

On the Creative Use of Language: the Form of Lexical Resources

David D. McDonald
14 Brantwood Road
Arlington, MA 02174, USA
mcdonald@cs.brandeis.edu

Federica Busa
Computer Science Department
Brandeis University 02254, USA
federica@cs.brandeis.edu

Abstract

We introduce the problem of referential creativity: how it is that a person can give a word or phrase a new denotation even though she has never heard or used in that way before. Using real examples that we have collected, we focus on the case of semantic type coercion, where a phrase of a given type is used in a textual context that requires a completely different type yet the intended message is perfectly well understood.

We frame our account as the problem of what form the linguistic resources available to a speaker must have such that she can appreciate the opportunity for creative phrasings. What systematic relationships exist in the speaker's lexicon that enable a phrase to convey something quite different than it normally would, and why should this ever occur during the generation process? We draw on a new theory of how information is associated with a word—Pustejovsky's Generative Lexicon—and we embed our account in a theory of generation as an incremental process that makes use of a rich model of the situation in which the utterance occurs.

Keywords: creativity, lexical choice, content determination, incremental generation

1. The problem

People are always saying new things. Usually the reason is obvious: we continually find ourselves in new situations, and the differences in these situations—our goals, the participants, the activities, the props, the locations, the recent history, etc.—naturally lead to differences in what we say. Other cases are not so easily explained. Differences in the situation will not account for those occasions when our choice of phrasing is inconsistent with the selectional restrictions of its context; when we make a reference that on its face should not be able to communicate our intended meaning yet in the situation in which it is uttered is both natural and completely understood.

Consider texts 1 and 2, which were uttered by the authors and transcribed immediately afterwards. Both were completely understood by their hearers.

- (1) *"I would really like to have you guys over for dinner, so let me know whether for you it is better before or after Florida."*
- (2) *"Just give me a few bites"*

In (1), a geographical location, Florida, provides information about a time interval. In (2), "a few bites (of food)" denotes a quantity, yet it is readily interpreted as a duration.¹ If hearers made their interpretations strictly

compositionally, these would be nonsense to them. If speakers strictly respected selectional restrictions, they would never generate such sentences. What is it that explains this ability? What do people know about their language and its relationship to situations that prompts them to produce these apparently type-inconsistent utterances and while still being understood?

We will argue in this paper that this ability follows from the use of a *generative lexicon*, one based on the dynamic operation of general rules that define what realizations are possible through interactions between the lexical semantics of our linguistic resources and the situation, rather than through the customary fixed enumeration of cases and decision criteria. Coupled with a theory of content selection as an incremental walk through a highly structured mental model of the speaker's 'relevant portion of the situation', and with an architecture for generation with multiple levels of (progressively more linguistic) representation that actively constrains what realizations are possible, this view of the lexicon can explain how speakers are led to these unusual 'creative' constructions and why they are understood by their hearers.

1.1 Related phenomena

There is every reason to believe that these two utterances and the others we will present later are new and unique. They were semantic patterns of expression that we had never uttered before nor ever heard anyone else say.

¹ The normal way to express the message in (2) would be "just give me a little while longer", meaning that the speaker wants to wait longer and, in this case, to eat a little bit more before doing something his interlocutor

has requested. Example (1) involves an event that takes place in Florida. It is discussed below.

Nevertheless, they are similar in kind to other turns of phrase that have become quite commonplace:

(3) “*Do you take Visa?*”

(4) “*With encouragement from Moscow, Bosnian Serbs appear to be complying with the UN ultimatum to remove their heavy weapons from Sarajevo.*”

(4) is an example of a now conventional meronymic construction where a part of an object, here the capital city Moscow, is able to stand in for the whole. This is a natural move that is probably prompted by the salience of the ‘part’ in the speaker’s mind: Moscow is, after all, the actual site of the governmental decision making even though we take the act as being done by the Russian Republic as a whole. Meronymy and other kinds of metonymic substitutions have been extensively studied (e.g. Cruise 1986, Leher & Kittay 1992 and their references) our contribution will be to couch some aspects of these phenomena in a computational context.

Example (3) comes closer to the problems of lexical semantics that concern us in this paper. Here the reconstruction that the hearer must make is not part-whole but based on the object’s purpose—what has been called its ‘telic’ role. It also reflects the ‘polymorphism’ that typifies most lexical knowledge, whereby one word, the name of the credit card, is able to convey both the physical aspect of the card as a piece of plastic and its role as a means of commercial exchange. Different situations will involve different aspects of the credit card, yet they all can be conveyed by the very same word.

