Chapter 3

Embodied Construction Grammar

What is an idea? It is an image that paints itself in my brain.
— Voltaire

My very photogenic mother died in a freak accident (picnic, lightning)...
— from Lolita, Vladimir Nabokov

Our current enterprise requires a theory of language structure that is compatible with the constraints set forth so far: it must be consistent with the notion of embodiment in language use; it must be formally described and it must be learnable. Embodied Construction Grammar (ECG) is a computational formalism intended to fill this tall order. This chapter provides a general introduction to the formalism, beginning with an overview in Section 3.1 of its role in supporting a broader simulation-based model of language understanding. Section 3.2 and Section 3.3 provide more detailed descriptions of the two basic ECG primitives, *schemas* and *constructions*.

As a relative newcomer in the world of grammatical theories, the ECG formalism remains under active development. Since its original formulation (primarily in Bergen & Chang 2005, with a shorter description in Chang *et al.* 2002a), it has been extended to address various (cross)linguistic, computational and psychological challenges.[1] The version presented here is motivated primarily by the demands of the language acquisition model and thus does not attempt to include all of these

[1] See `http://ecgweb.pbwiki.com` for current developments and publications.

extensions; rather, I focus on the core features needed for representing constituent structure and relational constraints, both within and across the paired domains of form and meaning. This version is sufficient for our current focus on English constructions appearing in early child-directed speech, and it also supports the detailed models of language understanding (Chapter 4) and language learning (Chapter 5) to follow.

3.1 Overview

Embodied Construction Grammar is a grammar formalism designed to support an embodied model of language use. We take as starting point many insights from the construction-based family of approaches outlined in Chapter 2 (Goldberg 1995; Kay & Fillmore 1999; Lakoff 1987; Langacker 1991; Croft 2001). Foremost among these is the observation that linguistic knowledge of all sizes and degrees of specificity, from morphemes to multi-word idioms, can be characterized as **constructions**, or pairings of form and meaning. These form-meaning pairings crosscut traditional levels of linguistic analysis (e.g., phonological, morphological, syntactic, semantic and pragmatic). Along with other construction grammarians, we assume that language users exploit constructions to communicate: during comprehension, to discern from a particular utterance a collection of interrelated conceptual structures; and during production, to generate a surface form that expresses some communicative intent.

ECG is distinguished by its focus on precisely *how* constructions facilitate communication. From this **process-oriented** perspective, it is not sufficient to specify constructional mappings between form and meaning; we must also describe how such mappings interact with world knowledge and the surrounding communicative context to support processes of meaningful language use. These structures and processes are further required to be computationally explicit; that is, they must be described formally enough to support computational implementation.

The current discussion is restricted to language comprehension, although the basic ideas of the approach taken here have analogues in language production. The key functional criterion for successful comprehension is that it allows the hearer to react appropriately to an utterance, whether with language (*e.g.*, by answering a question or responding to a comment), some other kind of action (*e.g.*, by complying with an order or request), or some set of updates to the current belief state. A hearer must therefore glean both the basic propositional content of the utterance and the speaker's intended meaning in context. The difficulties of this task have been amply documented: speakers imply, hedge and underspecify; they obey (and flout) Gricean convention and take poetic license; and even the most prosaic of utterances may be rife with ambiguity and polysemy. While such phenomena fall well beyond the scope of our immediate ambitions, they effectively illustrate how meaning arises from the integration of disparate kinds of information, including constructional mappings, embodied and world knowledge, social and communicative goals, and the dynamically evolving situational context.

The Simulation Hypothesis given in Chapter 1 provides a means of characterizing the interaction between specifically linguistic information (encoded by constructions) and domain-general structures and processes. Figure 3.1 provides an overview of the model:

- Linguistic knowledge consists of a repository of **constructions** that express generalizations linking the domains of **form** (typically, phonological schemas) and **meaning** (conceptual schemas). The meaning domain encompasses a variety of cognitive structures motivated by perceptual, motor, social and other kinds of embodied experience.

- Each **utterance** can be seen as instantiating a set of constructions (and their associated forms and meanings) in a particular **communicative context**. In comprehension, a **constructional analysis** process takes an utterance and its context and determines which constructions the utterance instantiates. The product of analysis is the **semantic specification** (or **semspec**), which specifies which embodied schemas are evoked by the utterance and how they are related. The semspec may undergo a further **contextual resolution** process that links its semantic components to specific contextually available referents.

- The **simulation** process takes as input the (resolved) semspec exploits embodied representations to simulate (or enact) the specified events, actions, objects, relations and states with respect to the current context. The inferences resulting from simulation shape subsequent processing and provide the basis for the language user's response.

The simulation-based architecture has some specific representational consequences, chief among which is that constructions (*i.e.*, "pure" linguistic knowledge) need not bear the entire inferential burden alone, but rather need specify only enough information to run a simulation using sensorimotor and cognitive structures. This division of labor reflects a fundamental distinction between conventionalized, schematic meanings that are directly associated with linguistic constructions, and indirect, open-ended inferences that result from detailed simulation. In effect, constructions provide a limited interface through which the discrete tools of symbolic language can approximate and exploit the multidimensional, continuous world of action and perception.

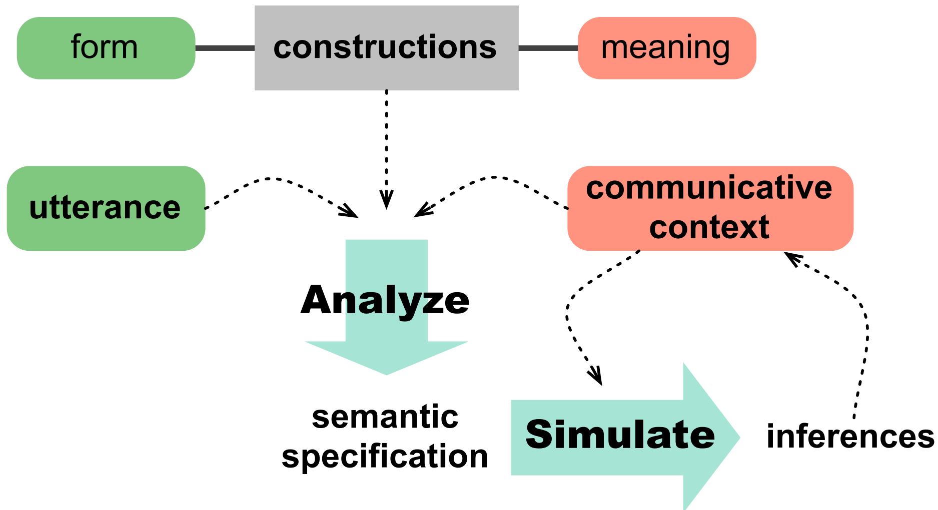

Figure 3.1. Overview of the simulation-based language understanding model: Constructions play a central role as the bridge between phonological and conceptual knowledge, supporting processes of constructional analysis, contextual resolution and embodied simulation.

The remainder of this section introduces the ECG formalism and its use as an interface between surface form and mental simulation, focusing on the representational requirements for understanding a simple example, *Harry ran home*. Section 3.1.1 casts the task in terms of the input sentence and corresponding output simulation, and Section 3.1.2 defines ECG structures that capture basic form, meaning and constructional information relevant to the example. The motivating concern here is to illustrate as simply as possible how the ECG formalism satisfies the broader architectural requirements laid out above. The structures defined are thus intended not to support the most general or linguistically nuanced investigation of the example sentence, but rather to demonstrate how a particular set of linguistic claims might be expressed in ECG. Section 3.1.3 addresses the representational challenges posed by more complex examples; the formal tools available in ECG for meeting these challenges are summarized in Section 3.1.4.

3.1.1 From sentence to simulation

Consider the simple English sentence *Harry ran home.* What surface cues are available to the addressee, and what inferences might be drawn from these? Let us assume that the sentence is pre-segmented into words (regardless of whether it is written or spoken). The surface cues available include, minimally, the individual word forms (*Harry*, *ran* and *home*) and the order in which they appear. We may also include some indication of intonational contour; in its written form the sentence ends with a period, suggesting a neutral or declarative contour for its spoken counterpart.

