A Neural Theory of Language and Embodied Construction Grammar

Jerome Feldman, Ellen Dodge, John Bryant

University of California, Berkeley and ICSI

Keywords: construction grammar, embodied, neural, best-fit, compositionality, schemas, cognitive, simulation,

Abstract: This chapter outlines an explicitly neural theory of language and a construction grammar formalism based on this theory. The formalism, ECG, combines deep insights from cognitive linguistics with advanced techniques of neural computation. We illustrate the theory and the formalism with detailed automatically generated semantic analyses of several related examples.

Recent developments in neuroscience and the behavioral sciences suggest approaching language as a cornerstone of Unified Cognitive Science. One such integrative effort has been underway for two decades in Berkeley. The NTL (Neural Theory of Language) project studies language learning and use as an embodied neural system using a wide range of analytic, experimental, and modeling techniques. The basic motivation for NTL and its relation to ongoing experimental work is discussed in several places[1]. The core idea is to take all the constraints seriously and to build explicit computational models that demonstrate the theoretical claims. At one level, NTL continues the tradition of Cognitive Linguistics (CL) represented by several chapters in this volume. But explicit computational modeling demands greater precision than is possible with the pictorial diagrams that remain standard in most CL work.

CL and related approaches to language stress the continuity of language with the whole mind and body and with society. Statistical considerations and incremental learning and adaptation are also deemed essential parts of the capacity for language. The challenge is to develop a methodology that honors the inseparability of language use while being sufficiently rigorous to support formal and computational analysis. The NTL approach is to postulate distinct levels of description, explicitly mirroring the levels in the natural sciences such as biology, chemistry, and physics. The discussions in this chapter will focus on *computational level* descriptions of fairly complex language phenomena. In other work (Feldman 2006), we suggest how such descriptions can be reduced to a *structured connectionist level* and then to postulated brain structures. There is now a fair body of behavioral and biological experimentation exploring these models (Boroditsky 2000; Gallese 2005; Hauk 2004).

Within the computational level, the NTL approach separates language understanding into two distinct phases called *analysis* and *enactment*. Schematically, analysis is a process that takes an utterance in context and produces the best fitting intermediate structure called the Semantic

---

[1] The ECGweb wiki can be found at http://ecgweb.pbwiki.com/. A web search using ECG NTL ICSI will also work.

Specification or *semspec* (cf. Figure 6). The semspec is intended to capture all of the semantic and pragmatic information that is derivable from the input and context. As we will see, this is at a rather deep level of embodied semantics. The semspec is used to drive inference through mental simulation or as we call it, enactment. Within NTL, enactment is modeled using executing networks or **x-nets** which model the aspectual structure of events, and support dynamic inference (Narayanan, 1999).

The grammar formalism of NTL is called Embodied Construction Grammar (ECG). It is a notation for describing language that is being used in a wide range of theoretical and applied projects. The ECG formalism is designed to be all of the following:

1) A descriptive formalism for linguistic analysis
2) A computational formalism for implementing and testing grammars
3) A computational module for applied language tasks
4) A cognitive description for reduction and consequent experiments
5) A foundation for theories and models of language acquisition

Embodied Construction Grammar is our evolving effort to define and build tools for supporting all five of these goals.

This chapter will focus on three related features of ECG: deep semantics, compositionality, and best-fit. The NTL theory behind ECG highlights two aspects of neural embodiment of language – deep semantics and neural computation. One can formalize deep semantics using ECG *Schemas.* This includes ideas such as goals and containers, which have been at the core of Cognitive Linguistics from its origins (Lakoff 1987). As we will see, ECG schemas such as the EventDescriptor (Figure 3) can also describe much broader concepts. ECG schemas can also be used to represent the linguistically relevant parameters of actions (as we will see below), which are packaged in the semspec.

NTL posits that the key to language analysis lies in getting the conceptual primitives right, which in turn depends on evidence from biology, psychology, etc. As linguists, we evaluate putative primitives by their ability to capture general linguistic phenomena. ECG provides a mechanism for expressing and (with the best-fit analyzer) testing linguistic explanations. But isolated phenomena do not suffice for eliciting powerful primitives; we need to examine how a range of cases can be treated compositionally. The tools provided by the ECG system become increasingly important as the size of the grammar increases, as they facilitate testing a wide range of examples and thus help greatly in the cyclic process of hypothesizing linguistic primitives and using these to model complex phenomena.

The core of the chapter is a detailed analysis of a set of related constructions covering purposeful action, with an emphasis on compositionality. The examples illustrate the notation and central ideas of the ECG formalism, but hopefully will also convey the underlying motivations of NTL deep semantics and conceptual composition.

The position of NTL and ECG in the current study of language

The general NTL effort is independent of any particular grammar formalism, but it is strongly aligned with integrated approaches to language including several chapters in this Handbook: Bybee (this volume), Caffarel (this volume), Fillmore (this volume), Langacker (this volume), and Michaelis (this volume). Jackendoff (this volume) presents a different perspective, preserving the separation of form and meaning, but linking them more tightly than earlier generative theories.

NTL also suggests that the nature of human language and thought is heavily influenced by the neural circuitry that implements it. This manifests itself in the best-fit ECG constructional analyzer that is described later in this chapter. An important aspect of both NTL and of the ECG analyzer is the dependence on a quantitative best-fit computational model. This arises from the computational nature of the brain and shares this perspective with statistical (Bod, this volume) and Optimality (Gouskava; de Swart, this volume) approaches to grammar.

One way to characterize the ECG project is as formal cognitive linguistics. ECG is a grammar formalism, methodology, and implementation that is designed to further the exploration and application of an integrated, embodied approach to language. The explicit simulation semantics of NTL plays an important role in ECG, because the output of an ECG analysis (cf. Figure 6) is a semspec.

At a technical level, ECG is a unification-based grammar, like HPSG (http://hpsg.stanford.edu/) and LFG (Asudeh, this volume) in which the mechanisms of unification and binding are extended to deep embodied semantics, discourse structure, and context, as we will show. A unique feature of the ECG notation is the **evokes** primitive, which formalizes Langacker's idea of a profile-base relation and models one aspect of spreading activation in the brain.

ECG is a kind of Construction Grammar (Goldberg 1995; Michaelis, this volume) because it takes as primitive explicit form-meaning pairs called *Construction*s. Both the schemas and constructions are organized in an inheritance lattice, similar to that described by Michaelis (this volume). ECG is called embodied because the meaning pole of a construction is expressed in terms of deep semantic schemas, based on postulated neural circuits and related to the image schemas of CL (Lakoff 1987). An explicit limitation is that no symbolic formalism, including ECG, can capture the spreading activation and contextual best-fit computations of the brain. In the conclusions, we will briefly discuss how NTL tries to unify ECG with neural reality.

Compositionality

Before introducing the technical details of the ECG formalism and illustrating their application in a more detailed case study, it may be helpful to first look at some of the challenges of compositional analysis and some ways these might be addressed.

One challenge is presented by the fact that a given verb will often exhibit more than one pattern of argument realization. For instance, *slide* appears in sentences such as: *The chair slid*; *Jack slid the chair*; and *Jack slid her the chair*. The verb's 'slider' role is expressed in each of these sentences (*the chair*), but in each case is associated with a different grammatical role. Moreover, these sentences describe different types of events, differing as to the presence/absence of causation and transfer of possession, and express different numbers and types of semantic roles.

Using argument structure constructions (Michaelis, this volume), one can handle examples such as these without necessarily positing different verb senses for each pattern. The argument structure construction provides semantic roles associated with the basic type of scene being described, and specifies relations between these roles and grammatical arguments. Verbs have semantic roles associated with them, but do not have to specify how these are linked to grammatical arguments. A given verb can potentially combine with different argument structure constructions, each of which may describe a different type of scene. When a verb unifies with a specific argument structure construction, some or all of their semantic roles will unify.

The meanings and semantic roles associated with the verb and argument structure constructions are clearly a central component of such an approach, but the semantic representations commonly used are often inadequate in several respects. To fully support a compositional analysis of a broad range of examples, constructional meaning representations should meet the following criteria:

- Since it is not entirely predictable, the exact pattern of how the roles of each construction compose with one another should be explicitly specified [rather than relying on the analyst's semantic intuitions]. For instance, to analyze *The chair slid*, there needs to be a specification of the fact that the 'sliding thing' verb role unifies with the argument structure construction's 'Theme' role.
- Because a given argument structure construction will unify with many verbs, its specifications should not be lexically specific. We don't, for example, want to have to list every verb-specific role that an argument structure's 'Theme' will unify with. Instead, more general specifications that capture the meaning common to these different verb roles are needed.
- Verb meaning should be represented so that it explicitly indicates the semantic similarities that motivate the unification of verb roles with the roles of argument structure constructions. This is not possible if verb meaning is defined too generally. For instance, we could use a 'Theme' role to represent the sliding thing, motivating its composition with an argument structure construction's 'Theme' role. But, this would not explain what motivates it to compose with the 'Patient' role in *He slid the chair*. On the other hand, if the argument structure construction for this second example is defined using a Theme role instead of a Patient role, the semantic distinction between transitive and intransitive constructions is obscured.
- The results of composition should be something more than just the conjunction of two role names (such as slider-patient or theme-patient); ideally, complex semantic roles should be defined using separately motivated conceptual structures.
- Meaning representations should capture both differences and similarities of meaning. Roles that are only defined at a very general level, such as thematic-type roles, support recognition

of semantic similarity, but obscure differences. More specific roles can make semantic differences more apparent, but need to be defined in such a way that it is also possible to recognize semantic similarities that motivate the composition of verb and argument structure constructions.

Each of these goals presents some challenges. Taken as a whole, it is clear that semantic representations need to include more than just very general thematic-type roles and very specific lexically-defined roles.

