a desired update to the conversational record and that can be recognized by the hearer. We go on to illustrate how SPUD's declarative processing provides a natural framework for addressing sentence planning subtasks like referring expression generation, lexical and syntactic choice and aggregation, and how it supports a concrete methodology for building grammatical resources for specific generation problems.

3 Grammar Organization

In SPUD, a grammar consists of a set of SYNTACTIC CONSTRUCTIONS, a set of LEXICAL ENTRIES, and a database of MORPHOLOGICAL RULES.

3.1 Syntactic Constructions

Syntactic constructions are specified by four components in SPUD:

- (17) a a NAME, an identifier under which other parts of the grammar refer to the construction;
- b a set of PARAMETERS, open variables for referential indices in the definition (which are instantiated to discourse referents in a particular use of the construction);
- c a PRAGMATIC CONDITION, which expresses a constraint that the construction imposes on the discourse context in terms of its parameters; and
- d a SYNTACTIC STRUCTURE, which maps out the linguistic form of the construction.

The syntactic structure is represented as a tree of compound nodes. Internal nodes in the tree bear the following attributes:

- (18) a a CATEGORY, such as NP, V, etc.;
- b INDICES, a list of the parameters that the node refers to and that additional syntactic material combined with this node may describe;

- c a TOP FEATURE STRUCTURE, a list of attribute-value pairs (including variable values shared with other feature structures elsewhere in the tree) which describes the syntactic constraints imposed on this node from above; and
- d a BOTTOM FEATURE STRUCTURE, another such list of attribute-value pairs which describes the syntactic constraints imposed on this node from below.

Leaves in the tree fall into one of four classes: SUBSTITUTION SITES, FOOT NODES, GIVEN-WORD NODES and lexically-dependent word nodes or ANCHOR NODES. Like internal nodes, substitution sites and foot nodes are loci of syntactic operations and are associated with categories and indices. Any tree may have at most one foot node, and that foot node must have the same category and indices as the root. A given-word node includes a specific lexeme (typically a closed-class or function item) which appears explicitly in all uses of the construction. An anchor node is associated with an instruction to include a word retrieved from a specific lexical entry; trees may have multiple anchors and lexical entries may contain multiple words. In addition, all leaves are specified with a single feature structure which describes the constraints imposed on the node from above. Note that, in the case of anchor nodes, these constraints must be satisfied by the lexical items retrieved for the node.

(19) shows the tree structure for the zero definite noun phrase required in (2) for *coupling nut* and *sealing ring*.

Evidently such structures, and the full specifications associated with them, can be quite involved. For exposition, henceforth we will generally suppress feature structures. We will write internal nodes in the form CAT(INDICES); anchors, in the form CAT $\Diamond$ N (for the Nth token of a lexical item, a word of category CAT); substitution nodes, in the form CAT(INDICES) $\downarrow$; foot nodes, in the form CAT(INDICES)*; and given-word nodes just by the words associated with them.

With these conventions, the syntactic entry for the zero definite construction (associated with *sealing-ring* for example) is given in (20).

- (20)
- a NAME: zerodefnpmtree
 - b PARAMETERS: U
 - c PRAGMATICS: $zero\text{-}genre \wedge def(U)$

Observe that (19) appears simply as (20d).

3.2 Lexical Entries

SPUD lexical entries have the following structure.

- (21) a a NAME, a list of the lexemes that anchor the entry (most entries have only one lexeme, but entries for idioms may have several);
 b a set of PARAMETERS, open variables for referential indices in the definition (which are instantiated to discourse referents in a particular use of the entry);
 c a TARGET, an expression constraining the category and indices of the node in a syntactic structure at which this lexical entry could be incorporated, and indicating whether the entry is added as a complement or as a modifier;
 d a CONTENT CONDITION, a formula specifying a constraint on the parameters of the entry that the entry will assert when the entry is used to update the conversational record;
 e PRESUPPOSITION, a formula specifying a constraint on the parameters of the entry that the entry must presuppose;
 f PRAGMATICS, a formula specifying a constraint on the status in the discourse of parameters of the entry;
 g an ANCHORING FEATURE STRUCTURE, a list of attribute-value pairs that constrain the anchor nodes where lexical material from this entry is inserted into a syntactic construction; and
 h a TREE LIST, specifying the trees that the lexical item can anchor by name and parameters (note that the tree list in fact determines what the target of the entry must be).

(22) gives an example of such a lexical item: the entry for *sealing-ring* as used, among other ways, with the zero definite noun phrase illustrated in (20).

- (22) a NAME: sealing-ring
 b PARAMETERS: N
 c TARGET: $\text{NP}(N)$ [complement]
 d CONTENT: $sr(N)$
 e PRESUPPOSITION: —
 f PRAGMATICS: —
 g ANCHOR FEATURES: [NUMBER : SINGULAR]
 h TREE LIST: $\text{zerodefnptree}(N), \dots$

3.3 Lexico-grammar

The basic elements of grammatical derivations are lexical entries used in specific syntactic constructions. These elements are declarative combinations of the two kinds of specifications presented in

Sections 3.1 and 3.2. Abstractly, the combination of a lexical entry and a syntactic construction requires the following steps.

- (23)
- a The parameters of the lexical entry are instantiated to suitable discourse referents.
 - b The parameters of the construction are instantiated to discourse referents as specified by the tree list of the lexical entry.
 - c Anchor nodes in the tree are replaced by corresponding given-word nodes constructed from the name of the lexical entry; and the top feature structures of anchor nodes are unified with the anchor features of the lexical entry to give the top features of the new given-word nodes.
 - d The assertion and the presupposition of the combined entry are determined, in one of two possible ways. In one possible case, the content condition of the lexical entry provides the assertion while the presupposition of the lexical entry provides the presupposition of the combined element. In the other, the content condition and any presupposition of the lexical entry are conjoined to give the presupposition of the combined element; in this case the element carries no assertion.
 - e The pragmatics of the syntactic construction is conjoined with the pragmatics of the lexical entry.

Thus, abstractly, we can see the syntactic construction of (20) coming together with the lexical entry of (22) to yield the particular lexico-grammatical option described in (24).

- (24)
- a TREE:

$$\begin{array}{c} \text{NP}(R) \\ | \\ \text{N}'(R) \\ | \\ \text{sealing-ring} \end{array}$$
 - b TARGET: $\text{NP}(R)$ [complement]
 - c ASSERTION: —
 - d PRESUPPOSITION: $sr(R)$
 - e PRAGMATICS: $def(R) \wedge zero\text{-}genre$

(Again, feature structures are suppressed here, but note that feature sharing ensures that each of the nodes in the tree is in fact marked with singular number.) This is the entry for *sealing-ring* which is used in deriving the communicative intent of Figure 5.

3.4 Morphological rules

We have seen that lexico-grammatical entries such as (24) contain not specific surface word-forms but merely lexemes labeled with features. This allows feature-values to be propagated through grammatical derivations. In this way, the derivation can select an appropriate realization for an underlying lexeme as a function of agreement processes in the language.

A database of morphological rules accomplishes this selection. Each lexeme is paired with a list of feature-realization patterns. To determine the form to use in realizing a given lexeme at a node with given features F in a grammatical derivation, SPUD scans this list until the feature structure in a pattern subsumes F ; SPUD uses the realization associated with this pattern.

For example, we might use (25) to determine the realization of *sealing-ring* in (24) as “sealing ring”.

Figure 8: Substitution of T_1 into T_2 .

- (25) a LEXEME: sealing-ring; PATTERNS:
 b [NUMBER : SINGULAR] $\rightarrow$ sealing ring
 c [NUMBER : PLURAL] $\rightarrow$ sealing rings

4 Grammatical Derivation and Communicative Intent

To assemble communicative intent, SPUD deploys lexico-grammatical entries like (24) one by one, as depicted in Figure 7. As Section 2 suggested, these steps involve both grammatical inference to link linguistic structures together and contextual inference to link linguistic meanings to domain-specific representations. We now describe the specific form of these inferential processes in SPUD.

4.1 Grammatical Inference

In SPUD’s grammar, the trees of entries like (24) describe a set of elementary structures for a feature-based lexicalized tree-adjoining grammar, or LTAG (Joshi et al., 1975; Vijay-Shanker, 1987; Schabes, 1990). In all TAG formalisms, entries can be combined into larger trees by two operations, called SUBSTITUTION and ADJOINING. Elementary trees without foot nodes are called INITIAL trees and can only substitute; trees with foot nodes are called AUXILIARY trees, and can only adjoin. The trees that these operations yield are called DERIVED trees; we regard the computation of derived trees as an inference about a complex structure that follows from a declarative specification of elementary structures. In a grammar with features, derived trees are completed by unifying the top and bottom features on each node.

In substitution, the root of an initial tree is identified with a leaf of another elementary or derived structure, called the SUBSTITUTION site. The top feature structure of the substitution site is unified with the top feature structure of the root of the initial tree. Figure 8 schematizes this operation.

Adjoining is a more complicated splicing operation, where an elementary structure DISPLACES some subtree of another elementary or derived structure. The node in this structure where the replacement applies is called the ADJUNCTION SITE; the excised subtree is then substituted back into the first tree at the distinguished FOOT node. As part of an adjoining operation, the top feature structure of the adjunction site is unified with the top feature structure of the root node of the auxiliary

Figure 9: Adjunction of T_1 into T_2

tree; the bottom feature structure of the adjunction site is unified with the feature structure of the foot node. After an adjoining operation, no further adjoining is possible at the foot node. This is schematized in Figure 9.

In substitution, the substitution site and the root node of the substituted tree must have the same category; likewise, in adjoining, the root node, the foot node and the adjunction site must all have the same category. Moreover, as our trees incorporate indices labeling the nodes, there is the further requirement that any nodes that are identified through substitution or adjoining must carry identical indices.

The identification of indices in trees determines the interface between syntax and semantics in SPUD. SPUD adopts an ontologically promiscuous semantics (Hobbs, 1985), in the sense that each entry used in the derivation of an utterance contributes a constraint to its overall semantics. Syntax determines when the constraints contributed by different grammatical entries describe the same variables or discourse anaphors. For example, take the phrase *slide the sleeve quickly*. Its lexical elements contribute constraints describing an event e in which agent x *slides* object y along path p ; describing an individual z that is a *sleeve*; and describing an event e' that is *quick*. The syntax–semantics interface provides the guarantee that $y = z$ and $e = e'$ (i.e., that the sleeve is what is slid and that the sliding is what is quick). It does so by requiring that the index y of the object NP substitution site of *slide* unify with the index z of the root NP for *sleeve*, and by requiring that the index

e of the VP adjunction site for *slide* unify with the index e' of the VP foot node for *quickly*. (See (Hobbs, 1985; Hobbs et al., 1993) for more details on ontologically promiscuous semantics.)

Note that this strategy contrasts with other approaches to LTAG semantics, such as (Candito and Kahane, 1998), which describe meanings primarily in terms of function-argument relations. (It is also possible to combine both function-argument and constraint semantics, as in (Joshi and Vijay-Shanker, 1999; Kallmeyer and Joshi, 1999).) Like Hobbs, we use semantic representations as a springboard to explore the relationships between sentence meaning, background knowledge and inference—relationships which are easiest to state in terms of constraints. In addition, the use of constraints harmonizes with our perspective that the essential microplanning task is to construct extended descriptions of individuals (Stone and Webber, 1998; Webber et al., 1999).

Let us illustrate the operations of grammatical inference by describing how the structure for *fuel-line* can combine with the structure for *sealing ring* by adjoining. *Fuel-line* will be associated with a combined lexico-syntactic realization as in (26).

- (26) a TREE:
- $$\begin{array}{c}
 N'(R) \\
 \swarrow \quad \searrow \\
 NP(F) \quad N'(R)* \\
 | \\
 N'(F) \\
 | \\
 \text{fuel-line}
 \end{array}$$
- b TARGET: $N'(R)$ [modifier]
c ASSERTION: —
d PRESUPPOSITION: $fl(F) \wedge nn(R, F, X)$
e PRAGMATICS: $def(F)$

We can adjoin (26a) into (24a) using the $N'(R)$ node as the adjunction site, to obtain the structure in (27).

- (27)
- $$\begin{array}{c}
 NP(R) \\
 | \\
 N'(R) \\
 \swarrow \quad \searrow \\
 NP(F) \quad N'(R) \\
 | \quad | \\
 N'(F) \quad \text{sealing-ring} \\
 | \\
 \text{fuel-line}
 \end{array}$$

When we put together entries by TAG operations, we can represent the meaning of the combined structure as the component-wise conjunction of the meanings of its constituents. In the case of (24) and (26) this would yield:

- (28) a ASSERTION: —
b PRESUPPOSITION: $fl(F) \wedge nn(R, F, X) \wedge sr(R)$
c PRAGMATICS: $def(F) \wedge def(R) \wedge zero-genre$

(As explained in the next section, we can also directly describe the joint interpretation of combined elements, in terms of intended links to the conversational record and intended updates to it.)

In addition to explicitly setting out the structure of a TAG derived tree as in (27), we can also describe a derived tree implicitly in terms of operations of substitution and adjoining which generate the derived tree. Such a description is called a TAG DERIVATION TREE (see (Vijay-Shanker, 1987) for a formal definition and discussion of TAG derivation trees). Each node in a derivation tree represents an elementary tree that contributes to the derived tree. Each edge in a derivation tree specifies a mode of combination: the child node is combined to the parent node by a specified TAG operation at a specified node in the structure. For example, (29) shows the derivation tree corresponding to (27).

Derivation trees indicate the decisions required to produce a sentence and outline the search space for the generation system more perspicuously than do derived trees. This makes derivation trees particularly attractive structures for describing an NLG system; for example, we can represent a TAG derivation tree for utterance (2) with a structure isomorphic to the the dependency tree (3).

4.2 Contextual Inference

SPUD assembles structures and meanings such as (27)–(29) to exploit connections between linguistic meanings and domain-specific representations. For example, the presupposition (28b) connects the meaning of the constituent *fuel-line coupling nut* with shared referents *f4* and *r11* in the aircraft domain; SPUD might use the connection to identify these referents to the user.

SPUD’s module for contextual inference determines the availability of such connections. The main resource for this module is a domain-specific knowledge base, specified as logical formulas. This knowledge base describes both the private information available to the system and the shared information that characterizes the state of the conversation. Tasks for contextual inference consult this knowledge base: SPUD first translates a potential connection between meaning and context into a theorem-proving query, and then confirms or rejects the connection by using a logic programming search strategy to evaluate the query against the contextual knowledge base. When the inferential connection is established, SPUD can record the inference as a constituent of its communicative intent.⁴

We now describe SPUD’s knowledge base, SPUD’s queries, and the inference procedure that evaluates them in more detail. SPUD’s knowledge base is specified in first-order modal logic. First-order modal logic extends first-order classical logic by the addition of MODAL OPERATORS; these operators can be used to relativize the truth of a sentence to a particular time, context or information-state. We will use modal operators to refer to a particular body of knowledge. Thus, if *p* is a formula

⁴Note that this strategy is strongly monotonic: SPUD’s inference tasks are deductive and the links SPUD adds to communicative intent cannot be threatened by the addition of further information. Previous researchers have pointed out that much inference in interpretation is nonmonotonic (Lascarides and Asher, 1991; Hobbs et al., 1993). We take it as future work to extend SPUD’s contextual inference, communicative-intent representations, and search strategy to this more general case.

and $\Box$ is a modal operator, then $\Box p$ is a formula; $\Box p$ means that p follows from the body of knowledge associated with $\Box$.

