

# Scaling Understanding up to Mental Spaces

**Eva Mok, John Bryant, Jerome Feldman**  
International Computer Science Institute  
1947 Center Street Suite 600,  
Berkeley, CA 94704  
{emok, jbryant, jfeldman}@icsi.berkeley.edu

## Abstract

Mental Space Theory (Fauconnier, 1985) encompasses a wide variety of complex linguistics phenomena that are largely ignored in today’s natural language processing systems. These phenomena include conditionals (e.g. *If* sentences), embedded discourse, and other natural language utterances whose interpretation depends on cognitive partitioning of contextual knowledge. A unification-based formalism, Embodied Construction Grammar (ECG) (Chang et al., 2002a) took initial steps to include space as a primitive type, but most of the details are yet to be worked out. The goal of this paper is to present a scalable computational account of mental spaces based on the Neural Theory of Language (NTL) simulation-based understanding framework (Narayanan, 1999; Chang et al., 2002b). We introduce a formalization of mental spaces based on ECG, and describe how this formalization fits into the NTL framework. We will also use English Conditionals as a case study to show how mental spaces can be parameterized from language.

## 1 Introduction

There are two dimensions to scalability: improving system performance (e.g. speed and size) for a fixed task, and expanding the range of tasks that the system can handle. Today’s natural language processing (NLP) systems are not very scalable in the latter dimension. They tend to ignore a wide range of cognitive linguistic phenomena, notably those associated with mental spaces, which are key to understanding any non-trivial piece of natural language. **Mental spaces** (Fauconnier, 1985) are partial cognitive structures built up during discourse that keep track of entities and relations in dif-

ferent contexts. Hypothetical reasoning, depictions (e.g. stories, paintings or movies) and reasoning about other minds are but a few examples where new mental spaces are required. Mental spaces provide an important partitioning of contextual knowledge that allows scalable reasoning in a partitioned large knowledge base.

However, the literature on Mental Space Theory does not address how these cognitive structures are specified compositionally, let alone offering formalizations of mental space representations or computational realizations. We thus seek to scale up the complexity of natural language understanding (NLU) systems by proposing a computational method of handling mental spaces. In this paper we give a brief introduction to Mental Space Theory (Section 2), and then explain how Mental Space Theory can be incorporated into the existing Neural Theory of Language (NTL) simulation-based understanding framework (Section 3). In this framework, each mental space is a separate thread of simulation. Each thread of simulation comprises a dynamic structure combining knowledge, event models and beliefs that evolve over time.

The use of a simulation-based framework imposes constraints on how mental spaces are represented, and we introduce a formalization of mental spaces based on Embodied Construction Grammar (ECG). As a case study (Section 4), we will walk through the formalization of English Conditionals (Dancygier and Sweetser, 2004) using mental space analysis. Through this case study we illustrate how mental spaces are parameterized compositionally by language, and capture this compositionality succinctly using construction grammar. Then with these formal tools, we address the issue of inference in mental spaces (Section 5), which is at the core of the scaling up understanding.

## 2 Mental Space Theory

Mental spaces refer to the partial cognitive structures built up, usually through discourse, that provide a partitioning of contextual as well as world knowledge. This partitioning in turn affects what inferences can be

drawn. In traditional mental space analysis, certain linguistic constructions called **space builders** may open a new mental space or shift focus to an existing space. Examples are *in this picture...*, *Nancy thinks...*, *if it rains...*, *back in the '50s...* (Fauconnier, 1997). Consider the following sentence:

- (1) *In Harry's painting of Paris, the Eiffel Tower is only half-finished.*

Harry's painting creates a new Depiction-Space. The Eiffel Tower is the only entity local to this Depiction-Space, and it maps to the physical Eiffel Tower. However, only the Eiffel Tower in the Depiction-Space has an additional attribute half-finished, and one should not be led to think that the real Eiffel Tower is also half-done. This kind of space-building is traditionally illustrated in diagrams similar to Figure 1.




Figure 1. The painting opens a Depiction-Space where the Eiffel Tower is half-finished.