In addition to verb and argument structure constructions, a compositional analysis of sentence meaning requires that various phrase-level constructions be defined. To support a compositional analysis, these constructions need to be defined in such a way that they will unify with other constructions instantiated in a sentence, with the result that the composed meanings of these constructions specify the sentence's meaning. Therefore, phrasal constructions need to be coordinated with the definition of argument structure and other constructions.

What phenomena do such constructions need to deal with? In the case study presented later in this chapter, we will analyze a set of sentences which describe the same type of events and actions, but which exhibit other, sometimes subtle, differences in meaning. As a prelude, we will briefly describe some of the challenges presented by sentences such as these. To start, consider *She kicked the table*, *Her foot kicked the table,* and *The table was kicked,* sentences which can all be used to describe the same scene, but which differ in terms of scene perspective, or 'participant profiling' While the same set of semantic roles are conceptually present in each case, there are differences as to how (and if) they are expressed. Consequently, such sentences may instantiate different, but related, argument structure constructions.

The sentences *You kicked the table, Did you kick the table?,* Which *table did you kick?, and Kick the table!* can all be used to describe the same type of kicking event. The kicking is in each case described from the same perspective (the kicker), which suggests that these sentences may all instantiate the same argument structure construction. However, they clearly differ as to their discourse purpose, as well as differing as to topic, presence/absence of auxiliary, and in other ways, suggesting the need for different phrase-level constructions specifying these differences. Ideally, the phrase-level constructions should be defined in such a way that they will unify with a range of argument structure constructions. So, for instance, they could also be used in the analysis of sentences such as *Did her foot kick the table? Which table was kicked?*, etc.

Ultimately, the goal of a compositional constructional analysis of an utterance is one of conceptual composition: when the constructions instantiated in a sentence unify, their meanings should compose, and this composition should represent the conceptual structure associated with the utterances as a whole. Consequently, it is critical not just to get the constructional 'decomposition' of utterances right, but also their conceptual decomposition. But how can we get this right? How can we carve up conceptual structure along the proper 'joints'?

Cognitive linguistics research has yielded many insights into the nature of the conceptual structures conveyed by language. Research on the meanings of spatial relations terms has shown that they can be analyzed in terms of combinations of primitive elements (e.g. Talmy 1972, 1983; Langacker 1976, 1987, this volume). These primitives include such things as bounded regions, paths, contact, etc. Cognitive linguists have also observed that primitives such as these recur across many different experiences, including but not limited to language. They have described various image schemas, each of which includes a relatively small numbers of parts or **role**s, and which have an internal structure that supports inferences (Johnson 1987). Later research supports the idea that such schemas are "embodied" neural structures (Regier 1996; Dodge & Lakoff 2005). Moreover, many types of culturally-specific experiences also exhibit schema-like structures, or frames (Fillmore 1982).

ECG builds upon these ideas, representing meaning using **schemas**, which are partial representations of neurally-based conceptual structures. Importantly, these schemas are defined independent of specific linguistic constructions. For example, we can define a Container schema, which can be used to represent the meaning of *in*. But, this schema's structure is also present in many non-linguistic experiences, such as experiences of putting things into and taking them out of various containers. Moreover, this same schema can be used in the meaning representations of many different constructions, such as those for *inside, inner, out,* and *outside*. Several examples are given in Figure 1, in a formal notation to be defined below.

Many complex conceptual structures can be 'decomposed' and be represented as combinations of more 'primitive' conceptual structures. For instance, going into a bounded region can be decomposed into a spatial relations component (a relation to a container) and a motion component. Crucially, the conceptual gestalt is a particular combination of these parts, not just some sort of random juxtaposition.

The meaning of many constructions can also be decomposed and represented using combinations of more primitive schemas. For example, the meaning of *into* can be represented as a composition of Container and SourcePathGoal schemas. And, when constructions compose, their meanings will also compose, generally resulting in more complex conceptual structure. Consequently, the meaning of the unified constructions can also be represented as a composition of schemas. Different constructional combinations will be associated with different complex conceptual structures, represented by different compositions of schemas.

To represent the meanings of a wide range of constructions, a relatively large, comprehensive inventory of schemas is needed. In addition to defining schema primitives that are used in descriptions of spatial relations, we also need schemas that represent other types of concepts, especially those that are directly relevant to argument realization, such as action, force, causation, motion, and change.

But, it is not a simple matter to develop such an inventory. To do so requires that we figure out what sorts of more 'primitive' conceptual structures might exist, as well as how other, more complex structures might be analyzed as involving productive compositions of these more primitive structures.

Our strategy is to examine a broad range of situations evidencing similar types of conceptual structures. By comparing them, we can gain insights into the types of distinctions and similarities that schemas representing such structure should capture. Further insights can be gained by using more than one kind of evidence to make these comparisons. We will briefly illustrate this strategy by examining some different types of motor-control actions, and discussing the sorts of schemas that seem to be needed to represent the conceptual structures associated with them. The following discussion uses linguistic examples to illustrate these comparisons, but there is also other cognitive and neural evidence to support these decompositions. A more detailed analysis is presented in Dodge (2009).

We start by considering motor control actions and the different types of 'roles' they may include. For example, actions such as sneezing, walking, or smiling involve an actor and (typically) some part of the actor's body, but not necessarily another entity. And, verbs describing such actions typically occur in utterances which express this actor role, e.g. *He walked/sneezed/smiled*. Many actions also involve another entity, but differ in other important respects. Actions such as looking and pointing don't necessarily involve contact or transfer of force to the other entity. And verbs for these actions typically occur in utterances that express this other 'target entity' role, but in such a way that does not indicate any physical affectedness, e.g. *He looked/pointed at her*. In actions such as kicking, squeezing, and pushing, the actor contacts and transfers force to another entity. In many cases, forceful actions result in some physical change to the entity being acted upon, suggesting an important distinction in affectedness between forceful and non-forceful actions. These distinctions suggest the need for a basic motor-control schema with an actor role, and at least two other related schemas (for forceful and non-forceful actions) that each contain some kind of role for an additional entity.

Next, let us look a little more closely at the relation between forceful actions and cause-effect. Forceful action verbs often occur in utterances that express the 'acted upon' entity in a way that indicates it is affected in some way, e.g. *He squeezed/pulled/kicked the bottle*. But, the same action can potentially produce many different effects. For instance, kicking may make a leaf move, break a pane of glass, or cause someone pain. And, in some cases, it may not cause any perceptible effect, e.g. *He kicked at the door/pulled on the rope, but it wouldn't budge*. Furthermore, in many cases, any "effects" that do occur are ones that can also occur independent of the action (and, in some cases independent of any readily observable 'causer'). For example, a leaf may fall off a tree, a window may break in a storm, and we may feel a sudden pain in our leg and not know what caused it. These observations suggest that schemas for forceful actions, as well as those for possible effects that are caused by such actions (such as motion and change of state), should be defined independent of causation. However, these same schemas can also serve as parts within more complex schemas, such as those for cause-effect events that involve a causal relation between a forceful action and the motion or change of some other entity.

The remainder of this Chapter formalizes these notions and shows how they can be combined with innovative computational tools to support deep semantic analysis of complex utterances.

ECG Notation and Primitives

In ECG, construction grammars are specified using two basic primitives: **constructions** and **schemas**. Constructions are paired form constraints and meaning constraints. ECG is different from other construction grammar formalisms because the meaning constraints are defined in terms of embodied semantic schemas, such as those in Figure 1.

There are four ways to specify relations between ECG primitives: roles, sub-typing (through the **subcase of** keyword), evoking a structure (through the **evokes** keyword), and constraints (co-indexation and typing). A **role** names a part of a structure, and the **subcase of** keyword relates the construction/schema to its type lattice, allowing for structure sharing through (partial) inheritance.

Evoking a structure makes it locally available without imposing a part-of or subtype relation between the evoking structure and the evoked structure. Using Langacker's standard example, the concept hypotenuse only makes sense in reference to a right triangle, but a hypotenuse is not a kind of a right triangle, nor is the right triangle a role of the hypotenuse. The evokes operator is used to state the relationship between the hypotenuse and its right triangle.

Like other unification-based formalisms, ECG also supports constraints on roles (features). The double-headed arrow operator is used for co-indexing roles (↔). Roles can be assigned an atomic value using the assignment operator (⟵). A type constraint (specified with a colon) constrains a role to only be filled by a certain type of filler.

The specific grammar described here will be called EJ1. Figure 1 shows a set of EJ1 semantic schemas ranging over conventional image schemas (TL and SPG), embodied processes (Process, ComplexProcess, and MotorControl), and motion schemas (Motion, TranslationalMotion and EffectorTranslationalMotion). The TL schema has roles for a trajector and a landmark. The SPG schema inherits the trajector and landmark roles by subcasing TL, and adds roles for describing a path including source, path, and goal. As we will see in Figure 6, the embodied semantics (semspec) of ECG consists entirely of schemas with bindings between their roles.

