Extending ECG to Communities, Mental Spaces, and Maps

8/12/2009 5:31 PM

Contents

1. Introduction

<footnote, this part of a volume see wiki or search on ntl ecg icsi>

NTL for decades ….

We should separate the tenets of the NTL project into three categories; *essential, scientific* and *methodological*. The most fundamental idea behind NTL is the direct neural realization and *continuity* of language and thought. We take this to be established - it alone does not suggest particular methodology, but it does entail a commitment to parallel *spreading activation* as opposed to serial processing. It is also simply true that a language community necessarily shares some skeletal beliefs, including shared grammar. Equally fundamental, although not as unequivocally established, is the notion of *simulation* semantics – language understanding involves much of the circuitry involved in perception, motor control, emotion, etc. A further core assumption is that language understanding involves a *best-fit* analysis phase that activates the knowledge and bindings used in simulation. Symmetrically, production is a best-fit process mapping communicative intent to surface structure, including gesture, etc. Language learning and change involve further best-fit modifications of grammar and other knowledge to changing needs.

Given these paradigmatic assumptions and a commitment to unification with the scientific enterprise, there are a wide variety of detailed questions that are addressable by existing experimental techniques and some of these are being explored. We are part of the multidisciplinary effort to understand the biological, psychological, and social reality of language and thought independent of particular theories. Of course, the tenets of NTL and related theories suggest experiments that might not emerge from other perspectives and, crucially, detailed theories and models greatly facilitate focused experimentation.

But any such concrete work requires much more precise specification and that, for NTL, is largely described in terms of ECG, Embodied Construction Grammar and the developing theory of simulation semantics based on CPRM, Coordinated Probabilistic Relation Models(ref). Neither the formalism itself, an ECG grammar for some language, analysis of a particular case, nor a demonstration system embodies any scientific claim. In contrast with the dominant tradition in linguistics, *the notation is not the theory*. We do try to constrain specific grammars, models, etc. by all the best established scientific results, but ECG notation is just a way of describing postulated neural structures. Some central tenets of NTL theory, particularly spreading activation cannot be expressed in ECG or any descriptive formalism.

More specifically, ECG is a *methodological* choice on how to explore the basic NTL principles in a way that is sufficiently expressive and tractable. Similarly, an ECG treatment of some language phenomenon should be seen as one way of explicating the hypothesized regularities. There will always be alternative ways of describing the same phenomenon and the question of what is realized in the brain is an (often approachable) scientific question, as Lakoff discusses in Section 5.

Embodied Construction Grammar is a formalism for specifying grammars that is being designed to simultaneously serve the following six functions.

a. Specify the shared grammar and conceptual skeleton of a belief community

b. A technical tool for linguistic analysis.

c. A computer specification for implementation of linguistic theory.

d. A front end system for applied language understanding tasks.

e. A representation for models and theories of language acquisition.

f. A high level functional description for biological and behavioral experiments.

The second item, linguistic analysis is covered in the chapter by Dodge and Bryant and implementation is described in Bryant's chapter and in considerably more detail in his thesis (ref). The use of ECG in studies of language acquisition is considered in some detail in the accompanying chapters by Chang and by Mok. The most comprehensive discussion of application is in Sinha's 2008 thesis (ref). The last of these roles will be discussed in Section 5 by George Lakoff. The Bryant chapter has some behavioral results and the ECG wiki <> has pointers to current experimental efforts. Some of the motivation for the developments described in this chapter comes from a more careful analysis of what is required for ECG to fulfill these multiple roles.

Most of the effort and essentially all of the implementation through 2008 has been focused on a collection of basic issues in language and has backgrounded some central issues that we know to be fundamental. This chapter will describe the current state of design and implementation of the most important of these considerations. First we will extend the domain of concern from individual communication and understanding to discourse involving *communities*. This has implications not only for intellectual and practical applications, but also for the role of grammar in general.

From its beginnings over two decades ago, the NTL effort has recognized the importance of multiple perspectives (Mental Spaces, ref) and mappings like metaphor (ref) and correspondence links between mental spaces (ref). This chapter presents the current state of design of the two additional ECG primitives that seem to be required to cover these phenomena: *Situations* and *Maps*. It also introduces some additions to the paradigm that make it more suitable for modeling communication among groups in addition to just individuals. But first we restate some of the basic motivations for ECG which, we have learned, have not been sufficiently clear.

One ambitious claim of the NTL/ECG effort is a redefinition of "grammar". The claim is that previous definitions were either too narrow or too broad. Of course it is not the word that matters but the scientific concept. It is clearly too restrictive to study only a notion of linguistic form with no attached meaning. It is also inadequate to work on only an abstract notion of meaning (e.g. logic) that has no grounding in human experience. More generally, the goal of linguistics should be to explain language acquisition, use, and change and how these integrate with science in general. This goal has been construed as incompatible with formalization, leading to an almost complete split between formal and cognitive linguistics.

Notwithstanding decades of great insights in cognitive linguistics and related fields, there have been no previous cognitive grammar formalizations or analysis programs. There are a number of non-technical reasons for this, but there are also technical issues. Cognitive linguistics is based on the unity of language and thought and it is certainly true that there is no separation of language and thought in the brain. So, how could we specify a scientifically accurate and also tractable notion of grammar? Is the entire conceptual system and embodied experience required for grammar?