For specifications in NLG, we use four such operators: [S] represents the private knowledge of the generation system; [U] represents the private knowledge of the other party to the conversation, the user; [CR] represents the content of the conversational record; and finally [MP] (for MEANING POSTULATES) represents a body of semantic information that follows just from the meanings of words. We regard the four sources of information as subject to the eleven axiom schemes presented in (30):

- (30) a $[S]p \supset p$. $[U]p \supset p$. $[CR]p \supset p$. $[MP]p \supset p$.
 b $[S]p \supset [S][S]p$. $[U]p \supset [U][U]p$. $[CR]p \supset [CR][CR]p$. $[MP]p \supset [MP][MP]p$.
 c $[MP]p \supset [CR]p$. $[CR]p \supset [S]p$. $[CR]p \supset [U]p$

The system's information, the user's information, the conversational record and the background semantic information are all accurate, according to the idealization of (30a). The effect of (30b) is that hypothetical reasoning with respect to a body of knowledge retains access to all the information in it. Finally, (30c) ensures that semantic knowledge and the contents of the conversational record are in fact shared. (Stone, 1998) explores the relationship between this idealization of conversation implicit in these inference schemes and proposals for reasoning about dialogue context by Clark and Marshall (1981) and others. For current purposes, note that inferences using the schemes in (30) are not intended to characterize the explicit beliefs of participants in conversation veridically. Instead, the inferences contribute to a data structure, communicative intent, whose principal role is to support conversational processes such as plan recognition, coordination and negotiation.

In this paper, we consider specifications of domain knowledge and queries of domain knowledge that can be restricted to the logical fragment involving *definitions* of category D and *queries* of category Q defined by the following, mutually-recursive rules:

- (31) $D ::= Q \mid Q \supset D \mid \forall x D$
 $Q ::= [CR]D \mid [S]D \mid [U]D \mid Q \wedge Q \mid A$

A schematizes over any atomic formula; x schematizes over any bound variable. We use the notation $?K \longrightarrow q$ to denote the task of proving a Q -formula q as a query from a knowledge base K consisting of a set of D -formulas; we indicate by writing $K \longrightarrow q$ that this task results in the construction of a proof, and thus that the query succeeds.

This fragment allows for the kind of clauses and facts that form the core of a logic programming language like Prolog. In addition, these clauses and facts may make free use of modal operators; they may have nested implications and nested quantifiers in the body of rules, provided they are immediately embedded under modal operators. There have been a number of proposals for logic programming languages along these lines, such as (Fariñas del Cerro, 1986; Debart et al., 1992; Baldoni et al., 1998). Our implementation follows (Stone, 1999), which also allows for more general specifications including disjunction and existential quantifiers. For a discussion of NLG inference using the more general modal specifications, see (Stone, 2000b).

SPUD's knowledge base is a set of D formulas. These formulas provide all the information about the world and the conversation that SPUD can draw on to construct and to evaluate possible communicative intent. Concretely, for SPUD to construct communicative intent, the knowledge base must support any assertions, presuppositions and pragmatics that SPUD decides to appeal to in its

utterance. Thus, the knowledge base should explicitly set up as system knowledge any information that SPUD may assert; if some intended update relates by inference to an assertion, the knowledge base must provide, as part of the conversational record, rules sufficient to infer the update from the assertion. Moreover, the knowledge base must provide, as part of the conversational record, formulas which entail the presuppositions and pragmatic conditions that SPUD may impose. Meanwhile, for SPUD to assess whether the hearer will interpret an utterance correctly, the knowledge base must describe the context richly enough to characterize not just the intended communicative intent for a provisional utterance, but also any potential alternatives to it.

For the communicative intent of Figure 5, then, the knowledge base must include the specific private facts that underlie the assertion in the instruction, as in (32):

- (32) [S]*move*(*a1*, *h0*, *n11*, *p*(*l*(*on*, *j2*), *l*(*on*, *e2*))).
 [S]*next*(*a1*).
 [S]*purpose*(*a1*, *a2*).
 [S]*uncover*(*a2*, *h0*, *r11*).

(Recall that, in words, (32) describes the *next* action, a *move* event which takes the nut along a specified path and whose *purpose* is to *uncover* the sealing-ring.) For this communicative intent, no further specification is required for the links between assertions and updates. Updates are expressed in the same terms as meanings here, so the connection will follow as a matter of logic.

At the same time, the knowledge base must include the specific facts and rules that permit the presuppositions and pragmatics of the instruction to be recognized as part of the conversational record. (33a) spells out the instances that are simply listed in the conversational record; (33b) describes the rules and premises that allow the noun-noun compound and the spatial presuppositions to be interpreted by inference as in (12) and (13).

- (33) a [CR]*partic*(*s0*, *h0*). [CR]*obl*(*s0*, *h0*).
 [CR]*surf*(*p*(*l*(*on*, *j2*), *l*(*on*, *e2*))). [CR]*zero-genre*.
 [CR]*cn*(*n11*). [CR]*def*(*n11*).
 [CR]*el*(*e2*). [CR]*def*(*e2*).
 [CR]*sr*(*r11*). [CR]*def*(*r11*).
 [CR]*fl*(*f4*). [CR]*def*(*f4*).
 b [CR]*for*(*r11*, *f4*). [CR] $\forall ab(\text{for}(a, b) \supset nn(a, b, \text{for}))$.
 [CR]*loc*(*l*(*on*, *j2*), *n11*)
 [CR] $\forall loe(\text{loc}(l, o) \supset \text{start-at}(p(l, e), o))$. [CR] $\forall se(\text{end-on}(p(s, l(\text{on}, e)), e))$

(Again, with our conventions, (33a) spells out such facts as that *s0* and *h0* are the speaker and hearer *participating* in the current conversation, and that *s0* is empowered to impose *obligations* on *h0*. Likewise, (33b) indicates that the ring is *for* the fuel-line, and that *for* is the right kind of relationship to interpret a noun-noun compound; that a path that starts where an object is *located starts at* the object; and that any path whose endpoint is *on* an object *ends on* the object.)

Of course, the knowledge base cannot be limited to just the facts that figure in this particular communicative intent. SPUD is designed to be supplied with a number of other facts, both private and shared, about the discourse referents evoked by the instruction. This way SPUD has substantive lexical choices that arise in achieving specified updates to the state of the conversation. SPUD also expects to be supplied with additional facts describing other discourse referents from the context.

This way SPUD can consult the specification of the context to arrive at meaningful assessments of ambiguities in interpretation. For instance, the knowledge base must describe any other fuel lines and other sealing rings to settle whether there are any referential ambiguity in the phrase *fuel-line sealing ring*. For exposition, we note only the bare-bones alternatives required for SPUD to generate (2) given the task of describing the upcoming uncovering motion:

- (34) a [CR]*sr*(a_{r11})
 b [CR]*surf*($p(l(on, j2), l(a_{on}, a_{e2}))$)

There must be another sealing ring a_{r11} for SPUD to explicitly indicate $r11$ as the *fuel-line* sealing ring; and there must be another path to slide $n11$ along, for SPUD to explicitly describe the intended path as *onto elbow*.⁵

Now we consider the steps involved in linking grammatical structures such as (24) or (27)–(28) to domain-specific representations. As described in Section 4.1, the grammar delivers an assertion A , a presupposition P and pragmatics Q for each derivation tree. Links to domain-specific representations come as SPUD constructs a communicative intent for this derivation tree by reasoning from the context.

In doing this, SPUD must link up P and Q in a specific way with particular referents and propositions from the conversational record. We introduce an assignment σ taking variables to terms to indicate the correspondence between anaphors and intended referents. (We write out assignments as lists of the form $\{\dots V_i \leftarrow t_i \dots\}$ where each variable V_i is assigned term t_i as its value; for any structure E containing variables, and any assignment σ of values to those variables, we use $E\sigma$ to indicate the result of replacing the occurrences of variables in E by the terms assigned by σ .) In addition, SPUD must link up the assertion A with particular open questions in the discourse in virtue of the information it presents about particular individuals. We schematize any such update as a condition U .

These links between A , P and Q and the context constitute the presumptions that SPUD makes with its utterance; SPUD explicitly records them in its representation of communicative intent. Since these links are inferences, constructing them is a matter of proof. In SPUD, these proof tasks are carried out using logic programming inference and a modal specification of context.

- Checking that the intended instance of the assertion A is true corresponds to the proof task:

$$?K \longrightarrow [S]A\sigma$$

That is, does some instance of $A\sigma$ follow from the information available to the speaker? As usual in logic programming, if σ leaves open the values of some variables, then the proof actually describes a more specific instance $[S]A\sigma'$ where the substitution σ' possibly supplies values for these additional variables.

- Checking that the intended instance of the assertion A leads to the update U corresponds to the proof task:

$$?K \longrightarrow [CR]([CR]A\sigma \supset [CR]U)$$

⁵SPUD's greedy search also requires that this alternative path not end *on* anything, but instead end perhaps *around* or *over* its endpoint. The explanation for this depends on the results of Section 4.4 and Section 5, but briefly, SPUD will adjoin the modifier *onto* only if *onto* by itself rules out some path referents (and thus by itself helps the hearer to interpret the instruction).

That is, considering only the content of the conversational record, can we show that when $A\sigma$ is added to the conversational record, U also becomes part of the conversational record? Note that $[CR]([CR]p \supset [CR]p)$ is a valid formula of modal logic, for any p . Such a query always succeeds, regardless of the specification K .

- Checking that a presupposition P is met for an intended instance corresponds to the proof task:

$$?K \longrightarrow [CR]P\sigma$$

That is, does $P\sigma$ follow from the conversational record? More generally, determining the potential instances under which the presupposition P is met corresponds to the proof task:

$$?K \longrightarrow [CR]P$$

Each proof shows how the context supports a specific resolution σ' of underspecified elements in the meaning of the utterance, by deriving an instance $P\sigma'$. Such instances need not be just the one that the system intends. Checking that pragmatic conditions Q are met for an intended instance also corresponds to a query $?K \longrightarrow [CR]Q\sigma$.

Our logic programming inference framework allows queries and knowledge bases to be understood operationally as instructions for search, much as in Prolog; see (Miller et al., 1991). For example, a query $\Box p$ is an instruction to move to a new possible world and consider the query p there; a query $\forall x p$ is an instruction to consider a new arbitrary individual in place of x in proving p . A query $p \supset q$ is an instruction to assume p temporarily while considering the query q ; a query $p \wedge q$ is an instruction to set up two subproblems for search: a query of p and a query of q . Logical connectives in knowledge-base clauses, meanwhile, are interpreted as describing matches for predicates, first-order terms, and possible worlds in atomic queries, and as setting up subproblems with additional queries of their own. Overall then, each theorem-proving problem initiates a recursive process where the inference engine breaks down complex queries into a collection of search problems for atomic queries, backward-chains against applicable clauses in the knowledge base to search for matches for atomic queries, and takes on any further queries that result from the matches.

As in Prolog, the course and complexity of the proof process can be determined from the form of the queries and the knowledge-base. Thus, when necessary, performance can be improved by astute changes in the representation and formalization of domain relationships. Proof search is no issue with (2), for example; inspection of the clauses in (32), (33) and (34) will confirm that logic programming search explores the full search space for generation queries for this instruction without having to reason recursively through implications.

4.3 Concrete Representations of Communicative Intent

We can now return to the communicative intent of Figure 5 to describe the concrete representations by which SPUD implements it. For reference, we repeat Figure 5 as Figure 10 here.

The grammar delivers a TAG derivation whose structure is isomorphic to the tree-structure of Figure 10. That derivation is associated with a meaning that we represent as the triple of conditions of (35a)–(35c); (35d) spells out the instantiation σ under which this meaning is to be linked to the communicative context:

$$(35) \quad a \quad \text{Assertion: } move(a1, H, N, P) \wedge next(a1) \wedge purpose(a1, a2) \wedge uncover(a2, H, R)$$

Figure 10: Communicative intent for (2). The grammar specifies meanings as follows: For *slide*, assertions *move* and *next*; for the bare infinitival adjunct, *purpose*; for *uncover*, *uncover*. For *slide*, presuppositions *partic*, *start-at* and *surf*; for *coupling-nut*, *cn*; for *onto*, *end-on*; for *elbow*, *el*; for *sealing-ring*, *cn*; for *fuel-line*, *fl* and *nn*. For *slide*, pragmatics *obl*; for other nouns, pragmatics *def* and *zero-genre*. The speaker's presumptions map out intended connections to discourse referents as follows: the speaker *S*, *s0*; the hearer *H*, *h0*; the nut *N*, *n11*; the path *P*, $p(l(on, j2), l(on, e2))$; the elbow *E*, *e2*; the ring *R*, *r11*; the fuel-line *F*, *f4*; the relation *X*, *for*. The fuel-line joint is *j2*.

- b Presupposition: $partic(S, H) \wedge start-at(P, N) \wedge surf(P) \wedge cn(N) \wedge end-on(P, E) \wedge el(E) \wedge sr(R) \wedge fl(F) \wedge nn(R, F, X)$
- c Pragmatics: $obl(S, H) \wedge def(N) \wedge def(E) \wedge def(R) \wedge def(F) \wedge zero-genre$
- d Instance: $\{H \leftarrow h0, S \leftarrow s0, N \leftarrow n11, P \leftarrow p(l(on, j2), l(on, e2)), R \leftarrow r11, E \leftarrow e2, F \leftarrow f4, X \leftarrow for\}$

(We abbreviate the assertion (35a) by M ; and abbreviate the instance (35d) by σ .)

SPUD connects these meanings with domain-specific representations as schematized by the inference notation of Section 2.1 and as formalized by the modal logic queries described in Section 4.2. For example, an inference schematized in (8), repeated as (36), is required to justify the assertion-instance $move(a1, h0, n11, p(l(on, j2), l(on, e2))) = move(a1, H, N, P)\sigma$ and to link it with one of the system's goals for the instruction.

$$(36) \quad \boxed{move(a1, h0, n11, p(l(on, j2), l(on, e2)))} \\ \quad \quad \quad | \\ \quad \quad \quad move(a1, H, N, P)$$

Concretely this corresponds to two proofs which we obtain from the knowledge base K :

- (37) a $K \longrightarrow [S]move(a1, H, N, P)\sigma$
- b $K \longrightarrow [CR]([CR](M\sigma \supset [CR]move(a1, h0, n11, p(l(on, j2), l(on, e2))))$

The proof (37a) shows that the speaker knows about this motion; the proof (37b) shows that the overall assertion of the sentence will add the description of this motion to the conversational record. Note that (37b) relates the overall assertion of the utterance to the update achieved by a particular word. In general, we anticipate the possibility that a single domain-specific fact may be placed on the conversational record by combining the information expressed by multiple words. For example, one word may both provide an inference on its own and complete a complex inference in combination with words already in a sentence. We return to this possibility in Section 6.5.

Each conjunct of the assertion in (35a) contributes its inference to the system's communicative intent. In each case, SPUD represents the inference portrayed informally as a tree in Figure 10 as a pair of successful queries from K , as in (37).

Next, consider a presupposition, such as the general form $nn(R, F, X)$ and its concrete instance $nn(r11, f4, for) = nn(R, F, X)\sigma$. Corresponding to the informal inference of (38) we have the proof indicated in (39).

$$(38) \quad \begin{array}{c} nn(R, F, X) \\ | \\ \underline{for(r11, f4)} \end{array}$$

$$(39) \quad K \longrightarrow [CR]nn(R, F, X)\sigma$$

The proof of (39) proceeds by backward chaining using the axiom $[CR]\forall ab(for(a, b) \supset nn(a, b, for))$ and grounds out in the axiom $[CR]for(r11, f4)$; hence the correspondence with (38).

Each conjunct of the presupposition and each conjunct of the pragmatics requires a link to the shared context—an inference as in (38)—and in each case SPUD represents this link by a successful query as in (39).

Appendix A gives a grammar fragment sufficient to generate (2) in SPUD. By reference to the trees of this grammar, SPUD's complete representation of communicative intent for (2) is given in Figure 11.