Based on this limited set of surface cues, a speaker of English is likely to infer that a running event took place sometime prior to speech time; that the runner is an individual named Harry (by default, a human male); and that the goal of running is someone's home (by default, the runner's).

Many other inferences may also be available, though potentially less certain or salient: the runner presumably has legs (at least two), expends energy to move his legs (along with the rest of his body), and starts out at some location other than home. These inferences are not specifically linguistic, but depend rather on general knowledge about runners and running events and homes. Moreover, they are contingent upon the particular referents chosen for *Harry* and *home*: specific interpretive contexts (say, the stipulation that Harry is a baseball player, or that Harry is a horse) might alter the most likely referents of these words, the inferred relationship between them, and the nature of the running motion itself.

Within the framework described here, the bulk of these inferences arise from a mental simulation involving the motor schema for running. Computationally, we represent this simulation using the executing schema formalism described in Section 2.3.3, shown in Figure 3.2 for the relevant running event. The RUN x-schema is similar to the WALK x-schema shown earlier: it is an iterative, energy-consuming process with various parameters, preconditions and effects. Given a particular set of parameters, a simulation based on this x-schema yields a slew of detailed inferences about the event's temporal, causal and interactional structure.

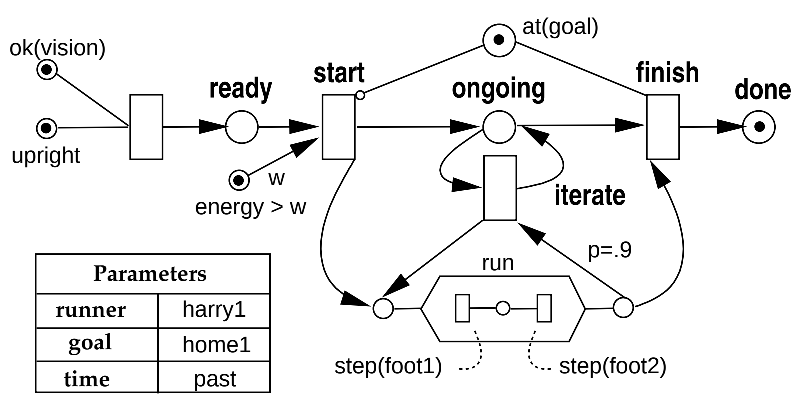

Parameters	
runner	harry1
goal	home1
time	past

Figure 3.2. An executing schema for running representing a simulation of *Harry ran home*. Tokens in the places corresponding to the conditions done and at(goal) indicate that the goal location has been reached. The parameters harry1 and home1 refer to contextually chosen referents of *Harry* and *home*, respectively, while past indicates that the event takes place before speech time.

We can thus narrow the specifically linguistic portion of language understanding to the following task: given a set of surface cues (corresponding to an utterance) and discourse context (corresponding to its meaning), produce a set of simulation parameters (*i.e.*, the semspec). As noted above, this task divides further into two interrelated processes: constructional analysis, which finds a set of constructions (and associated semspec) that accounts for the input forms; and contextual resolution, which finds a set of contextually appropriate referents for use in the simulation. In this case, the analysis and resolution processes should ideally produce a semspec for a running event with a runner (a referent of *Harry*), a goal location (a referent of *home*), and the time of running (in the past, *i.e.*, prior to speech time); the particular referents resulting from resolution are denoted in Figure 3.2 by the simulation parameters harry1 and home1.

Both of these processes must cope with ambiguity and uncertainty in the input: a given grammar may license many constructional analyses of an utterance, offering competing constituent structures, or varying in degree of constructional specificity; these analyses may in turn afford many contextually licensed resolutions. They thus require some means of evaluating and choosing among candidate analyses and resolutions. Chapter 4 will describe these processes and their associated evaluation strategies in detail. For the purposes of illustrating the ECG formalism, the focus here will be on the representational requirements for supporting a single, simple and unambiguous analysis of the example sentence. The structures defined below thus correspond to an analysis that has minimal internal structure, with no constituent corresponding to the predicate *ran home*; and that is partially lexically specific, treating *home* as a fixed element of the expression, but allowing variation in its other constituents (thus generalizing to sentences like *Theodore sauntered home* and *The dragon is flying home*).[2] More general candidate analyses are considered in Section 3.1.3.

3.1.2 Running example in ECG

Our example analysis, depicted graphically in Figure 3.3, involves three lexical constructions, one for each word in the sentence, along with one phrasal construction, the MOTION-HOME construction. These are shown in the center column of the figure, with arrows from the phrasal construction to the lexical constructions indicating constituency, similar to that in a traditional syntactic parse tree. Each construction also links the form domain (on the left) with the meaning domain (on the right). Here, each lexical construction links a word with a particular conceptual schema, while the phrasal construction links a word order relations (indicated as arrows on the dotted schematic time line) with a set of semantic identity relations (indicated as double-headed arrows).

Such a pictorial representation necessarily elides much detail. Each of the structures shown — including forms, meanings and constructions — is not just an abstract symbol, but rather corre-

[2]This treatment is motivated in part by the idiosyncratic behavior of *home* as a locative noun that can also serve as a path expression (compare with **Harry ran school*).

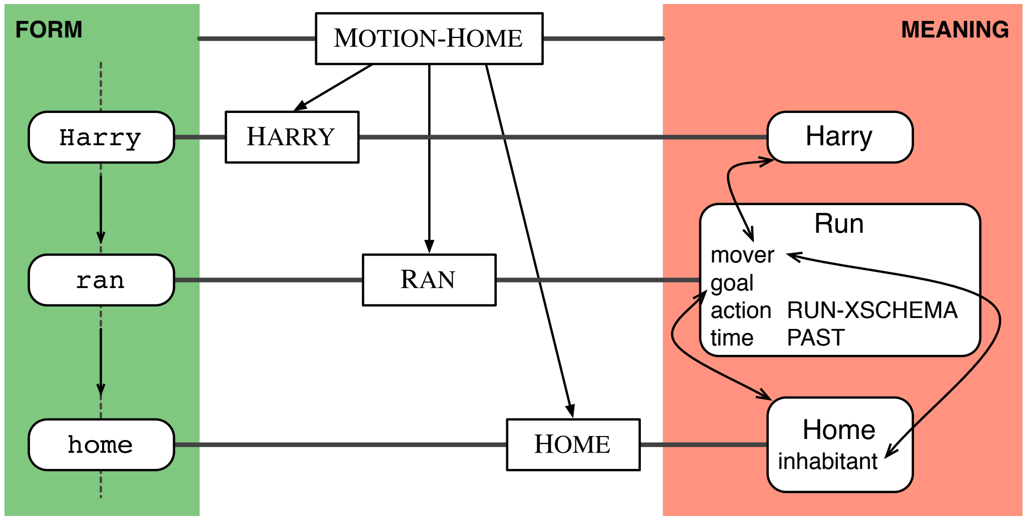

Figure 3.3. A lexically specific analysis of *Harry ran home*. The central boxes depict the constructs involved, each linking schemas in the domains of form (left) and meaning (right). The MOTION-HOME construction asserts constraints on its constituents, including ordering constraints on their forms and identity relations on their meanings.

sponds to a coherent chunk of knowledge within a broader, interconnected (and interdefined) inventory of structures needed to support language use. Below these structures and processes are defined more precisely using the ECG formalism. Our goal is to provide an informal guided tour of the **schema** formalism (for forms and meanings) and the **construction** formalism (for form-meaning maps); see Section 3.1.4 for a detailed summary of the formal tools available in ECG.

The schemas most relevant to our analysis are shown in Figure 3.4. In brief, schemas may have **roles**, corresponding to features or semantic roles (participants, locations, etc.), which may be constrained to be of a specific type (using a colon) or have a particular value (using the $\longleftarrow$ notation). Schemas may also be **subcases** of other schemas, inheriting their roles and constraints.