It might be useful to consider the analogy of laws and courts. There are several sets of laws (local, state, federal, etc.) that are written down and serve as the basis for court decisions. But any court action involves an inseparable interaction between the (possibly conflicting) applicable laws and the facts and personalities involved in the particular case. The analogy could be extended to include the role of salient examples (case law), diachronic change, etc. There is nothing exotic about having both cultural conventions and individual behavior that is only partly constrained by these conventions. The relation between grammatical convention and individual behavior shares motivation with the old linguistic competence/performance distinction, but leads to very different detailed theory.

ECG includes an explicit attempt to provide tools for describing the cultural conventions of grammar and skeletal beliefs for a language *community*. The full conceptual system is different for each person. Despite this, people are able to communicate in some language because there is a shared conceptual belief *skeleton* that is largely consistent within that language community. This skeleton is richly structured and complex, but is still a small subset of the all the knowledge and experience in your head.

The shared skeleton necessarily includes more than conventions of grammar – there must also be enough of a shared conceptual system to support communication. There are a lot of (embodied) universal concepts shared by all people; these are viewed as the core of the structure of all belief communities. More technically, ECG postulates a lattice of BDA (Belief, Desire, Action) communities and models communication among them. For example, most English speakers can communicate fairly well using the general of idea of game, but we each have our own BDA structure for games.

A person is normally a member of several BDA communities, for example at home, at work, by nationality, etc. Successful communication requires that the speaker and the audience share an adequate BDA community. Grammar is an important part of any BDA community and the one of most direct concern here. But an ECG grammar explicitly maps from linguistic form to deep conceptual structure so skeletal shared beliefs are part of grammar. These include both (postulated) conceptual universals such as image schemas (ref) and cultural frames (FN ref) such as games.

Suppose we say that the job of grammatical analysis is to extract the communally shared skeletal structure from an utterance in context. This is now a well formed and tractable problem, subject to formalization, implementation, and experiment. The basic conceptual skeleton information that is conveyed by language is represented formally by ECG *schemas* and *bindings.* ECG schemas include the standard image schemas and cultural frames, but also linguistic structures like causal links, discourse relations, etc. The subcase lattice over schemas model the skeletal organization of conceptual knowledge. Schemas have many uses in ECG; importantly they constitute the meaning poles of ECG *constructions*. ECG constructions are form-meaning pairs in which the meaning is described in terms of the conceptual skeleton. This is the basis for the E~ Embodied in ECG. As we will see in Sections 3 and 4 of this chapter, the formal definition of ECG schemas is also extended to the two new primitives: *Situations* and *Maps*.

Another key element in the approach involves best-fit constructional analysis. As mentioned, ECG *constructions* are form-meaning pairs, where the meaning pole is an ECG skeletal schema, with additional bindings and constraints. The implemented analyzer (Bryant, this volume) while very useful, is not part of the theory. But the notion of integrated *best-fit analysis* is central. Grammatical analysis inherently involves determining the collection of construction instances (aka constructs) and bindings that best matches an utterance in context. For ECG, the best match must include factors for the semantic and contextual fit as well as the standard constituent structure form fit. As with schemas, there is a subcase lattice over ECG constructions, which helps in organizing the compositional structure of a grammar.

In addition, the meaning pole of an ECG construction specifies what modifications to the currently evolving SemSpec are sanctioned if that construction is included in the best-fit analysis. These modifications can include evoking additional skeletal schemas and constraining and linking schema roles.

So, this redraws the boundaries between grammatical analysis and full comprehension in a significantly different manner than any previous theory of language understanding. It makes the grammatical analysis problem simpler and more realistic in several ways. Using best-fit, there is no reason to require or expect a categorical grammar and the notion of well-formedness becomes contextual. The role of grammar and parsing is reduced to the natural one of producing a SemSpec that captures the schemas and links that are specified by the utterance in context. These schemas and bindings are part of the shared conceptual skeleton, but they are only links to the full experience of the understander.

There are additional bindings that are not explicitly specified constructionally, such as anaphors, definites, and frame inferences. The process of trying to fit implicit bindings is called *resolution* in ECG and is discussed in the Mok chapter. The goal of ECG and of the analysis process is to specify and match all of the schemas and bindings that are available from the utterance in context. Simulation uses the resulting SemSpec towards the goals of the understanding agent. Production is the converse process – mapping from a SemSpec and communicative goals to surface structure, but we have not yet implemented this.

2. Simulation Semantics and the SemSpec

Another foundation of our approach comes from the NTL formulation of *simulation semantics*. NTL postulates that the full meaning of an utterance is the result of active imaginative simulation. The idea here is that the skeletal information conveyed by language triggers, in the hearer, an active simulation that elaborates the message wrt his full conceptual experience, goals, etc. More technically, we postulate that grammatical analysis produces a *Semantic Specification* (SemSpec), which consists of conceptually linked skeletal schemas. The SemSpec is then (in programs or in people) elaborated by simulation. Since this volume is oriented towards computation, we will look more carefully at the SemSpec and its use in computer systems.