Assert:	
$K \longrightarrow [S]move(a1, H, N, P)\sigma$	$K \longrightarrow [S]next(a1)\sigma$
$K \longrightarrow [CR]([CR]M\sigma \supset [CR]move(a1, h0, n11, p(l(on, j2), l(on, e2))))$	
$K \longrightarrow [CR]([CR]M\sigma \supset [CR]next(a1))$	
$K \longrightarrow [S]purpose(a1, a2)\sigma$	$K \longrightarrow [S]uncover(a2, H, R)\sigma$
$K \longrightarrow [CR]([CR]M\sigma \supset [CR]purpose(a1, a2))$	$K \longrightarrow [CR]([CR]M\sigma \supset [CR]uncover(a2, h0, r11))$

Presuppose:		
$K \longrightarrow [CR]partic(S, H)\sigma$	$K \longrightarrow [CR]start-at(P, N)\sigma$	$K \longrightarrow [CR]surf(P)\sigma$
$K \longrightarrow [CR]cn(N)\sigma$	$K \longrightarrow [CR]end-on(P, E)\sigma$	$K \longrightarrow [CR]el(E)\sigma$
$K \longrightarrow [CR]sr(R)\sigma$	$K \longrightarrow [CR]fl(F)\sigma$	$K \longrightarrow [CR]nm(R, F, X)\sigma$

Pragmatics:		
$K \longrightarrow [CR]obl(S, H)\sigma$	$K \longrightarrow [CR]def(N)\sigma$	$K \longrightarrow [CR]def(E)\sigma$
$K \longrightarrow [CR]def(R)\sigma$	$K \longrightarrow [CR]def(F)\sigma$	$K \longrightarrow [CR]zero-genre\sigma$

Figure 11: SPUD's representation of the communicative intent in Figure 10. Note two abbreviations for the figure:

$$M := move(a1, H, N, E) \wedge next(a1) \wedge purpose(a1, a2) \wedge uncover(a2, H, R)$$

$$\sigma := \{H \leftarrow h0, S \leftarrow s0, N \leftarrow n11, P \leftarrow p(l(on, j2), l(on, e2)), R \leftarrow r11, E \leftarrow e2, F \leftarrow f4, X \leftarrow for\}$$

Note also that K refers to the knowledge base specified in (32) and (33).

4.4 Recognition of Communicative Intent

Recall from Section 2.1 that structures such as that of Figure 11 represent not only the interpretations that speakers intend for utterances but also interpretations that hearers can recognize for them; in the ideal case, an utterance achieves the updates to the conversation that the speaker intends because the hearer successfully recognizes the speaker’s communicative intent. In generating an utterance, SPUD anticipates the hearer’s recognition of its intent by consulting a final, inferential model.

This model incorporates some simplifications that reflect the constrained domains and the constrained communicative settings in which NLG systems are appropriate. Each of these assumptions represents a starting point for further work to derive a more systematic and more general model of interpretation.

- We assume that the hearer can identify the intended lexical elements as contributing to the utterance, and can reconstruct the intended structural relationships among the elements. That is, we assume successful parsing and word-sense disambiguation. On this assumption, the hearer always has the correct syntactic structure for an utterance and a correct representation of its assertion, presupposition and pragmatics. For example, for utterance (2) as in Figure 11, the hearer gets the syntactic structure of the figure and the three conditions of meaning from (35a)–(35c).
- We assume that each update that the utterance is intended to achieve must either be an instance of an open question that has been explicitly raised by preceding discourse, or correspond to an assertion that is explicitly contributed by one of the lexical elements in the utterance itself. Once the hearer identifies the intended instance of the assertion $M\sigma$, the hearer can arrive at the intended update-inferences by carrying out a set of queries of the form $[CR]([CR]M\sigma \supset [CR]Q)$. Our assumption dictates that the set of possible formulas for Q is finite and is determined by the hearer’s information; we make the further assumption that the domain inferences are sufficiently short and constrained that the search for each query is bounded (of course, the generator requires this to design its utterances—whether or not it assesses the hearer’s interpretation). The two assumptions justify counting all updates as successfully recognized as long as the hearer can recognize the intended instance σ of the assertion.
- We assume that the hearer attempts to resolve the presupposition according to a shared ranking of SALIENCE. This ranking is formalized using the notion of a CONTEXT SET. Each REFERENT, e , comes with a context set $D(e)$ including it and its distractors; the context set for e determines all the referents that a hearer will consider as possible alternatives in resolving a variable X that the speaker intends to refer to e . This can represent a ranking because we can have $a \in D(b)$ without $b \in D(a)$; in this case a is more salient than b . During the reference resolution process, then, the hearer might have to run through the context set for a before expanding the search to include the context set for b . In practice, we simply assume that the hearer must recognize the context set successfully. That means that the hearer will consider a set of potential resolutions where variables are instantiated to elements of appropriate context sets; we represent this set of potential resolutions as a set of substitutions $D(\sigma)$ defined as follows:

$$(40) \quad \sigma' \in D(\sigma) \text{ if and only if for each variable } X \text{ that occurs in the presupposition of the utterance, } \sigma'(X) \in D(\sigma(X))$$

To make this assumption reasonable we have made limited use of gradations in salience.

- We assume that the hearer does not use the pragmatic conditions in order to determine the speaker's intended substitution σ . The hearer simply checks, once the hearer has resolved σ using the presupposition, that there is a unique inference that justifies the corresponding instance of the pragmatics.

It follows from these assumptions that interpretation is a constraint-satisfaction problem, as in (Mellish, 1985; Haddock, 1989; Dale and Haddock, 1991). In particular, the key task that the hearer is charged with is to recognize the inferences associated with the presupposition of the utterance. That presupposition is an open formula P composed of the conjunction of the individual presupposition formulas P_i contributed by lexical elements. The resolutions compatible with the hearer's information about the utterances are the instances of P that fit the conversational record and the attentional state of the discourse. Formally, we can represent this as Σ' defined in (41).

$$(41) \quad \Sigma' := \{\sigma' \in D(\sigma) : K \longrightarrow [\text{CR}]P\sigma'\}$$

Each of the formulas P_i determines a relation R_i on discourse referents that characterizes instances that the speaker may have intended; SPUD computes this relation by querying the knowledge base as in (42), and represents it compactly in terms of the free variables that occur in P_i .

$$(42) \quad R_i = \{\sigma' \in D(\sigma) : K \longrightarrow [\text{CR}]P_i\sigma'\}$$

SPUD then uses an arc-consistency constraint-satisfaction heuristic on these relations to solve for Σ' (Mackworth, 1987). (This is a conservative but efficient strategy for eliminating assignments that are inconsistent with the constraints.) SPUD counts the inferences for the presupposition as successfully recognized when the arc-consistency computation leaves only a single possibility, namely the intended resolution σ .

5 Microplanning as a Search Task

The preceding sections have been leading up to a characterization of microplanning as a formal search task (Nilsson, 1971). We argued in Section 2 that a generator must represent the interpretation of an utterance as a data structure which records inferences that connect the structure of an utterance with its meaning, ground the meaning of an utterance in the current context, and draw on the meaning of the utterance to register specified information in the conversational record. In Section 3, we described the grammatical knowledge which defines the structure and meaning of utterances; in Section 4.2, we described the inferential mechanisms which encode the relationships between utterance meaning and an evolving conversational record. With these results, we obtain the specific data structure that SPUD uses to represent communicative intent, in the kinds of records schematized in Figure 11; and the concrete operations that SPUD uses to derive representations of communicative intent, by the steps of grammatical composition and contextual inference described in Sections 4.1, 4.2, and 4.4. Thus, we obtain a characterization of the microplanning problem as a SEARCH, whose RESULT is an appropriate communicative-intent data structure, and which PROCEEDS by steps of grammatical derivation and contextual inference.

5.1 A Formal Search Problem

In SPUD, the specification of a microplanning search problem consists of the following components:

- (43) a a background specification of a GRAMMAR G describing the system's model of language (as outlined in Section 3) and a KNOWLEDGE BASE K describing the system's model of its domain, its user and the conversational record (as outlined in Section 4.2);
- b a set of formulas, UPDATES, describing the specified facts that the utterance must add to the conversational record;
- c a specification of the ROOT NODE of the syntactic tree corresponding to the utterance. This specification involves a syntactic category; variables specifying the indices of the root node; a substitution σ_0 describing the intended values that those variables must have; and a top feature structure, indicating syntactic constraints imposed on the utterance from the external context; cf. (18).

For instance, we might specify the task of describing the sliding action $a1$ by an instruction such as (2) as follows.

- (44) a The GRAMMAR G outlined in Appendices A and B; the knowledge base outlined in (32), (33), and (34).
- b Four UPDATES: $move(a1, h0, n11, p(l(on, j2), l(on, e2)))$; $next(a1)$; $purpose(a1, a2)$; $uncover(a2, h0, r11)$.
- c A root node $S \downarrow (E)$ with intended instance $\{E \leftarrow a1\}$.

The grammar and knowledge base of (43a) determine the search space for the NLG task. States in the search space are data structures for communicative intent, as argued for in Section 2 and as illustrated in Section 4.3. In particular, each state involves:

- (45) a a syntactic structure T derived according to G and paired with a meaning $\langle A, P, Q \rangle$ giving the assertion, presupposition and pragmatics of T (respectively);
- b a substitution σ determining the discourse referents intended for the variables in A , P , and Q ;
- c inferences $K \longrightarrow [S]A\sigma$, $K \longrightarrow [CR]P\sigma$, and $K \longrightarrow [CR]Q\sigma$ —such inferences show that the context supports use of this utterance to describe σ ;
- d inferences of the form $K \longrightarrow [CR]([CR]A\sigma \longrightarrow [CR]F)$ where F is an update—such inferences witness that the utterance supplies needed information;
- e a constraint network approximating $\Sigma' := \{\sigma' \in D(\sigma) : K \longrightarrow [CR]P\sigma'\}$ —this network represents the hearer's interpretation of reference resolution.

The INITIAL STATE for search is given in (46).

- (46) a a syntactic structure consisting of a single substitution site matching the root node of the problem specification (43c) and paired with an empty meaning;
- b the specified intended resolution σ_0 of variables in this syntactic structure;
- c no inferences—a record that suffices to justify the empty meaning of the initial state but which shows that this state supplies no needed information;
- d an unconstrained network realizing $\Sigma' := \{\sigma' \in D(\sigma_0)\}$.

A GOAL STATE for search is one where the three conditions of (47) are met.

- (47) a The syntactic structure of the utterance must be complete: top and bottom features of all syntactic nodes must agree, and all substitution sites must be filled.
- b For each update formula F , the communicative intent must include an update inference that establishes a substitution instance of F . More formally, on the assumption that M is the assertion of the utterance and that σ is the intended instance of M , the requirement is that the communicative intent include an inferential record of the form $K \longrightarrow [\text{CR}]([\text{CR}]M\sigma \supset [\text{CR}]F\sigma')$.
- c The arc-consistency approximation to the key presupposition-recognition problem the hearer faces for the communicative intent, as defined in Section 4.4, identifies uniquely the intended substitution of knowledge-base discourse referents for discourse-anaphor variables in the utterance.

The requirements of (47) boil down simply to this: the generator's communicative intent must provide a complete sentence (47a) that says what is needed (47b) in a way the hearer will understand (47c). Observe that the communicative intent of Figure 11 fulfills the conditions in (47) for the microplanning problem of (44).

To derive a new state from an existing state as in (45) involves the steps outlined in (48).

- (48) a Construct a lexico-grammatical element L , according to the steps of (23).
- b Apply a syntactic operation combining L with the existing syntactic structure T (cf. Section 4.1); the result is a new structure T' and a new meaning $\langle A \wedge A', P \wedge P', Q \wedge Q' \rangle$ that takes into account the contribution $\langle A', P', Q' \rangle$ of L .
- c Ensure that the use of this element is supported in context, by proving $K \longrightarrow [\text{S}]A'\sigma$, $K \longrightarrow [\text{CR}]P'\sigma$ and $K \longrightarrow [\text{CR}]Q'\sigma$; the result is a refined substitution σ' describing the intended instantiation not just of T but also of L .
- d Record the communicative effects of the new structure in any inferences $K \longrightarrow [\text{CR}]([\text{CR}](A \wedge A')\sigma' \longrightarrow [\text{CR}]F)$ for outstanding updates F .
- e Refine the constraint network to take into account the new constraint P' .

Any state so derived from a given state is called a NEIGHBOR of that state.

Because such searches begin at an initial substitution site and derive neighbors by incorporating single elements into the ongoing structure, this characterization of microplanning in terms of search builds in SPUD's head-first derivation strategy. On the other hand, it is compatible with any search algorithm, including brute-force exhaustive search, a traditional heuristic search method such as A* (Hart et al., 1968), or a stochastic optimization search (Mellish et al., 1998).

5.2 A Greedy Search Algorithm

We chose to implement a greedy search algorithm in SPUD. Greedy search applies iteratively to update a single state in the search space, the CURRENT STATE. In each iteration, greedy search first obtains all the neighbors of the current state. Greedy search then ranks the neighbors by a heuristic evaluation intended to assess progress towards reaching a goal state. The neighbor with the best heuristic evaluation is selected. If this state is a goal state, search terminates; otherwise this state becomes the current state for the following iteration.

In developing SPUD, we have identified a number of factors that give evidence of progress towards obtaining a complete, concise, natural utterance that conveys needed information unambiguously.

1. How many update formulas the utterance has conveyed. Other things being equal, if fewer updates remain unrealized, then fewer steps of lexical derivation will be required to convey this further required information.
2. How many alternative values the hearer could consider for each free variable which the system must resolve. Other things being equal, the fewer values remain for each variable, the fewer steps of lexical derivation will be required to supply content that eliminates the ambiguity for the hearer. The concrete measure for this factor in SPUD is a sorted list containing the number of possible values for each ambiguous variable in the constraint network; lists are compared by the lexicographic ordering.
3. How SALIENT the intended values for each free variable are. Other things being equal, an utterance referring to salient referents may prove more coherent and easier for the hearer to resolve (irrespective of its length). Again, the concrete measure for this factor in SPUD is a sorted list of counts, compared lexicographically; the counts here are the sizes of context sets for each intended referent.
4. How many FLAWS remain in the syntactic structure of the utterance. Flaws are open substitution sites and internal nodes whose top and bottom features do not unify. Each flaw can only be fixed by a separate step of grammatical derivation. Other things being equal, the fewer flaws remain, the fewer further syntactic operations will be required to obtain a complete grammatical utterance. We also prefer states in which an existing flaw has been corrected but new flaws have been introduced, over a structure with the same overall number of flaws but where the last step of derivation has not resolved any existing flaws.
5. How SPECIFIC the meanings for elements in the utterance are. In general, an element with a more specific assertion offers a more precise description for the hearer; an element with a more specific presupposition offers more precise constraints for identifying objects; an element with a more specific pragmatic conditions fits the context more precisely. We assess specificity off-line using the semantic information associated with the operator [MP]. If the query $?K \longrightarrow [MP](M \supset N)$ succeeds, we count formula M as at least as specific as N . We prefer words with more specific pragmatics; then (other things being equal) words with more specific presuppositions; then (other things being equal) words with more specific content; then (other things being equal) words in constructions with more specific pragmatics.

In our implementation of SPUD, we use all these criteria, prioritized as listed, to rank alternative options. That is, SPUD ranks option S ahead of option S' if one of these factors favors S over S' and all factors of higher priority are indifferent between S and S' .⁶

In designing SPUD with greedy search, we drew on the influential example of (Dale and Hadcock, 1991), which used greedy search in referring expression generation; and on our own experience using greedy algorithms to design preliminary plans to achieve multiple goals (Webber et al.,

⁶It happens that this is also the treatment of ranked constraints in optimality theory (Prince and Smolensky, 1997)!

1998). As described in Sections 6 and 7, we believe that our experience with SPUD supports our decision to use a sharply constrained search strategy; consistent search behavior makes it easier to understand the behavior of the system and to design appropriate specifications for it. However, we do NOT claim that our experience offers a justification for the specific ranking we used beyond two very general preferences—a primary preference for adding lexical elements that make some progress on the generation task over those that make none (on syntactic, informational or referential grounds); and a secondary preference based on pragmatic specificity. In general, the relationships between search algorithms, specification development and output quality for microplanning based on communicative intent, remains an important matter for future research.