In the mental space literature, the transfer of assumptions between spaces is guided by the **presupposition float principle**. The presupposition float principle states that any presupposed structure in the parent space can float to a child space unless it conflicts with structure already in that space. In the above example, any attributes about the real Eiffel Tower can be assumed in the Depiction-Space, as long as they do not depend on it being finished. However, this account of assumption and inference transfer is incomplete. Specifically, it is incorrect to assume that different types of mental spaces obey the same presupposition float principle. For example, if we are having a conversation right now about Harry's painting, very little of what is currently happening should transfer into the Depiction-Space. On the other hand, if we are having a conversation about our plans for tomorrow, our current situation is very relevant to our actions tomorrow, and this information should carry over to the future space. The other key piece that is missing from the presupposition float account is how inference is drawn across spaces in general. Any computational account of mental spaces must address this inference process precisely and supply a formalized representation that inference can operate

with. We will outline such a computational solution in the next two sections of this paper.

### 3 Simulation-Based Understanding

A central piece of a scalable computational treatment of mental spaces is a robust language understanding framework. Our work relies on the NTL simulation-based understanding paradigm (Narayanan, 1999; Chang et al., 2002b), and extends the model in a conceptually straightforward way. The simulation-based understanding paradigm stipulates that, in addition to constructional analysis of the surface form, language understanding requires active simulation.

The constructional analysis is based on Embodied Construction Grammar (Chang et al., 2002a) which contains four primitive types: schemas, constructions, maps and mental spaces. Schemas are the basic ECG unit of meaning, capturing embodied concepts such as image schemas, actions, and events. Constructions are the basic linguistic unit, pairing meaning schemas with representations of linguistic form (words, clauses, etc.). Maps and mental spaces are the subject of this paper, and will be discussed in detail in the next section. It is worth noting that in the ECG formalism, in addition to support of an inheritance hierarchy (with the keyword subcase of), there is also an evokes relation that makes an outside structure accessible to a schema through a local name. The evokes relation is neither a subcase-of or part-of relation, but is analogous to spreading activation in the neural sense.

During analysis, a **Semantic Specification (Sem-Spec)** is created from the meaning poles of the constructions, and is essentially a network of schemas with the appropriate roles bounded and filled in. Crucially, within this network of schemas are **executing schemas** (or **X-schemas**), which are models of events. They are active structures for event-based asynchronous control that can capture both sequential flow and concurrency.

**Simulation** is a dynamic process which includes executing the X-schemas specified in the SemSpec and propagating belief updates in a belief network. This mechanism is used for metaphor understanding in (Narayanan, 1999), and is being generalized to Coordinated Probabilistic Relational Models (CPRM) in current efforts (Narayanan, submitted). The CPRM mechanism is discussed in more detail in Section 5.

Within a simulation-based understanding paradigm, each mental space involves a new thread of simulation, with its own separate belief network and simulation trace. This is necessary for keeping track of possibly contradictory beliefs, such as the alternative scenarios where it is sunny or rainy tomorrow. Each alternative scenario exists within its own mental space, and in many situations, there can be a large number of alternatives. However, not only is it computationally expensive

to create a new thread of simulation, but cognitive capacity also constrains the number of concurrently open spaces. We need both a cognitively plausible and computationally feasible theory of how mental spaces are manipulated.

The insight in addressing this problem is that at any given level of granularity, not all spaces need to be opened at the same time. Harry’s painting, in example (1), may be represented at different granularity depending on the context. If it is discussed simply as a wall-hanging, the Depiction-Space need not be expanded and the painting should be treated schematically as an object. However, once the contents of the painting are under discussion (e.g., the trip to Paris during which the painting was done), inference in a separate mental space is required, and the Depiction-Space needs to be built.

As illustrated by this example, the simulation process dictates the actual building of mental spaces. The analysis process is responsible for supplying all the necessary parameterization of the spaces and their corresponding maps in case they need to be built. As a result, each potential space-builder is represented at two levels of granularity – as an object in its schematic form and as a full mental space.