Of the five functions of ECG, the third and fourth - applications and acquisition models, involve significant interpretation of the SemSpec. As of now, the linguistic and cognitive science aspects of NTL have not made extensive use of the SemSpec, except in the metaphor modeling (ref). A critical initial point is that we expect a grammar and the SemSpec resulting from the analysis of an utterance to be independent of the beliefs and goals of the audience, as formalized by its BDA structure. By contrast, the *interpretation* (or simulation) of a SemSpec depends entirely on the audience. This follows from the fact that an ECG grammar, analysis process and resulting SemSpec are intended to model the linguistic and cultural knowledge shared by the speaker and audience.

There have been two fairly ambitious systems built that focus on the interpretation of a SemSpec and it will be useful to review these. The first is the model of children's Mandarin grammar acquisition, discussed in Chapter xx of this book by Eva Mok. Nancy Chang's earlier model (Chapter yy) uses similar ideas but predated the current formalization. As with any research effort, each of these simplifies certain aspects of the problem to concentrate on others. Mok's work uses minimal child language and a conventional AI model of simulation, and focuses on the double loop over language understanding, context updating, and grammar learning.

< more from Eva>

The other major implementation of SemSpec based simulation involves applications to understanding news stories and related material. Many of the original ideas for the SemSpec (as well as the metaphor maps of Section 4) came from earlier work of Narayanan (1999) on simulation based understanding, which pre-dated ECG. In contrast to Mok's work (Ch xx), which completely formalizes a simple child language situation, Sinha studied much more complex language and only implemented certain key modules.

Reasoning about event structure is a fundamental research problem in Artificial Intelligence. Event scenarios and procedures are inherently about change of state. To understand them and answer questions about them requires a means of describing, simulating and analyzing the underlying processes, taking into account preconditions and effects, the resources they produce and consume, and their interactions with each other. We propose a novel comprehensive event schema that covers many of the parameters required and has explicit links to language through FrameNet. Based on the event schema, we have implemented a dynamic model of events capable of simulation and causal inference. We describe the results of applying this event reasoning platform to question answering and system diagnosis, providing responses to questions on justification, temporal projection, ability and 'what-if' hypotheticals, as well as complex problems in diagnosis of systems with incomplete knowledge.

< segue to situations?>

Of course, none of this solves the central questions of language understanding. Linguists still need to determine what schematized conceptual information is conveyed by a specific language and what lexical and structural mechanisms are used for conveying it. What the ECG system provides is a mechanism for writing, testing, and exploiting theories about universals and the structure of particular languages, according to the paradigm outlined above.

These considerations can also guide the choice of what level of detail to include in the SemSpec. It should include ways of expressing any distinction that can be expressed in the surface form. The SemSpec should not require specification of any distinction that can be omitted from the surface form. Some degree of under-specification is implicit in all human communication.

Even for basic literal language there are additional considerations. Morphology provides a rich set of tools for conveying schematic conceptual structure and the ECG/best-fit model has been extended to cover a lot of this. And language is not the only means of communication; mirror neuron studies show that gesture and intonation convey direct emotional messages largely independent of the symbolic structures discussed here. Even within symbolic language, the story presented above is oversimplified – our understanding is not literal and not always against the same background and context.

It currently appears that only two additional primitives are required to extend the ECG design to handle the full range of symbolic language. These have not yet been fully implemented and have received much less effort, but they do seem to fit quite nicely into the paradigm outlined above.

In ECG, Mental Spaces are loci of simulation. Again we need to be careful to distinguish *definitions* of Mental Spaces from *instances* that will play a major role in SemSpecs. As before, there is a simulation process that takes a Semspec, context, and goals and carries out mental actions, updating the agent's situational beliefs.

The individual enactment or simulation of an utterance in context inherently depends on the BDA structure and current goals and situation of the individual. This requires an analysis process that activates the schemas and bindings specified by the utterance. But this simulation is a distinct, deeply individual, process and is modeled in NTL by CPRM. The combined analysis (ECG+Analyzer) and inference (CPRM + simulation) has been formalized as a model of language understanding and is extensible to production and acquisition.

3. Situations and Mental Spaces.

The assumption that a single SemSpec can always characterize an utterance is unrealistic. A great deal of discourse describes events at various times and places, counterfactuals, thoughts of others, etc. These are traditionally studied as "mental space" phenomena (refs). The core idea for treating theses in ECG is basically simple; specify in the SemSpec a separate simulation specification for each situation.

But we also need to model enactment of language (etc.) in real physical spaces. The fundamental difference is that mental simulation is totally internal to an agent while acting in the physical or social world inherently entails input and output consequences that cannot be completely modeled by the agent. For a variety of reasons, ECG treats external interactions as an alternative subcase under the higher type, "Situation" and thus ECG has distinct mental and extra-personal space types.

As a practical matter, some modeling tasks will inherently require the system to maintain a model of physical (or social) space that is distinct from the mental spaces in the minds of any agents involved. For example, take Eva Mok's example of a mother and child in a room (Chapter XX). If something is dropped, it falls to the floor independent of whether or not anyone observed it and this must be simulated in the system. Some dialog only makes sense against a background of (simulated) ongoing events. For the philosophically burdened, this does not imply that the system/model has privileged access to the true nature of the world – if the physical or social assumptions in the world model are inappropriate, it is a bad model.