The Process and ComplexProcess schemas are general descriptions of actions and events in which a single participant is profiled using the protagonist role. A ComplexProcess is made up of two sub-processes called process1 and process2. The ComplexProcess schema shows how roles can be bound (co-indexed) -- required to have the same filler. The ComplexProcess's

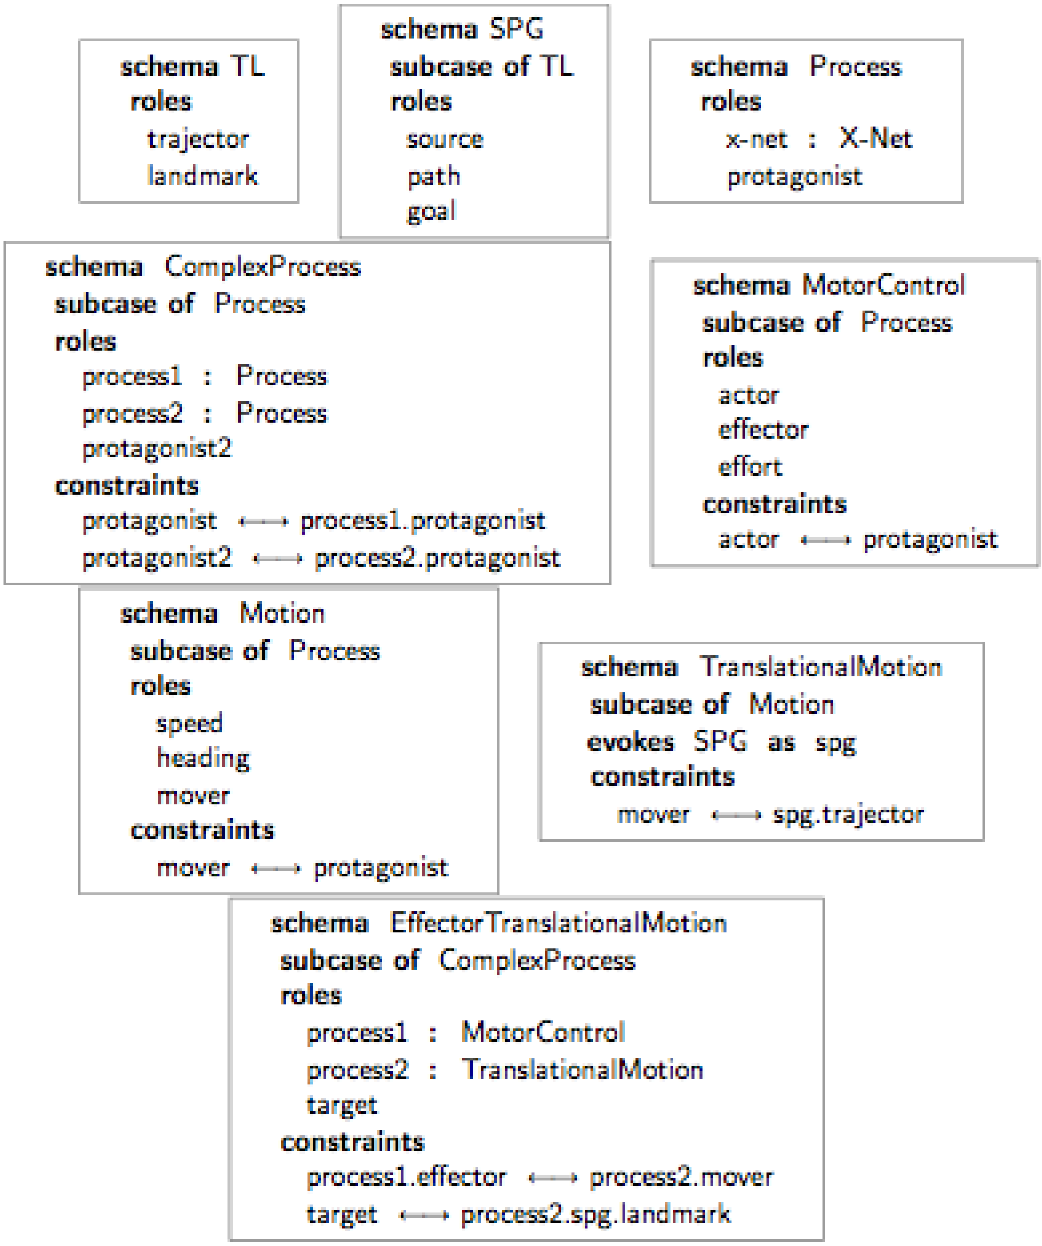

```
schema TL
  roles
    trajector
    landmark

schema SPG
  subcase of TL
  roles
    source
    path
    goal

schema  Process
  roles
    x-net  :  X-Net
    protagonist

schema  ComplexProcess
  subcase of  Process
  roles
    process1  :  Process
    process2  :  Process
    protagonist2
  constraints
    protagonist  ←→  process1.protagonist
    protagonist2  ←→  process2.protagonist

schema MotorControl
  subcase of  Process
  roles
    actor
    effector
    effort
  constraints
    actor  ←→  protagonist

schema  Motion
  subcase of  Process
  roles
    speed
    heading
    mover
  constraints
    mover  ←→  protagonist

schema  TranslationalMotion
  subcase of  Motion
  evokes  SPG  as  spg
  constraints
    mover  ←→  spg.trajector

schema  EffectorTranslationalMotion
  subcase of  ComplexProcess
  roles
    process1  :  MotorControl
    process2  :  TranslationalMotion
    target
  constraints
    process1.effector  ←→  process2.mover
    target  ←→  process2.spg.landmark
```

Figure 1: ECG representations of image schemas (TL and SPG), embodied processes (Process, ComplexProcess, and MotorControl), and motion (Motion, TranslationalMotion and EffectorTranslationalMotion).

primary protagonist role (inherited from Process) is co-indexed (using ↔) with the protagonist role of process1, and the secondary protagonist2 role is bound to the protagonist of process2.

The Motion schema describes a process in which the mover is the protagonist and has a speed and heading. The TranslationalMotion schema is a subcase of Motion and adds the constraint that the motion is conceptualized as occurring along a path. The TranslationalMotion schema shows an example of the evokes keyword; the path is represented by the evoked SPG schema. In both schemas, the mover is defined as the protagonist using a co-indexation constraint. Each process has a role called x-net that further specifies the kind of action that is modeled by the schema. For example, the Motion process can describe walking, crawling and many other methods of motion. Specific aspects of particular actions are represented by the filler of the x-net role. As a consequence, the Motion schema acts as an abstraction over all the different motion x-nets. This is one crucial requirement for compositionality.

The MotorControl schema has a special significance in the grammar. It is the semantic root of embodied, controlled processes. It adds roles for an actor, effector, and effort. The actor is the embodied protagonist, the effector is the controlled body part, and the effort is the energy expenditure.

The schema EffectorTranslationalMotion puts all of these schemas together to represent the idea of an entity controlling the motion of an effector. EffectorTranslationalMotion is a ComplexProcess in which process1 has a MotorControl type constraint and process2 is constrained to be a TranslationalMotion. The schema also adds a role for the target towards which the effector is moving.

Figure 2 shows the schemas that are combined to define the meaning of verbs of impact such as “hit”, “slap”, “kick” etc. The ForceTransfer schema describes a transfer of a particular amount of force between any kind of supplier and recipient. The ForceApplication schema describes MotorControl actions in which force is applied, and thus evokes the ForceTransfer schema. The ForceApplication schema adds roles for an instrument and an actedUpon entity. The constraints block of the schema then binds the appropriate roles. As we discussed earlier, a judicious choice of embodied schemas enables us to capture conceptual regularities and the ECG formalism supports this.

ECG Constructions

Constructions are pairings of form and meaning, and in ECG, this pairing is represented by a form block (defined by the **form** keyword) and a **meaning** block. Both the form pole and meaning pole of a construction can be typed. In this chapter, we simply constrain the form pole of HitPastTense and SlapPastTense to be a word using the Word schema (not shown). Form blocks can also have form constraints, and in these simple lexical constructions, the form constraint specifies that the orthography of the HitPastTense construction is “hit” by binding the string “hit” to self.f.orth. The slot chain self.f.orth uses the **self** keyword to refer to the construction itself, and the keyword **f** to refer to the construction's form pole. The Word schema has a role called orth to represent the orthography of a word.

Figure 2 also shows how past tense lexical constructions for hit and slapped can be defined in terms of the ForcefulMotionAction schema. Like schemas, constructions are arranged into a subcase lattice and the HitPastTense and SlapPastTense constructions subcase a general PastTense construction (not shown) that specifies facts about the tense and aspect of PastTense verbs.

The meaning block of these two lexical constructions is typed as ForcefulMotionAction. The meaning block of a construction is quite similar to a semantic schema, and thus it also allows for semantic constraints as well as evoking structure. In the case of the lexical HitPastTense and SlapPastTense constructions, the (inherited) x-net role in ForcefulMotionAction is assigned the appropriate X-Net. Using the general ForcefulMotionAction schema to represent the meaning of "hit" and "slap" provides a useful level of generalization over broad range of verbs, and as we will see below, it enables the definition of a simple transitive argument structure construction to cover this semantic class of verbs.

ECG argument structure constructions must also provide guidance about how a scene should be simulated. For example, active and passive provide differing perspectives on the same scene, and such a perspective shift must be communicated to the simulator. For this the EJ1 grammar uses a general VerbPlusArguments construction and its associated abstraction over scenes (events) called the EventDescriptor schema (both shown in Figure 3) to represent perspectivized facts about a scene.

The central importance of the EventDescriptor schema extends the central function of predication in grammar as proposed in Croft's Radical Construction Grammar (2001). We suggest that an utterance has two primary construals available to it: the scene provided by the argument structure construction and a particular process provided by the verb. Often these two processes are the same, but they are not required to be the same.