Summarizing the information expressed by the schemas in Figure 3.4:

- Word has a role phon that is a Phonological-String and a role orth that is an Orthographic-String.
- Harry is a kind of Human whose (inherited) sex role is specified as MALE.
- Home is a kind of Location. Home has a role inhabitant that is Animate.
- SelfMotion has a role mover that is an Entity, a role goal that is a Location, a role path that is an SPG (Source-Path-Goal schema, defined below), a role action that is an X-schema, and a role time that is a Time. Its goal role and mover roles identified, respectively, with the path role's subsidiary goal and trajector.
- Run is a SelfMotion whose (inherited) action role is specified as RUN-XSCHEMA.

The schema formalism is flexible, in that the same notations are deployed to represent, in Figure 3.4(a), general knowledge about words (here, that they are associated with orthographic and

(a)

```
schema Word
  roles
    phon : Phonological-String
    orth : Orthographic-String
```

(b)

```
schema Harry
  subcase of Human
  constraints
    sex ⟵ MALE
```

```
schema Home
  subcase of Location
  roles
    inhabitant : Animate
```

```
schema SelfMotion
  roles
    mover : Entity
    goal : Location
    path : SPG
    action : Xschema
    time : Time
  constraints
    goal ⟷ path.goal
    mover ⟷ path.trajector
```

```
schema Run
  subcase of SelfMotion
  constraints
    action ⟵ RUN-XSCHEMA
```

Figure 3.4. ECG representations of the schemas used in *Harry ran home*, capturing (a) a category for word forms; and (b) conceptual categories contributing to the meaning of the sentence.

phonological forms); and, in Figure 3.4(b), a variety of conceptual structures evoked by the sentence. The formalism is also concise, since many schemas are subcases of others and need only specify any new (non-inherited) roles and constraints. Thus, though the Human schema is not shown, we can assume that its definition is responsible for licensing the inferences that instances of its subcase Harry schema can move, act as a source of energy, have intentions, experience feelings and do whatever else humans do. Likewise, since the Run schema is defined as a subcase of SelfMotion (self-propelled motion), it need only specialize its action role with the appropriate x-schema.

We turn next to the construction definitions needed for the example, beginning with the lexical constructions shown in Figure 3.5. These exploit similar notational devices as the schema definitions, but they include additional structures that reflect their function of mapping form and meaning. The blocks labeled **form** and **meaning** stand for their two linked domains, or **poles** (corresponding to the left and right arrows in Figure 3.3).. These poles list the roles and constraints (if any) within each domain; they should also be considered special roles of the construction, which can themselves be type-constrained. The Ran construction can thus be summarized as follows:

- Ran is a kind of SelfMotion-Verb.
- Ran has a form pole of type Word with its orth role specified as the orthographic string "ran".
- Ran has a meaning pole of type Run with its time role specified as PAST.

The other two constructions likewise link their form poles (words with specialized orthographic or phonological values) to meaning poles that specialize one of the schemas defined above. Both are also subcases of the REF-EXPR (referring expression) construction, a general construction corresponding to the traditional noun phrase. (Details of this and the SELFMOTION-VERB construction are not important for current purposes; see Section 3.3.)

```
construction HARRY
  subcase of REF-EXPR
  form : Word
    orth ⟵ "Harry"
    phon ⟵ /heriʸ/
  meaning : Harry
```

```
construction RAN
  subcase of SELFMOTION-VERB
  form : Word
    orth ⟵ "ran"
  meaning : Run
    time ⟵ PAST
```

```
construction HOME
  subcase of REF-EXPR
  form : Word
    orth ⟵ "home"
  meaning : Home
```

Figure 3.5. ECG representations of the lexical constructions used in *Harry ran home*, each linking a specific form with embodied meanings and constraints.

The final construction needed for the analysis, shown in Figure 3.6, combines the lexical constructions into a cohesive gestalt. Like the lexical constructions already described, the MOTION-HOME construction has blocks for describing its form and meaning poles. It also, however, has a block labeled **constituents** for describing its internal constituent structure (as depicted in by the vertical arrows between constructions in Figure 3.3). This block, analogous to the block labeled **roles** in schema definitions, lists its constituent constructions by their local names and associated type constraints. Each of these constituents is itself a construction (or, more precisely, an instance of a construction), with form and meaning poles that are accessible for reference within the larger construction using a subscripted f and m on the relevant constituent name. Moreover, those poles themselves are (instances of) schemas, whose internal roles are also accessible for reference, where dotted notation (x.y) allows reference to accessible structures through a chain of structures. Additional notations are used to impose the form and meaning relations depicted in Figure 3.3: the `before` relation constrains constituent word order; the **identification** operator (indicated by the double-headed arrow $\longleftrightarrow$) asserts semantic identity relations; and the keyword **self** allows reference to the structure being defined (here, the MOTION-HOME construction; the subscripted version **self**$_m$ thus refers to the meaning pole of this construction).

The construction definition in Figure 3.6 can be restated as follows:

- MOTION-HOME has a constituent r that is a REF-EXPR, a constituent s that is a SELFMOTION-VERB, and a constituent h that is a HOME.
- The form pole of r comes before the form pole of s, which comes before the form pole of h.
- MOTION-HOME has a meaning pole of type SelfMotion, which is identified with the meaning pole of s. This structure also has a role mover that is identified with the meaning pole of m, and a role goal that is identified with the meaning pole of h.

construction MOTION-HOME
constituents
r : REF-EXPR
s : SELFMOTION-VERB
h : HOME
form
r_f `before` s_f
s_f `before` h_f
meaning : SelfMotion
self$_m$ $\longleftrightarrow$ s_m
s_m.mover $\longleftrightarrow$ r_m
s_m.goal $\longleftrightarrow$ h_m
h_m.inhabitant $\longleftrightarrow$ r_m

Figure 3.6. ECG representations of the clausal MOTION-HOME construction, associating word order relations with a self-motion scene.

- The meaning pole of h has a role inhabitant that is identified with the meaning pole of r.

Together, these schema and construction definitions convey the basic information needed to analyze our example sentence. The notational devices provided by the ECG formalism make it possible to express both the relevant domain-specific structures (in form and meaning schemas) and their cross-domain associations (in constructions), making explicit the structures and processes shown in Figure 3.3. Besides elaborating the internal structures of the various schemas and constructions, the formal definitions codify inheritance relations among the structures and allow each structure to specify its locally relevant constraints.

These relations and constraints are combined during the analysis process to produce the constructional graph and its associated set of interrelated schemas. The Run schema illustrates how information from disparate structures is combined: as shown in Figure 3.3, it includes roles inherited from SelfMotion; its action role is specified as RUN-XSCHEMA by the Run schema itself; its time role is specified as PAST by the RAN construction; and its mover and goal roles are bound by the MOTION-HOME construction to the Harry and Home schemas, respectively. This set of bound schemas — *i.e.*, the semspec — supplies the parameters for the requisite simulation using the run x-schema shown in Figure 3.2. The rich set of inferences, encompassing not just the semantic relations expressed by the semspec but also the detailed, context-dependent motor and perceptual consequences of the simulated running event, constitute the meaning of *Harry ran home*.

3.1.3 Beyond *Harry ran home*

We have now illustrated how some simple forms, meanings and constructions are defined using the ECG formalism. These structures support one possible analysis (and corresponding embodied simulation) of our example sentence. This analysis, though pedagogically convenient, is by no means the only one possible, nor necessarily the best. Alternative analyses might involve purely cosmetic variations (perhaps with a more perspicuous choice of names for the relevant schemas, roles and constituents); or they might require more substantive changes to constructional level of specificity or to inheritance relations among structures. For example, we could replace the MOTION-HOME construction with the even more lexically specific HARRY-RAN-HOME construction (with constituents typed as instances of HARRY, RAN and HOME). Such a construction, while of more limited utility than the MOTION-HOME construction defined earlier, would result in the same set of schemas and bindings; *i.e.*, it would be *informationally equivalent* with respect to simulation.