So, physical (and later social) actions can be captured by their effects on a situation model that is updated according to putative regularities of the external world. If an ECG system were embodied in a robot, the robot would need to have an internal model that was updated by perception. But if we wanted to simulate the robot's behavior before placing it in physical danger, there would still need to be a model of the physical environment distinct from the robot's own beliefs. For people, this corresponds to the difference between an initial guess about the effects of an action and the result of thinking through (simulating) its consequences.

The suggested ECG primitive should encompass both physical and mental environments, and is called a *Situation*. As with the other ECG primitives, there is a lattice of Situation *definitions* and we will use Situation *instances* in building Semspecs. Mental Spaces and (modeled) Physical/Social Situations are subcases. A fragment of the Situation lattice with TemporalSpace highlighted is shown in the left pane of Figure 5.1, which is in the format described in [BG]. We will not discuss it here, but there is also

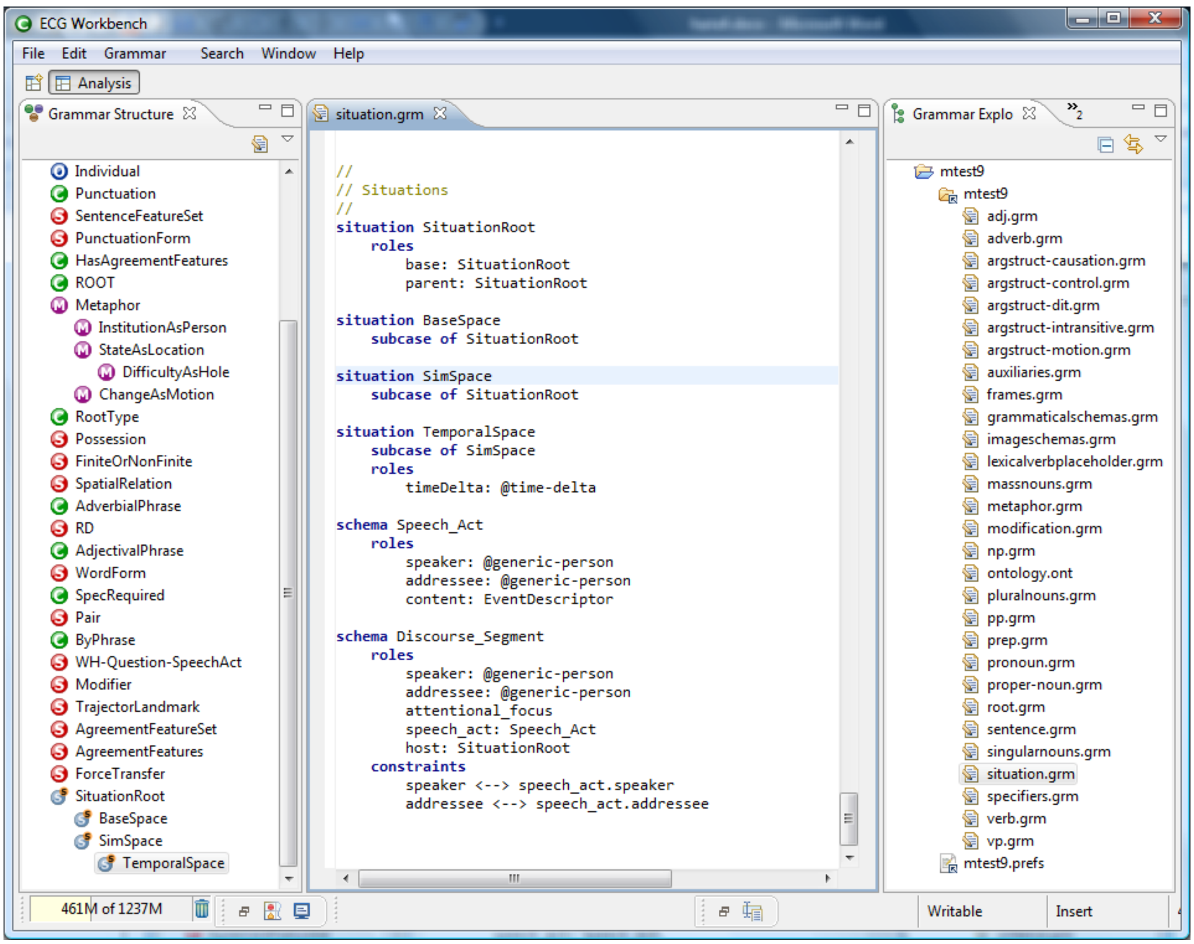

Figure 5.1

the third level involving some evolving model of the current state of either a mental space or a physical/social world model. Until we involve physical robots, any physically situated task will use a model of the physical situation, which will be updated according to some model of how the world works (cf. [Mok]).

The definition for TemporalSpace is shown in the middle panel of Figure 5.1. In so far as possible, the ECG notation for situations follows the conventions of the Mental Space literature (refs). This has never before been formalized and some terms have been used in multiple incompatible ways. We can also see in Figure 5.1 that any situation has links to two predecessor situations: *parent*, the immediate predecessor and *base*, the starting situation for the current episode.

A full treatment of Mental Spaces and other Situations is beyond the scope of this chapter and, in fact, is not feasible in the current implementation. All of the previous ECG examples in this book involve analyzing isolated sentences without explicit consideration of the discourse or situational context. This is traditional in linguistics and the specific examples are sufficiently rich to allow consideration of many interesting questions.