6 Solving NLG tasks with SPUD

In this section, we support our claims that decision-making based on communicative intent provides a uniform framework by which SPUD can simultaneously address all the subtasks of microplanning. We further argue that such a framework is essential for generating utterances that are EFFICIENT, in that they exploit the contribution of a single lexico-grammatical element to multiple goals and indeed to multiple microplanning subtasks. Throughout the section, we illustrate how SPUD’s grammatical resources, inference processes, and search strategy combine to solve these problems together for instruction (2). Additional examples of using SPUD in generation can be found in (Bourne, 1998; Cassell et al., 2000); we also investigate these issues from the perspective of designing specifications for SPUD in Section 7.

6.1 Referring Expressions

The problem of generating a referring expression for a simple (i.e., non-event) discourse referent a is to devise a description that can be realized as a noun phrase by grammar G and that uniquely identifies a in context K . Such a problem can be posed to SPUD by the problem specification of (49).

- (49) a the grammar G and context K
b no updates to achieve
c an initial node $\text{NP} \downarrow (X)$ and an initial substitution $\sigma_0 = \{X \leftarrow a\}$

By the criteria of (47), a solution to this task is a record of communicative intent which specifies a complete grammatical noun phrase and which determines a constraint-satisfaction network that identifies a unique intended substitution, including the assignment $X \leftarrow a$.

The following example demonstrates the close affinity between SPUD’s strategy and the algorithm of (Dale and Haddock, 1991). In Figure 12, we portray a context K which supplies a number of salient individuals, including a rabbit $r1$ located in a hat $h1$; K records each individual with visual properties such as kind, size, and location. We consider the problem of generating a referring expression to identify $r1$.

With a suitable grammar, K allows us to construct the communicative intent schematized in Figure 13 for (50).

- (50) the rabbit in the hat

SPUD’s model of interpretation, like Dale and Haddock’s, predicts that the hearer successfully recognizes this communicative intent, because the context supplies a unique pair of values for variables

Figure 12: A representation of the context for a referring expression generation task

Structure:		
rabbit (definite) in (noun postmodifier) hat (definite)		
Assert:		
(none)		
Presuppose:		
<i>rabbit(R)</i>	<i>in(R, H)</i>	<i>hat(H)</i>
<u><i>rabbit(r1)</i></u>	<u><i>in(r1, h1)</i></u>	<u><i>hat(h1)</i></u>
Pragmatics:		
<i>def(R)</i>	<i>def(H)</i>	
<u><i>def(r1)</i></u>	<u><i>def(h1)</i></u>	

Figure 13: Communicative intent for *the rabbit in the hat*.

R and H such that R is a rabbit, H is a hat, and R is in H . Thus, (50) represents a potential solution to the reference task both for SPUD and for Dale and Haddock.

In fact, in deriving *the rabbit in the hat*, the two algorithms would use parallel considerations to take comparable steps. SPUD's derivation, like Dale and Haddock's, consists of three steps in which specific content enriches a description: first *rabbit*, then *in* and finally *hat*. For both algorithms, the primary consideration to use these steps of derivation is that each narrows the domain of values for variables more than the available alternative steps.

We note three important contrasts between SPUD's approach and Dale and Haddock's, however. First, SPUD typically formulates referring expressions not in isolated subtasks as suggested in (49) but rather as part of a single, overall process of sentence formulation. SPUD's broader view is in

fact necessary to generate instructions such as (2)—a point we return to in detail in Section 6.5.

Second, SPUD’s options at each step are determined by grammatical syntax, whereas Dale and Haddock’s must be determined by a separate specification of possible conceptual combinations. For example, SPUD directly encodes the syntactic requirement that a description should have a head noun using the NP substitution site; for Dale and Haddock this requires an *ad hoc* restriction on what concepts may be included at certain stages of description.

Third, Dale and Haddock adopt a fixed, depth-first strategy for adding content to a description. Particularly since (Dale and Reiter, 1995), such fixed (and even domain-specific) strategies have become common for referring expressions made up of properties of a single individual. It is difficult to generalize a fixed strategy to relational descriptions, however. Indeed, Horacek (Horacek, 1995) challenges fixed strategies with examples that show the need for modification at multiple points in an NP, such as (51).

(51) the table with the apple and with the banana

In SPUD, the order of adding content is flexible. An LTAG derivation allows modifiers to adjoin at any node at any step of the derivation. This places descriptions such as (51) within SPUD’s search space. (SPUD’s flexibility also contrasts with a top-down derivation in a context-free grammar, where modifiers must be chosen before heads and there is a resulting tension between providing what the syntax requires and going beyond what the syntax requires. See (Elhadad and Robin, 1992) for discussion of the resulting difficulties in search.)

6.2 Syntactic Choice

The problem of syntactic choice is to select an appropriate grammatical construction in which to realize a given lexical item. For example, for English noun phrases, the problem is to select an appropriate determiner from among options including the indefinite marker *a*, the definite marker *the* and the demonstrative markers *this* and *that*. With main verbs in English sentences, the problem involves such decisions as the appropriate use of active or passive voice, and the appropriate fronting or preposing of marked argument constituents.

For SPUD, alternatives for such syntactic choices are represented as alternative states which SPUD’s greedy search must consider at some stage of generation. All alternative syntactic entries whose pragmatic conditions are supported in the context will be available. Since these syntactic alternatives share a common lexical specification, their interpretations differ only by the contribution of the distinct pragmatic conditions. Recall that the pragmatics contributes neither to the updates that an utterance achieves nor to the resolution of referential ambiguity, in SPUD’s model of interpretation. Accordingly, SPUD’s ranking of these alternatives is based only on the specificity of the pragmatic conditions. SPUD’s strategy for syntactic choice is to select a licensed form whose pragmatic condition is maximally specific.

As an illustration of this strategy, consider the syntactic frame for the verb *slide* in instruction (2). The instruction exhibits the imperative frame *slide* NP. Recall that we associate this frame semantically with the condition that a sliding is the next action that the hearer should perform; we associate it with the pragmatic condition that the speaker is empowered to impose obligations for action on the hearer. This pragmatic condition distinguishes *slide* NP from other possible descriptions of this action. One such possibility is *you should slide* NP; we would represent this as a neutral alternative with an always true pragmatic condition. Thus, when SPUD considers both alterna-

tives, it favors *slide* NP because of its specific pragmatics. (In (Stone and Doran, 1997), we consider choice of a topicalized frame, represented with the pragmatic conditions proposed for topicalization in (Ward, 1985), over an unmarked frame; we describe how the generation of *the syntax book*, *we have* follows from this specification under SPUD's preference for specificity.)

Syntactic frames for the noun phrases provide a similar illustration. Noun phrases in our aircraft maintenance manuals are realized in one of two frames: a zero definite realization for a unique referent, as in *coupling nut*, and a realization with an explicit numeral, used in the other cases (plural referents, such as *two coupling nuts*, and indefinite singular referents, such as *one coupling nut*). We associate the zero definite realization with a pragmatic condition, as in (20), requiring a definite referent and an appropriate linguistic genre; the realization with the explicit numeral is a default whose pragmatic conditions are always satisfied for this genre. The zero definite is chosen whenever applicable, by specificity. More generally, whichever of the two entries, the zero-definite noun phrase or the numerical noun phrase, best applies to a referent in the maintenance domain, SPUD will prefer that entry to the corresponding ordinary definite (*the*) or indefinite (*a*) noun-phrase entry. The genre-restricted entry carries a pragmatic condition on genre which the ordinary entry lacks; thus the genre-restricted entry is selected as more specific.

We credit to systemic linguistics the idea that choices in syntactic realization should be made incrementally, by consulting a model of the discourse and a specification of the functional consequences of grammatical choices. (Mathiessen, 1983) is a classic implementation for generation, while (Yang et al., 1991) explores the close connection between systemic linguistics and TAG. However, SPUD departs from the systemic approach in that pragmatic conditions are associated with individual constructions rather than linguistic systems; this departure also necessitates SPUD's criterion of specificity. Inspiration for both of these moves can be found in such recent research on the discourse function of syntactic constructions as (Prince, 1986; Hirschberg, 1985; Ward, 1985; Gundel et al., 1993; Birner, 1992). More generally, as hinted in our contrast of zero-definite noun phrases versus *the* noun phrases, we hypothesize that pragmatically-conditioned constructions, selected in context by specificity, make for grammars that can incorporate general defaults in realization while also modeling the tendency of specific genres or sublanguages to adopt characteristic styles of communication (Kittredge et al., 1991). This hypothesis merits further detailed investigation.

6.3 Lexical Choice

Problems of lexical choice arise whenever a microplanner must apportion abstract content onto specific lexical items that carry this content (in context). Our model of this problem follows (Elhadad et al., 1997). According to this approach, in lexical choice, the microplanner must select words to contribute several independently-specified conditions to the conversational record. Some of these conditions characteristically "float", in that they tend to be realized across a range of syntactic constituents at different linguistic levels, and tend to be realized by lexical items that put other needed information on the record. We agree with the argument of Elhadad et al. that a solution to such problems depends on declarative conceptual and linguistic descriptions of lexical items and accurate assessments of the contribution of lexical items to interpretation. (We agree further that this lexical choice cannot be solved as an isolated microplanning subproblem, and must be solved concurrently with such other tasks as syntactic choice.)

Elhadad et al.'s example is (52); the sentence adopts an informal and concise style to describe

an AI class for an academic help domain.

(52) AI requires six assignments.

The choice of verb *requires* here responds to two generation goals. First, it conveys simply that the AI class involves a given set of assignments. The generator has other lexical alternatives, such as *y has x* or *there are x in y*, that do the same. In addition, *requires* conveys that the assignments represent a significant demand that the class places on its students. This second feature distinguishes *requires* from alternative lexical items and accounts for the generator's selection of it.

Both for Elhadad et al. and for SPUD, the selection of *requires* for (52) depends on its lexical representation, which must spell out the two contributions the verb can make. In SPUD, these contributions can be represented as assertions made when using *require* to describe a state S associating a class C with assignments A , as in (53).

(53) Assertion: $involve(S, C, A) \wedge demand(A)$

Meanwhile, a microplanning task might begin with goals to convey two specific instances about the AI class, $c1$; its assignments, $a1$; and an eventuality, $s1$, as in (54).

(54) a $involve(s1, c1, a1)$
b $demand(a1)$

In a context which supplies the information in (54), SPUD can add an instance of *require* as in (53) to augment a sentence about $s1$; the instantiation σ has $\{S \leftarrow s1, C \leftarrow c1, A \leftarrow a1\}$. Using M to abbreviate the *require* assertion from (53), SPUD's assessment of interpretation now records the completed inferences in (55).

(55) a $[CR](M\sigma \supset involve(s1, c1, a1))$
b $[CR](M\sigma \supset demand(a1))$

Thus, SPUD recognizes the opportunistic dual contribution of *require*, and will therefore prefer *require* to other lexical alternatives that do not make a similar contribution.

Despite the high-level similarity, SPUD's mechanisms for grammatical and contextual inference are quite different to those of (Elhadad et al., 1997). Elhadad et al. achieve flexibility of search by logic-programming constructs that allow programmers to state meaningful dependencies and alternatives in the generator's decisions in constructing a context-free phrase structure by top-down traversal. For SPUD, dependencies and alternatives are represented using the extended domain of locality of LTAG; SPUD's strategy for updating decisions about the linguistic realization of floating constraints thus depend on its LTAG derivation and incremental interpretation.

Moreover, because SPUD's model of interpretation is broader, we account for more diverse interactions in microplanning; we explore this in more detail in Section 6.5 and explore its consequences for the design of SPUD specifications for lexical choice in Section 7.3.

6.4 Aggregation

The microplanning process of aggregation constructs complex sentences in which assemblies of lexical items achieve multiple simultaneous updates to the conversational record. Instruction (2) represents a case of aggregation because the combination of *slide*, a bare infinitival purpose clause,

and *uncover* conveys four updates to the conversational record with a single sentence: the *next* event is a *sliding* whose *purpose* is an *uncovering*.

Aggregation is so named because many microplanners produce complex sentences through syntactic operations that combine together, or aggregate, specifications of simple linguistic structures (Reiter and Dale, 2000). For example, such a system might derive instruction (2) by stitching together specifications for these simple sentences: *slide the coupling nut to the elbow*; *the sliding has a purpose*; *the purpose is uncovering the sealing ring*. Each of these sentence specifications directly corresponds to a single given update. The specifications can be combined by describing transformations that create embedded syntactic structures under appropriate syntactic, semantic and pragmatic conditions.

In SPUD, aggregation is not a distinct stage of microplanning that draws on idiosyncratic linguistic resources; instead, aggregation arises as a natural consequence of the incremental elaboration of communicative intent using a grammar. Initial phases of lexicalization leave some updates unexpressed; for example, after SPUD's selection in (2) of the imperative transitive verb *slide*, SPUD still has the goals of updating the conversational record to the event's *purpose*, of *uncovering*. These lexical and syntactic decisions also trigger new grammatical entries that adjoin into SPUD's provisional linguistic structure and augment the provisional communicative intent. Such entries provide the grammatical resources by which SPUD's subsequent lexicalization decisions can directly contribute to complex sentences that achieve multiple communicative goals.

For example, in (2), *slide* introduces a VP_{purp} node indexed by the sliding event *a1* and its agent *h0*. This is a site where the lexico-syntactic entry in (56) could adjoin.

- (56) a TREE:
- $$\begin{array}{c} VP_{purp}(A1, H) \\ \swarrow \quad \searrow \\ VP_{purp}(A1, H)^* \quad S_i(A2, H) \downarrow \end{array}$$
- b TARGET: $VP_{purp}(A1, H)$
- c ASSERTION: $purpose(A1, A2)$
- d PRESUPPOSITION and PRAGMATICS: —

(56) is a declarative description of the form and meaning of an English bare infinitival purpose construction, expressed in the general terms required for reasoning about the interpretation of assemblies of linguistic constructions in context. Specifically, (56a) assumes that the purpose clause modifies a specific VP node and subcategorizes for an infinitive *s*.⁷

At the same time, (56) also has an operational interpretation for generation, as a pattern of possible aggregation: (56) describes when and how a description of an event can be extended to include a characterization of the purpose of the event. This operational interpretation provides a complementary motivation for each of the constituents of (56). An aggregation pattern must indicate how new material can be incorporated into an existing sentence; this is the role of the target in (56b). And it must indicate what updates are realized by the addition; this is the role of the assertion in (56c).

More generally, an aggregation pattern must indicate how the syntactic realization of aggregated material depends on its subordination to or coordination with other linguistic structure. Languages

⁷Lexicalization purists could add a covert subordinating conjunction to head the tree in (56a), but SPUD does not require it.

generally offer lightweight constructs, such as participles and prepositional phrases, which augment a sentence with less than another full clause. Syntactic trees such as that in (56a) provide a natural specification of these constructs. Finally, the pattern must characterize the idiosyncratic interpretive constraints that favor one aggregated realization over another. Not all realizations are equally good; alternatives may require specific informational or discourse relationships, such as the inferrability between events that some adjuncts demand (Cheng and Mellish, 2000). As an aggregation pattern, (56) represents such characterizations of requirements on context by appropriate pragmatic conditions or presuppositions.

Selecting entry (56) is SPUD’s analogue of an aggregation process; by using it, SPUD derives a provisional sentence including *slide* and requiring a further infinitive clause. SPUD substitutes *to uncover* for the infinitive sentence in the purpose clause in a subsequent step of lexicalization. This grammatical derivation results in a single complex sentence that achieves four updates to the conversational record.