The space of possible analyses grows more daunting as we look beyond our introductory example. Consider the minimal variants shown in (3–1):

(3–1) a. {The boy / He / My oldest son} ran home.
b. Harry {is running / will walk / has flown} home.
c. Harry ran {out / here / *house / to the store / around the house}.

These sentences could be analyzed as having the same high-level constituent structure as our example. In fact, given some additional constructions to license the wider range of expressions for the relevant constituents, the variants in (3–1a) (referring expressions) and (3–1b) (tense and aspect constructions) could make use of the MOTION-HOME construction. The variants in (3–1c), however, necessitate a more general treatment of path expressions, with subcases corresponding to directional particles (like *out*), spatial prepositional phrases (like *to the store* or *around the house*), and locative expressions that express a path (like the locative pronoun *here*).[3] (A number of the requisite constructions are defined in Section 3.3.)

Different communicative functions may also be expressed with respect to our main event, as exemplified in (3–2). These include questions about the veracity of the event as a whole or its individual role-filler values, as well as commands and focus-manipulating constructions:

(3–2) a. Did Harry run home?
b. Who ran home?
c. Where did Harry run?
d. Run home, Harry!
e. It was Harry who ran home.

Venturing past the confines of Harry's running activities, we observe instances of some other basic clausal types in English, corresponding to the main argument structure constructions discussed by Goldberg (1995):

[3] The idiosyncratic behavior of *home* (compared with other common nouns) may justify its inclusion in this last category.

(3–3) a. Mary threw a ball. *(transitive scene)*
b. Mary threw a ball to me. *(caused motion scene)*
c. Mary threw me a ball. *(transfer scene)*

All of these variations raise issues of linguistic representation deliberately sidestepped in our initial example. In particular, while many potential grammars might be informationally equivalent with respect to these examples (in the sense mentioned above of producing the same semspec), they may vary along many other dimensions. A standard criterion employed in linguistic analysis favors grammars that capture shared structure and (relatedly) generalize to further examples; for example, a general SUBJECT-PREDICATE construction is typically assumed to capture commonalities across the sentences in (3–3), such as possible constraints on agreement and word order. In some circumstances, however, one might prefer a grammar that involves fewer levels of analysis (*e.g.*, one in which the examples in (3–3) are analyzed using three wholly independent clausal constructions).

The ECG formalism is, by design, expressive enough to support a wide range of linguistic claims and analyses, but agnostic about which of these should be considered optimal. Indeed, as noted earlier, nothing in the formalism precludes multiple competing or compatible analyses. In principle, the best analysis in a given situation may depend on a host of factors, including dynamic aspects of the analysis and simulation context; the applicability of crosslinguistic, psychological, aesthetic or practical considerations; and, most relevantly for the current work, criteria based on processes of language learning. ECG's theoretical commitments thus apply mainly to the structural level — that is, to the *kind* of information that must be represented to capture the interdependence of linguistic structure, process and context inherent in our approach to language understanding.

3.1.4 The formalism at a glance

This section briefly summarizes the notational devices available in the ECG formalism; many of these have been introduced informally in the preceding sections, and all are illustrated with more examples in Section 3.2 and Section 3.3. See Chapter 4 for more formal definitions.

Computationally, the formalism draws most directly from constraint-based grammars of the kind discussed in Section 2.3.1 (Shieber 1986; Pollard & Sag 1994): the underlying representation is a typed feature structure with unification constraints and feature inheritance. But, as illustrated in the last section, the formalism bears a closer surface resemblance to object-oriented programming languages like Java, with special keywords and symbols for expressing relations and constraints among structures and allowing self-reference and reference to non-local structures.

The version of the formalism presented here simplifies many aspects of linguistic representation. As discussed in Section 2.2.1, constructions are like other complex radial categories in exhibiting graded membership and internal structure; the multiple inheritance hierarchy assumed here captures some but not all aspects of constructional category structure. In particular, graded category membership might be better expressed in a probabilistic framework (as discussed again in Chapter 9). The simpler inheritance-based organization used here, however, suffices for our current focus on structural relations among constructions.

Representational primitives

ECG includes formalisms for several representational **primitives** needed for capturing a range of linguistic phenomena; only the two most relevant of these are discussed here:

- A **schema** is the basic structured unit of representation for both the form and meaning domains. Each schema specifies relationships among a set of interdefined participant **roles**. Roles can be instantiated by particular values (or *fillers*). Form schemas provide information relevant to surface form (*e.g.*, associated phonological or orthographic strings, intonational information, temporal ordering), while meaning schemas help to specify parameters for embodied simulations.
- A **construction** is the basic linguistic unit that pairs elements and constraints in the form and meaning domains, or **poles**. Each construction has a **form** pole and a **meaning** pole, which can be constrained to instantiate specific form and meaning schemas, respectively. Some constructions also have internal **constituents** that are themselves constructions. Constructions may also have features encoding properties of the construction that do not reside solely in the form or meaning domain.

The other primitives are *spaces* and *maps*; see (Chang *et al.* 2002a; Mok *et al.* 2004) for discussion of how these are used to account for metaphorical mappings and mental space phenomena.

Relations among primitives

Structures may be related in several ways:

- **subcase**: Schemas and constructions are organized in multiple inheritance hierarchies, each a partial order induced by the **subcase** relation between a structure (schema or construction) and its more general **base** structure (or set of base structures), notated using the tag **subcase of** x. The subcase relation is similar to the *is-a* link used in semantic networks. The roles (and constituents, in the case of constructions) of each base structure are accessible to its subcases, and its constraints apply.
- **constituency**: Complex schemas and constructions may combine a set of simpler subsidiary structures. A schema may contain *roles* instantiating other schemas, and a construction may contain *constituents* instantiating other constructions. This relationship corresponds to a *part-whole* relation, allowing one schema or construction to include another.
- **evocation**: A schema may *evoke* (or activate) an instance of another schema x with the local identifier y, without implying either inheritance or constituency, using the notation **evokes** x

as y. This underspecification provides needed flexibility for building semantic specifications. In combination with the **self** notation (see below), it also allows one structure to be defined or raised to prominence against a background set of structures, formalizing the notion of *profiling* used in frame semantics (Fillmore 1982) and Cognitive Grammar (Langacker 1991).

- **containment**: In addition to constituency, a construction may be seen as containing its form and meaning poles, as well as any additional roles it introduces in the form, meaning or constructional domains.

Accessible structures

ECG includes notations for expressing several kinds of constraints. Arguments to these constraints must be *accessible* structures within the relevant definition, *i.e.*, one of the following:

- the structure itself, expressed using the keyword **self**;
- locally defined roles, constituents and evoked structures;
- inherited roles, constituents and evoked structures (*i.e.*, any structures accessible via the subcase relation);
- roles and constituents recursively accessible through other accessible structures, using a dotted "slot chain" notation to refer to a role y of a structure x as $x.y$; and
- the form and meaning poles of any accessible construction, including those of the structure itself or any of its constituents, using a subscripted f or m.

Constraints

The following constraint types are allowed:

Type (or **category**) constraints (indicated with a colon, as $x : y$) restrict x to be filled by an instance of schema y.

Binding constraints: ECG has two constraints that correspond to standard unification or coindexation. **Identification** constraints (indicated with a double-headed arrow, as $x \longleftrightarrow y$) cause fillers to be shared between x and y, thus indicating how roles of different structures involved in simulation are aligned. **Filler** constraints (indicated with a single-headed arrow, as $x \longleftarrow y$) indicate that the role x is filled by the element y (a constant value).

Ordering constraints: Temporal relations among form segments are notated using form constraints. In principle, the formalism can express any binary relation between intervals, including sequence, overlap, contiguity and others specified in Allen's (1984) Interval Algebra. Of these, the most common relations are **before** (for precedence) and **meets** (for immediate precedence). In the absence of any explicit order constraint, a weaker co-occurrence constraint holds among the forms of different constituents of the same construction.