But many of the phenomena addressed by Mental Space theory explicitly involve discourse, alternative possibilities, etc. The ECG Situation and Map primitives provide a mechanism for formalizing much of this material, but there are fundamental limitations on what the existing system can do. We will start with a simplified treatment of the basic temporal example sentence:

"Last year, he slid."

```
schema Speech_Act
  roles
    speaker: @generic-person
    addressee: @generic-person
    content: EventDescriptor

    addressee <--> speech_act.addressee
```

```
schema Discourse_Segment
  roles
    speaker: @generic-person
    addressee: @generic-person
    attentional_focus
    speech_act: Speech_Act
    host: SituationRoot
  constraints
    speaker <--> speech_act.speaker
    addressee <--> speech_act.addressee
```

Figure 5.2

This is an example of a TemporalSpace as shown in Figure 5.1, which in turn is a subcase of SimSpace, a situation that can be simulated.

Our example sentence is also an instance of an explicit speech act. In particular the semantics needs to capture the difference between the time of utterance and the event being described in the main clause. Simple ECG schemas for modeling discourse elements and speech acts are depicted in Figure 5.2. All of the previous illustrations in the book are also examples of discourse segments, but that was not explicitly included.

Similarly, any ECG analysis is with respect to a Base Situation, but this was also not explicitly brought out in the previous examples, because nothing depended upon this. Recall that in ECG, the semantics of any clause is given by an Event Descriptor (ED). For the more general treatment presented here, we should think of the ED as having an additional role: situation, which was not displayed in the earlier examples.

<more>

A basic set of constructions for analyzing sentences like the example are presented in Figure 5.3. One essential new item *focus* has been added to ECG; this corresponds fairly well with the traditional use of the term in the Mental Space literature. Recall that the analysis process builds SemSpec for subsequent simulation. In temporal space examples, there need to be separate simulations for the different times involved – many things will be different. In our example, he is a year older at speech time than when he slipped. The ECG analyzer builds partitioned SemSpecs for differing times (or other situations) - the *focus* variable specifies which mental space is being referenced in each linguistic expression. Specialized constructions are used to recognize and model such focus shifts. The “space builders” of Mental Space theory are prototypical examples of this.

```
schema TemporalModifier
  subcase of Modifier
  roles
    profiledInterval
    timeDelta
```

```
general construction TemporalPhrase
  subcase of ModifierPhrase
  meaning: TemporalModifier
```

```
construction TemporalPhrase3
  subcase of TemporalPhrase
  constructional
    constituents
      ta: TemporalAdj
      np: SingularNoun
  form
    constraints
      ta.f meets np.f
  meaning: TemporalModifier
    constraints
      self.m ↔ ta.m
```

```
general construction TemporalAdj
  subcase of TemporalPhrase, Word
```

```
construction Last
  subcase of TemporalAdj
  form
    constraints
      self.f.orth ⟵ "last"
  meaning
    constraints
      self.m.profiledInterval ⟵ "past"
      self.m.timeDelta ⟵ "1"
```

Figure 5.3

The representation of and updating of beliefs and/or the world model is outside the scope of ECG. The system described in earlier chapters (x, y) uses a partitioned relational data base, over items from the ontology, for both internal beliefs and its model of the room. This is oversimplified in a number of ways, but is good enough for a number of additional studies. Conceptually, we have been thinking about partitioned CPRM for capturing belief spaces, extending Narayanan's (ref) thesis design. Leon Barrett (ref) is working on the details of linking beliefs to action and perception wrt their internal model of the speaker, which may differ from their own beliefs.

A preliminary version of Situations (and Mental Spaces) has been implemented within the Bryant ECG analyzer (Ch XX). We will describe this proof of concept version and then discuss some serious barriers to a full implementation. As was mentioned, the key semantics of Situations is that each situation involves a separate simulation, possibly with different beliefs, etc. The basic problem for the grammar and the analyzer is to build a SemSpec that specifies which networks of schemas go with each simulation situation. The following example discusses a simple example:

“Last year, he slid”