Similarly, the American speaker’s proclivity to ‘verb’ nouns is well-known:²

(5) “*You can Mastercard your pledge.*”

(6) “*I’ll sprig the grapes.*”

What these and other creative uses of words have in common is that they are only uttered in what we call ‘high salience’ situations—situations where the activity under discussion is completely clear to all the interlocutors and the interlocutors are only attending a few, mutually known, entities. But this class of utterances as well, however, follows patterns that we all appreciate as completely conventional, though at one time they too were new.

(7) “*Roy irons his own shirts*”

(8) “*Bob vacuumed the floor without being asked*”

Here the verbs are the same word as the nouns that name the instruments that are prototypically used to perform the actions: ‘iron’, ‘vacuum (cleaner)’. As with credit cards, these words are taking on a role that relies, polymorphically, on their telic uses as instruments for certain prototypical acts.

² Example (5) was spoken by a public radio announcer during an on-air fund drive. (6) was spoken by one of the authors and referred to the operation of taking a large bunch of grapes and cutting it into small, hand-sized portions (‘sprigs’) as part of getting ready for a dinner party.

The capability these examples illustrate is a new problem for research in natural language generation. It points to a productive capacity that people have that enables us to exploit our knowledge of the language to construct new uses for established terms, in the process providing shorter, more cohesive realizations for content that could otherwise be quite awkward to express. It is an important problem because it is part of the scientific problem of understanding what language is and how people employ it, and also because having an account it should aid us in the engineering problem of designing machine speakers by providing insights into representations and procedures that will make them more natural and fluent.

2. Referential Creativity

Every new phenomenon needs a name, and our choice is “referential creativity”. We call it *creative* because the speaker has never before heard or used that precise pattern of semantic elements and grammatical relations. The principles and representational structures involved are surely an established part of our knowledge of language, but their use is the application of productive principles in the face of a new situation, and not the instantiation of a well-trod template or mapping.

Creativity is taken by many to be a central property of the human use of language. Chomsky (1966) introduced it into modern linguistics, and took the notion back to ideas implicit in the works of von Humboldt in the early 1800s. He used it to revolutionize the machinery of syntax. We apply it here to the process of generation, building on its application to lexical semantics by Pustejovsky (1991).

We call this phenomenon *referential* creativity because in it a speaker uses a new, ostensibly type-inconsistent referent in a linguistic context that would normally be filled with a referent of a quite different type. In (1) a geographical entity was used instead of the end-points of an event; in (2) a measurement was used instead of a duration.

The referent that is used and the one that it invokes in the mind of the hearer are always logically related. We can see this in (7), which is one of the less abbreviated ways to realize the content expressed in (1).

(7) “... *before you leave for your vacation in Florida or after you get back*”

Here the reference to Florida appears as a component unit within the larger events of leaving for (and returning from) a vacation there. The two transitional events, the end points of the event of taking a vacation, fit the selectional restrictions imposed by “before” (“after”), and so constitute a normal, uncreative, realization.

Referentially creative texts are not compositional. The New Englander’s world knowledge that one takes winter vacations in Florida notwithstanding, a hearer cannot reconstruct the intended referent of “Florida” in (1) in isolation from the immediately prior linguistic context. The meaning of the whole involves the *interaction* of the

meaning of the parts and not just their sum, a phenomenon called *co-compositionality*, (Pustejovsky (1994).

We see co-compositionality as part of the speaker's processing as well. At any given moment we believe there will typically be several units of content vying for expression which, while distinct, are interconnected and even redundant. At the point of the referentially creative choice in (1) these will include at least the event (the vacation) and its location. The choice between the two is not made in isolation based solely on the information they contain, but is strongly influenced by the choice made just before, namely to use the semi-fixed phrase "... before or after". The speaker's commitment to that phrasing imposes a constraint on what follows that interacts with the potential content and influences what happens.

As the speaker looks for the next unit of content to select, her goal is simply to communicate the two moments, not to deploy any particular unit to which she was committed long before reaching that point in the utterance. Any unit that will convey the information directly or indirectly can be used. Such a lack of prior commitment is the whole point of an incremental model of generation, which we are now exploring at the level of content selection rather than just the customary level of realization and surface form (see Kempen & Hoenkamp 1987, DeSmedt 1990, McDonald 1980).