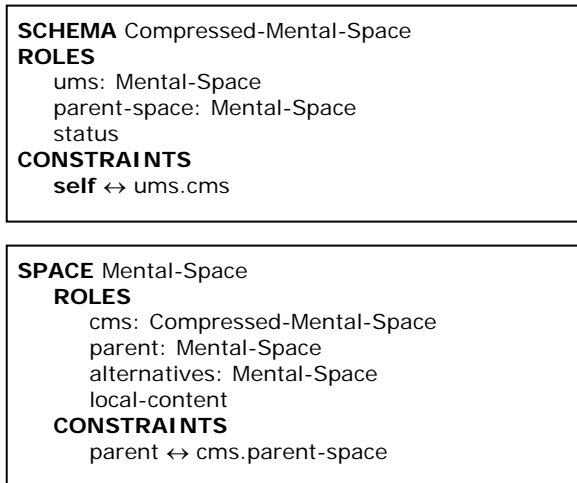


Figure 2. ECG notation for Compressed and Uncompressed Mental Spaces

Formalizing this idea in ECG, mental spaces are represented in two ways: as a compressed mental space and an uncompressed version. In the ECG notation in Figure 2, the Compressed-Mental-Space is just a schema, and Mental-Space is of the space primitive type. In each version there is pointer to its counterpart, ums and cms respectively. The role parent-space points to the parent of this space. The uncompressed mental space contains the list of alternatives and the local-content. Alternatives are scenarios that are different and cannot co-exist, such as the different activities one might be doing tomorrow at noon. Local-content

at noon. Local-content provides the local semantics of mental spaces, maintaining a list of predications that are true in this space, *ceteris paribus*. Each predication contains a role local-to that denotes the space to which it belongs. The predication is then automatically added to the local-content of the space when this role is assigned. Figure 3 shows an example of the Cause-Effect schema with the cause, effect, and local-to role.




Figure 3. Cause-Effect is a predication that contains a pointer to a Mental-Space

In the next section, we will demonstrate the use of the above formalization with a case study on conditional sentences in English. In Section 5, we will discuss how this representation supports the needed inference.

#### 4 Case Study: English Conditionals

One of the most common classes of space-building expressions is the **predictive conditional**. Predictive conditionals are sentences like

(2) *If it rains tomorrow, the game will be cancelled.*

They are space-builders, setting up a primary conditional space and an alternative space<sup>1</sup> (Dancygier and Sweetser, 2004). As shown in Figure 4, in the primary conditional space for tomorrow, it rains, and therefore the game is cancelled. In the alternative space, it does not rain, and the game is not cancelled.

This case study will focus on predictive conditionals in English. An English predictive conditional is characterized by tense backshifting in the *if* clause, i.e., the use of the present tense for a future event. On the meaning side, the condition and the conclusion are related by some causal or enablement relationship.

In this section we will gradually build up to the Conditional-Prediction construction by introducing the relevant schemas and smaller constructions. It is important to stress how construction grammar succinctly captures how mental spaces can be parameterized in a compositional way. The larger constructions supply information about the spaces that is not contained in any of its smaller constituents. Through compositionality, the grammar becomes much more scalable in handling a wide range of linguistic variations.

<sup>1</sup> Readers interested in the mental space analysis of other types of English conditionals should refer to (Dancygier and Sweetser, 2004), as well a technical report for the formalization (Bryant and Mok, 2003).

If it rains tomorrow, the game will be cancelled.




Figure 4. A predictive conditional sets up two spaces

#### 4.1 Conditions

The condition, often preceding the conclusion in a conditional statement, sets up or locates the space in which the conclusion is to be placed. Therefore the Conditional-Schema, given below, is a subcase of the Compressed-Mental-Space. In addition to the inherited roles, the Conditional-Schema has roles for a condition, a premise, and a conclusion. The condition role stands for the condition P as expressed, and premise can be P or  $\sim P$ , depending on the conjunction. Finally, epistemic-stance is the degree of commitment that the speaker is making towards the condition happening, as indicated by the choice of verb tense. For example, a sentence such as *If you got me a cup of coffee, I'd be grateful*

**SCHEMA** Conditional-Schema  
**SUBCASE OF** Compressed-Mental-Space  
**ROLES**  
 epistemic-stance  
 condition: Predication  
 premise: Predication  
 conclusion: Predication  
 ums: Conditional-Space  
**CONSTRAINTS**  
 epistemic-stance  $\leftrightarrow$  ums.epistemic-stance  
 premise  $\leftrightarrow$  ums.premise  
 conclusion  $\leftrightarrow$  ums.conclusion

**SPACE** Conditional-Space  
**SUBCASE OF** Mental-Space  
**ROLES**  
 cms: Conditional-Schema  
 epistemic-stance  
 premise: Predication  
 conclusion: Predication  
**CONSTRAINTS**  
 premise.local-to  $\leftarrow$  self  
 conclusion.local-to  $\leftarrow$  self

Figure 5. Conditional-Space has roles for premise, conclusion and epistemic-stance

*forever* (Dancygier and Sweetser, 2004) has a negative epistemic stance. The speaker makes a low commitment (by using the past tense *got*) so as to not sound presumptuous, even though she may think that the addressee will very likely bring her coffee. On the other hand, a counterfactual such as *If I'd've known you were coming, I'd've stayed home* has a very negative epistemic stance. The speaker, in this case, implies that he did not know the addressee was coming.