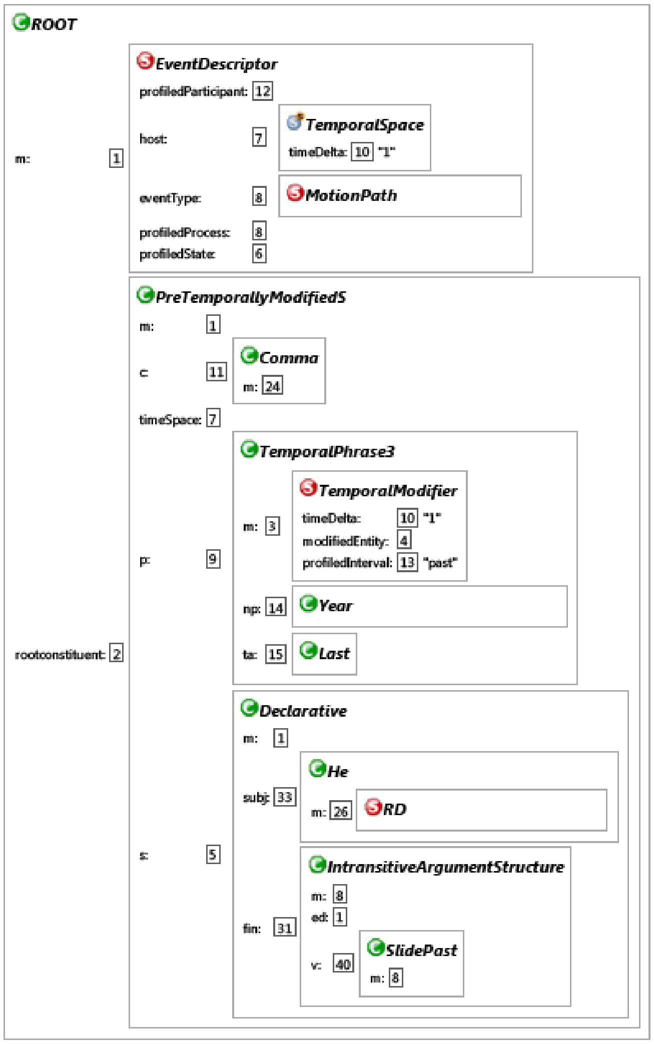

Figure 5.4

This involves a mental space situation a year before the utterance time and a link to some sliding event in the speech context, which of course was somewhat different in the past. This is pretty much the minimal example, but still presents many challenges.

At the basic system level, there needs to be some mechanism to group the SemSpec into appropriate chunks. The current method uses a global system variable *focus* that can be set to different situations. We see in Figure 5.5 that the TempClause (XXX: **TemporallyModifiedS?)** construction resets focus to a past situation, in our example.

<CxN for example, Semspec also>

Even our simple example brings in many of the issues of Mental Space Theory (ref). We need to make explicit the Situation of the discourse itself; traditionally called the *base*. The focus variable always starts in the base situation.

<More, show a SemSpec shot, etc.>

One crucial function in understanding mental space language is recognizing when reference is being made to different spaces. The traditional MS literature discusses “space builder” constructions (ref), but there is a lot more involved. Dialog frequently involves complex interactions between multiple spaces. Part of modeling this requires the fourth ECG primitive, the *Map*. Among other things, these formalize and implement what have been called “correspondence links” in MS research.

```
general construction TemporallyModifiedS
  subcase of ModifiedS
  constructional
    constituents
      p: TemporalPhrase
  meaning
    evokes TemporalSpace as timeSpace
        constraints
          timeSpace.timeDelta ↔ p.m.timeDelta
          self.m.host ↔ timeSpace
```

```
construction PreTemporallyModifiedS
  subcase of TemporallyModifiedS
  constructional
    constituents
      optional c: Comma
  form
    constraints
      p.f before s.f
      p.f before c.f
      c.f before s.f
```

Figure 5.5

4. Maps: Metaphors and Correspondence Links

The fourth and (and currently final) ECG primitive is the *map*. One major use of ECG maps is to specify conceptual relations between different situations or mental spaces. This formalizes the traditional diagrams in the mental space literature. Following the general paradigm above, analysis of discourse involving multiple situations results in a partitioned SemSpec where the elements of various partitions are linked by instances of maps. As we saw, in the example above there should be maps linking both a person of the present to the past. The full interpretation of the utterance would be done, as always, by the simulation process. Individuals and cultures differ significantly on what inferences project across such a map. One major difference is the relative importance placed on inherent traits versus context in human behavior (refs).

ECG maps have another major function in the treatment of figurative language. In contrast with the maps linking situations described above, ECG metaphor maps characteristically link two schemas (often cultural frames). For example, an Event Structure Metaphor map could be defined linking a general Goal schema to the basic Self-Motion schema, which is used in describing action descriptions. Again, analysis of a metaphorical usage will result in a SemSpec that contains instances of the two schemas and the metaphor map. As always, the full interpretation of a metaphoric utterance will be determined by the simulation process and depends on the beliefs and goals of the understander.

For concreteness we will work through the complete analysis of a simple example sentence:

“France slid into recession”.