Two other constraint types are allowed, as described elsewhere (Chang *et al.* 2002a; Chang *et al.* 2002b): *predicates*, a general means of expressing open-ended relational constraints; and *dynamic constraints* that allow other constraints to be asserted as holding only within a specific setting (temporal, spatial, etc.). These are not exploited in the current work.

* * *

The next two sections elaborate on the formalism's expressive capacity. For expository ease (and maximal relevance to the learning model), these focus on the representational demands posed by crosslinguistic patterns of lexicalization and grammaticization common in early child language.

3.2 Schemas

Embodied schemas generalize over a long tradition of representations in linguistic analysis. These include feature-based representations for phonological, morphological and syntactic variation, as well as the role-filler semantic relationships associated with case-marking and thematic roles. All of these are schematic structures based on structured relations over variable values. ECG schemas are distinguished from other feature-based representations in two key respects: (1) ECG schemas are explicitly associated with richer underlying structures: schema roles are viewed as supplying both parameters to the associated grounded representations and points of access for other schemas and constructions. Thus, they can function not only as symbols that name relationships (as features often do) but, more importantly, as an interface between structures at different levels of granularity and across different domains. (2) ECG schemas include additional notational devices that make it convenient to express both traditional role-filler bindings and the more complex relationships that arise within interdefined clusters of concepts.

The sections below describe the main varieties of embodied schemas relevant to the current work. These include schemas from the basic domains of form (Section 3.2.1) and meaning (Section 3.2.2), as well as a special class of schemas for capturing communicative context (Section 3.2.3).

3.2.1 Form

The canonical domain of linguistic form is sound. The crosslinguistically most common sound patterns are described abstractly in terms of phonological sequences, tone and stress, or more concretely as the acoustic patterns and the articulatory gestures that produce them; written languages also employ orthographic representations. Form in its most general sense, however, may extend to any kind of signifier, including manual and facial gestures, pictorial icons like those taught to non-human primates, and tactile patterns like those in Braille.

Relationships among forms vary by domain. Many of these include a temporal dimension (or a spatial dimension with a temporal correlate), and thus naturally afford certain relations among

```
schema Schematic-Form
```

```
schema Word
  subcase of Schematic-Form
  roles
    phon : Phonological-String
    orth : Orthographic-String
```

```
schema Complex-Word
  subcase of Word
  roles
    stem
    prefix
    suffix
  constraints
    prefix before stem
    stem before suffix
```

Figure 3.7. Form schemas: The Schematic-Form schema is the most general form schema; its (simplified) subcase Word schema has roles for phonological and orthographic representations.

temporal segments or intervals: *e.g.*, one segment may precede one another, either immediately or with intervening segments; segments may include other segments, possibly as initial or final segments; two segments may be temporally coincident (Allen 1984). Such interval relations are sufficient to describe many of those associated with concatenative morphology (namely, affixation) and syntax (word order). Other phenomena, like templatic morphology, vowel shifting and vowel harmony, may require domain-specific similarity- and variation-based relations, or even the coordination of multiple forms in different domains (such as the linked manual and facial gestures in some sign languages).

The ECG schema formalism can accommodate any of the forms and relations above. The current work makes use of a small subset of the possible forms and relations, focusing on orthographic information, as in the introductory example. Figure 3.7 shows some example form schemas, including the abstract Schematic-Form schema of which all other form schemas are subcases; a Word schema (repeated from Section 3.1.1) with roles for phonological and orthographic strings; and a Complex-Word schema that might be useful for languages with concatenative morphology.

3.2.2 Meaning

We take the scope of linguistic meaning to include all conceptualizable notions, from concrete entity, relation and action categories to the more abstract categories usually associated with grammar. It is thus appropriate that embodied schemas for meaning generalize over many proposed semantic representations. The basic notion — that of a schematic structure capturing relations among a set of roles that can be filled by variable values — is similar to the *semantic schema* assumed in Langacker's (1987) Cognitive Grammar and compatible with Lakoff's (1987) *idealized cognitive model* and Fillmore's (1982) *frame* (as discussed further below).

Figure 3.8 shows schemas representing some standard examples from the literature. As depicted here, both concrete physical categories like those in (a) and specific named individuals like those in (b) are represented as schemas, but the latter are also defined as subcases of a NamedEntity category. The action schemas in (c) specialize the SelfMotion schema defined earlier, and (d) includes some commonly grammaticized image schemas, such as Trajector-Landmark schema (spatial or attentional asymmetry), Container (enclosed or partially enclosed regions) and SPG (motion along a path) (Johnson 1987; Lakoff & Johnson 1980).[4]

(a)

```
schema Human
  subcase of Animate
  roles
    sex
    status
```

```
schema Girl
  subcase of Human
  constraints
    sex ⟵ FEMALE
    status ⟵ CHILD
```

```
schema Cafe
  subcase of Location, Business
  constraints
    product ⟵ COFFEE
```

(b)

```
schema Harriet
  subcase of Girl, NamedEntity
```

```
schema Starbucks
  subcase of Cafe, NamedEntity
```

(c)

```
schema Walk
  subcase of SelfMotion
  roles
    mover : Animate
  constraints
    action ⟵ WALK-XSCHEMA
```

```
schema Hop
  subcase of SelfMotion
  roles
    mover : Animate
  constraints
    action ⟵ HOP-XSCHEMA
```

(d)

```
schema Trajector-Landmark
  roles
    trajector
    landmark
    relation
```

```
schema Container
  roles
    interior
    exterior
    portal
    boundary
```

```
schema SPG
  roles
    source : Location
    path : Path
    goal : Location
    trajector : Entity
```

Figure 3.8. Meaning schemas: (a) entities; (b) named entities; (c) actions; (d) image schemas.

A key feature of semantic representation, associated most closely with the study of *frame semantics* (Fillmore 1982), is that many concepts are defined against a background set of related participants and props (or *frame*). The classic example of a culturally specific frame is that of the commercial transaction, *i.e.*, an exchange of goods and money between a buyer and seller. Different aspects of this inferentially rich backdrop are highlighted, or *profiled*, by lexical items like *buy*, *sell*, *price*, *cost* and *discount*. Langacker (1987) provides a geometric example exhibiting the same representational pattern: the concept associated with the word *hypotenuse* cannot be defined independently of

[4] As noted in 2.2.2, a variety of terms appear in the literature for related concepts. Alternative terminology more in line with Talmy's (2000) schematization might include the Figure-Ground schema (with roles figure and ground) instead of Trajector-Landmark and a Path schema (with roles figure, ground, path and manner) instead of SPG. Our purpose here is not to choose among schematizations but to illustrate that the formalism can accommodate either.

its base concept of a right triangle. The ECG schema formalism expresses this profile-base relation using the **evokes** notation, which allows a schema to evoke a background structure (corresponding to the background frame or base schema). Roles of this evoked structure can then be bound as appropriate; in particular, they may be bound to the current structure (denoted by the **self** notation). Schemas for Hypotenuse and its evoked schema Right-Triangle are shown in Figure 3.9. Note that this approach to profiling can also be used within the meaning pole of a construction definition, as illustrated earlier by the HOME construction in Figure 3.5.

```
schema Right-Triangle
 roles
   leg1 : LineSegment
   leg2 : LineSegment
   hyp : LineSegment
   angle1 : Angle
   angle2 : Angle
   angle3 : Angle
 constraints
   angle3 ⟵ 90
```

```
schema Hypotenuse
 evokes Right-Triangle as rt
 constraints
   self ⟵ rt.hyp
```

Figure 3.9. Profiling by evoking in ECG: The Hypotenuse schema **evokes** the Right-Triangle and binds itself to (or profiles) one of its roles.

Complex interactional patterns combine the notational devices illustrated above. For example, the word *into* involves a dynamic spatial relation in which one entity moves from the exterior to the interior of another. This situation is depicted informally in Figure 3.10; the corresponding ECG schema is shown in Figure 3.11. In brief, Into is defined as a particular kind of trajector-landmark relationship, one whose landmark is an instance of the Container schema and whose trajector is also the trajector of an evoked instance of the SPG schema. Identification constraints assert that the trajector's path takes it from the exterior to the interior of the container. (The same evoked schemas with a different set of bindings would be needed to express the meaning of *out of*.)