6.5 Interactions in Microplanning

SPUD is capable of achieving specified behavior on isolated microplanning tasks, but a key strength of SPUD is its ability to model INTERACTIONS among the requirements of microplanning. Different requirements can usually be satisfied in isolation by assembling appropriate syntactic constituents—for example, by identifying an individual using a noun phrase that refers to it or by communicating a desired property of an action using a verb phrase that asserts it. However, many sentences exhibit an alternative, more efficient strategy which we have called TEXTUAL ECONOMY: the sentences satisfy some microplanning objectives implicitly, by exploiting the hearer’s (or reader’s) recognition of inferential links to material elsewhere in the sentence that is there for independent reasons (Stone and Webber, 1998). Such material is therefore *overloaded* in the sense of (Pollack, 1991).⁸

The main clause of (2), repeated as (57), is in fact illustrative of textual economy that exploits interactions among problems of referring expression generation and lexical choice within a single clause.

(57) Slide coupling nut onto elbow.

Consider the broader context in which (57) will be used to instruct the action the depicted in Figure 2. Given the frequent use of coupling nuts and sealing rings to join vents together in aircraft, we cannot expect this context to supply a single, unique coupling nut. Indeed, diagrams associated with instructions in our aircraft manuals sometimes explicitly labeled multiple similar parts. Allocating tasks of verb choice and referring expression generation to independent constituents in such circumstances would therefore lead to unnecessarily verbose utterances like (58).

(58) Slide coupling nut that is over fuel-line sealing ring onto elbow.

Instead, it is common to find instructions such as (57), in which these parts are identified by abbreviated descriptions; and such instructions seem to pose no difficulty in interpretation. Intuitively, the hearer can identify the intended nut from (57) because of the choice of verb: one of the semantic features of the verb *slide* is the constraint that its object (here, the coupling nut) moves

⁸Pollack used the term *overloading* to refer to cases where a single intention to act is used to wholly or partially satisfy several of an agent’s goals simultaneously.

in contact along a surface to reach its destination (here, the elbow). Identifying the elbow directs the hearer to the coupling nut on the fuel line, since that coupling nut alone lies along a common surface with the elbow.

The formal representation of communicative intent in Figure 11 implements this explanation. It associates the verb *slide* with proofs $K \rightarrow [\text{CR}]surf(P)$, $K \rightarrow [\text{CR}]start-at(P, N)$ and $K \rightarrow [\text{CR}]end-on(P, E)$ which together require the context to establish that the nut lie on a common surface with the elbow. Accordingly, the constraint-network model of communicative-intent recognition described in Section 4.4 uses this requirement in determining candidate values for N and E . The network will heuristically identify coupling nuts that lie on a common surface with an elbow. In this case, the constraints suffice jointly to determine the arguments in the action. Thus, when SPUD constructs the communicative intent in Figure 11, it models and exploits an interaction between the microplanning tasks of referring expression generation and lexical choice.

In (Stone and Webber, 1998), we make a similar point by analyzing the instruction (59) in the context depicted in Figure 12.

(59) Remove the rabbit from the hat.

From (59), the hearer should be able to identify the intended rabbit and the intended hat—even though the context supplies several rabbits, several hats, and even a rabbit in a bathtub and a flower in a hat. The verb *remove* presupposes that its object (here, the rabbit) starts out in the source (here, the hat), and this distinguishes the intended rabbit and hat in Figure 12 from the other ones.

Where instructions such as (57) exploit interactions between referring expression generation and lexical choice, instructions exhibiting PRAGMATIC OVERLOADING exploit interactions between aggregation and lexical choice (Di Eugenio and Webber, 1996). DiEugenio and Webber characterize the interpretation of instructions with multiple clauses that describe complex actions, such as (60).

(60) a Hold the cup under the spigot—
b —to fill it with coffee.

Here, the two clauses (60a) and (60b) are related by enablement, a kind of purpose relation. Because of this relation, the description in (60b) forms the basis of a constrained inference that provides additional information about the action described in (60a). That is, while (60a) itself does not specify the orientation of the cup under the spigot, its purpose (60b) can lead the hearer to an appropriate choice. To fill a cup with coffee, the cup must be held vertically, with its concavity pointing upwards. As noted in (Di Eugenio and Webber, 1996), this inference depends on the information available about the action in (60a) and its purpose in (60b). The purpose specified in (61) does not constrain cup orientation in the same way:

(61) Hold the cup under the faucet to wash it.

In a representation of communicative intent, the pragmatic overloading of (60) manifests itself in an update to the conversational record that is achieved by inference. Suppose that we represent the cup as $c1$, the action of holding it under the spigot as $a1$, and the needed spatial location and orientation as $o1$; at the same time, we may represent the filling as action $a2$, and the coffee as liquid $l1$. We contribute by inference that the orientation is upright—*upright*($o1$)—because we assert that $a1$ is an action where the hearer $h1$ holds $c1$ in $o1$ —*hold*($a1, h1, c1, o1$)—whose purpose is the

action $a2$ of filling $c1$ with $l1$ — $\text{purpose}(a1, a2) \wedge \text{fill}(a2, h1, c1, l1)$; and because we count on the hearer to recognize that an event in which something is held to be filled must involve an upright orientation—in symbols:

$$(62) \quad [\text{CR}] \forall e' x c o l [\text{hold}(e, x, c, o) \wedge \text{purpose}(e, e') \wedge \text{fill}(e', x, c, l) \supset \text{upright}(o)]$$

The notation of Section 2.1 records this inference as in (63), a constituent of the communicative intent for (60).

Because SPUD assesses the interpretations of utterances by looking for inferential possibilities such as (63), it can recognize the textual economy in utterances such as (60). Moreover, because SPUD interleaves reasoning for aggregation and lexical choice (and referring expression generation), SPUD can orchestrate the lexical content of clauses in order to take advantage of inferential links like that of (63).

Thus, suppose that SPUD starts with the goal of describing the holding action in the main clause, describing the filling action, and indicating the purpose relation between them. For the holding action, SPUD’s goals include making sure that the sentence communicates where the cup will be held and how it will be held (i.e., *upright*). SPUD first selects an appropriate lexico-syntactic tree for imperative *hold*; SPUD can choose to adjoin in the purpose clause next, in an aggregation move, and then to make the appropriate lexico-syntactic choice of *fill*. After this substitution, the semantic contributions of the sentence describe an action of *holding an object* which *can bring about* an action of *filling that object*. As shown in (Di Eugenio and Webber, 1996), and as formalized in (62), these are the premises of an inference that the object is held upright during the filling. When SPUD assesses the interpretation of this utterance, using logical queries about the updates it could achieve, it finds that the utterance has in fact conveyed how the cup is to be held. SPUD has no reason to describe the orientation of the cup with additional content.

7 Building specifications

We have seen how SPUD plans sentences not by a modular pipeline of subtasks, but by general reasoning that draws on detailed linguistic models and a rich characterization of interpretation. While this generality makes for an elegant uniformity in microplanning, it also poses substantial obstacles to the development of SPUD specifications. Because of SPUD’s general reasoning, changes to any lexical and syntactic entry have far-reaching and indirect consequences on generation results.

In response to this challenge, we have developed a methodology for constructing lexicalized grammatical resources for generation systems such as SPUD. Our methodology involves guidelines for the construction of syntactic structures, for semantic representations and for the interface between them. In this section, we describe this methodology in detail, and show, by reference to a case study in a specific instruction-generation domain, how this methodology helps ensure that SPUD deploys its lexical and syntactic options as observed in a corpus of desired output. In the

future, we hope that this methodology can serve as a starting point for automatic techniques of specification development and validation from possibly paired corpora of syntactic and semantic representations—a problem that has begun to draw attention from the perspective of interpretation as well (Hockenmaier et al., 2001).

The basic principle behind all of our guidelines is this: THE REPRESENTATION OF A GRAMMATICAL ENTRY MUST MAKE IT AS EASY AS POSSIBLE FOR THE GENERATOR TO EXPLOIT ITS CONTRIBUTION IN CARRYING OUT FURTHER PLANNING. This principle responds to two concerns. First, SPUD is currently constrained to greedy or incremental search for reasons of efficiency. At each step, SPUD picks the entry whose interpretation goes furthest towards achieving its communicative goals. As the generator uses its grammar to build on these greedy choices, our principle facilitates the generator in arriving at a satisfactory overall utterance. More generally, we saw in Section 6 many characteristic uses of language in which separate lexico-syntactic elements jointly ensure needed features of communicative intent. This is an important way in which any generator needs to be able exploit the contribution of an entry it has already used, in line with our principle.

7.1 *Syntax*

Our first set of guidelines describes the elementary trees that we specify as syntactic structures for lexical items (including lexical items that involve a semantically-opaque combination of words).

1. The grammar must associate each item with its observed range of complements and modifiers, in the observed orders. This constraint is common to any effort in grammar development; it is sufficiently well-understood to allow induction of LTAGs from treebanks (Chen and Vijay-Shanker, 2000; Sarkar, 2001).
2. All syntactically optional elements, regardless of interpretation, must be represented in the syntax as modifiers, using the LTAG operation of adjunction. This allows the generator to select an optional element when it is needed to achieve updates not otherwise conveyed by its provisional utterance. Recall that, in LTAG, a substitution site indicates a constituent that must be supplied syntactically to obtain a grammatical sentence; we call a constituent so provided a SYNTACTIC ARGUMENT. The alternative is to rewrite a node so as to include additional material (generally optional) specified by an auxiliary tree; we call material so provided a SYNTACTIC ADJUNCT. If optional elements are represented as syntactic adjuncts, it is straightforward to select one whenever its potential benefit is recognized. With other representations—for example, having a set of syntactic entries, each of which has a different number of syntactic arguments—the representation can result in artificial dependencies in the search space in generation, or even dead-end states in which the grammar does not offer a way to more precisely specify an ambiguous reference. To use this representation successfully, a greedy generator such as SPUD would have to anticipate how the sentence would be fleshed out later in order to select the right entry early on.
3. The desired linear surface order of complements and modifiers for an entry must be represented using hierarchies of nodes in its elementary tree. In constructions with fixed word-order (the typical case for English), the nodes we add reflect different semantic classes which tend to be realized in a particular order. In constructions with free word-order (the typical case in many other languages), node-ordering would instead reflect the information-structure sta-

tus of constituents. Introducing hierarchies of nodes to encode linear surface order decouples the generator's search space of derivations from the overt output word-order. It allows the generator to select complements and modifiers in any search order, while still realizing the complements and modifiers with their correct surface order. This is important for SPUD's greedy search; alternative designs—representing word-order in the derivation itself or in features that clash when elements appear in the wrong order—introduce dependencies into the search space for generation that make it more difficult for the generator to build on its earlier choices successfully. However, for a generator which explores multiple search paths, the more flexible search space will offer more than one path to the same final structure, and additional checks will be required to avoid duplicate results.

Because of strong parallels in natural language syntax across categories (see for example (Jackendoff, 1977)), we anticipate that these guidelines apply for all constructions in a similar way. Here we will illustrate them with verbs, a challenging first case that we have investigated in detail; other categories, particularly complex adjectives, adverbials and discourse connectives, merit further investigation.

We collected occurrences of the verbs *slide*, *rotate*, *push*, *pull*, *lift*, *connect*, *disconnect*, *remove*, *position* and *place* in the maintenance manual for the fuel system of the American F16 aircraft. In this manual, each operation is described consistently and precisely. Syntactic analysis of instructions in the corpus and the application of standard tests allowed us to cluster the uses of these verbs into four syntactic classes; these classes are consistent with each verb's membership in a distinct Levin class (Levin, 1993). Differences among these classes include whether the verb lexicalizes a path of motion (*rotate*), a resulting location (*position*), or a change of state (*disconnect*); and whether a spatial complement is optional (as with the verbs just given) or obligatory (*place*). The sentences from our corpus in (64) illustrate these alternatives.

- (64) a Rotate valve one-fourth turn clockwise. [Path]
 b Rotate halon tube to provide access. [Path, unspecified]
 c Position one fire extinguisher near aircraft servicing connection point. [Resulting location]
 d Position drain tube. [Resulting location, unspecified]
 e Disconnect generator set cable from ground power receptacle. [Change of state, specified source]
 f Disconnect coupling. [Change of state, unspecified source]
 g Place grommet on test set vacuum adapter. [Resulting location, required]

We used our guidelines to craft SPUD syntactic entries for these verbs. For example, we associate *slide* with the tree in (65). The structure reflects the optionality of the *path* constituent and makes explicit the observed characteristic order of three kinds of modifiers: those specifying path, such as *onto elbow*, which adjoin at VP_{path}; those specifying duration, such as *until it is released*, which adjoin at VP_{dur}; and those specifying purpose, such as *to uncover sealing ring*, which adjoin at VP_{purp}.

The requirements of generation in SPUD induce certain differences between our trees and other LTAG grammars for English, such as the XTAG grammar (Doran et al., 1994; The XTAG-Group, 1995), even in cases when the XTAG trees do describe our corpus. For example, the XTAG grammar represents *slide* simply as in (66).

The XTAG grammar does not attempt to encode the different orders of modifiers, nor to assign any special status to path PPs with motion verbs.

7.2 Semantic Arguments and Compositional Semantics

Recall that, to express the semantic links between multiple entries in a derivation, we associate each node in a syntactic tree with indices representing individuals. When one tree combines with another, and a node in one tree is identified with a node in the other tree, the corresponding indices are unified. Thus, the central problem of designing the compositional semantics for a given entry is to decide which referents to explicitly represent in the tree and how to distribute those referents as indices across the different nodes in the tree. (Of course, these decisions also inform subsequent specification of lexical semantics.)

We refer to the collection of all indices that label nodes in an entry as the SEMANTIC ARGUMENTS of the entry. This notion of semantic argument is clearly distinguished from the notion of syntactic argument that we used in Section 7.1 to characterize the syntactic structure of entries. Each syntactic argument position corresponds to one semantic argument (or more), since the syntactic argument position is a node in the tree which is associated with some indices: semantic arguments. However, semantic arguments need not be associated with syntactic argument positions. For example, in a verb entry, we do not have a substitution site that realizes the eventuality that the verb describes. But we treat this eventuality as a semantic argument to implement a Davidsonian account of event modifiers, cf. (Davidson, 1980). Because we count these implicit and unexpressed referents as semantic arguments, our notion is broader than that of (Candito and Kahane, 1998) and is more similar to Palmer’s essential arguments (Palmer, 1990).

Our strategy for specifying semantic arguments is as follows. We always include at least one implicit argument that the structure as a whole describes; these are the MAJOR ARGUMENTS of the structure. (This is common in linguistics, e.g. (Jackendoff, 1990), and in computational linguistics, e.g. (Joshi and Vijay-Shanker, 1999).) Moreover, since complements require semantic arguments,

we have found the treatment of complements relatively straightforward—we simply introduce appropriate arguments.

The treatment of optional constituents, however, is more problematic, and requires special guidelines. Often, it seems that we might express the semantic relationship between a head h and a modifier m in two ways, as schematized in (67).

- (67) a $h(R, A) \wedge m(A)$
 b $h(R) \wedge m(R, A)$

In (67a), we represent the head as relating its major argument R to another semantic argument A ; we interpret the modifier m as specifying A further. In this case, we must provide A as an index at the node where m adjoins. In contrast, in (67b), we interpret the modifier m as relating the major argument R of the head directly to A . In this case, A need not be a semantic argument of h , and we need only provide R as an index at the node where m adjoins.

We treat the case (67b) as a default, and we require specific distributional evidence before we adopt a representation such as (67a). If a class of modifiers such as m passes any of the three tests below, we represent the key entity A as a semantic argument of the associated head h , and include A as an index of the node to which m adjoins.

1. The PRESUPPOSITION TEST requires us to compare the interpretation of a sentence with a modifier m , in which the head h contributes an update, to the interpretation of a corresponding sentence without the modifier. If the referent A specified by the modifier can be identified implicitly as discourse-bound—so that the sentence without the modifier can have the same interpretation as the sentence with the modifier—then the modifier must specify A as a semantic argument of the head h . In fact, A must figure in the presupposition of h . This is only a partial diagnostic, because semantic arguments need not always be presupposed.