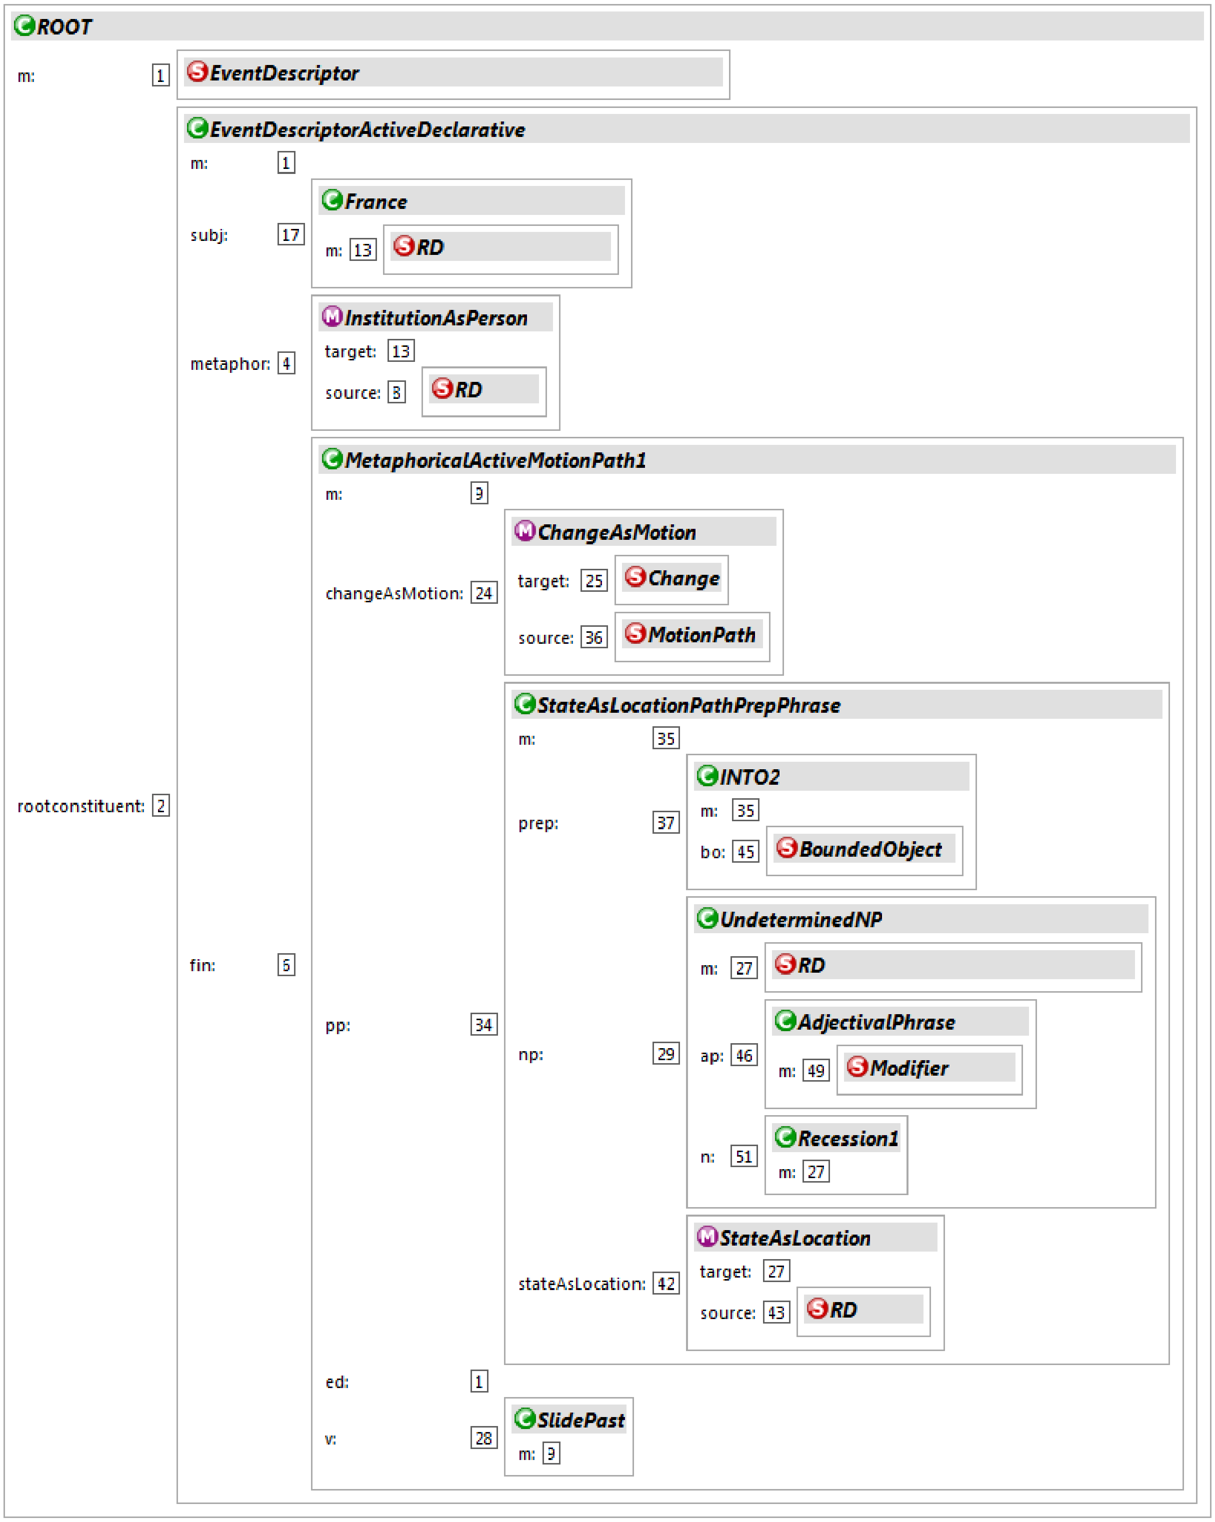

Figure 5.6

Recall that metaphors are modeled by ECG maps. Figure M1 presents the three maps

```
schema Pair
  roles
    first
    second
```

```
general map Metaphor
  roles
    target
    source
```

```
schema StateAsLocationPair
  subcase of Pair
  roles
    first: @state
    second: @location
```

```
map StateAsLocation
  subcase of Metaphor
  roles
    target: RD
    source: RD
    stateAsLocation: StateAsLocationPair
  constraints
    stateAsLocation.first ↔ target.ontological-category
    stateAsLocation.second ↔ source.ontological-category
```

```
construction StateAsLocationPathPrepPhrase
  subcase of PathPrepPhrase1
  meaning
    evokes StateAsLocation as stateAsLocation
    constraints
      ignore self.m.landmark ↔ np.m.referent
      stateAsLocation.source.referent ↔ self.m.landmark
      stateAsLocation.target ↔ np.m
      ignore np.m.ontological-category ⟵ @concrete-entity
      np.m.ontological-category ⟵ @abstract-entity
```