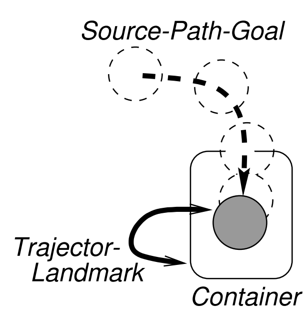

Figure 3.10. An iconic representation of the meaning of *into*, combining the simpler schemas Container, Trajector-Landmark and Source-Path-Goal.

```
schema Into
  subcase of Trajector-Landmark
  evokes SPG as spg
  roles
    trajector : Entity
    landmark : Container
  constraints
    trajector ⟷ spg.trajector
    spg.source ⟷ landmark.exterior
    goal ⟷ landmark.interior
```

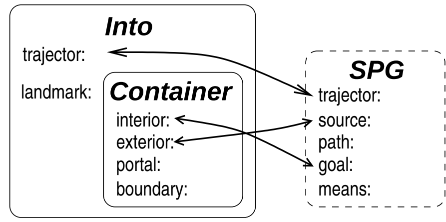

Figure 3.11. The Into schema, defined using the ECG formalism (left) and informally depicted as a linked set of schemas (right). The Into schema is defined here as a subcase of Trajector-Landmark that evokes an instance of the SPG schema (shown with a dashed boundary at right).

3.2.3 Function

Utterances are rooted in specific situational and discourse contexts and fulfill specific discourse functions: speakers pick out individuals, indicate whether they have been previously mentioned, make assertions and ask questions about them, and so forth (as demonstrated by the examples in (3–2) from Section 3.1.3). The ECG formalism defines several special schemas that allow explicit mention of the current context and invoke the basic communicative functions of *reference* and *predication*, as shown in Figure 3.12:

- The current discourse context, or **discourse space**, can be characterized in terms of a set of interlocutors and other salient participants, and the surrounding attentional and activity context (as in the DiscourseSpace schema).

- **Reference** is a basic communicative function in which the speaker evokes or directs attention to specific entities and events (the **referent** or set of referents). Languages differ in precisely what kinds of information speakers supply; as shown in the Referent schema, these typically include the referent's ontological category (*e.g.*, human, ball, picture), number and gender; its default level of *accessibility* (Lambrecht 1994) in the current context (active, accessible, inactive, unidentifiable, etc.); and any restrictions and attributions that apply to its open-class characteristics (such as size or color). The last role shown, the resolved-referent, is a special role filled by the actual resolved entity described by the schema.

- **Predication** is the relational counterpart to reference. Speakers make attributions and assert relations as holding of particular entities; and they locate, or ground, these relations (in time and space) with respect to the current speech context. The Predication schema provides roles for an overall scene (typically set by an argument structure construction) and its central means (typically set by verbal constructions); the predicated event's protagonist (main participant), event-structure (aspectual profile) and temporal or spatial setting; and any associated speech act type.

```
schema DiscourseSpace
  roles
    speaker
    addressee
    joint-attention
    activity
```

```
schema Referent
  roles
    category
    number
    gender
    accessibility
    restrictions
    attributions
    resolved-referent
```

```
schema Predication
  roles
    scene
    means
    event-structure
    time
    place
    speech-act
```

Figure 3.12. Schemas for describing communicative function

These structures play a central structuring role in grounding schemas with respect to the current context (by constraining the contextual resolution process described further in Chapter 4) and setting parameters for the simulation to be executed. But, as shown here, they can be schematized and represented using the ECG schema formalism, thus allowing them to mingle freely with other constructional meanings.

* * *

The examples above demonstrate the expressive capacity of the ECG schema formalism. Again, we make no strong claims about the particular schemas defined here, which might be modified or extended, for example, to support crosslinguistic variation, or reflect other task-specific assumptions. The formal tools they illustrate are, however, adequate for describing a wide range of form, meaning and function schemas, including those most relevant in early language learning.

3.3 Constructions

Constructions serve as the crucial symbolic link between the domains of form and meaning. Each construction is associated with two *domain poles*, representing these two domains. These poles are formally similar to ECG schemas: they may be typed as instances of existing schema types, and they may be further specialized with additional constraints. Constructions also resemble schemas in that they are defined within a multiple inheritance hierarchy, and they may have additional constituent structure, analogous to subsidiary schema roles. The ECG construction formalism thus employs many of the notations used for schemas, along with additional mechanisms for describing cross-domain associations. This section surveys the varieties of cross-domain mapping permitted in linguistic constructions (Section 3.3.1) and summarizes some basic construction types (Section 3.3.2 and Section 3.3.3).

3.3.1 Cross-domain mapping types

As previously described, constructions cut across traditional linguistic divisions: they vary in both size (from morphological inflections to intonational contours) and level of concreteness (from lexical items and idiomatic expressions to clausal units and argument structure patterns). Formally, however, the most salient difference among constructions is based on the presence or absence of internal constituent structure. A **simple** construction consists of only one cross-domain mapping. That is, though both the form and meaning domains may contain arbitrarily complex sets of schemas, the mapping between them itself lacks internal structure. In contrast, a **structured** (or **complex**) mapping includes at least one subsidiary cross-domain mapping, corresponding to a constituent.

Simple and structured mappings correspond canonically to lexical and phrasal (or clausal) constructions, respectively. Our introductory example in Section 3.1.1 included three simple, lexical constructions (for *Harry*, *ran* and *home*) and one structured mapping combining them. Figure 3.13 illustrates a similar pattern using constructions that license the phrase *into Starbucks*. The two lexical constructions in (a) correspond to the preposition *into* and the proper noun *Starbucks*. Each of these links a single form schema (a word with an orthographic constraint) with a single meaning schema (taken from Figure 3.8). In (b) we show the schematic PATH-EXPR (pairing a form of unspecified type with a path description) and its subcase PATH-PHRASE construction, an ECG version of the familiar prepositional phrase. Separate constituents denote the particular spatial relation involved (PATH-PREP) and a landmark referent; the form poles of these constituents are constrained to come in a particular order, and their meanings are aligned as indicated. Specifically, the Trajector-Landmark relationship denoted by the prepositional constituent rel has a landmark role (rel_m.landmark) that is identified with the referent of the other constituent (lm_m). The construction's overall meaning is also bound to the (evoked) spg role of the prepositional constituent (rel_m.spg), an instance of the SPG schema (as defined in Figure 3.11).

The distinction between simple and structured mappings applies to constructions of all levels

(a)

construction INTO-CXN
  subcase of PATH-PREP
  form : Word
    orth ⟵ "into"
  meaning : Into

construction STARBUCKS
  subcase of REF-EXPR
  form : Word
    orth ⟵ "Starbucks"
  meaning : Starbucks

(b)

construction PATH-EXPR
  form : Schematic-Form
  meaning : SPG

construction PATH-PHRASE
  subcase of PATH-EXPR
  constituents
    rel : PATH-PREP
    lm : REF-EXPR
  form : Phrase
    rel_f `before` lm_f
  meaning : SPG
    rel_m.landmark ⟷ lm_m
    $\textbf{self}_m$ ⟷ rel_m.spg

Figure 3.13. Simple and structured constructional mappings: (a) simple lexical constructions for INTO and STARBUCKS; (b) the PATH-PHRASE construction, a subcase of PATH-EXPR with constituents capturing a structured mapping.

of abstraction and size. Lexical items may have internal (morphological) constituents, and multi-word phrases may be linked *in toto* by a simple mapping to an associated (and perhaps frozen idiomatic) meaning. In fact, both kinds of mappings for the same construction can co-exist: the preposition *into*, defined as a simple map in Figure 3.13, could also be described as a compound with separate semantic contributions from *in* and *to*; the phrase *out of* offers a similarly rich set of possible analyses. Multiple mappings may also be useful for capturing partial idiomaticity, like that in the idiom *kick the bucket*: the phrase as a whole maps non-compositionally to a simple meaning like that of the word *die*, but a separate simple mapping for the verb *kick* could license variation based on tense-aspect-modality constructions. As we will see in Chapter 6, this ability to represent multiple kinds of mappings allows initially simple representations to be reanalyzed and enriched over the course of development.