(68) illustrates an application of the presupposition test for the locative modifier of the verb *disconnect*.

- (68) a (Find the power cable.) Disconnect it from the power adaptor.
 b (The power cable is attached to the power adaptor.) Disconnect it.

In (68b), it is understood that the power cable is to be disconnected *from the power adaptor*; the modifier in (68a) makes this explicit. Thus *disconnect* and *from the power adaptor* pass the presupposition test.

The motivation for the presupposition test is as follows. In SPUD, implicit discourse-bound references can occur in an entry h used for an update, only when the presupposition of h evokes a salient referent from the conversational record, as suggested by (Saeboe, 1996). In (68b), for example, this referent is the power adaptor and the presupposition is that the power cable is connected to it. The representation of such presuppositions must feature a variable for the referent—we might have a variable A for the adaptor of (68b). Accordingly, in SPUD's model of interpretation, the speaker and hearer coordinate on the value for this variable (that A is the power adaptor, say) by reasoning from the presupposed constraints on the value of this variable. To guarantee successful interpretation (again using greedy search), SPUD needs to be able to carry out further steps of grammatical derivation that add additional constraints

on these variables. (For example, SPUD might derive (68a) from (68b) by adjoining *from the power adaptor* to describe *A*.) But this is possible only if the variable is represented as a semantic argument.

2. The CONSTITUENT ELLIPSIS TEST looks at the interpretation of cases of constituent ellipsis—certain anaphoric constructions that go proxy for a major argument of the head *h*. If modifiers in the same class as *m* cannot be varied across constituent ellipsis, then these modifiers must characterize semantic arguments other than the major argument of *h*.

For verbs, *do so* is one case of constituent ellipsis. The locative PPs in (69a) pass the constituent ellipsis test for *do so*, as they cannot be taken to describe Kim and Chris’s separate destinations; the infinitivals in (69b), which provide different reasons for Kim and Sandy, fail the constituent ellipsis test for *do so*:

- (69) a *Kim ran quickly to the counter. Chris did so to the kiosk.
 b Kim left early to avoid the crowd. Sandy did so to find one.

A successful test with *do so* suggests that *m* contributes a description of a referent that is independently related to the event—in other words, that *m* specifies some semantic argument. Its meaning should therefore be represented in the form $m(A)$. For (69a), for example, we can use a constraint $to(P, O)$ indicating that the path *P* (a semantic argument of the verb) goes to the object *O*.

A failed test with *do so* suggests that *m* directly describes a complete event. Its meaning should therefore be represented in the form $m(R, A)$, where *m* is some relational constraint and *R* is an event variable. For (69b), for example, we can use the constraint $purpose(E, E')$, which we have already adopted to describe bare infinitival purpose clauses.

A theoretical justification for the constituent ellipsis test depends on the assumption that material recovered from context in constituent ellipsis is invisible to operations of syntactic combination. (For example, the material might be supplied atomically as discourse referent, as in (Hardt, 1999), where *do so* recovers a property or action discourse referent that has been introduced by an earlier predicate on events.) Then a phrase that describes the major argument *R* can combine with the ellipsis, but phrases that describe any other implicit referent *A* cannot; these implicit referents are syntactically invisible.

3. The TRANSFORMATION test looks at how modifiers are realized across different syntactic frames for *h*; it is particularly useful when *m* is headed by a closed-class item. If some frames for *h* permit *m* to be realized as a discontinuous constituent with an apparent “long-distance” dependency, then the modifier *m* specifies a semantic argument. (Note that failure of the transformation test would be inconclusive in cases where syntax independently ruled out the alternative realization.)

For verbs, *wh*-extraction constructions illustrate the transformation test:

- (70) a What did you remove the rabbit from? (A: the hat)
 b *What did you remove the rabbit at? (A: the magic show)

In these cases, a modifier is realized effectively in two parts: *what...from* in (70a) and *what...at* in (70b). Intuitively, we have a case of extraction of the NP describing *A* from within *m*. When this is grammatical, as in (70a), it suggests that *m* specifies *A* as a semantic argument of the head; when it is not, as in (70b), the test fails.

In LTAG, a transformation is interpreted as a relation among trees in a tree family that have essentially the same meaning and differ only in syntax. (In one formalization (Xia et al., 1998), these relationships between trees are realized as descriptions of structure to add to elementary trees.) A transformation that introduces the referent *A* in the syntax–semantics interface and relates *A* to the available referent *R* in the semantics cannot be represented this way. However, if some semantic argument *A* is referenced in the original tree, the transformed analogue to this tree can easily realize *A* differently. If we describe the source location as the semantic argument *A* in (70a) for example, the new realization involves an initial *wh*-NP substitution site describing the source *A*, and the corresponding stranded structure of the PP *from t*.

Of course, these tests are not perfect and have on occasion revealed difficult or ambiguous cases; here too, further research remains in adapting these tests to categories of constituents that did not require intensive investigation in our corpus.

We have combined these tests to designing the syntax–semantics interface for verbs in our generation grammar. In the case of *slide*, these tests show that the path of motion is a semantic argument but a syntactic modifier. (71) presents our diagnostics: extraction is good, *do so* substitution is degraded, and *slide* can make a presupposition about the path of motion that helps to identify both the object and the path.

- (71) a What did you slide the sleeve onto?
 b *Mary slid a sleeve onto the elbow and John did so onto the pressure sense tube.
 c Slide sleeve onto elbow [acceptable in a context with many sleeves, but only one connected on a surface with the elbow].

Suppose we describe an event *A* in which *H* slides object *O* along path *P*. We label the nodes of (65) with these indices as in (72).

- (72) a subject NP: *H*
 b object NP: *O*
 c S, VP_{dur}: *A*
 d VP_{purp}: *A, H*
 e VP_{path}: *A, O, P*

This labeling is motivated by patterns of modification we observed in maintenance instructions. In particular, the index *H* for (72d) allows us to represent the control requirement that the subject of the purpose clause is understood as the subject of the main sentence; meanwhile, the indices *O* and *P* for (72e) allows us to represent the semantics of path particles such as *back*; *back* presupposes an event or state preceding *A* in time in which object *O* was located at the endpoint of path *P*.

7.3 Lexical Semantics

To complete a SPUD specification, after following the methods outlined in Sections 7.1 and 7.2, we have only to specify the meanings of individual lexical items. This task always brings potential

difficulties. However, the preceding decisions and the independent effects of SPUD's specifications of content, presupposition and pragmatics greatly constrain what needs to be specified.

By specifying syntax and compositional semantics already, we have determined what lexicalized derivation trees the generator will consider; this maps out the search space for generation. Moreover, our strategy for doing so keeps open as many options as possible for extending a description of an entity we have introduced; it allows entries to be added incrementally to an incomplete sentence in any order, subject only to the constraint that a head must be present before we propose to modify it. Syntactic specifications guarantee correct word order in the result, while the syntax–semantics interface ensures correct connections among the interpretations of combined elements. Thus, all that remains is to describe the communicative intent that we associate with the utterances in this search space.

The communicative intent of an utterance is made up of records for assertion, presupposition and pragmatics that depend on independent specifications from lexical items. The content condition determines the generator's strategy for contributing needed information to the hearer; the presupposition determines, *inter alia*, reference resolution; the pragmatics determines other contextual links. Thus we can consider these specifications separately and base each specification on clearly delineated evidence. In what follows we will describe this process for the motion verbs we studied.

We begin with the content condition. We know the kind of relationship that this condition must express from the verb's syntactic distribution (i.e., for *slide*, the frames of (64) that lexicalize an optional path of motion), and from the participants in the event identified as semantic arguments of the verb (i.e., *slide*, the event itself and its agent, object and path). To identify the particular relationship, we consider what basic information we learn from discovering that an event of this type occurred in a situation where the possibility of this event was known. For verbs in our domain, we found just four contrasts:

- (73) a Whether the event merely involves a pure change of state, perhaps involving the spatial location of an object but with no specified path; e.g., *remove* but not *move*.
b Whether the event must involve an agent moving an object from one place to another along a specified path; e.g., *move* but not *remove*.
c Whether the event must involve the application of force by the agent; e.g., *push* but not *move*.
d Whether the event must be brought about directly through the agent's bodily action (and not through mechanical assistance or other indirect agency); e.g., *place* but not *position*.

Obviously, such contrasts are quite familiar from such research in lexical semantics as (Talmy, 1988; Jackendoff, 1990); they have also been explored successfully in action representation for animation (Badler et al., 1999; Badler et al., 2000)

Many sets of verbs are identical in content by these features. One such set contains the verbs *move*, *slide*, *rotate* and *turn*; these verbs contribute just that the event involves an agent moving an object along a given path. Note that when SPUD assesses the contribution of an utterance containing these verbs, it will treat the agent, object and path as particular discourse referents that it must and will identify. This is why we simply assume that the path is given in specifying the content condition for these verbs. Of course, the verbs do provide different path information; we represent this separately, as a presupposition.

To specify the presupposition and pragmatics of a verb, we must characterize the links that the

verbs impose between the action and what is known in the context about the environment in which the action is to be performed. In some cases, these links are common across verb classes. For instance, all motion verbs presuppose a current location for the object, which they assert to be the beginning of the path traveled. In other cases, these links accompany particular lexical items; an example is the presupposition of *slide*, that the path of motion maintains contact with some surface.

In specifying these links, important evidence comes from the uses of lexical items observed in a corpus. The following illustration is representative. In the aircraft vent system, pipes may be sealed together using a sleeve, which fits snugly over the ends of adjacent pipes, and a coupling, which snaps shut around the sleeve and holds it in place. At the start of maintenance, one *removes* the coupling and *slides* the sleeve away from the junction between the pipes. Afterwards, one *(re-)positions* the sleeve at the junction and *(re-)installs* the coupling around it. In the F16 corpus, these actions are always described using these verbs.

This use of verbs reflects not only the motions themselves but also the general design and function of the equipment. For example, the verb *position* is used to describe a motion that leaves its object in some definite location in which the object will be able to perform some intended function. In the case of the sleeve, it would only be IN POSITION when straddling the pipes whose junction it seals. Identifying such distinctions in a corpus thus points to the specification required for correct lexical choice. In this case, we represent *position* as presupposing some “position” where the object carries out its intended function.

These specifications now directly control how SPUD realizes the alternation. To start, SPUD’s strategy of linking the presupposition and pragmatics to a knowledge base of shared information restricts what verbs are applicable in any microplanning task. For example, when the sleeve is moved away from the junction, we can only describe it by *slide* and not by *position*, because the presupposition of *position* is not met.

At the same time, in contexts which support the presupposition and pragmatics of several alternatives, SPUD selects among them based on the contribution to communicative intent of presupposition and pragmatics. We can illustrate this with *slide* and *position*. We can settle on a syntactic tree for each verb that best fits the context; and we have designed these trees so that either choice can be fleshed out by further constituents into a satisfactory utterance. Similarly, these items are alike in that their assertions both specify the motion that the instruction must convey to the hearer.⁹ The syntax, the syntax–semantics interface, and the assertion put *slide* and *position* on an equal footing, and only the presupposition and pragmatics could distinguish the two.

With differences in presuppositions come differences in possible resolutions of discourse anaphors to discourse referents; the differences depend on the properties of salient objects in the common ground. The fewer resolutions that there are after selecting a verb, the more the verb assists the hearer in identifying the needed action. This gives a reason to prefer one verb over another. In general, we elect to specify a constraint on context as a presupposition exactly when we must model its effects on reference resolution.

In our example, general background indicates that each sleeve only has a single place where it belongs, at the joint; meanwhile, there may be many “way points” along the pipe to slide the sleeve

⁹Note that if the assertions were different in some relevant respect, the difference would provide a decisive reason for SPUD to prefer one entry over another. SPUD’s top priority is to achieve its updates. For example, SPUD would prefer an entry if its assertion achieved a specified update by describing manner of motion and alternative entries did not.

to. This makes the anaphoric interpretation of *position* less ambiguous than that of *slide*; to obtain an equally constrained interpretation with *slide*, an additional identifying modifier like *into its position* would be needed. This favors *position* over *slide*—exactly what we observe in our corpus of instructions. The example illustrates how SPUD’s meaning specifications can be developed step by step, with a close connection between the semantic distinctions we introduce in lexical entries and their consequences for generation.

With differences in pragmatics come differences in the fit between utterance and context. The more specific the pragmatics the better the fit; this gives another reason to prefer one verb over another. We did not find such cases among the motion verbs we studied, because the contextual links we identified all had effects on reference resolution and thus were specified as presuppositions. However, we anticipate that pragmatics will prove important when differences in meaning involve the perspective taken by the speaker on an event, as in the contrast of *buy* and *sell*.

Appendix B details our results for the ten verbs we studied; (74) presents the final sample entry for *slide*. The tree gives the syntax for one element in the tree family associated with *slide*, with its associated semantic indices; the associated formulas describe the semantics of the entry in terms of presuppositions and assertions about the individuals referenced in the tree.

(74) a Syntax and syntax–semantics interface:

b Assertion: $move(A, H, O, P)$

c Presupposition: $start-at(P, O) \wedge surf(P)$

Of course, the corresponding entries (75) and (81) that we used in assembling concrete communicative intent for (2) in Figure 11 refine (74) only in adopting the specific syntactic and semantic refinements of an imperative use of the verb. The entries are provided as (75) and (81) in Appendix A.

8 Previous Work

In the discussion so far, we have been able to contrast SPUD with a range of research from the sentence planning literature. As first observed in Section 2.3 and substantiated subsequently, SPUD’s representations and algorithms, and the specification strategies they afford, greatly improve on prior proposals for communicative-intent–based microplanning such as (Appelt, 1985; Thomason and Hobbs, 1997). Meanwhile, as catalogued in Section 6, SPUD captures the essence of techniques for referring expression generation, such as (Dale and Haddock, 1991); for syntactic choice, such as (Mathiessen, 1983; Yang et al., 1991); for lexical choice, such as (Nogier and Zock, 1991; Elhadad et al., 1997; Stede, 1998); and for aggregation, such as (Dalianis, 1996; Shaw, 1998).

At the same time, SPUD goes beyond these pipelined approaches in modeling and exploiting interactions among microplanning subtasks, and SPUD captures these efficiencies using a uniform

model of communicative intent. In contrast, other research has succeeded in capturing particular descriptive efficiencies only by specialized mechanisms. For example, Appelt's planning formalism includes plan-critics that can detect and collapse redundancies in sentence plans (Appelt, 1985). This framework treats subproblems in generation as independent by default; and writing tractable and general critics is hampered by the absence of abstractions like those used in SPUD to simultaneously model the syntax and the interpretation of a whole sentence. Meanwhile, in (McDonald, 1992), McDonald considers descriptions of events in domains which impose strong constraints on what information about events is semantically relevant. He shows that such material should and can be omitted, if it is both syntactically optional and inferentially derivable:

FAIRCHILD Corporation (Chantilly VA) Donald E Miller was named senior vice president and general counsel, succeeding Dominic A Petito, who resigned in November, at this aerospace business. Mr. Miller, 43 years old, was previously principal attorney for Temkin & Miller Ltd., Providence RI.

Here, McDonald points out that one does not need to explicitly mention the position that Petito resigned from in specifying the resignation sub-event, since it must be the same as the one that Miller has been appointed to. Whereas McDonald adopts special-purpose module to handle this, we regard it as a special case of pragmatic overloading.

More generally, like many sentence planners, SPUD achieves a flexible association between the content input to a sentence planner and the meaning that comes out. Other researchers (Nicolov et al., 1995; Rubinoff, 1992) have assumed that this flexibility comes from a mismatch between input content and grammatical options. In SPUD, such differences arise from the referential requirements and inferential opportunities that are encountered.