Part of the speaker's knowledge, given the particular circumstances in which this actually attested utterance was made, is that just saying "Florida" will communicate the event of the vacation and through it the two time points, thereby achieving her communicative goal. We hold that this knowledge is mediated by the lexical entry for Florida, whose content and form follow a theory of lexical semantics known as the Generative Lexicon ("GL"), the subject of Section Five. Before going into that, however, we have to establish the context in which this lexical knowledge is deployed.

3. Situational Cohesion

Because a referentially creative expression seems to violate the selectional restrictions of its context, its use requires the hearer to apply a generative lexical operation called 'type coercion' in GL theory (see §5). The hearer must first apply knowledge of the lexical entry for the apparently inappropriate referent plus an appreciation of the linguistic context and the situation in order to decode the expression and 'coerce' the semantic type of the given expression to the type that the context requires, and then use that new type information to reconstruct the intended referent.

This effort by the hearer is surely not without some cost, which prompts us to ask the question of why a speaker should make the listener go to the extra work that the 'decoding' of type coercion implies. Any cost that the speaker imposes on the hearer must presumably be offset

by some corresponding benefit or else people would stop doing it.

While we do not yet have an understanding of what this benefit might be, we can at this point narrow its character by identifying it with another set of phenomena, which, while also having no real explanation as yet in terms of their utility (but see Granville 1990 or Valduví 1992), seem to us to be phenomena of the same kind. These are the textual resources of the language known as 'cohesion', as that term is used by Halliday and Hasan (1976).

Cohesion, to Halliday and Hasan, is the defining property of a text qua text—what it is that makes a text a semantic unit rather than a jumble of unconnected phrases. Sophistication in the use of cohesive linguistic devices is what distinguishes a fluent computer speaker from an awkward one, e.g. the ability to correctly use pronouns, ellipsis, substitutions of generals for specifics, elision of information because it is redundant in content, etc. All of these linguistic resources are typified by an omission or substitution that a listener can only resolve by consulting earlier parts of the text. What we are proposing here as the benefit of referential creativity is that it involves an omission that is only resolved by consulting the *situation* in which the utterance occurred, hence our name for the phenomenon: "situational cohesion".

Selecting a unit of content that is referentially creative—one that cannot be interpreted in its textual context directly but requires the listener to make a connection from it to some larger event or relationship that she knows to be salient in both her own and the speaker's model of the situation—creates a tie from the text to the unspoken but mutually appreciated situation in which the utterance is taking place. In just the same way that a normal cohesive device creates a link to an earlier part of the text and thereby reinforces the fact that the two points are part of the same semantic unit and not disconnected ideas, we can see referential creativity—a device that creates situational cohesion—as knitting the act of making the utterance that much closer to the situation in which it occurs and thereby reinforcing the link between the interlocutors by explicitly tying them to the same situation as a precondition for the speaker being understood.

The phenomenon of situational cohesion is one more piece of evidence of the central place of the situation in the speaker's processes, and with that the need on the part of the generation research community to pay increased attention to it and how it may be formally characterized.³ As a possible step in that direction, we will lay out our, so-

³ Which is by no means to say that the situation has been ignored in past work, as the efforts in functional theories of linguistics and some lines of work in philosophy of language will attest; see, e.g., Halliday 1978, Bateman 1985; or Barwise & Perry 1983, Devlin 1991., but only that it needs a strong computational treatment that we have yet to see evidence of.

far only partially implemented, conception of how generation is {situated} in the speaker's situation, and from that move on to describing elements of our generative lexicon and its use in a rational reconstruction of an attested referentially creative utterance.

4. The processing framework

All actions by a human agent take place in the context of the agent's mental model of the situation. We assume by hypothesis that this model is highly structured, compositional, and organized according to a vast number of reified 'situation types' (Barwise & Perry 1983, Halliday 1978 pg. 29) that provide the prototypical assumptions and categorizations by which a person makes sense of the situation she perceives.

In McDonald, Meteer & Pustejovsky 1987, we proposed a general model of the generation process as a sequence of incremental mappings to successive, progressively more linguistic and more concrete representations of the utterance. These begin with the identification by the speaker of what we call "the Relevant Portion of the Situation", or "RPS".⁴ This is a subset of the total situation, that picks out just those elements to which the speaker is attending given her intentions as she commences to speak.

Much of the organization of the RPS—the roles and relative salience assigned to the entities in the external context and to the speaker's mental entities—derives from the situation-type(s) governing the speaker's actions. It is the situation type, as an orchestrator of the speaker's intentions and attitudes, that we see as ultimately in control of the generation process. The situation type may be very simple, as when answering the phone or saying hello to a colleague, or very complex, as when being interviewed for a job or going out on a third date.