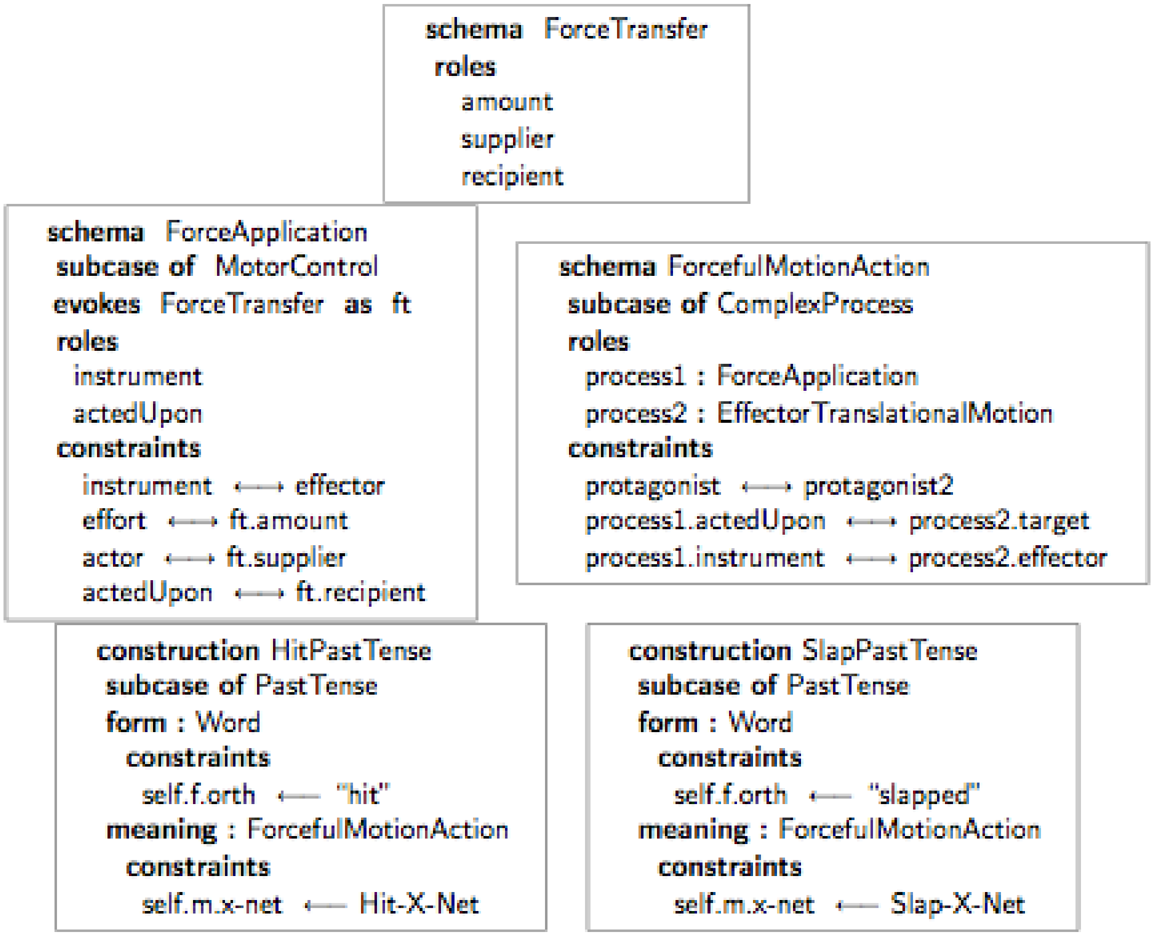

```
schema  ForceTransfer
 roles
   amount
   supplier
   recipient
```

```
schema  ForceApplication
 subcase of  MotorControl
 evokes  ForceTransfer  as  ft
 roles
  instrument
  actedUpon
 constraints
   instrument  ⟷ effector
   effort  ⟷ ft.amount
   actor  ⟷ ft.supplier
   actedUpon  ⟷ ft.recipient
```

```
schema ForcefulMotionAction
 subcase of ComplexProcess
 roles
  process1 : ForceApplication
  process2 : EffectorTranslationalMotion
 constraints
  protagonist  ⟷ protagonist2
  process1.actedUpon  ⟷ process2.target
  process1.instrument  ⟷ process2.effector
```

```
construction HitPastTense
 subcase of PastTense
 form : Word
   constraints
    self.f.orth  ⟵  "hit"
 meaning : ForcefulMotionAction
   constraints
    self.m.x-net  ⟵  Hit-X-Net
```

```
construction SlapPastTense
 subcase of PastTense
 form : Word
   constraints
    self.f.orth  ⟵  "slapped"
 meaning : ForcefulMotionAction
   constraints
    self.m.x-net  ⟵  Slap-X-Net
```

Figure 2: ECG representations of processes related to force application including ForceTransfer, ForceApplication, and ForcefulMotionAction. These schemas represent the meaning of verbs of impact such as “hit”, “strike”, “kick”, “slap” etc.

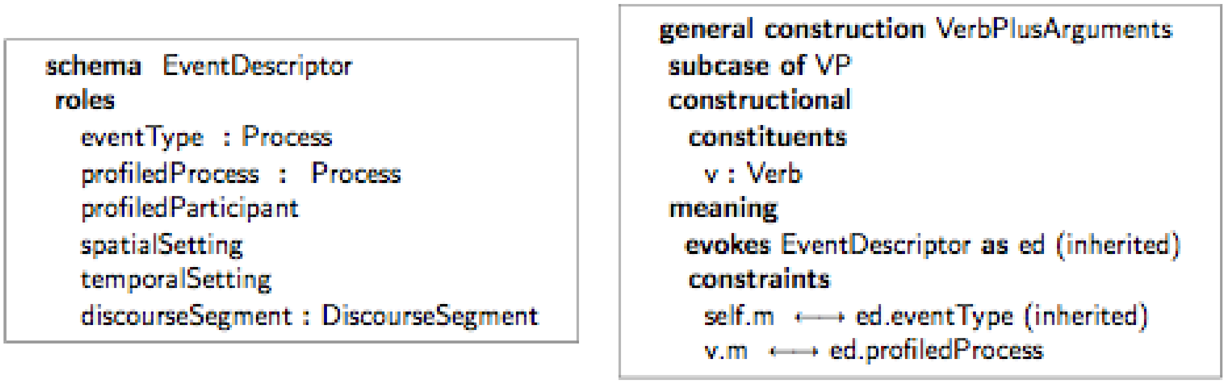

```
schema  EventDescriptor
 roles
   eventType  : Process
   profiledProcess  :  Process
   profiledParticipant
   spatialSetting
   temporalSetting
   discourseSegment : DiscourseSegment
```

```
general construction VerbPlusArguments
 subcase of VP
 constructional
   constituents
    v : Verb
 meaning
   evokes EventDescriptor as ed (inherited)
   constraints
    self.m  ⟷ ed.eventType (inherited)
    v.m  ⟷ ed.profiledProcess
```

Figure 3: The EventDescriptor Schema and a general VerbPlusArguments construction that functions as the root of the argument structure hierarchy.

In the EventDescriptor schema:

- The eventType role is bound to the Process that represents the scene being described. An argument structure construction supplies the filler of this role.
- The profiledProcess role is bound to the subprocess in the scene that is being profiled. The verb supplies the filler of this role.
- The profiledParticipant role is bound to the participant in the scene that is being profiled. This role can be thought of as the semantic correlate of subject and it is bound to different roles in a scene depending on whether the utterance is active or passive voice.
- Roles temporalSetting and locativeSetting are bound to the time and location.
- The discourseSegment role is typed to a simplified DiscourseSegment schema, which has roles for the speechAct of the utterance and the topic of the utterance. In this chapter, the speechAct role will be bound to a simple atomic value such as ``Declarative" or ``WH-Question". The topic role will be bound to the topic specified by each finite clause. Table 1 shows how the roles of the EventDescriptor are co-indexed for some examples.

The profiledParticipant role provides a lot of leverage in the grammar. It allows for a simple semantic distinction between active and passive sentences and makes it straightforward to implement control (described below). As a consequence, the subject constituent is just like any other constituent in the grammar, and has no special status apart from the fact that its meaning is bound to the profiledParticipant role. Thus a construction like the imperative without a subject is not problematic in that the profiledParticipant is just bound to the addressee.

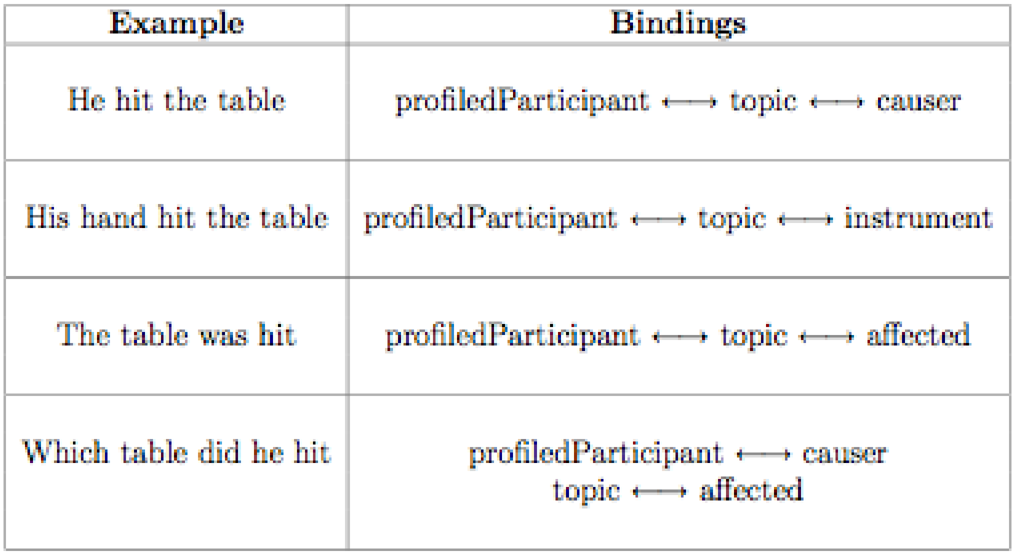

Example	Bindings
He hit the table	profiledParticipant ⟷ topic ⟷ causer
His hand hit the table	profiledParticipant ⟷ topic ⟷ instrument
The table was hit	profiledParticipant ⟷ topic ⟷ affected
Which table did he hit	profiledParticipant ⟷ causer topic ⟷ affected

Table 1: Different ways the EventDescriptor.profiledParticipant, DiscourseSegment.topic and the semantic roles of CauseEffectAction are co-indexed for four different example sentences.

The notion of topic is distinct from both the subject constituent and the profiledParticipant role. While a simple declarative construction (described below) would co-index the profiledParticipant with the topic role, this is not a requirement. A WH-Question with a fronted NP constituent binds the meaning of the fronted NP to the topic role. Again, the ECG EventDescriptor formalizes and extends several insights from CL. By formalizing theoretical findings from CL, ECG allows for conceptual compositions that model human language use.