Figure M1

needed for our example. The current implementation of metaphor maps uses a tiny schema called *pair,* also shown in Figure M1, that links a target item to its source domain counterpart. We need two such pairs: InstitutionAsPersonPair and StateAsLocationPair. The top of Figure M1 illustrates that Metaphor is one subcase of the map primitive. The first metaphor example, InstitutionAsPerson, is a prototypical case that just maps one target item to its counterpart. StateAsLocation has the same structure. More typically, the ChangeisMotion map links multiple elements in the target and source domains, using two distinct pairs. As we will see in Figure M3, instances of the maps become part of the SemSpec For this proof-of-concept implementation, we need to add specific additional grammatical constructions for all metaphorical uses. The basics of ECG constructions are discussed in Chapter X and the analysis process in Chapter Y of this volume. The extra constructions required for our example are shown in Figure M2. The first of these, StataeAsLocation-PP, is the metaphorical version of the standard Spatial-PP described on p.? of Chapter ??, again using the StateAsLocation map. The **ignore** statement overrides the action of the non-metaphorical version. Instead, the referent and landmark roles of the prepositional phrase are assigned to the target and source roles of the evoked StateAsLocation metaphor map. This represents the metaphorical usage in the SemSpec and also works for a literal subject as in: “John fell into a depression”.

In the current implementation, we also need specific metaphorical constructions at the phrasal and clausal level, also shown in Figure M2. The MetaphoricallyActiveTranslationalMotion construction is an extension of the standard ActiveTranslationalMotion described in Chapter Y. It evokes an instance of the InstitutionAsPerson metaphor and, as before, assigns items to metaphorical roles. Here, the ProfiledParticipant (France) is the target and mover role is unified with the source domain item. This represents the fact that the agent of the motion is also metaphorical in our example.

Finally, the MetaphoricallyActiveDeclarative (XXX: **EventDescriptorActiveDeclarative?)** construction of Figure M2 describes how the whole example sentence fits together. It differs from the standard active declarative (Chapter X) in requiring its finite constituent (fin) to be metaphorical. It evokes the ChangeisMotion metaphor and resets the speech act type in the EventDescriptor to be metaphorical. The referent of the sentence subject (France) is set to mover ???.

```
construction EventDescriptorActiveDeclarative
  subcase of Declarative
  constructional
    constituents
      fin: MetaphoricalActiveMotionPath1
  meaning // this is an ED
    evokes InstitutionAsPerson as metaphor
    constraints
      ignore subj.m.referent ↔ self.m.profiledParticipant
      ignore subj.m.ontological-category ⟵ @concrete-entity
      subj.m.ontological-category ⟵ @abstract-entity
      metaphor.target ↔ subj.m
      metaphor.source.referent ↔ self.m.profiledParticipant
```

Figure M2

However, this entire discussion only covered an ECG treatment of what might be called “static” maps – ones that are known in advance and do not change during a discourse. Novel metaphors, like novel words and grammatical constructions are assumed in ECG/NTL to involve separate learning processes like those discussed in the Chapters by Chang and Mok. We assume that the analysis and simulation of a particular input relies only on pre-existing linguistic knowledge.

This all works fine for many purposes, but is not sufficient for correspondence mappings between ECG Situations, of which Mental Spaces are a subcase. Even the most basic Mental Space constructions involve “dynamic” maps between spaces that are established by the discourse itself and thus cannot be set up in advance. For example, the phrase “When I was one and twenty” specifies a mental space earlier in time than the discourse and also a map from the speaker to his younger counterpart. If the next clause introduces a wise man and his saying into the earlier space, this must also be mapped forward to the current discourse. ECG needs a clean way of treating these “dynamic maps”. There are proposed designs, but none has been implemented because dynamic synthesis and use of ECG Maps would require massive redesign of Bryant’s analyzer as described in Chapter ZZ of this book. This is in addition to the modifications described above that would be required for the analyzer to recognize and compile SemSpec for unanticipated instances of metaphors.

< footnote "When I was one-and-twenty..." by A. E. Housman (1859-1936) >

5. Neural Level

6. Further Extensions and Conclusions

The three pillars of ECG grammatical analysis: shared skeletal concepts, simulation semantics, and contextual best fit constitute a new paradigm. By restricting the goal of grammatical analysis to producing an intermediate network of shared skeletal concepts, we integrate form, meaning, and context in a tractable and psychologically plausible way. This allows the insights of cognitive linguistics and related fields to be expressed and tested rigorously, for the first time. Compared to contemporary semantic parsers, the full contextual best fit algorithm realizes a long standing goal. This is only possible because the deep conceptual semantics of ECG provides a formalism in which all the relevant factors can be combined.

. Independently of all this, the ECG workbench provides tools for organizing and testing grammars that far surpass those developed explicitly for this purpose. This is because the workbench adapts the Eclipse <link> framework for software development, which benefits from a vastly larger community effort.

So, how can we separate out the problem of detecting the meaning bearing elements and their (grammatical) relations in an utterance?