We see the RPS as comprised of a structured heap of independent units of information of various sizes whose epistemological character follows the lines common to today's KR systems. The amount of information in a unit, particularly for the minimal units from which complex ones are composed, is a matter of conjecture. The way we approach this question is to work backwards from the phrasings and variations that we observe in language and to design hypothetical units with enough articulation to produce those texts while trying to work out an aesthetic by which to evaluate alternative design decisions.

We further hypothesize that the RPS, being language oriented, is restricted in what it can contain.⁵ Specifically, a

⁴ The idea of giving special status to a 'relevant' portion of the situation, often under the construal of a particular situation type, is also present in the work of Halliday and no doubt others. We see our contribution as trying to give it computational consequences by tying it to a means of ensuring expressibility.

⁵ There are plainly things one can perceive, categorize, and think about that one cannot talk about in language. Imagine trying to communicate what a headache really feels like or a sunset really looks like. One can get another person to think along similar lines or can, with some deliberate effort, work out a nearly private sublanguage to provide lexicalizations for perceptions, as when at a wine tasting, but there are always limits.

mental unit is only allowed in the RPS if it is *expressible* in the sense of Meteer 1992. We identify this property with whether or not the unit has a lexical entry. The lexical entry dictates the possibilities for the unit's realization, and governs its mappings.

We view the generation process as fundamentally one of traversing of the units in the RPS under the control of a linguistic schema associated with the situation type. The network of interconnections among the units defines the paths available to the generator and the basis of judgments about the content that should be used. As a unit is reached and selected it is mapped to the first linguistic level of representation, Text Structure (Meteer 1990, 1992).

The mapping to Text Structure involves three simultaneous and inter-constrained actions: the choice of the specific unit (possibly a sub-unit of the unit reached by the traversal), the choice or delimitation of its linguistic realization, and the choice of the position where it is to appear in the Text Structure. This simultaneous constraint plays an important role in our framework as we will see in §8, as it is the locus of the operations that lead to referential creativity.

Text Structure is the level where the selected content is first rendered into an abstract linguistic form that dictates its approximate sequential order and hierarchical structure. Working off the Text Structure, later processes incrementally construct a linguistic specification that is then mapped to the surface structure level of representation where the grammatical form of the utterance is established. The surface structure in turn is incrementally mapped to a stream of words for presentation as written text. These later procedures have been thoroughly described in earlier papers (Meteer 1992, Meteer et al. 1987).

5. The Generative Lexicon

Our account of referential creativity builds on observations made within the framework of Generative Lexicon theory ("GL"). This framework defines a method of lexical decomposition where the semantic information in the lexical representation is highly structured. The meaning of a lexical item is taken as a relation (or a set of relations) between all its components, with the important consequence that the lexical entry will include all the parameters that give rise to what appear to be distinct (but related) senses of the same lexical item when projected to surface structure:

(8) a. *John tripped and broke the bottle.* (container of liquid)

b. *Let's drink this bottle, it should be good.* (liquid in container)

Besides the advantage of reducing the size of the lexicon, a generative theory of lexical semantics captures the underlying relatedness among different surface realizations by means of generative mechanisms operating over the lexical structures. One such mechanism is *type coercion* which we take to be crucial to model our notion

of referential creativity. Informally, type coercion is a mechanism that allows the expression of a lexical item of a different type than the one expected in a particular context, only if that lexical item has available in its representation a denotation of the expected type:

(9) a. *Bob enjoys reading detective novels.*

b. *Bob enjoys detective novels.*

The generalization provided in Pustejovsky (1991) is that those predicates that like "enjoy" have multiple subcategorization frames on the surface select for only one deep semantic type, thus in (9a) the 'deep' type selected by the verb is an event, as is that in (9b) where the complement does not directly denote an event.

The theory of type coercion, without making any commitment to what actual event will result from the coercion, states that the object in complement position has to have an event denotation in its lexical representation in order for the sentence to be acceptable. Thus the default interpretation of (9b) is that of 'reading detective novels'.⁶

This view reduces the semantic load of the predicate by distributing it across the surrounding lexical items whose type specifications reduce the set of possible interpretations—a phenomenon known in GL theory as co-composition.