Also shown in Figure 3 is the phrasal VerbPlusArguments construction. *Phrasal* constructions share many properties with *lexical* constructions. The primary difference between the lexical and phrasal constructions presented in this chapter is the presence of a **constructional** block. A phrasal construction's **constituents** are defined in the constructional block.

The VerbPlusArguments construction is the root of the argument structure hierarchy. Because it is marked with the **general** keyword, it is not used directly for interpretation or production, but instead represents a generalization over all its subtypes. The generalization that the VerbPlusArguments captures is that all argument structure constructions (in this EJ1 grammar) have a verb constituent called v.

The VerbPlusArguments construction has no additional form constraints to add, and thus the form block is omitted. In its meaning block, the VerbPlusArguments construction inherits an evoked EventDescriptor from general construction VP (not shown) as well as the constraint that the meaning of the construction itself (specified by self.m) is bound to the EventDescriptor's eventType role. It then adds the semantic constraint that the meaning of the verb is bound to the EventDescriptor's profiledProcess role. This constraint cashes out the intuition described above that a verb profiles a particular process associated with the scene being described.

Specific subtypes of VerbPlusArguments are shown in Figure 4. These argument structure constructions define transitives with a salient causer. The meaning of these transitives is defined using the CauseEffectAction schema. CauseEffectAction is a complex process in which process1 (the cause) is a ForceApplication, and process2 is the effect. It also adds roles for a causer and an affected participant, and uses co-indexations to bind the causer and affected roles to the appropriate roles in process1 and process2.

The TransitiveCEA construction represents transitive VPs with causal verbs with a force application component such as “cut”, “chop”, “crush” etc. It defines an additional NP constituent to represent its object, and uses the form block to add a form constraint requiring that the verb's form (specified by v.f) comes before the form of the np (specified by np.f). The meaning of the construction is defined to be the CauseEffectSchema.

In the semantic constraints block of TransitiveCEA, the construction specifies how its semantic roles relate to its constituents. Its semantic roles are referred to using the slot chain self.m, which in this case refers to a CauseEffectAction schema. Thus self.m.causer refers to its semantic causer role. The first constraint in the meaning block of TransitiveCEA co-indexes the causer with the profiledParticipant, and the second constraint co-indexes the affected participant role with the meaning of the NP constituent. The final constraint unifies the meaning of the construction itself (self.m) with the meaning of the verb. The first consequence of this co-indexation is that it limits the kinds of verbs that can co-occur with the TransitiveCEA construction to those that have meaning poles that can unify with CauseEffectAction. The second consequence is that it imports the meaning of the verb into the meaning pole of the construction via unification.

```
schema CauseEffectAction
 subcase of ComplexProcess
 roles
  process1 : ForceApplication
  causer
  affected
 constraints
  protagonist ⟷ causer
  protagonist2 ⟷ affected
  process1.actedUpon ⟷ affected
```

```
construction TransitiveCEA
 subcase of  VerbPlusArguments
 constructional
  constituents
    np  :  NP
 form
  constraints
    v.f before np.f
 meaning : CauseEffectAction
  constraints
   self.m.causer ⟷ ed.profiledParticipant
   self.m.affected ⟷ np.m
   self.m ⟷ v.m
```

```
construction TransitiveCEAProfiledCause
 subcase of TransitiveCEA
 meaning : CauseEffectAction
  evokes ForcefulMotionAction as fma
  constraints
   ignore: self.m ⟷ v.m
   v.m ⟷ fma
   self.m.process1 ⟷ fma.process1
```

```
construction TransitiveCEAProfiledInstrument
 subcase of TransitiveCEAProfiledCause
 meaning : CauseEffectAction
  constraints
   ignore: self.m.causer ⟷ ed.profiledParticipant
   self.m.process1.instrument ⟷ ed.profiledParticipant
```

Figure 4: A selection of argument structure constructions that model transitives like *he cut the steak* (TransitiveCEA), *he hit the table* (TransitiveCEAProfiledCause), and *his hand hit the table* (TransitiveProfiledInstrument).

Analysis Examples

Having introduced some of the schemas and constructions in EJ1, we now show how they are used to support a compositional analysis of various types of sentences. Our emphasis here is on the mechanisms of ECG; the linguistic analysis of these examples and many more can be found in Dodge (2009).

We assume the following points: (1) A sentence not only instantiates lexical constructions, but also phrasal and other types of constructions. All of these constructions are meaningful, with meaning represented using schemas. (2) When these constructions unify, their meanings compose in a way that is consistent with the constraints specified in each construction. (3) In ECG this composed meaning is represented as a semantic specification for a simulation (semspec).

The same construction can be instantiated in many different sentences, and should therefore compose with a variety of other constructions. For each sentence, the instantiated constructions should unify to produce a semspec that is consistent with our intuitions about that sentence's meaning. Similarities and differences in sentence meaning should be reflected in their semspecs. In addition, the ECG lattice of constructions facilitates expressing generalizations across constructions.

We will first present an in-depth analysis of the simple declarative sentence *He hit the table*. Then, we will look at sentences which are similar in many respects, but which present some challenges to linguistic analysis, such as instrument subjects, passives, and questions. For each, we will describe how the EJ1 grammar supports an analysis involving many of the same or similar constructions composed in different ways. Crucially, differences in meaning are captured as a few key differences in the semspecs that result from these different compositions.

He hit the table instantiates several different constructions, whose meanings are unified to produce the semspec shown in Figure 6. To understand which elements of the semspec each of these constructions provide, how they unify, and what type of sentence meaning this semspec represents, let's look more closely at the instantiated constructions. The key ones described here are: Declarative (shown in Figure 5), TransitiveCEAProfiledCause (an argument structure construction, shown in Figure 4); and HitPastTense (a verb construction shown in Figure 2). In addition, there are nominal constructions for 'he' and 'the table'.

Declarative identifies its meaning with an EventDescriptor schema (Figure 3), indicating that this type of construction is used to describe some kind of event. Declarative inherits a subj constituent and the constraint that this constituent's meaning is bound to the EventDescriptor's profiledParticipant role. In this way, Declarative indicates that the event should be simulated from the perspective of the entity referred to by the subj constituent. Constituent subj, in the current example, unifies with HE, providing information that the entity that fills the profiledParticipant role is, in this case, MALEANIMATE.

Declarative does not, however, specify what type of event is being described, nor does it specify which event-related semantic role the profiledParticipant is associated with; this information is instead supplied by whichever argument structure construction Declarative unifies with, in this case TransitiveCEAProfiledCause.

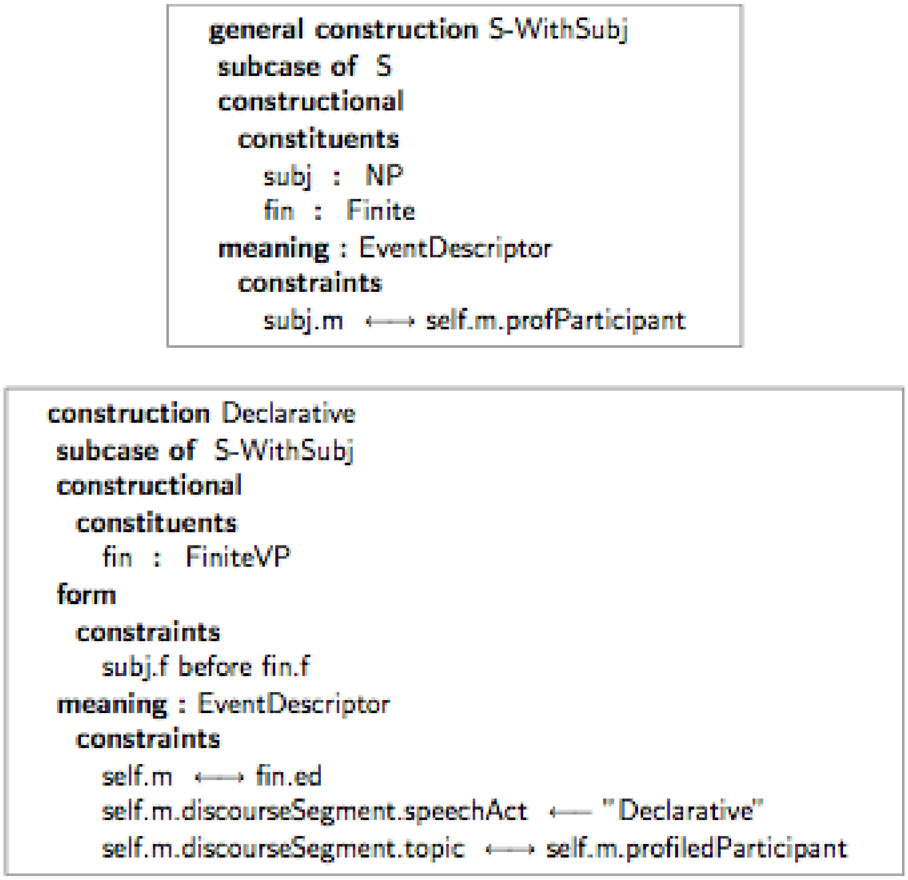

```
general construction S-WithSubj
 subcase of  S
 constructional
   constituents
      subj  :   NP
      fin   :   Finite
 meaning : EventDescriptor
   constraints
      subj.m   ←→ self.m.profParticipant
```

```
construction Declarative
 subcase of  S-WithSubj
 constructional
   constituents
      fin   :   FiniteVP
 form
   constraints
      subj.f before fin.f
 meaning : EventDescriptor
   constraints
      self.m   ←→ fin.ed
      self.m.discourseSegment.speechAct   ←  "Declarative"
      self.m.discourseSegment.topic   ←→ self.m.profiledParticipant
```

Figure 5: Speech act constructions that set the profiledParticipant role and the topic roles of the EventDescriptor and DiscourseSegment schemas, respectively. The Declarative construction covers basic declarative sentences.