Previous authors (McDonald and Pustejovsky, 1985; Joshi, 1987) have noted that TAG has many advantages for generation as a syntactic formalism, because of its localization of argument structure. (Joshi, 1987) states that adjunction is a powerful tool for elaborating descriptions. These aspects of TAGs are crucial for us; for example, lexicalization allows us to easily specify local semantic and pragmatic constraints imposed by the lexical item in a particular syntactic frame.

Various efforts at using TAG for generation (McDonald and Pustejovsky, 1985; Joshi, 1987; Yang et al., 1991; Danlos, 1996; Nicolov et al., 1995; Wahlster et al., 1991) enjoy many of these advantages. They vary in the organization of the linguistic resources, the input semantics and how they evaluate and assemble alternatives. Furthermore, (Shieber et al., 1990; Shieber, 1991; Prevost and Steedman, 1993; Hoffman, 1994) exploit similar benefits of lexicalization and localization. Our approach is distinguished by its declarative synthesis of a representation of communicative intent, which allows SPUD to construct a sentence and its interpretation simultaneously.

9 Conclusion

Most generation systems pipeline pragmatic, semantic, lexical and syntactic decisions (Reiter, 1994). With the right formalism—an explicit, declarative representation of COMMUNICATIVE INTENT—it is easier and better to construct pragmatics, semantics and syntax simultaneously. The approach elegantly captures the interaction between pragmatic and syntactic constraints on descriptions in a sentence, and the inferential interactions between multiple descriptions in a sentence. At the same time, it exploits linguistically motivated, declarative specifications of the discourse functions of syntactic constructions to make contextually appropriate syntactic choices.

Realizing a microplanner based on communicative intent involves challenges in implementation and specification. In the past (Appelt, 1985), these challenges may have made communicative-intent-based microplanning seem hopeless and intractable. Nevertheless, in this paper, we have described an effective implementation, SPUD, that constructs representations of communicative intent through top-down LTAG derivation, logic-programming and constraint-satisfaction models of interpretation, and greedy search; and we have described a systematic, step-by-step methodology for designing generation grammars for SPUD.

With these results, the challenges that remain for the program of microplanning based on communicative intent offer fertile ground for further research. SPUD's model of interpretation omits important features of natural language, such as plurality (Stone, 2000a), discourse connectivity (Webber et al., 1999) and such defeasible aspects of interpretation as presupposition-accommodation (Lewis, 1979). SPUD's search procedure is simplistic, and is vulnerable to stalled states where lookahead is required to recognize the descriptive effect of a combination of lexical items. (Gardent and Striegnitz, 2001) illustrate how refinements in SPUD's models of interpretation and search can lead to interesting new possibilities for NLG. At the same time, the construction of lexicalized grammars for generation with effective representations of semantics calls out for automation, using techniques that make lighter demands on developers and make better use of machine learning.

Acknowledgments

This paper has been brewing for a long time—during which the investigation has been supported by: NSF and IRCS graduate fellowships, a RuCCS postdoctoral fellowship, NSF grant NSF-STC SBR 8920230, NSF Research Instrumentation award 9818322, ARPA grant N6600194C6-043, and ARO grant DAAH04-94-G0426. We are grateful for discussion and comments from Gann Bierner, Betty Birner, Julie Bourne, Justine Cassell, Aravind Joshi, Alistair Knott, Ellen Prince, Owen Rambow, Joseph Rosenzweig, Anoop Sarkar, Mark Steedman, Mike White and Hao Yan.

References

- Appelt, D. (1985). *Planning English Sentences*. Cambridge University Press, Cambridge England.
- Bach, E. (1989). *Informal Lectures on Formal Semantics*. State University of New York Press, Albany, NY.
- Badler, N., Bindiganavale, R., Allbeck, J., Schuler, W., Zhao, L., Lee, S.-J., Shin, H., and Palmer, M. (2000). Parameterized action representation and natural language instructions for dynamic behavior modification of embodied agents. In *Agents*.
- Badler, N., Palmer, M., and Bindiganavale, R. (1999). Animation control for real-time virtual humans. *Communications of the ACM*, 42(7):65–73.
- Baldoni, M., Giordano, L., and Martelli, A. (1998). A modal extension of logic programming: Modularity, beliefs and hypothetical reasoning. *Journal of Logic and Computation*, 8(5):597–635.
- Birner, B. (1992). *The Discourse Function of Inversion in English*. PhD thesis, Northwestern University.

- Bourne, J. (1998). Generating effective instructions: Knowing when to stop. PhD Thesis Proposal, Department of Computer & Information Science, University of Pennsylvania.
- Bratman, M. E. (1987). *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge, MA.
- Cahill, L. and Reape, M. (1999). Component tasks in applied NLG systems. Technical Report ITRI-99-05, ITRI, University of Brighton.
- Candito, M. and Kahane, S. (1998). Can the TAG derivation tree represent a semantic graph? an answer in the light of Meaning-Text Theory. In *TAG+4*.
- Carberry, S. and Lambert, L. (1999). A process model for recognizing communicative acts and modeling negotiation subdialogues. *Computational Linguistics*, 25:1–53.
- Cassell, J. (2000). Embodied conversational interface agents. *Communications of the ACM*, 43(4):70–78.
- Cassell, J., Stone, M., and Yan, H. (2000). Coordination and context-dependence in the generation of embodied conversation. In *First International Conference on Natural Language Generation*, pages 171–178.
- Chen, J. and Vijay-Shanker, K. (2000). Automated extraction of TAGs from the Penn treebank. In *Proceedings of the 6th International Workshop on Parsing Technologies*, Trento, Italy.
- Cheng, H. and Mellish, C. (2000). An empirical analysis of constructing non-restrictive NP components to express semantic relations. In *First International Conference on Natural Language Generation*, pages 108–115.
- Clark, H. H. (1996). *Using Language*. Cambridge University Press, Cambridge, UK.
- Clark, H. H. and Marshall, C. R. (1981). Definite reference and mutual knowledge. In Joshi, A. K., Webber, B. L., and Sag, I., editors, *Elements of Discourse Understanding*, pages 10–63. Cambridge University Press, Cambridge.
- Dale, R. (1992). *Generating Referring Expressions: Constructing Descriptions in a Domain of Objects and Processes*. MIT Press, Cambridge MA.
- Dale, R. and Haddock, N. (1991). Content determination in the generation of referring expressions. *Computational Intelligence*, 7(4):252–265.
- Dale, R. and Reiter, E. (1995). Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science*, 18:233–263.
- Dalianis, H. (1996). *Concise Natural Language Generation from Formal Specifications*. PhD thesis, Royal Institute of Technology, Stockholm. Department of Computer and Systems Sciences.
- Danlos, L. (1996). G-TAG: A formalism for Text Generation inspired from Tree Adjoining Grammar: TAG issues. Unpublished manuscript, TALANA, Université Paris 7.

- Davidson, D. (1980). The logical form of action sentences. In *Essays on actions and events*, pages 105–148. Clarendon Press, Oxford.
- Debart, F., Enjalbert, P., and Lescot, M. (1992). Multimodal logic programming using equational and order-sorted logic. *Theoretical Computer Science*, 105:141–166.
- Di Eugenio, B. and Webber, B. (1996). Pragmatic overloading in natural language instructions. *International Journal of Expert Systems*, 9(2):53–84.
- Doran, C., Egedi, D., Hockey, B. A., Srinivas, B., and Zaidel, M. (1994). XTAG System - a wide coverage grammar for English. In *Proceedings of COLING*.
- Elhadad, M., McKeown, K., and Robin, J. (1997). Floating constraints in lexical choice. *Computational Linguistics*, 23(2):195–240.
- Elhadad, M. and Robin, J. (1992). Controlling content realization with functional unification grammars. In Dale, R., Hovy, E., Rösner, D., and Stock, O., editors, *Aspects of Automated Natural Language Generation: 6th International Workshop on Natural Language Generation*, Lecture Notes in Artificial Intelligence 587, pages 89–104. Springer Verlag, Berlin.
- Fariñas del Cerro, L. (1986). MOLOG: A system that extends PROLOG with modal logic. *New Generation Computing*, 4:35–50.
- Gardent, C. and Striegnitz, K. (2001). Generating indirect anaphora. In *Proceedings of IWCS*, pages 138–155.
- Grice, H. P. (1957). Meaning. *The Philosophical Review*, 66:377–388.
- Gundel, J. K., Hedberg, N., and Zacharski, R. (1993). Cognitive status and the form of referring expressions in discourse. *Language*, 69(2):274–307.
- Haddock, N. (1989). *Incremental Semantics and Interactive Syntactic Processing*. PhD thesis, University of Edinburgh.
- Hardt, D. (1999). Dynamic interpretation of verb phrase ellipsis. *Linguistics and Philosophy*, 22(2):187–221.
- Hart, P. E., Nilsson, N. J., and Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on SSC*, 4:100–107. Correction in *SIGART Newsletter* 37:28–29.
- Heeman, P. and Hirst, G. (1995). Collaborating on referring expressions. *Computational Linguistics*, 21(3).
- Hirschberg, J. (1985). *A Theory of Scalar Implicature*. PhD thesis, University of Pennsylvania.
- Hobbs, J., Stickel, M., Appelt, D., and Martin, P. (1993). Interpretation as abduction. *Artificial Intelligence*, 63:69–142.
- Hobbs, J. R. (1985). Ontological promiscuity. In *Proceedings of ACL*, pages 61–69.

- Hobbs, J. R., Stickel, M., Appelt, D., and Martin, P. (1988). Interpretation as abduction. In *Proceedings of ACL*, pages 95–103.
- Hockenmaier, J., Bierner, G., and Baldridge, J. (2001). Providing robustness for a CCG system. Submitted.
- Hoffman, B. (1994). Generating context-appropriate word orders in Turkish. In *Proceedings of the Seventh International Generation Workshop*.
- Horacek, H. (1995). More on generating referring expressions. In *Proceedings of the Fifth European Workshop on Natural Language Generation*, pages 43–58, Leiden.
- Jackendoff, R. (1977). *$\bar{X}$ Syntax: A Study of Phrase Structure*. MIT.
- Jackendoff, R. S. (1990). *Semantic Structures*. MIT Press, Cambridge, MA.
- Joshi, A. and Vijay-Shanker, K. (1999). Compositional semantics with lexicalized tree-adjoining grammar (LTAG). In *International Workshop on Computational Semantics*, pages 131–145.
- Joshi, A. K. (1987). The relevance of tree adjoining grammar to generation. In Kempen, G., editor, *Natural Language Generation*, pages 233–252. Martinus Nijhoff Press, Dordrecht, The Netherlands.
- Joshi, A. K., Levy, L., and Takahashi, M. (1975). Tree adjunct grammars. *Journal of the Computer and System Sciences*, 10:136–163.
- Kallmeyer, L. and Joshi, A. (1999). Factoring predicate argument and scope semantics: underspecified semantics with LTAG. In *12th Amsterdam Colloquium*.
- Kamp, H. and Rossdeutscher, A. (1994). DRS-construction and lexically driven inference. *Theoretical Linguistics*, 20:97–164.
- Kittredge, R., Korelsky, T., and Rambow, O. (1991). On the need for domain communication knowledge. *Computational Intelligence*, 7(4):305–314.
- Lascarides, A. and Asher, N. (1991). Discourse relations and defeasible knowledge. In *Proceedings of ACL 29*, pages 55–62.
- Levelt, W. J. M. (1989). *Speaking*. MIT, Cambridge, MA.
- Levin, B. (1993). *English Verb Classes and Alternations: A preliminary investigation*. University of Chicago, Chicago.
- Lewis, D. (1979). Scorekeeping in a language game. In *Semantics from Different Points of View*, pages 172–187. Springer Verlag, Berlin.
- Mackworth, A. (1987). Constraint Satisfaction. In Shapiro, S., editor, *Encyclopedia of Artificial Intelligence*, pages 205–211. John Wiley and Sons.

- Mathiessen, C. M. I. M. (1983). Systemic grammar in computation: the Nigel case. In *Proceedings of EACL*, pages 155–164.
- McDonald, D. (1992). Type-driven suppression of redundancy in the generation of inference-rich reports. In Dale, R., Hovy, E., Rösner, D., and Stock, O., editors, *Aspects of Automated Natural Language Generation: 6th International Workshop on Natural Language Generation*, Lecture Notes in Artificial Intelligence 587, pages 73–88. Springer Verlag, Berlin.
- McDonald, D. D. and Pustejovsky, J. D. (1985). TAG’s as a grammatical formalism for generation. In *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics*, pages 94–103, Chicago, IL.
- Mellish, C., O’Donnell, M., Oberlander, J., and Knott, A. (1998). An architecture for opportunistic text generation. In *9th International Workshop on Natural Language Generation*.
- Mellish, C. S. (1985). *Computer Interpretation of Natural Language Descriptions*. Ellis Horwood, Chichester, UK.
- Miller, D., Nadathur, G., Pfenning, F., and Scedrov, A. (1991). Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic*, 51:125–157.
- Moore, J. (1994). *Participating in Explanatory Dialogues*. MIT Press, Cambridge MA.
- Moore, J. D. and Paris, C. L. (1993). Planning text for advisory dialogues: capturing intentional and rhetorical information. *Computational Linguistics*, 19(4):651–695.
- Nicolov, N., Mellish, C., and Ritchie, G. (1995). Sentence generation from conceptual graphs. In G. Ellis, R. Levinson, W. R. and Sowa, F., editors, *Conceptual Structures: Applications, Implementation and Theory (Proceedings of Third International Conference on Conceptual Structures)*, pages 74–88. Springer.
- Nilsson, N. (1971). *Problem-solving Methods in Artificial Intelligence*. McGraw-Hill, New York.
- Nogier, J. and Zock, M. (1991). Lexical choice as pattern-matching. In Nagle, T., Nagle, J., Gerholz, L., and Eklund, P., editors, *Current Directions in Conceptual Structures Research*. Springer.
- Palmer, M. (1990). *Semantic Processing for Finite Domains*. Cambridge University Press.
- Pereira, F. C. N. and Shieber, S. M. (1987). *Prolog and Natural Language Analysis*. CSLI, Stanford CA.
- Pollack, M. (1991). Overloading intentions for efficient practical reasoning. *Noûs*, 25:513–536.
- Pollack, M. E. (1992). The uses of plans. *Artificial Intelligence*, 57:43–68.
- Prevost, S. and Steedman, M. (1993). Generating contextually appropriate intonation. In *Proceedings of the Sixth Conference of the European Chapter of ACL*, pages 332–340, Utrecht.
- Prince, A. and Smolensky, P. (1997). Optimality: From neural networks to universal grammar. *Science*, 275:1604–1610.

- Prince, E. (1986). On the syntactic marking of presupposed open propositions. In *Proceedings of the 22nd Annual Meeting of the Chicago Linguistic Society*, pages 208–222, Chicago. CLS.
- Rambow, O. and Korelsky, T. (1992). Applied text generation. In *Applied Natural Language Processing Conference*, pages 40–47.
- Reiter, E. (1994). Has a consensus NL generation architecture appeared, and is it psycholinguistically plausible? In *Seventh International Workshop on Natural Language Generation*, pages 163–170.
- Reiter, E. and Dale, R. (2000). *Building Natural Language Generation Systems*. Cambridge.
- Rich, C., Sidner, C. L., and Lesh, N. (2001). COLLAGEN: applying collaborative discourse theory to human-computer interaction. *AI Magazine*. to appear.
- Rubinoff, R. (1992). Integrating text planning and linguistic choice by annotating linguistic structures. In Dale, R., Hovy, E., Rösner, D., and Stock, O., editors, *Aspects of Automated Natural Language Generation: 6th International Workshop on Natural Language Generation*, Lecture Notes in Artificial Intelligence 587, pages 45–56. Springer Verlag, Berlin.
- Saeboe, K. J. (1996). Anaphoric presuppositions and zero anaphora. *Linguistics and Philosophy*, 19(2):187–209.
- Sarkar, A. (2001). Applying co-training methods to statistical parsing. In *Proceedings of the North American Association for Computational Linguistics*.
- Schabes, Y. (1990). *Mathematical and Computational Aspects of Lexicalized Grammars*. PhD thesis, Computer Science Department, University of Pennsylvania.
- Shaw, J. (1998). Clause aggregation using linguistic knowledge. In *Ninth International Workshop on Natural Language Generation*, pages 138–148.
- Shieber, S., van Noord, G., Pereira, F., and Moore, R. (1990). Semantic-head-driven generation. *Computational Linguistics*, 16:30–42.
- Shieber, S. M. (1991). A uniform architecture for parsing and generation. In *ICLP*, pages 614–619.
- Stede, M. (1998). A generative perspective on verb alternations. *Computational Linguistics*, 24(3):401–430.
- Steedman, M. (1997). Temporality. In van Benthem, J. and ter Meulen, A., editors, *Handbook of Logic and Language*, pages 895–935. Elsevier.
- Stone, M. (1998). *Modality in Dialogue: Planning, Pragmatics and Computation*. PhD thesis, University of Pennsylvania.
- Stone, M. (1999). Indefinite information in modal logic programming. Technical Report RUCCS Report 56, Rutgers University.