#### 4.2 The If construction

The abstract construction Conditional-Conjunction is a supertype of lexical constructions *If* and other conditionals conjunctions. Each conjunction leads to a different way of parameterizing the spaces, therefore a Conditional-Conjunction does not have meaning on its own. Instead, it EVOKES two copies of the Conditional-Schema, one as primary and one as alternative.

Given the Conditional-Conjunction and the Conditional-Schemas it evokes, the *If* construction only needs to hook up the correct premise and set the epistemic-stance to neutral. The condition itself is not filled in until a larger construction uses the Conditional-Conjunction as a constituent.

**CONSTRUCTION** Conditional-Conjunction  
**MEANING**  
**EVOKES** Conditional-Schema AS cs  
**EVOKES** Conditional-Schema AS alt

lexical **CONSTRUCTION** *If*  
**SUBCASE OF** Conditional-Conjunction  
**FORM:** Word  
 self.f.orth  $\leftarrow$  "if"  
**MEANING**  
 cs.premise  $\leftrightarrow$  cs.condition  
 cs.epistemic-stance  $\leftarrow$  neutral

Figure 6. Conditional-Conjunctions evokes two Conditional-Schemas. The *If* construction identifies the premise with the condition in the primary space

#### 4.3 The Condition Construction

The Condition construction forms a Subordinate-Clause from a Conditional-Conjunction and a Clause, such as *If it rains tomorrow* from our game cancellation example in (2). The most important aspect of this construction is that it identifies its meaning pole (a Conditional-Schema) with the Conditional-Schema that is evoked by the Conditional-Conjunction, thereby preserving all the constraints on the premise, epistemic-stance and status that the conjunction sets up.

In addition, it also fills in the content of the condition role with the meaning of the Clause. It sets the parent-space to the current focus-space.

**CONSTRUCTION** Condition  
**SUBCASE OF** Subordinate-Clause  
**CONSTRUCTIONAL**  
 conj: Conditional-Conjunction  
 cl: Clause  
**FORM**  
 conj meets cl  
**MEANING:** Conditional-Schema  
**self.m** ↔ conj.m.cs  
**self.m.condition** ↔ cl.m  
**self.m.parent-space** ↔ focus-space

Figure 7. The Condition construction fills in the content of the condition role

#### 4.4 Predictions

Predictions are not space builders in and of themselves. Instead, they are simply events that are marked as future. The Prediction-Schema is therefore not a subcase of Compressed-Mental-Space. It includes a predicted-event, which is a predication of category Event (denoted through the evokes statement and the binding), and a basis-of-prediction. A predicted event also has an associated probability of happening, which may be supplied linguistically (through hedges like *perhaps* or *probably*). This probability directly affects what inferences we draw based on the prediction, and is captured by the likelihood-of-predicted-event role.

**SCHEMA** Prediction-Schema  
**EVOKES** Event **AS** e  
**ROLES**  
 predicted-event: Predication  
 likelihood-of-predicted-event  
 basis-of-prediction: Predication  
**CONSTRAINTS**  
 predicted-event.category ↔ e  
 predicted-event.time-location ← future

Figure 8. Predictions are not space-builders by themselves.

#### 4.5 The Prediction Construction

The Prediction construction is a Clause that evokes a Prediction-Schema in its meaning, such as *the game will be cancelled*. The meaning pole of a Clause is a Predication. The nature of prediction requires that the time reference be future with respect to the viewpoint-space, in this case, today. The meaning of the Prediction construction is itself the predicted-event.