Declarative also specifies discourse-related information. This is represented here in simplified fashion by specifying that the discourse segment's speechAct role has the atomic value "Declarative". Additionally, the topic of the discourse segment is bound to the profiledParticipant role. Thus, the profiledParticipant, the topic of the discourse segment, and the meaning of the subj constituent are all bound to one another, as shown by their sharing boxed number 2 in Figure 6.

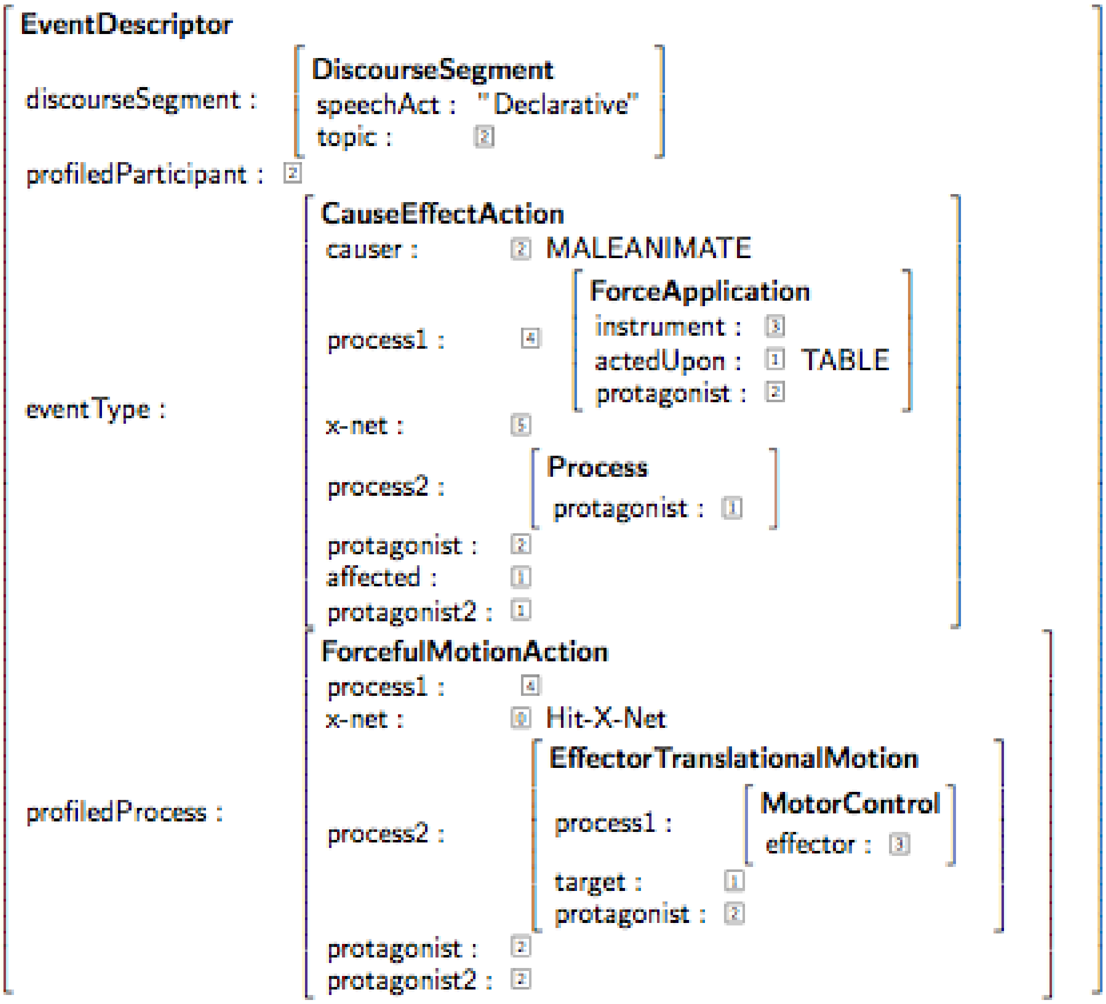

Figure 6: The semspec for the sentence *He hit the table*.

Declarative has a second constituent, fin (a type of FiniteVP, not shown) which, in this example, unifies with TransitiveCEAProfiledCause. This argument structure construction unifies with verbs such as “hit”, “slap”, and “kick”. Declarative also specifies that the EventDescriptor evoked by this argument structure construction is to be identified with that of Declarative, indicating that both constructions are describing the same event.

As with other argument structure constructions, TransitiveCEAProfiledCause provides information about the type of scene that is being described. This is specified through the (inherited) constraint that its meaning is bound to the eventType role of the EventDescriptor. TransitiveCEAProfiledCause is a member of a family of transitive argument structure constructions, all of which identify their meaning with a CauseEffectAction schema (Figure 4). This schema defines a complex process in which one process, a ForceApplication, has a causal relation to another process. Two key participant roles are defined by this schema: an animate causer who performs the action, and an affected entity, which is acted upon and (potentially) affected in some way by this action. Thus, this family of constructions reflects the causal semantics prototypically associated with transitivity, and includes semantic roles similar to those of Agent and Patient.[2] In ECG, both the representation and underlying conceptualization of these roles is semantically complex, and they are defined relative to embodied schemas/gestalts, rather than just being names whose meaning is left unspecified.

In EJ1 grammar, argument structure constructions not only specify what type of event is being described, but also specify which event participant is 'profiled', i.e. from whose perspective the event should be simulated. TransitiveCEAProfiledCause is a subcase of the TransitiveCEA, and inherits its constraint that the profiledParticipant is the causer. Therefore in the semspec, causer is co-indexed with profiledParticipant. As described later, other argument structure constructions may specify that the profiledParticipant is bound to a different event role.

As with other argument structure constructions, TransitiveCEAProfiledCause inherits a verb constituent, the meaning of which serves to elaborate a process or subprocess related to the event as a whole. In the semspec, the meaning of the verb is bound to the profiledProcess of the EventDescriptor. Its parent construction, TransitiveCEA, defines a central case situation in which the verb constituent meaning is identified with the same schema as the argument structure construction, indicating a very close correspondence in meaning between the two constructions. TransitiveCEAProfiledCause represents an extension to the central case, a situation in which the verb constituent provides information about the causal process of CauseEffectAction (process1), but does not elaborate the effect (process2). This is specified within the construction by: (1) constraining verb constituent meaning to be ForcefulMotionAction and specifying that the inherited constraint that verb meaning is the same as argument structure construction meaning should be ignored, and (2) binding the causal subprocess, ForceApplication, to the verb constituent meaning. Because constructional meaning is specified using conceptual primitives, this argument structure construction is not lexically-specific, covering all verbs which identify their meaning with ForcefulMotionAction (such as "punch", "pat", and "tap"). The unification of the verb and argument structure constructions results in the co-indexation of many different roles.

In Figure 6, the causer (boxed 2) is also protagonist of CauseEffectAction, ForceApplication, ForcefulMotionAction and EffectorTranslationalMotion. The affected (boxed 1) is:

- CauseEffectAction.protagonist2,
- CauseEffectAction.process2.protagonist,
- ForceAppplication.actedUpon, and
- EffectorTranslationalMotion.target.

TransitiveCEAProfiledCause's verb constituent is unified with the HitPastTense verb construction, whose meaning (ForcefulMotionAction) meets the constraints specified by the argument structure construction. HitPastTense specifies a particular type of X-net (a hitting routine). The NP constituent of TransitiveCEAProfiledCause is bound to the affected role. It provides information about the affected role of the CauseEffectAction. In this particular example, this constituent is unified with an NP construction whose N constituent is TABLE.

---

[2] Note, though, that not all argument structure constructions of transitive form necessarily share this same meaning.

The semspec for this sentence thus supports a simulation of an event in which a male causal actor performs a particular kind of forceful action (hitting) on a table, and this table is (at least potentially) affected in some way. Neither the argument structure construction nor the verb specify what type of effect this is; the particular effect will depend on the particular fillers of the causer and affected roles. Compare: The baby/weightlifter hit the table/wine glass/policeman. The simulation of effect will also depend on the specific ForcefulMotionAction described by the verb. For instance, substitute *patted*, which specifies a low amount of force, in the examples above.

We have now shown how a particular sentence can be analyzed as instantiating several constructions, whose meanings compose to produce a semspec for that sentence. Next, we will show how sentences similar in form and meaning to *He hit the table* can be analyzed as compositions of many – but not all-- of these same constructions.

First, consider the sentence *The hammer hit the table*. Under an instrumental reading[3] of this sentence, the event being described is one in which a person used a hammer to hit the table – roughly, someone hit the table with a hammer. The meaning of this sentence is very similar to *He hit the table*, but with an important difference in participant profiling: *The hammer hit the table* foregrounds the instrument and its motion and contact with the table, and backgrounds the actor who is wielding this instrument.

In ECG, the particular 'event perspective' or 'participant profiling' of an utterance is specified in the semspec through a binding between profiledParticipant and a particular event participant role. In the case of *He hit the table*, profiledParticipant is bound to causer. But, for *The hammer hit the table,* this role is bound to an instrument role instead [see Table 1].

This distinction is specified in EJ1 by using a different, but closely related, argument structure construction. *He hit the table*, instantiated TransitiveCEAProfiledCause, which specifies that the profiledParticipant is bound to the causer role. *The hammer hit the table* is analyzed as instantiating a subcase of the argument structure construction: TransitiveCEAProfiledInstrument [see Figure 4]. This subcase **ignore**s the parent's specification that profiledParticipant is bound to causer, and instead specifies that it is bound to the instrument role of the ForceApplication action (i.e. the instrument the actor used to apply force to the actedUpon). But, these argument structure constructions are the same in all other respects.