The decomposition of the lexical meaning is given in terms of the *qualia structure* which organizes the distinct facets of the meaning of a lexical item based on the four Aristotelian 'modes of explanation': the AGENTIVE role picks out that which event brings about the lexical item; the CONSTITUTIVE role defines the relation between it and its parts; the FORMAL role distinguishes it in some larger domain; and the TELIC role gives its purpose and function.

These relationships not only account for regularities in a word's behavior, but they reflect the way in which we sort out what we know about the meaning of a word, such that the way we will use it in context will make direct reference to the distinct but tightly related aspects of such knowledge.

A system that, like GL, captures the expressiveness of language without resorting to multiple listings of the same lexical entry has the capacity to account for phenomena that, like referential creativity, can be best treated generatively. In particular, while keeping the number of senses fairly constant, generative mechanisms over highly organized lexical structures allow for an unbounded space of possible interpretations which reflect the unbounded set of situation types and their compositional properties.

⁶ This is what is taken to be the default reading. Consider instead the same sentence where Bob is a writer of detective novels, where the interpretation "writing" because of the sortal specification of the subject in determining the type of the object.

6. A Rational Reconstruction

Consider example 10 below, which was spoken by one of the authors while shoveling the snow off his sidewalk.

(11) “*When we get it down like that it will stay clean for a couple of inches.*”

The motivation behind (10) was to admonish the speaker’s snow-shoveling partner to continue to clean the sidewalk thoroughly. There had been a problem with ice forming under footprints when there was a lot of snow, and this utterance reflects the observation that if nearly all the snow was removed (“*when we get it down like that*”), that when the next snowstorm came at least two inches of snow would have to accumulate before the icing problem would arise again (“*it [the sidewalk] will stay clean*”).⁷

What is referentially creative about (10) is that it uses a measurement of depth—the amount of snow that would have to accumulate in the course of the next storm—in a phrasal context whose selectional restrictions require a time adverbial, i.e. a measurement of time: the duration during which the sidewalk will remain in the state of ‘staying clean’.

A ‘normal’ realization, one where the type of the adverbial fits the type expected by the predicate, might be:

(11) “... *it will stay clean for two hours.*”

though it would not be communicating the same information since the crucial datum is the depth of accumulated snow, not the amount of time it will take for that snow to fall.

Since the accumulation is a function of the rate at which the snow falls and the duration of the storm, the measurement of snow depth does in fact implicitly denote an amount of time, namely the length of time in which it takes that amount of snow to fall, and in so doing it makes that time information available through a process of type coercion.

As the listener hears (10) incrementally from left to right, she will have set up an expectation at the point when she hears “*stay clean for*” that she is about to encounter a time phrase and that they should construct an interpretation of it as a duration.⁸ Instead, she encounters a measurement phrase. She must now coerce that phrase into a duration, which is possible by considering the information available via the lexical semantics of the phrase—information in its lexical entry—and an appreciation of the elements in the

⁷ This utterance is obviously muddled and obscure when it comes to the fit between what the speaker was trying to communicate and what he actually said, but this does not have any impact on its value as an example of referential creativity. Such are the complications of dealing with real data.

⁸ That listeners can construct an interpretation that quickly is shown from the performance of ‘close shadowers’ in experiments done by Marslen-Wilson and Tylor (1980).

situation that holds at that time. We expect that this is precisely the same information, though used in a different manner, as the speaker employed to sanction the use of the phrase in that position. Let us look at what these elements probably were.

First the RPS. The overarching situation type is that of ‘shoveling out after a snowstorm’, something that had become all too common during the winter that this paper is being written. Within this, we have the speaker-internal situation of having made an observation that a certain benefit will accrue if some activity is continued. This gives us the following likely elements for the relevant portion of the speaker’s situation:

- the two interlocutors cum snow-shovelers,
- the activity of shoveling snow,
- the likely next snow storm (a few hours away),
- the sidewalk,
- the ongoing problem of ice forming under people’s footprints (‘stay clear’), and
- the thoroughness with which the fellow snow-shoveler was cleaning the sidewalk (‘get it down that far’),
- the observation that, given that thoroughness, a surprisingly large amount of snow would have to accumulate before ice was likely to become a problem again (‘about two inches’).

Next come rhetorical issues. How are the units of content in an utterance to be organized—what is their order, linguistic roles, prosodic tune—such that the intended message will be communicated. One basic rule, particularly for a topic-initial utterance like this one, is progression from ‘given’ to ‘new’: The utterance must start with an element that is mutually known given the situation, and should then progress to new elements that the hearer is probably unaware of. That rule provides a selection criterion to pick out one of the schemas available for communicating cause and effect (here ‘achievement’ and ‘benefit’) by fixing the order in which the two elements must occur.