**CONSTRUCTION** Prediction  
**SUBCASE OF** Clause  
**CONSTRUCTIONAL**  
 time-reference ← relative-future (viewpoint-space)  
**MEANING:**  
**EVOKES** Prediction-Schema **AS** ps  
 ps.predicted-event ↔ **self.m**

Figure 9. The Prediction construction makes itself the predicted-event in the Prediction-Schema

#### 4.6 Conditional Statements

Conditional-Statement is a very general construction that puts a Condition and a Clause together, in unrestricted order. The most important thing to notice is that this larger construction finally fills in the conclusion of the Condition with the meaning pole of the statement.

**CONSTRUCTION** Conditional-Statement  
**CONSTRUCTIONAL**  
 cond: Condition  
 statement: Clause  
**MEANING**  
 cond.conclusion ↔ statement.m

Figure 10. The Conditional-Statement construction puts together a Condition and a Clause

#### 4.7 Predictive Conditionals

To a first approximation, a Conditional-Prediction is just a special case of the Conditional-Statement where the statement has to be a Prediction. However, extra care has to be taken to ensure that the alternative spaces and cause-effect relations between the premises and conclusions are set up correctly.

Recall that the Conditional-Conjunction EVOKES two Conditional-Schemas, which are either partially filled in or not filled in at all. Intuitively, the goal of this construction is to completely fill out these two schemas (and their respective spaces), and put a Cause-Effect relation between the premise and conclusion in the local-content of each space.

A role alt is created with type Conditional-Schema to capture the fact that there is an alternative space parameterized by this construction. alt is then identified with the alternative Conditional-Schema evoked in the Conditional-Conjunction. This allows the unused alternative Conditional-Schema in the If construction to be filled in.

The complete filling out of both Conditional-Schemas are done by identifying the premise in the alternative schema with the negation of the premise in the

**CONSTRUCTION** Conditional-Prediction  
**SUBCASE OF** Conditional-Statement  
**CONSTRUCTIONAL**  
statement: Prediction  
statement.time-reference  $\leftarrow$  relative-future  
(condition.time-reference)

**MEANING:**  
**EVOKES** Cause-Effect **AS** ce-primary  
**EVOKES** Cause-Effect **AS** ce-alternative

alt: Conditional-Schema  
alt  $\leftrightarrow$  cond.conj.alt  
alt.premise  $\leftrightarrow$  not(cond.premise)  
alt.conclusion  $\leftrightarrow$  not(cond.conclusion)  
alt.epistemic-stance  $\leftarrow$  opposite  
(cond.epistemic-stance)

ce-primary.cause  $\leftrightarrow$  cond.premise  
ce-primary.effect  $\leftrightarrow$  cond.conclusion  
ce-primary.local-to  $\leftarrow$  cond.msp  
alt.msp.local-to  $\leftarrow$  cond.msp.alternatives

ce-alternative.cause  $\leftrightarrow$  not(cond.premise)  
ce-alternative.effect  $\leftrightarrow$  not(cond.conclusion)  
ce-alternative  $\leftarrow$  alt.msp  
alt.msp.alternatives  $\leftarrow$  cond.msp

ce-primary  $\leftrightarrow$  assertion.ps.predicted-  
event.ce1

Figure 11. The Conditional-Prediction construction fills out the parameterization of both the primary and alternative spaces

primary schema, and like-wise for the conclusion. The epistemic-stance of the alternative schema is set to the opposite of that in the primary. For example, the opposite of a neutral stance is neutral, and the opposite of a negative stance is a positive stance.

So far, the only things in the local-content of both spaces are the premise and conclusion, which are just predications without any relations between them. The next two sections assert a Cause-Effect relation between the respective premise and conclusion, and place that in the local-content. It also adds the other space to its list of alternatives.

Finally, the last statement identifies the primary cause-effect with ce1 of the predicted-event (in the Event schema), thereby filling in the cause of the predicted-event.

#### 4.8 Example

With the Conditional-Prediction construction and the smaller constructions in hand, along with the related schemas and spaces, we now return to example (2) in the beginning of this section: *If it rains tomorrow, the game will be cancelled.*

Figure 12 shows the schemas and mental spaces that result from the analysis of this sentence. The first half of

the sentence, *If it rains tomorrow*, is an instance of the Condition construction. The Conditional-Conjunction *If* evokes a primary and an alternative Conditional-Schema, and partially fills out the primary one. Specifically, the *If* construction sets the epistemic-stance to neutral, and identifies the premise with the condition. The Condition construction then fills in the condition role with the meaning of the Clause *it rains tomorrow*. For simplicity, the actual schemas for representing a rain event are omitted from the diagram.