The other constructions instantiated in this example are the same, with the exception of the particular NP that is bound to Declarative's subj constituent. Recall that Declarative specifies

---

[3] [Note alternative reading = non-agentive impact, e.g. The hammer fell and hit the table. This reading would also focus on the hammer's motion and contact, but the actor and actions would no longer be part of the conceptual picture. This reading would be analyzed as a different type of event, one which does not include an agentive causer.

that the meaning of its subj constituent is bound to profiledParticipant. Unification therefore results in co-indexation of the meaning of this NP construction with the semantic role bound to profiledParticipant, which in this case is the Instrument of ForceApplication thus noting that instrument is of type HAMMER, information not present in previous example. At the same time, we lose information that causer is MALEANIMATE, since this role is no longer explicitly expressed. However, the role itself is still part of the event description, indicating that the causer is still conceptually present.

The resultant semspec therefore would differ from Figure 6 in that profiledParticipant is co-indexed with an Instrument role rather than the causer role. Such a semspec indicates that the event should be simulated from the perspective of the instrument, not the causer.

Extensions

These examples just discussed differ as to which argument structure constructions they instantiate. But, in both cases, the argument structure construction composes with an instance of Declarative, indicating that in both cases the kind of speech act is the same, and the subject is also the discourse topic. Of course to apply to a broader range of linguistic data, these argument structure constructions must also compose with constructions that provide different discourse-related specifications.

Questions – Which table did he hit?

To illustrate how this can be done, consider the sentence *Which table did he hit?,* which differs from *He hit the table* in terms of the type of speechAct. Both sentences have the same subject, but differ as to discourse topic. However, the meanings of both sentences are similar with respect to the type of event being described and the perspective from which it is described.

Which table did he hit instantiates the same verb and argument structure constructions as *He hit the table*, as well as similar nominal constructions. When the instantiated constructions unify, the resultant semspec is very similar to that of *He hit the table.* Both specify the same eventType and profiledProcess, and in both profiledParticipants is co-indexed with causer. The key differences relate to the DiscourseSegment roles. Firstly, speechAct types are different [“wh-question” vs. “declarative”.] And secondly, the topic is co-indexed with different semantic roles in each case: for declarative, the topic is co-indexed with causer (and meaning of subj constituent) whereas in the wh-question the topic is co-indexed with affected (and the meaning of the extraposed constituent). The best-fit analyzer hypothesizes that the extraposed phrase can fill the affected role of the TransitiveCEAProfiledCause construction because there is a good form and meaning fit.

These analyses are possible, in part, because the EventDescriptor schema is defined such that the notion of topic is separate from both the subject constituent and the profiledParticipant role. In Declarative, these conceptual elements are all bound to one another, indicating a particular type of conceptual whole. But, because this particular combination is not represented by an atomic role it is also possible to write constructions in which these are not all bound, such as the question construction instantiated in *Which table did he hit*? EJ1 argument structure constructions are defined such that they can compose not only with declarative, but also question (and other types of speech act) constructions.

Passives -- The table was hit (by him)

The Passive is typically analyzed in relation to active, but the exact relation remains a topic of continuing linguistic research. In the analysis that we sketch here, actives and passives are treated as different families of constructions which are related through common semantics. The basic idea is to have passive constructions use the exact same schemas as their active counterparts, while inheriting their form constraints through the passive hierarchy.

Using EJ1, *The table was hit (by him)* is analyzed as instantiating a passive argument structure construction whose meaning is CauseEffectAction, and which specifies that the meaning of its verb constituent is one of ForcefulMotionAction. In these respects, this construction is the same as the TransitiveCEAProfiledCause argument structure construction that was instantiated in *He hit the table*. However, the passive argument structure construction specifies that the profiledParticipant role is bound to the affected role of the CauseEffectAction, not the causer role. Thus, both the passive and the active argument structure constructions specify the same type of event, and both have a verb constituent that elaborates the causal action within this event. But, they differ on which simulation perspective they specify, with active specifying that of the causer, and passive that of the affected.

Constructionally, the passive differs from active in terms of its constituents. Unlike TransitiveCEAProfiledCause, the passive argument structure construction does not have an NP constituent, but does have an optional prepositional phrase, whose meaning is bound to the causer role. In addition, passive has different verb constituent constraints, including the fact that the verb form is that of past participle.

The table was hit therefore instantiates a different (though semantically similar) argument structure construction than *He hit the table*. Most of the other instantiated constructions are the same for both examples, including Declarative, a HitPastTense verb construction, and an NP construction for ‘the table’. When these constructions are unified, the semspecs are also very similar, with the key difference that in *Which table did he hit?* the profiledParticipant and topic is the affected participant rather than the causer [see Table 1].

Control -- He wants to hit the table.

The strategy for handling control relations in ECG also relies on leveraging the power of the profiledParticipant role. The basic idea can be illustrated by a description of the strategy for analyzing the sentence *He wants to hit the table*, as follows.

First, define a set of control verbs, such as *want*, whose meaning can be defined as involving an additional 'event' role. The meaning of *want*, for example, would be represented as a 'wanting' process in which a wanter desires that some type of event take place.

Next, define an argument structure construction whose verb constituent is a control verb, and which adds another constituent that is itself an argument structure construction. Thus, there will be two EventDescriptors in the semspec, one associated with the control argument structure construction itself, and the other associated with the constituent argument structure construction. For the current example, this second argument structure construction would be TransitiveCEAProfiledCause, the same argument structure construction instantiated in *He hit the table.* In the resulting semspec, one for the control EventDescriptor describes the event of wanting something, and the second for the thing that is wanted, in this case the hitting event. In addition, the profiledParticipant role of the main EventDescriptor is bound to the control verb's protagonist role, which in the current example is the 'wanter' role.

For the subject control argument structure construction instantiated in this example, the profiledParticipant is also bound to the profiledParticipant of the constituent argument structure construction. Therefore, in the semspec for *He wanted to hit the table,* the profiledParticipants of each event descriptor are co-indexed, and are co-indexed with the wanter of the wanting event, and the causer of the 'hitting the table' event. And, because *He wanted to hit the table* also instantiates Declarative, profiledParticipant is also be bound to the meaning of the subj constituent, 'he'.

Analyzer/Workbench

The analyses described in this chapter are produced by a system called the constructional analyzer (Bryant 2008). Constructional analysis is the process of interpreting an utterance in context using constructions, and the analyzer maps an utterance onto an instantiated set of ECG constructions and semantic schemas. The design of the system is informed by the fields of construction grammar/functional linguistics, natural language processing and psycholinguistics, and the constructional analyzer is a cognitive model of language interpretation within the tradition of Unified Cognitive Science and NTL.

The power of the analyzer comes from combining constructions with best-fit processing. Best-fit is a term we use to describe any decision making process that combines information from multiple domains in a quantitative way. Thus best-fit constructional analysis is a process in which decisions about how to interpret an utterance are conditioned on syntactic, semantic and contextual information. Because constructions provide explicit constraints between form, meaning and context, they are well-suited to a best-fit approach (Narayanan and Jurafsky, 2001).

The best-fit metric computes the conditional likelihood of an interpretation given the grammar and the utterance and is implemented as a factored probabilistic model over syntax and semantics. The syntactic factor incorporates construction-specific preferences about constituent expression/omission and the kinds of constructional fillers preferred by each constituent. The semantic factor scores a semspec in terms of the fit between roles and fillers.

The constructional analyzer uses a psychologically plausible sentence processing algorithm to incrementally interpret an input utterance. Each partial (incremental) interpretation is a subset of the instantiated constructions and schemas that go into the final, intended interpretation. Intuitively, this means that there are a set of competing partial interpretations that are each trying to explain the prefix of the input that has been seen so far. The best-fit metric is used to focus the analyzer's attention on more likely partial interpretations.

The analyzer produces rich linguistic analyses for a range of interesting constructions including embodied semspecs for the various motion and force-application constructions designed by Dodge (2009). An array of syntactically interesting constructions are also easy to implement within the analyzer including constructions for passives, simple wh-questions, raising and radial category description of the ditransitive argument structure construction.

Although the English construction grammar is currently the most linguistically well-motivated grammar processed by the analyzer, the analyzer is not tied to English. It analyzes Mandarin child-directed utterances as well, using a Mandarin grammar. Productive omission is incorporated into the system and scored by the best-fit metric (Mok and Bryant 2006). Omitted arguments are resolved to a candidate set by a simple context model.

The analyzer also predicts differences in incremental reading time for reduced relative data. The factored syntactic and semantic model plays an important role in making the reading time predictions. The syntactic factor implements the syntactic bias for main verb interpretations over reduced relative interpretations, and the semantic model implements the good agent/good patient biases that differentiate the two experimental conditions (Bryant 2008). Crucially, the system is a model of deep semantic interpretation first, and it predicts reading time data as a byproduct of its design.

Conclusions

The main purpose of this Chapter is to introduce both the technical aspects of ECG and the scientific basis for the underlying NTL. By presenting detailed examples, we hope to convey how ECG and the related tools can be used for deeper linguistic analysis than is otherwise available.

Even from the fragments presented here, it is clear that ECG grammars employ a large number of constructions, contrary to traditional minimalist criteria for language description. This is partly a question of style, as one could define an equivalent formalism that had fewer, but more complicated (parameterized), constructions. More basically, we believe that constructional compositionality crucially depends on a deep semantics that captures the rich structure of human conceptual systems. The semantic poles of ECG constructions are based on our best understanding of the conceptual and communicative primitives. In addition, we suggest that the critical resource in language processing is time, not space. By having explicit, compositional, constructions for language variations we simplify grammar writing and analysis for the analyzer program and, we believe, this is true for human language processors as well.

The NTL project and the ECG grammar formalism are both undergoing continuous development and this chapter presents only a snapshot of one thread. Below, we outline some of the current areas of research. John Bryant's dissertation (Bryant 2008) contains more information on all of this work.

Constructions should be able to capture form-meaning relations at all levels of language use. We have extended ECG and the Analyzer to handle complex morphology, including Semitic, by constructing a parallel morphological analyzer that coordinates with the system described here. The idea of a common deep semantics linking various forms of time-locked input is, in principle, extendable to speech, intonation, and gesture.

There are now pilot implementations of two additional ECG primitives, *situations* and *maps*, which are needed to handle key CL mechanisms such as metaphors and mental spaces. As part of this extension we are also incorporating techniques for modeling language communities and social communication.