- Stone, M. (2000a). On identifying sets. In *First International Conference on Natural Language Generation*, pages 116–123.
- Stone, M. (2000b). Towards a computational account of knowledge, action and inference in instructions. *Journal of Language and Computation*, 1:231–246.
- Stone, M., Bleam, T., Doran, C., and Palmer, M. (2000). Lexicalized grammar and the description of motion events. In *TAG+: Workshop on Tree-Adjoining Grammar and Related Formalisms*, pages 199–206.
- Stone, M. and Doran, C. (1996). Paying heed to collocations. In *International Natural Language Generation Workshop*, pages 91–100.
- Stone, M. and Doran, C. (1997). Sentence planning as description using tree-adjoining grammar. In *Proceedings of ACL*, pages 198–205.
- Stone, M. and Webber, B. (1998). Textual economy through close coupling of syntax and semantics. In *Proceedings of International Natural Language Generation Workshop*, pages 178–187.
- Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, 12:49–100.
- The XTAG-Group (1995). A Lexicalized Tree Adjoining Grammar for English. Technical Report IRCS 95-03, University of Pennsylvania. Updated version available at <http://www.cis.upenn.edu/~xtag/tr/tech-report.html>.
- Thomason, R. H. (1990). Accommodation, meaning and implicature. In Cohen, P. R., Morgan, J., and Pollack, M. E., editors, *Intentions in Communication*, pages 325–363. MIT Press, Cambridge, MA.
- Thomason, R. H., Hobbs, J., and Moore, J. (1996). Communicative goals. In *ECAI Workshop on Gaps and Bridges: New Directions in Planning and Natural Language Generation*.
- Thomason, R. H. and Hobbs, J. R. (1997). Interrelating interpretation and generation in an abductive framework. In *AAAI Fall Symposium on Communicative Action*.
- van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics*, 9(2):333–377.
- Vijay-Shanker, K. (1987). *A Study of Tree Adjoining Grammars*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania.
- Wahlster, W., André, E., Bandyopadhyay, S., Graf, W., and Rist, T. (1991). WIP: The coordinated generation of multimodal presentations from a common representation. In Stock, O., Slack, J., and Ortony, A., editors, *Computational Theories of Communication and their Applications*. Berlin: Springer Verlag.
- Wanner, L. and Hovy, E. (1996). The HealthDoc sentence planner. In *Seventh International Workshop on Natural Language Generation*, pages 1–10.

- Ward, G. (1985). *The Semantics and Pragmatics of Preposing*. PhD thesis, University of Pennsylvania. Published 1988 by Garland.
- Webber, B., Knott, A., Stone, M., and Joshi, A. (1999). Discourse relations: A structural and pre-suppositional account using lexicalised TAG. In *Association for Computational Linguistics*, pages 41–48.
- Webber, B. L. (1988). Tense as discourse anaphor. *Computational Linguistics*, 14(2):61–73.
- Webber, B. L., Carberry, S., Clarke, J. R., Gertner, A., Harvey, T., Rymon, R., and Washington, R. (1998). Exploiting multiple goals and intentions in decision support for the management of multiple trauma: A review of the TraumAID project. *Artificial Intelligence*, 105:263–293.
- Xia, F., Palmer, M., Vijay-Shanker, K., and Rosenzweig, J. (1998). Consistent grammar development using partial tree-descriptions for lexicalized tree-adjoining grammars. In *TAG+4*.
- Yang, G., McCoy, K. F., and Vijay-Shanker, K. (1991). From functional specification to syntactic structures: systemic grammar and tree-adjoining grammar. *Computational Intelligence*, 7(4):207–219.

A Instruction Grammar Fragment

A.1 Syntactic Constructions

(75) a NAME: axnpVnpopp

b PARAMETERS: A, H, O, P, S

c PRAGMATICS: $obl(S, H)$

(76) a NAME: bvpPsinf

b PARAMETERS: $A1, H, A2$

c PRAGMATICS: —

(77) a NAME: anpxVinp

b PARAMETERS: A, H, O

c PRAGMATICS: —

- (78) a NAME zeroDefNP
b PARAMETERS: R
c PRAGMATICS: $zero\text{-}genre \wedge def(R)$

- (79) a NAME: bvpPnp
b PARAMETERS: E, O, P, R
c PRAGMATICS: $zero\text{-}genre \wedge def(R)$

- (80) a NAME: bNnn
b PARAMETERS: A, B
c PRAGMATICS: $def(A)$

A.2 Lexical Entries

- (81) a NAME: slide
b PARAMETERS: A, H, O, P, S
c CONTENT: $move(A, H, O, P) \wedge next(A)$
d PRESUPPOSITION: $start\text{-}at(P, O) \wedge surf(P) \wedge partic(S, H)$
e PRAGMATICS: —

- f TARGET: $S(A)$ [complement]
 - g TREE LIST: $\text{axnpVnpopp}(A, H, O, P, S)$
- (82)
- a NAME: $\langle \text{purpose} \rangle$
 - b PARAMETERS: $A1, H, A2$
 - c CONTENT: $\text{purpose}(A1, A2)$
 - d PRESUPPOSITION: —
 - e PRAGMATICS: —
 - f TARGET: $\text{VP}_2(A1, H)$ [modifier]
 - g TREE LIST: $\text{bvpPsinf}(A1, H, A2)$
- (83)
- a NAME: uncover
 - b PARAMETERS: A, H, O
 - c CONTENT: $\text{uncover}(A, H, O)$
 - d PRESUPPOSITION: —
 - e PRAGMATICS: —
 - f TARGET: $S_i(A, H)$
 - g TREE LIST: $\text{anpxVinp}(A, H, O)$
- (84)
- a NAME: sealing-ring
 - b PARAMETERS: N
 - c CONTENT: $\text{sr}(N)$
 - d PRESUPPOSITION: —
 - e PRAGMATICS: —
 - f TARGET: $\text{NP}(N)$ [complement]
 - g TREE LIST: $\text{zerodefnpintree}(N)$
- (85)
- a NAME: coupling-nut
 - b PARAMETERS: N
 - c CONTENT: $\text{cn}(N)$
 - d PRESUPPOSITION: —
 - e PRAGMATICS: —
 - f TARGET: $\text{NP}(N)$ [complement]
 - g TREE LIST: $\text{zerodefnpintree}(N)$
- (86)
- a NAME: onto
 - b PARAMETERS: E, O, P, R
 - c CONTENT: $\text{end-on}(P, R)$
 - d PRESUPPOSITION: —
 - e PRAGMATICS: —
 - f TARGET: $\text{VP}_{\text{path}(E, O, P)}$ [modifier]

- g TREE LIST: $\text{bvpPnp}(E, O, P, R)$
- (87) a NAME: elbow
b PARAMETERS: N
c CONTENT: $el(N)$
d PRESUPPOSITION: —
e PRAGMATICS: —
f TARGET: $\text{NP}(N)$ [complement]
g TREE LIST: $\text{zerodefnpintree}(N)$
- (88) a NAME: fuel-line
b PARAMETERS: N, R, X
c CONTENT: $fl(N) \wedge nn(R, N, X)$
d PRESUPPOSITION: —
e PRAGMATICS: —
f TARGET: $N'(R)$ [modifier]
g TREE LIST: $\text{bNnn}(N)$

B Motion Verb Entries

B.1 Pure Motion Verbs

The verbs *slide*, *rotate*, *turn*, *push*, *pull*, and *lift* all share a use in which they describe an event A in which some agent H moves an object O along a path P . Our analysis of this use was presented in detail in Section 7. (89) gives the syntactic frame for this class.

Semantically, *slide*, *rotate* and *turn* all assert simple motions; the verbs differ in that *slide* presupposes motion along a surface while *turn* presupposes a circular or helical path around an axis by which an object can pivot and *rotate* presupposes a circular path around an axis through the center of an object. (90) represents this.

- (90) a *slide*: assert $\text{move}(A, H, O, P)$; presuppose $\text{start-at}(P, O) \wedge \text{surf}(P)$
b *turn*: assert $\text{move}(A, H, O, P)$; presuppose $\text{start-at}(P, O) \wedge \text{around}(P, X) \wedge \text{pivot}(O, X)$
c *rotate*: assert $\text{move}(A, H, O, P)$; presuppose $\text{start-at}(P, O) \wedge \text{around}(P, X) \wedge \text{center}(O, X)$

The verbs *push*, *pull* and *lift* involve force as well as motion; they differ in presuppositions about the direction of force and motion: for *push*, it is away from the agent; for *pull*, it is towards the agent; *lift* has an upward component:

- (91) a push: assert *forced-move*(A, H, O, P); presuppose *start-at*(P, O) $\wedge$ *away*(P, H)
 b pull: assert *forced-move*(A, H, O, P); presuppose *start-at*(P, O) $\wedge$ *towards*(P, H)
 c lift: assert *forced-move*(A, H, O, P); presuppose *start-at*(P, O) $\wedge$ *upwards*(P)

B.2 Pure Change-of-state Verbs

This category of verbs describes an event A in which an agent H changes of state of an object O ; these verbs appeal to a single optional semantic argument U which helps to specify what the change of state is. Examples of this class are *remove* [*from* U], *disconnect* [*from* U] and *connect* [*to* U]; U is a landmark object and the change-of-state involves a spatial or connection relation between O and U .

Our diagnostic tests give a number of reasons to think of the parameter U as a semantic argument that is referenced in the tree but described by syntactic adjuncts. Here are illustrations of these tests for the case of *disconnect*. It is possible to extract from it, and impossible to supply it by *do so* substitution.

- (92) a What did you disconnect the cable from ϵ ?
 b ?Mary disconnected a coupling from system A, and John did so from system B.

It is possible to take the initial connection between O and U as presupposed, and to factor in this constraint in identifying O and U . Thus, with many systems and couplings, we might still find:

- (93) Disconnect the coupling from system A.

These considerations lead to the syntactic frame of (94).

Note that syntactic features can allow the verb to determine which preposition is used to specify the optional argument. That is, we can use lexical entries for verbs that indicate that they impose feature-value constraints on the syntactic features of the anchor $v \diamond$ node.

In order to characterize the semantics of change-of-state verbs, we introduce a predicate *caused-event*(A, H, O) indicating that A is an event in which H has a causal effect on O ; and an operator *result*(A, p) indicating that the proposition p holds in the state that results from doing A . (For more on this ontology, see (Steedman, 1997).) (95) uses this notation to describe *connect*, *disconnect* and *remove*.

- (95) a connect: assert *caused-event*(A, H, O) $\wedge$ *result*($A, connected(O, U)$); presuppose *free*(O, U)
 b disconnect: assert *caused-event*(A, H, O) $\wedge$ *result*($A, free(O, U)$); presuppose *connected*(O, U)

- c remove: assert *caused-event*(A, H, O) $\wedge$ *result*($A, \text{free}(O, U)$); presuppose *dependent*(O, U)

That is, *connecting* causes O to be connected to the optional argument U where O is presupposed to be presently spatially independent of, or *free* of, U ; *disconnecting*, conversely, causes O to be *free* of U , where O is presupposed to be connected to U . Finally, *remove* is more general than *disconnect*. It presupposes only that there is some *dependent* spatial relation between O and U ; O may be attached to U , supported by U , contained in U , etc.

B.3 Near-motion Verbs

Distinct from motion verbs and ordinary change-of-state verbs is a further class which we have called near-motion verbs: near-motion verbs are change-of-state verbs that encode a spatial change by evoking the final location where an object comes to rest. Semantically, they involve arguments A, H, O , and L —the fourth, spatial argument L represents a spatial configuration rather than a path (as in the case of motion verbs). The canonical near-motion verb is *position*; others are *reposition* and *install*. According to our judgments, *turn* and *rotate* can be used as near-motion verbs as well as genuine motion verbs, whereas *slide*, *push*, *pull* and *lift* cannot.

Now, whenever there is a change of location, there must be motion (in our domain); and whenever an object moves to a new place, there is a change of location. This semantic correspondence between motion verbs and near-motion verbs is mirrored in similar syntactic realizations with prepositional phrases that describe an final location. So we find both:

- (96) a Push the coupling on the sleeve.
b Position the coupling on the sleeve.

The difference between motion verbs and near-motion verbs is that motion verbs permit an explicit description of the PATH the object takes during the motion, while near-motion verbs do not:

- (97) a Push the coupling to the sleeve.
b *Position the coupling to the sleeve.

Another way to substantiate the contrast is to consider the interpretation of ambiguous modifiers. In (98a), *downward* modifies the path by describing the direction of motion in the event. In (98b), with the near-motion verb, this path interpretation is not available: the reading of *downward* instead is that it describes the final orientation of the object that is manipulated.

- (98) a Push handle downward.
b Position handle downward.

These readings are paraphrased in (99).

- (99) a Push handle in a downward direction.
b Position handle so that it is oriented downward.

The natural *wh*-questions associated with the two constructions are also different:

- (100) a { In which direction, *How } did you push the handle? Downward.

- b { *In which direction, How } did you position the handle? Downward.

(101) schematizes the syntax of near-motion verbs.

Like motion verbs, near-motion verbs share a common assertion—there is an event A of H acting on O whose result is that O is located at place L . The differences among near-motion verbs lie in their presuppositions: *position* presupposes that L is a position in which O will be able to perform its intended function, as in (102a); *reposition* further presupposes a state preceding A where O was located at L —we write this as *back*(A, O, L) in (102b); finally, *install* presupposes that the spatial position for O is one which fastens O tightly, as in (102c).

- (102) a position: assert *caused-event*(A, H, O) $\wedge$ *result*($A, loc(L, O)$); presuppose *position-for*(L, O)
 b reposition: assert *caused-event*(A, H, O) $\wedge$ *result*($A, loc(L, O)$); presuppose *position-for*(L, O) $\wedge$ *back*(A, O, L)
 c install: assert *caused-event*(A, H, O) $\wedge$ *result*($A, loc(L, O)$); presuppose *position-for*(L, O) $\wedge$ *fastening*(L, O)

B.4 Put Verbs

Closely related to the near-motion verbs are the *put* verbs. These differ from near-motion verbs only in that *put* verbs take the configuration PP as a syntactic complement—rather than as an optional syntactic modifier.

Verbs in this class include not only *put*, but also *place*.

- (104) a put: assert *caused-event*(A, H, O) $\wedge$ *result*($A, loc(L, O)$)
 b place: assert *body-caused-event*(A, H, O) $\wedge$ *result*($A, loc(L, O)$); presuppose *place-for*(L, O)

Note that a placement must be performed by hand; the presupposition that *L* be a *place* for *O* signifies that *O*'s specific location at *L* is required for the success of future actions or events. (*Place* contrasts with *position* in that places depend on the action of an agent on the object in a particular activity whereas positions are enduring regions that depend on the functional properties of the object itself; contrast *working place* and *working position*.)