This gives us a “When X, Y” schema, which we take to be an active player, setting up the ordering for how the RPS is to be traversed. With this process set in motion, units begin to be mapped to Text Structure and from there to surface structure (where their grammatical properties are manifest) and the first words of the utterance are spoken. We see no reason why all of the different stages of this pipeline wouldn’t be active simultaneously, and we further assume that all mappings are indelible: once made they cannot be retracted—all operations build monotonically on what has gone before.

7. Linguistic reasoning during content selection

We now pick up this process at the point where nearly all of the second clause has been introduced into text

structure. We are at the moment when the complement of its adjunct ‘for’ phrase is to be added, i.e.

“... *it will stay clear for* ___”

The eventuality denoted by this clause (the unit that is the clause’s source) is situated at the time of the next snow storm. We assume that that time was the direct source of the “*will*”.

Semantically the matrix is a persistent state and the adjunct is describing when it will cease to hold. They both derived from a complex unit in the RPS, which would have been queued up as a whole and its parts then traversed incrementally. For this unit the rhetorical considerations are probably overshadowed by the grammatical constraints on the order of the subunits. However, since it is the large depth of the to-be-accumulated snow that is what is significant (indeed, it is probably what prompted the whole utterance), there is a strong reason to make sure that the unit representing the depth of snow is positioned at the end of the utterance where it can receive tonic stress in the default prosodic tune.

The choice of “*for*” as the preposition reflects the particular way that the state has been conceptualized. Alternative conceptualization would have yielded different prepositions: ‘state persists *while* state/process’, ‘state persist(s) *until* event’. Here it is viewed as lasting throughout (“*for*”) some interval of time. We take it that this conceptualization was not a generation ‘choice’, per se, but just reflected the type of the source unit within the RPS. It is something that the generator must accommodate to in what follows, not something it was free to phrase in a different way.

We think that at the level of content selection the generation process is so incremental and the pipeline below it so fast (the ‘planning units’ so small) that the speaker will have committed to a formulation in terms of the matrix +*for* adjunct before reaching the next unit in its path, the snow. This means that the realization of this next unit must be accommodated to a site in Text Structure that is constrained to take only an interval of time.

On its face this is not possible, and the unit would ipso facto be ‘inexpressible’ and its realization blocked. However given a generative lexicon, the generator has at its disposal the linguistic resource of type coercion, and this makes new options available by in effect expanding on what constitutes the type of a unit.

The information carried by the ‘a couple of inches of snow’ unit goes considerably beyond what appeared at the surface. In words it is something like

‘the depth (about two inches) to which the snow from the next storm will have to accumulate before ice will form ...’

Obviously there is no simple lexical entry for a unit like this. The unit is a complex, highly interrelated object that makes reference to several other units (the impending storm, the sidewalk, etc.) that are all necessary concomitants of each other—the one will not exist without

the others. Consequently its lexical entry must have an equally elaborated type with equally interrelated elements. This, however, is actually common for words in a natural language.

To have a handle on it, we think of this entry as the lexical entry for “*snow*”, but that is only because that word appears so often in the realizations that pick out its various aspects, and because it has a more specific content than the other linked words such as “*fall*” or “*accumulate*”. Consider these excerpts from news articles:

(12) a. *The snow further snarled the plans of travelers.*
(process)

b. *Boston has received 24.7 inches of snow this month.* (state, result)

Following the symbolic presentation of Pustejovsky and Busa 1993, the relevant part of the entry for snow is the following. A full lexical entry also gives the temporal and aspectual relationships of the events, the linguistic properties of the words coordinated by this entry such as “*depth*”, and other qualia roles; here we just showed the rules that pick out units in the current RPS.

snow(e^P, e^S, x, y)

FORMAL: depth ($e^S, x, \text{location}$)

CONSTITUTIVE: flakes (x, y)

AGENTIVE: snow storm ($e^P, y, \text{location}$)

According to this entry, snow is polymorphic. It is brought about (the ‘agentive’) by a process, e^P , namely a snow storm at some location ‘y’. It is also the state, , of an accumulation of some depth of snow, ‘x’, at y.

As seen by the constraint mechanisms operating at Text Structure, the only relevant aspect of a unit in the RPS is its expressive semantic types, something we have identified with the unit’s lexical entry. Here the entry makes reference relations of three types: a ‘depth’ (which is how one conceptualizes an amount of snow on the ground before it starts to melt), a mass (the snowflakes), and a process (the storm). There is also reference to the types of the parameters in those relations. Of crucial interest here is of course the standard time parameter in the agentive, the duration of the storm. The Text Structure is imposing the constraint that the unit convey information of type duration; this parameter provides it.