This chapter, and the bulk of NTL work, has focused on language recognition; modeling *production* brings in a wide range of additional issues of audience modeling, etc. Interestingly, the best-fit analyzer already does analysis-by-synthesis and would not require major redesign to generate the best surface form, given metrics on the desiderata.

One of the most ambitious current projects involving ECG is an attempt to model in detail how children acquire their early grammatical constructions. Because of its explicit linking of embodied conceptual structure to linguistic form, ECG seems to provide a uniquely appropriate foundation for such studies.

All inductive learning is statistical, but the NTL work differs from purely statistical studies in postulating some conceptual and grammatical primitives as the hypothesis space for learning. The conceptual primitives include all of the embodied concepts (including emotional, social, etc.) that the child brings to language learning. The grammatical prior consists of three basic assumptions:

a) The child knows many meaning (conceptual) relations
b) The child can recognize relations in language form (e.g. word order)
c) Grammatical constructions pair form relations with meaning relations

Since the primitive relations in both form and meaning are bounded, the learning problem for the child (and our computer models) is not intractable. Ongoing work by Nancy Chang and Eva Mok (Chang and Mok 2006) demonstrates that ECG-based programs can learn complex grammatical constructions from labeled input, even for languages like Mandarin that allow massive omission.

We have said relatively little in this chapter about the neural realization of our Neural Theory of Language. A great deal of ECG-based linguistic analysis can be done without explicit neural considerations, just as much biology can be done without chemistry. But the neural perspective is crucial for many aspects including developing testable models and conceptual primitives. Our idea of how all the levels integrate is presented in (Feldman 2006), as part of Unified Cognitive Science.

As we suggested at the beginning, this growing Unified Cognitive Science presents opportunities of new possibilities for deep semantic grammars for theoretical, scientific and practical uses. When we add powerful tools, such as those described in this chapter, the future of linguistic analysis looks very promising.

References

Bailey, David. N. Chang, J. Feldman, S. Narayanan (1998). Extending Embodied Lexical Development. Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society COGSCI-98, Madison.

Barlow, Michael, and Suzanne Kemmer (eds.) 2000. *Usage-Based Models of Language*. Stanford: CSLI Publications.

Bergen, Benjamin and Nancy Chang. 2005. Embodied Construction Grammar in *Simulation-Based Language Understanding*. In Jan-Olastman and Miriam Fried (Eds.), Construction Grammars: Cognitive grounding and theoretical extensions.

Boroditsky, Lera. 2000. Metaphoric Structuring: Understanding time through spatial metaphors. Cognition, 75(1), 1-28.

Bryant, John 2008. *Best-fit Constructional Analysis*. Doctoral Dissertation, UC Berkeley.

Bybee, Joan, and Paul Hopper (eds.) 2001. *Frequency and the Emergence of Linguistic Structure*. Amsterdam and Philadelphia: John Benjamins. Typological Studies in Language 45.

Chafe, Wallace. 1998. Language and the Flow of Thought. In Michael Tomasello (ed.), *The New Psychology of Language: Cognitive and Functional Approaches to Language Structure*, 93-111. Mahwah, NJ and London: Erlbaum.

Chang, Nancy and Eva Mok. 2006. Putting Context in Constructions. *The Fourth International Conference on Construction Grammar* (ICCG4), Tokyo, Japan.

Croft, William. 2001. *Radical Construction Grammar: Syntactic Theory in Typological Perspective*. Oxford: Oxford University Press.

Dodge, Ellen. 2009. *Conceptual and Constructional Composition*. Doctoral Dissertation, UC Berkeley.

Dodge, Ellen, and George Lakoff. 2005. Image Schemas: From Linguistic Analysis to Neural Grounding. In Beate Hampe (ed), *From Perception to Meaning: Image Schemas in Cognitive Linguistics,* 57-91. Berlin and New York: Mouton de Gruyter.

Fauconnier, Gilles. 1985. *Mental Spaces: Aspects of Meaning Construction in Natural Language*. Cambridge, MA and London: MIT Press/Bradford.

Fauconnier, Gilles, and Mark Turner. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.

Feldman, Jerome. 2006.*From Molecule to Metaphor: A Neural theory of Language*. MIT Press

Feldman, Jerome. and Narayanan, Srini. (2004). Embodied Meaning in a Neural Theory of Language. *Brain and Language*, 89, 385-392.

Fillmore, Charles J. 1982. Frame Semantics. In Linguistic Society of Korea (eds) *Linguistics in the Morning Calm.* Seoul, Hanshin.

-----. 1988. The Mechanisms of "Construction Grammar". *Proceedings of the Annual Meeting of the Berkeley Linguistics Society* 14:35-55.

Fried, Mirjam, and Hans C. Boas (eds.) 2005. *Grammatical Constructions: Back to the Roots*. Amsterdam and Philadelphia: John Benjamins. Constructional Approaches to Language 4.

Gallese, Vittorio. and Lakoff, George. (2005). The Brain's Concepts: The Role of the Sensory-Motor System in Reason and Language. Cognitive Neuropsychology.

Glenberg, Arthur. & M. Kaschak (2002). Grounding language in action. *Psychonomic Bulletin & Review, 9,* 558-565.

Goldberg, Adele E. 1995. *Constructions: A Construction Grammar Approach to Argument Structure*. Chicago and London: University of Chicago Press.

-----. 2006. *Constructions at Work: The Nature of Generalizations in Language*. Oxford: Oxford University Press.

Hauk, O., I. Johnsrude, & F. Pulvermüller (2004). . Somatotopic representation of action words in human motor and premotor cortex. *Neuron, 41(2),* 301-7.

Johnson, Mark. 1987. *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. Chicago and London: University of Chicago Press.

Lakoff, George. 1987. *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. Chicago and London: University of Chicago Press.

Lakoff, George, and Mark Johnson. 1980. *Metaphors We Live By*. Chicago and London: University of Chicago Press.

-----. 1999. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. New York: Basic Books.

Lakoff, George. 2008. *The Political Mind*.

Lakoff, George, and Rafael E. Núñez. 2000. *Where Mathematics Comes From: How the Embodied Mind Brings Mathematics into Being*. New York: Basic Books.

Langacker, Ronald. 1976. Semantic Representations and the linguistic relativity Hypothesis. *Foundations of Language* 14:307-357.

-----. 1987. *Foundations of Cognitive Grammar*, vol. 1: Theoretical Prerequisites. Stanford: Stanford University Press.

-----. 2005. Construction Grammars: Cognitive, Radical, and Less So. In Francisco J. Ruiz de Mendoza Ibáñez and M. Sandra Peña Cervel (eds.), *Cognitive Linguistics: Internal Dynamics and Interdisciplinary Interaction*, 101-159. Berlin and New York: Mouton de Gruyter. Cognitive Linguistics Research 32.

-----. 2008. *Cognitive Grammar: A Basic Introduction*. New York: Oxford University Press.

Matlock, Teenie. 2004. The Conceptual Motivation of Fictive Motion. In Günter Radden and Klaus-Uwe Panther (eds.), *Studies in Linguistic Motivation*, 221-248. Berlin and New York: Mouton de Gruyter. Cognitive Linguistics Research 28.

Matlock, Teenie, Michael Ramscar, and Lera Boroditsky. 2005. The Experiential Link Between Spatial and Temporal Language. *Cognitive Science* 29:655-664.

Mok, Eva and John Bryant. 2006. A Best-Fit Approach to Productive Omission of Arguments. (*BLS-32*), Berkeley, CA.

Narayanan, Srini. (1999). Moving Right Along: A Computational Model of Metaphoric Reasoning about Events. *Proceedings of the National Conference on Artificial Intelligence AAAI-99*. Orlando, Florida.

Narayanan, S. and D. Jurafsky (2001). A Bayesian Model Predicts Parse Preferences and Reading Times in Sentence Comprehension, *Neural Information Processing Systems (NIPS2001)*.

Regier, T. 1996. The Human Semantic Potential: Spatial Language and Constrained Connectionism. Cambridge: MIT Press.

Talmy, Leonard. 1972. Semantic Structures in English and Atsugewi. Ph.D. diss., University of California, Berkeley.

-----. 1983. How Languages Structure Space. In H. Pick and L. Acredolo (eds) *Spatial Orientation: Theory, Research and Application*, 225-282. New York: Plenum Press.

-----. 1996. Fictive Motion in Language and “Ception”. In Paul Bloom, *et al.* (eds.), *Language and Space*, 211-276. Cambridge, MA and London: MIT Press/Bradford.

Tettamanti, M., Buccino, G., Saccuman, M.C., Gallese, V., Danna, M., Scifo, P., Fazio, F., Rizzolatti, G., Cappa, S.F. and Perani, D. (2005). Listening to action-related sentences activates fronto-parietal motor circuits. J Cogn. Neurosci. 17: 273-281.

Tomasello, Michael. 2003. *Constructing a Language: A Usage-Based Theory of Language Acquisition*. Cambridge, MA and London: Harvard University Press.

Ziemke, Tom, Jordan Zlatev, and Roslyn M. Frank (eds.) 2007. *Body, Language and Mind*, vol. 1, *Embodiment*. Berlin and New York: Mouton de Gruyter. Cognitive Linguistics Research 35.1.

Brief Biographies

Jerome A. Feldman is a Professor of Computer Science and Cognitive Science at the University of California and a research scientist at ICSI in Berkeley. He received the PhD from Carnegie-Mellon in 1964 and previously taught at Stanford and Rochester.

Ellen K. Dodge is a PhD candidate in Linguistics at the University of California, Berkeley. Her main area of research is in cognitive linguistics, and includes work on image schemas, meaning representation, construction grammar, and compositionality.

John Bryant is a researcher at the International Computer Science Institute. He received his PhD in computer science at the University of California at Berkeley in 2008. His research investigates precise, cognitively motivated implementations of language interpretation using constructions.