3.3.2 Reference and predication

So far we have shown scattered examples of constructions needed for simple examples. We now embark on a more organized tour of some basic English constructions. Though neither exhaustive nor definitive, this overview exemplifies the main relationships among form, meaning and function expressible by the ECG construction formalism. We analyze many traditional grammatical categories and relations as motivated by the primary propositional acts of reference and predication

(Croft 1990; Croft 1991; Croft 2001).[5] That is, they are defined in terms of their communicative functions, as represented by the Referent and Predication schemas defined in Section 3.2.3. Figure 3.14 shows the schematic REF-EXPR and PRED-EXPR constructions, which generalize over all referring and predicating expressions, respectively. Some constructions are defined as subcases of these, linked directly to the relevant communicative schemas; others use them as background frames against which specific roles are profiled. We illustrate some of the major construction types below.

construction REF-EXPR	construction PRED-EXPR
form : Schematic-Form	**form** : Schematic-Form
meaning : Referent	**meaning** : Predication

Figure 3.14. Schematic constructions for referring and predicating expressions

Referring expressions

A wide range of expressions draw on the notion of reference. Some of these function as full referring expressions — *i.e.*, they map a simple or complex form directly to the Referent schema. Figure 3.15 shows two referring expressions, both simple subcases of REF-EXPR:

Proper nouns pick out specific named entities (by default) and can typically be used with no prior mention. In the STARBUCKS construction (which elaborates the version in Figure 3.13), the resolved-referent role is directly bound to the specified schema, and only a minimal or inactive level of accessibility is required (*i.e.*, no prior mention is generally needed).

Pronouns pick out contextually available referents (*i.e.*, ones for which the interlocutors have active representations in the current discourse) and assert additional constraints on features like number and gender, as in the HE construction. In this example, we include a case constructional feature (in a **constructional** block) to distinguish the HE construction from the constructions for *him* (object case) and *his* (possessive case). Languages with wider use of case may define this feature for many other constructions.

More complex referring expressions require structured mappings, as well as auxiliary constructions that constrain specific aspects of the referent. Figure 3.16 illustrates constructions that license the phrase *the cafe*:

[5] Croft identifies *modification* as a third category, exemplified by adjectival and adverbial constructions. We analyze these instead as evoking the same structures as involved in reference and predication and asserting further constraints, but a general constructional category for such expressions could also be defined.

```
construction STARBUCKS (2)
  subcase of PROPER-NOUN
  form : Word
    orth ⟵ "Starbucks"
  meaning
    accessibility ⟵ inactive
    resolved-referent ⟷ Starbucks
```

```
construction HE
  subcase of PRONOUN
  constructional
    case ⟵ subject
  form : Word
    orth ⟵ "he"
  meaning
    accessibility ⟵ active
    number ⟵ singular
    gender ⟵ male
```

Figure 3.15. Referring expression examples: The proper noun STARBUCKS and pronoun HE constructions bind the Referent schema's resolved-referent role and set default levels of accessibility. The HE construction also constrains its case constructional feature.

```
construction DETERMINER
  form : Schematic-Form
  meaning
    evokes Referent as ref
```

```
construction COMMON-NOUN
  form : Schematic-Form
  meaning
    evokes Referent as ref
    self_m ⟷ ref.category
```

```
construction THE-CXN
  subcase of DETERMINER
  form : Word
    orth ⟵ "the"
  meaning
    ref.accessibility ⟵ accessible
```

```
construction CAFE-CXN
  subcase of COMMON-NOUN
  form : Word
    orth ⟵ cafe
  meaning : Cafe
```

```
construction DETERMINEDNOUN-EXPR
  subcase of REF-EXPR
  constructional
    det : DETERMINER
    com-noun : COMMON-NOUN
  form
    det_f before com-noun_f
  meaning
    self_m ⟷ com-noun_m.ref
    self_m ⟷ det_m.ref
```

Figure 3.16. Determined noun construction: Determiners and common nouns evoke and constrain a Referent schema; they are combined in the DETERMINEDNOUN-EXPR construction.

Determiners assert various constraints on a referent, for example on its number or contextual accessibility. The construction for *the* constrains the latter of these to be accessible; analogous constructions for *a* or *some* would constrain the referent to be inaccessible and singular or plural, respectively.

Common nouns denote categories; the COMMON-NOUN construction evokes a Referent and binds its meaning pole to the category role. Subcases like *cafe* or *boy* specialize their form and meaning poles with appropriate phonological or orthographic strings and semantic categories.

Determined nouns combine a determiner with a category-denoting expression (such as common nouns) to pick out an individual instance of the specified category. A simple case of this is a phrase like *the cafe*, referring to a specific member of the conceptual category of Cafe that has been previously mentioned or whose identity is otherwise inferrable.

This treatment of reference, though preliminary, has been successfully extended to modifiers, complements, quantifiers and measure phrases. While further research is necessary, we believe the basic strategy of evoking and constraining the Referent schema can scale to accommodate a wide full range of referential phenomena.

Predicating expressions

Constructions supporting the basic function of predication are represented using a strategy similar to that employed for reference: the Predication schema serves as the background frame whose roles specify what kind of action or relation is being asserted, commanded or queried. Various kinds of constructions can contribute to an overall predication by evoking this schema and binding it as appropriate; the unification of these evoked predications forces them to be understood as — *i.e.*, construed as — part of a single coherent event.

Verbs and verb complexes designate relations among entities. We define verbs as evoking a predication and binding to (or profiling) its central schema (or means). Subcases might specialize the action or event category (as do the SelfMotion-Verb and Walk constructions in Figure 3.17) and assert additional constraints on tense (locating the predication with respect to the current speech time, as in the Walked construction), aspect (event structure) or modality (relation to reality).

```
construction Verb
  form : Word
  meaning
    evokes Predication as pred
    self_m ←→ pred.means
```

```
construction SelfMotion-Verb
  subcase of Verb
  form : Word
  meaning : SelfMotion
```

```
construction Walk
  subcase of SelfMotion-Verb
  form
    orth ←− "walk"
  meaning : Walk
```

```
construction Walked
  subcase of Walk
  form
    orth ←− "walked"
  meaning
    pred.time ←− past
```

Figure 3.17. Basic verb constructions: The Verb construction and its subcases SelfMotion-Verb, Walk and Walked.

Argument structure constructions specify a basic pattern of interaction among a set of participants within a *scene* (Goldberg 1995). The ArgStruct construction shown in Figure 3.18 does this by evoking a predication and profiling its scene role; this predication must also unify with that evoked by its required verbal constituent. Subcases of ArgStruct may declare additional constituents, specialize existing constituents, and assert further constraints. The SelfMotion-Cxn, for example, adds a constituent corresponding to the mover entity, specializes the verbal constituent to require a Motion-Verb, and asserts constraints on case, word order and role bindings. Other argument structure constructions corresponding to the basic scenes discussed in Section 2.2.1 (*e.g.*, caused motion, transitive, transfer) may likewise add, specialize and constrain constituents as appropriate.

construction ArgStruct
constituents
v : Verb
form : Phrase
meaning
evokes Predication **as** pred
self$_m$ $\longleftrightarrow$ pred.scene
v$_m$.pred $\longleftrightarrow$ pred

construction SelfMotion-Cxn
subcase of ArgStruct
constituents
a : Ref-Expr
v : Motion-Verb
a.case $\longleftarrow$ subject
form
a$_f$ `before` v$_f$
meaning : SelfMotion
mover $\longleftrightarrow$ a$_m$
action $\longleftrightarrow$ v$_m$

Figure 3.18. Basic argument structure constructions: The ArgStruct construction and its subcase SelfMotion-Cxn.