The second half of the sentence, *the game will be cancelled*, is an instance of the Prediction construction. The basic job that the Prediction construction performs is to fill in the predicted-event with the actual prediction, i.e., game cancellation tomorrow.

At this point, given a Condition with *if* and a Prediction, an instance of the Conditional-Prediction construction is formed, and the diagram in Figure 12 is completed. The predicted-event is filled into the conclusion in the primary Conditional-Schema. The alternative Conditional-Schema, previously untouched by *If*, now gets the negated premise and conclusion. A primary Cause-Effect (ce-Primary) is evoked and placed into the local-content of the primary Conditional-Space, and likewise for the alternative space. The two spaces are then linked as alternatives.

## 5 Inference and Scalability

As we discussed in Section 3, the ECG simulation semantics approach to NLU involves both dynamic simulations and extensive belief propagation. This approach leads to systems that are scalable in semantic depth. For such systems to be practical, we also need them to be scalable in size. In recent work, (Narayanan, submitted) has shown how the dynamic simulations and belief propagation techniques can be tightly coupled in a highly scalable formalism called CPRM, Coordinated Probabilistic Relation Models.

This same formalism also provides the tools for a systematic treatment of inference across mental spaces, which correspond to separate threads of simulation. Returning to example (2) in the last section, *If it rains tomorrow, the game will be cancelled*, we can now make additional inference about the hypothetical scenario and the actions of the participants involved. One can ask, what if the game is cancelled and the participants need to plan for other activities? How will these activities affect their actions today?

The CPRM mechanism elegantly handles the transfer of assumptions and inferences needed to answer such questions. While the varying types of mental spaces have different rules of inference, all of the couplings are of only two different types, both of which are handled nicely by CPRM. Any two mental spaces will be related either by some shared assumptions, some