9. Summary

We have made a rational reconstruction of an actually attested utterance that exhibited what appears to us to have been a completely new pattern of semantic types, making it an instance of what we have called referential creativity. We have used this reconstruction to illustrate our design for an incremental content selection process. This process selects conceptual entities or ‘units of information’ from a region of the speaker’s situation that is established at the moment she decides to say something, what we call the Relevant

Portion of the Situation. Simultaneous with a unit's selection is its assignment to a location in the text plan as it exists up to that point and the determination of the linguistic resource to which it is to be mapped. The text plan is Meteer's Text Structure level of representation—the level at which units must be shown to be expressible—and consequently there is a constraint on the selection process that the mapping chosen must be consistent with the syntactico-semantic constraints of the target location.

Since judging the consequences of a choice of realization could involve a great deal of linguistic reasoning if the vocabulary in which the choice was deliberated were too detailed or too dependent on facts about surface syntax, we base the decisions on abstract representations of the facts in a unit's lexical entry, the Lexical-Conceptual Paradigms of Generative Lexicon theory, instead of on the entries themselves. In addition, we draw on GL's qualia structure representation of a word's meaning to indicate when a phrase will communicate more information than it appears to do looking just at its surface content. This permits the selection of units that would otherwise be inconsistent with the Text Structure location to which they are assigned, resulting in the need for the listener to do what GL calls 'type coercion' to recover the meaning that the phrase is intended to indicate. We noted that in cases like this, the omission of explicit information creates a cohesive link to the situation in which the utterance occurs, thus reinforcing the integrity of the text as a semantic unit.

Our work suffers from the lack of a real, implemented machine speaker that could be in real situations, have real intentions that lead to the identification of an RPS, and thereby provide a real check on whether the kinds of linkages of units that we have posited are coherent in a system that has more to do than just produce texts. We have accepted this limitation because it is the only way that we could look at language as it is really used by people, and with that be able to investigate a process whose complexity, we believe, is commensurate with what is happening in people's minds.

Indeed, the only realistic reason to propose an incremental processing architecture like this one is to formulate hypotheses about the nature of language processing by people, not machines. If the goal were simply to have computers produce the best possible texts to fit the content, then we would surely use multi-pass, revision-driven architectures such as the one developed by Gabriel (1981, 1988), which, after all, passed the ultimate test of a machine generation system by writing three paragraphs that were folded into the exposition in the journal article about the system and were completely undetectable as machine-text.

As a hypothesis about people, an architecture should offer some explanation for phenomena that we can observe, and should make predictions about aspects that we had not thought to look for before. We have seen the first signs of such explanations when we looked closely at the contexts that led to referential creativity. For example, had the

utterance about the amount of snow been based on a different phrasing for the temporal adjunct, then a normal realization would have been possible:

(13) "...it will stay clean until there are a couple inches of snow on it."

Similarly, if example 1 had not involved the fixed phrase "before or after" then a normal realization of the whole event would have been available as discussed earlier.

We expect that in part it is the difficulties that such phrasal choices create for the expressibility of the normal following content that contribute to the recourse to referentially creative utterances, where the difficulties no doubt combine with the centrality and salience of the type-inappropriate unit to make the unit more available in the RPS than it might otherwise be.

This suggests that referential creativity is an escape route from a context in Text Structure that would otherwise block any content from being selected. As a hypothesis about human processing, this may be testable through a close examination of a class of speech errors, a known to be fruitful source of insight into human processing (Levelt 1989, Garrett 1980). This class is the pattern of 'block and restart' errors that are so common, even ubiquitous, in casual unpracticed speech that we accept them just part of the process of talking and don't even think of them as "errors". Here, for example, is such a block and restart transcribed off the radio during an interview. The person is speaking about the flooding in Mississippi and how people are dealing with it:

(14) "*Some of these emergency workers have not had a // have not been to bed for three days.*"

If we allow ourselves to speculate about what this person might have said had he continued at the point of the "//" rather than restarting, a likely possibility is that he would have said that the workers had not had "a nap", or "a chance to sleep". This possible content is high frequency locution, just like the phrase "before or after", but it probably does not fit the facts of the situation. That is, we speculate that the intended completion was "have not had a nap in three days".