This treatment allows a more general analysis of *Harry ran home*, along the lines discussed in Section 3.1.3. The DirectedMotion-Cxn construction in Figure 3.19 is defined as a subcase of SelfMotion-Cxn with an additional argument expressing the path.[6] Unlike the lexically specific Motion-Home construction in Figure 3.5, the DirectionMotion-Cxn allows a variety of path expressions to specify the motion's direction. Note, however, that both kinds of constructions can coexist in a grammar, and neither is intrinsically better than the other: the more general construction accounts for more variations, but the more specific construction is more predictive of those variants it licenses. It is precisely this tradeoff that drives the construction learning model to be described.

The examples given here do not do justice to the varieties of predicating constructions posited in the literature, such as those that capture generalizations about subject-predicate agreement, active-passive voice and question formation. These fall beyond the scope of our current focus on child language learning.

[6]The analysis in Section 3.1.1 is also simplified in that the relevant constraints contributed by the verb (action and time) and clausal construction (scene) are asserted directly on the SelfMotion scene, in lieu of an evoked Predication.

```
construction DirectedMotion-Cxn
  subcase of SelfMotion-Cxn
  constituents
    b : Path-Expr
  form
    v_f before b_f
  meaning
    path ⟷ b_m
```

Figure 3.19. The DirectedMotion-Cxn, a subcase of SelfMotion-Cxn

3.3.3 Other constructions

The constructions shown so far cover only basic aspects of reference and predication. The principles of linguistic representation they exemplify can, however, be extended to other constructional domains. Here we briefly review a few of those most relevant to early English constructions.

Spatial relations

Constructions related to the domain of spatial relations are among the most commonly grammaticized and earliest acquired crosslinguistically, both in the single-word stage (*e.g.*, English *up*) and in the earliest word combinations. Semantic representations related to this domain have thus figured prominently in the literature: the case grammars of the kind proposed by Fillmore (1968) typically include *theme*, *goal* and *source* as roles in the thematic hierarchy; Jackendoff's (1990)'s lexical conceptual semantics includes *GO* as a conceptual primitive; and Talmy (2000) offers a comprehensive catalogue of spatial relations and their many dimensions of linguistically marked variation. As noted in Section 3.1.1, these notions have been represented using *image schemas* (Johnson 1987; Lakoff & Johnson 1980), which in turn lend themselves very naturally to ECG formalization. The SPG (Source-Path-Goal) schema of Figure 3.4 and the Trajector-Landmark, Container and Path schemas in Figure 3.8(d) illustrate how some of the most common image schemas can be defined.

Languages vary in the formal devices they exploit to express spatial relations. As we have seen already, argument structure constructions (such as the Directed-Motion construction from Section 3.1.1) may evoke a path schema. In addition, verbs like *enter* and *leave* specifically denote motion along a path and thus directly convey image-schematic information.[7] English has several classes of constructions especially relevant to the spatial domain. We have already introduced ex-

[7] Languages display typological preferences for whether paths are expressed mainly by verbs or satellites. The formalism allows flexible means of capturing either alternative, though as yet has no explicit encoding of this typological preference.

amples for the specific class of path-related expressions; more generally, these expressions include the following:

Prepositions link a word form with a spatial relation schema, such as the Trajector-Landmark schema or one of its subcases. The INTO-CXN from Figure 3.13 has a meaning pole that is typed as an instance of the Into schema, which in turn evokes and asserts bindings on instances of the Container and SPG schemas. It is defined as a subclass of PATH-PREP(osition) in view of its connection to the SPG schema.

Prepositional phrases are structured mappings with two constituents, a preposition and a referring expression, where the latter furnishes a landmark for the former, as exemplified for path expressions by the PATH-PHRASE construction in Figure 3.13.

Particles are represented as a kind of truncated prepositional phrase, of which only the relation is expressed; the landmark is assumed to be inferrable in context. Figure 3.20 defines an IN-PARTICLE as a specific path-related particle that evokes the Into schema of Figure 3.11 along with a Referent bound to its landmark. (Note that the HOME construction of the *Harry ran home* example similarly expresses a path, though with different lexically specific constraints.)

```
construction IN-PARTICLE
  subcase of PATH-PARTICLE
  form
    orth ⟵ "in"
  meaning : SPG
    evokes Referent as r
    evokes Into as i
    i.lm ⟵ r
```

Figure 3.20. The IN-PARTICLE construction

Variations on these constructions can be defined to accommodate crosslinguistic diversity in how spatial (and other) relations are marked. An ADPOSITIONAL phrase construction, for example, may allow either prepositional or postpositional word order. For languages that employ morphological case marking to indicate spatial relationships, morphological constructions similar to the prepositional phrase but with word-internal form constraints (as well as possible phonological constraints) could be defined. The meaning domain could also be broadened to encompass other kinds of relations (such as temporal ordering and possession).

Context-dependent constructions

Many of the earliest childhood constructions refer to the current discourse or situational context. We can capture the relevant meanings by explicitly evoking the DiscourseSpace schema defined in Figure 3.12 and constraining it appropriately. Two simple examples are shown in Figure 3.21:

Speech act constructions provide explicit cues to the speaker's intent. The IMPERATIVE construction is a simple mapping between a specific intonational contour and speech act type.

Deictic pronouns are defined as pronouns that refer directly to participants in the current discourse space. The YOU construction, for example, is linked to the current addressee.

```
construction IMPERATIVE
  subcase of INTONATIONAL-CXN
  form : Intonational-Contour
    type ⟵ falling
  meaning
    evokes DiscourseSpace as ds
    ds.speechAct ⟵ command
```

```
construction YOU
  subcase of PRONOUN
  form : Word
    orth ⟵ "you"
  meaning
    evokes DiscourseSpace as ds
    self_m ⟷ ds.addressee
```

Figure 3.21. Contextually grounded constructions: The IMPERATIVE and YOU constructions both evoke and constrain the DiscourseSpace schema.

Scaling up

The domains illustrated so far are not intended to capture the full range of formal devices and meaningful distinctions available in English, let alone crosslinguistically; they aim merely to provide reasonable coverage of the early utterances most relevant to the learning model. Nonetheless, our interest in the learning model to be proposed rests in part on the assumption that its target grammar formalism can scale gracefully beyond these to more complex linguistic phenomena. To date, ECG has been extended to handle the formal and semantic demands of many (adult) constructions in several typologically distinct languages, including:

- *Morphological constructions:* Morphologically complex words can be represented as structured mappings with constituents, with word-internal analogues of ordering relations used to capture concatenative morphological constraints. ECG analyses have been proposed for English verbal suffixes, Russian and Georgian case marking, and Hebrew verbal morphology.
- *Argument structure constructions:* A broader system of constructions capturing English argu-

ment structure and argument structure alternations has been developed (Dodge In Preparation; Dodge & Lakoff 2005).

- *Referential expressions:* Modifications of the Referent schema have been used to capture a variety of English referential phenomena, including modifiers, measure phrases like *a cup of tea*, and quantifiers.

- *Predicational expressions:* These include directional specifiers (deictic particles in Mandarin and German), English aspectual constructions, and preliminary studies of conditionals and hypotheticals in English and Mandarin.

3.4 Summary

This chapter has provided a brief overview of Embodied Construction Grammar, a representational formalism that supports a construction-based approach to grammar and a simulation-based model of language understanding. Two primary formal structures have been described:

- schemas, capturing generalizations over experience in the domains of form or meaning; and

- constructions, capturing generalizations over form-meaning pairs.

Both structures are represented using a common formalism that facilitate the parameterization of simulations of the kind described in Section 2.3.3. Together they provide a formally precise means of representing constructions, in the symbolic, unified and gestalt sense laid out in Section 2.2.1. The discussion here has centered on the needs of the learning model, drawing examples mainly from early English constructions. The schema and construction formalisms are, however, sufficiently expressive to describe a wide range of adult and crosslinguistic constructions. More importantly, they provide the crucial information needed for driving the dynamic language understanding processes to which we turn next.