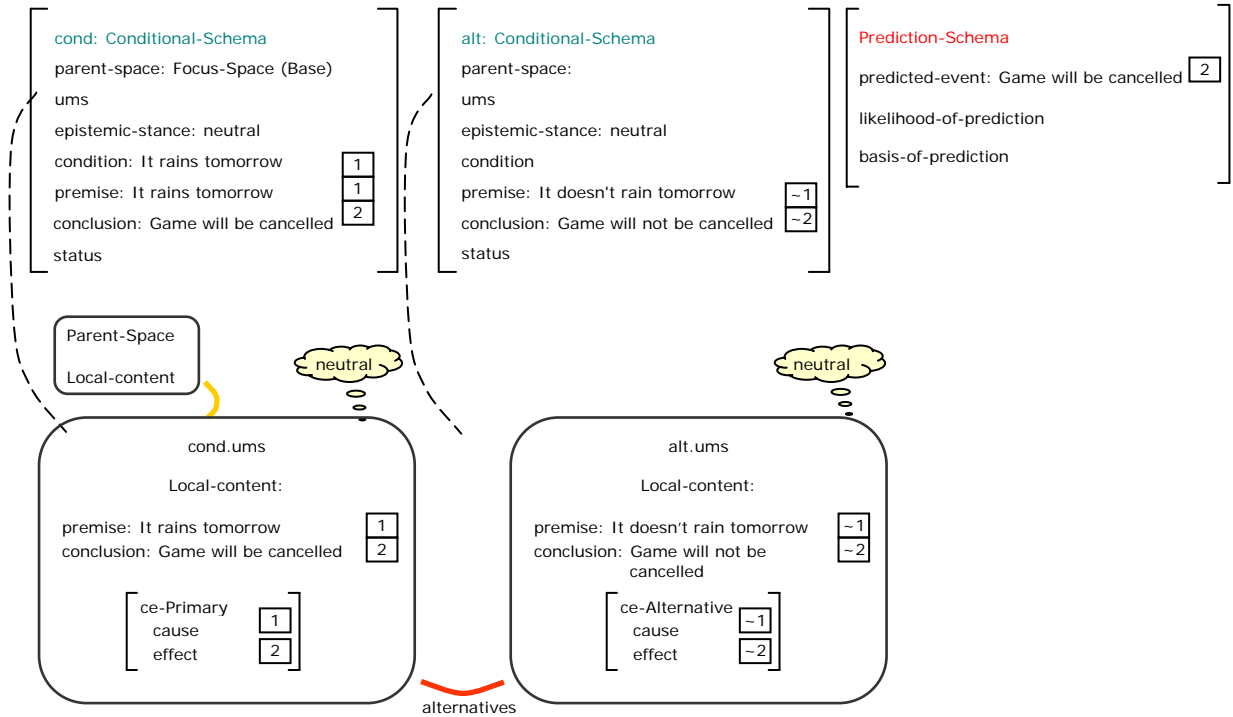


Figure 12. *If it rains tomorrow, the game will be cancelled.*

influence links, or both. Since the CPRM formalism is inherently nested, it is straightforward to have **shared spaces**. For example, if several people are watching a game and are talking about it, the progress of the game is a (dynamic) shared space that is common ground for all the speakers.

**Influence links** are the central primitive in all belief networks, including CPRM. These encode the effect of one attribute on the conditional probabilities of another attribute. In the mental space context, we employ explicit influence links to encode the dependencies of one space on attributes of another space. In the game cancellation example, the participants can choose to further plan for the scenario where it does rain. They might pick a backup plan, for example going to the theater. If this backup plan requires some resource (e.g. a discount card), there should be an influence link back to the present plan suggesting that the discount card be brought along.

More generally, each particular kind of mental space relation will have its own types of shared knowledge and influence links. Some of these can be evoked by particular constructions. For example: *Harry never agrees with Bob* would set up dependency links between our models of the minds of these two individuals. It should be feasible to use these mechanisms to formalize the informal insights in the Cognitive Linguistics litera-

ture and therefore significantly extend the range of NLU.

## 6 Conclusion

In this paper we have provided a computational realization of mental spaces. Within a simulation-based understanding framework, a mental space corresponds to a new thread of simulation, implementable using the Coordinated Probabilistic Relational Model formalism. Cognitive and computational constraints demand that each mental space be represented at two levels of granularity: as a schema (compressed-mental-space) and as a full space. The analyzer, using constructions like the ones shown in the case study, creates a Semantic Specification parameterizing the compressed and uncompressed versions of each mental space. Simulation determines the correct level of granularity to operate on, and builds new mental spaces only when it is necessary to perform inference in the new space. Once a new mental space is built, shared spaces and influence links between any two mental spaces can be defined to allow the transfer of inference between the spaces.

Our proposed formalization of mental spaces allows systems to be scalable in both size and semantic depth: (i) Our formalization makes explicit how mental spaces partition contextual knowledge into manageable chunks,

thereby providing significant computational advantage. (ii) By using Embodied Construction Grammar, our formalization provides a compositional approach to parameterizing mental spaces. Compositionality has the advantage of allowing a small grammar to handle a large degree of linguistic variation. (iii) During simulation, new threads of simulation are built only as needed, obeying cognitive capacity constraints as well as making mental spaces computationally tractable. (iv) CPRM provides a tightly coupled, scalable inference mechanism that handles the couplings between mental spaces. Our proposed mental space formalism thus provides a precise and scalable means for handling a rich body of complex linguistics phenomena beyond the reach of current NLU systems.

## References

- John Bryant and Eva Mok. 2003. Constructing English Conditionals: Building Mental Spaces in ECG. Technical Report.
- Nancy Chang, Jerome Feldman, Robert Porzel and Keith Sanders. 2002. *Scaling Cognitive Linguistics: Formalisms for Language Understanding*. First International Workshop on Scalable Natural Language Understanding (SCANALU 2002).
- Nancy Chang, Srinu Narayanan and Miriam R.L. Petruck. 2002. *From Frames to Inference*. First International Workshop on Scalable Natural Language Understanding (SCANALU 2002).
- Barbara Dancygier and Eve Sweetser. 2004. *Mental Spaces In Grammar: Conditional Constructions*. Cambridge University Press. In Press.
- Gilles Fauconnier. 1985. *Mental spaces: Aspects of meaning construction in natural language*. Cambridge: MIT Press.
- Gilles Fauconnier. 1997. *Mappings in Thought and Language*. New York: Cambridge University Press.
- Srinu Narayanan. 1999. *Moving Right Along: A Computational Model of Metaphoric Reasoning about Events*. Proceedings of the National Conference on Artificial Intelligence (AAAI '99), Orlando, Florida, July 18-22, 1999, pp 121-128, AAAI Press, 1999.