We expect that it may be very often be the case that a block and restart occurs because the content unit just before the block, the last one said or partially said (given that the process overall is a pipeline and the speaker can be expected to appreciate a failure of expressibility at an early stage), a high frequency standard phrasing, creates a context in which the next unit to be said cannot be expressed. Then, if there is no high salience, situationally cohesive unit available in the RPS that might implicitly convey the information and allow the the speaker to make a referentially creative choice, then the speaker will have no alternative other than aborting their plan and starting anew. If such an account for block and restart errors proves fruitful it will provide empirical support for the design of the content selection process that we have described in this paper.

10. References

- Bateman, John (1985) *Utterances in Context*, Ph.D. Dissertation, Univ. of Edinburgh.
- Barwise, Jon & John Perry (1983) *Situations and Attitudes*, MIT Press.
- Busa, Federica & James Pustejovsky (1993) "Unaccusativity and Event Composition", proc. Cortona Tense and Aspect meeting, Oct. 1993.
- Chomsky, Noam (1966) *Cartesian Linguistics*, Harper & Row.
- Cruse, D.A. (1986) *Lexical Semantics*, Cambridge Univ. Press
- Devlin, Keith (1991) *Logic and information*, Cambridge Univ. Press.
- Gabriel, Richard P. (1981) *An Organization for Programs in Fluid Domains*, Ph.D. Dissertation, Stanford University.
- Gabriel, Richard P. (1988) "Deliberate Writing", in McDonald & Bolç (eds.) *Natural Language Generation Systems*, Springer-Verlag, Berlin.
- Garrett, Merrill F. (1980) "Levels of processing in sentence production", in Butterword (ed.) *Language production: Vol. 1. Speech and Talk*, Academic Press, London.
- Granville, Robert A. (1990) "The Role of Underlying Structure in Text Generation", in the proceedings of the Fith International Workshop on Natural Language Generation, June 3–6, 1990, pp.105-111.
- Halliday, M.A.K. (1978) *Language as social semiotic*, Arnold.
- Halliday, M.A.K. & Ruqaiya Hasan (1976) *Cohesion in English* Longman, London.
- Kempen, Gerard & Eduard Hoenkamp (1987) "An incremental procedural grammar for sentence formulation", *Cognitive Science* 11, pp.201-258.
- Lehrer & Kittay (eds.) (1992) *Frames, Fields, and Contrasts*, Lawrence Erlbaum.
- Levelt, Willem J.M. (1989) *Speaking*, MIT Press.
- Marslen-Wilson, W.D. & L.K. Tyler (1980) "The temporal structure of spoken language understanding", *Cognition* 8: pp.1-71.
- McDonald, David D. (1980) "A Linear-time Model of Language Production: some psychological implications", in the proceedings of the 18th Annual Meeting of the Association for Computational Linguistics, June 19-22, Philadelphia, pp.55-56
- McDonald, David D., Marie Meteer, & James Pustejovsky (1987) "Factors contributing to efficiency in natural language generation", in Kempen (ed.) *Natural Language Generation*, Nijhoff, Dordrecht, The Netherlands, pp. 159-182.
- Meteer, Marie W. (1990) "Abstract Linguistic Resources for Text Planning", in the proceedings of the Fith International Workshop on Natural Language Generation, June 3–6, 1990, pp.62-69.
- Meteer, Marie W. (1992) *Expressibility and the Problem of Efficient Text Planning*, Pinter, London.
- Meteer, Marie W., David McDonald, Scott Anderson, David Forster, Linda Gay, Alison Huettner & Penelope Sibun, (1987) *Mumble-86: Design and Implementation*, TR #87-87 Dept. Computer & Information Science, UMass., 174 pgs.
- Pustejovsky, James P. (1991) "The Generative Lexicon", *Computational Linguistics* 17, pp.409-441.
- Pustejovsky, James P. (1994) *The Generative Lexicon: A Theory of Computational Lexical Semantics*, MIT Press, to appear.
- Reiter, Ehud (1990) "A New Model for Lexical Choice for Open-Class Words", in the proceedings of the Fith International Workshop on Natural Language Generation, June 3–6, 1990, pp.23-30.
- De Smedt, Koenraad J.M.J. (1990) *Incremental Sentence Generation: a computer model of grammatical encoding*, TR# 90-01, Nijmegen Institute for Cognition Research and Information Technology, Nijmegen, The Netherlands.
- Vallduví, Enric (1992) *The Informational Component*, Garland, New York.